

Roth, Duncan; Moffat, John D.; Garloff, Alfred

Book

Cohort size and labour-market outcomes

IAB-Bibliothek, No. 367

Provided in Cooperation with:

Institute for Employment Research (IAB)

Suggested Citation: Roth, Duncan; Moffat, John D.; Garloff, Alfred (2018) : Cohort size and labour-market outcomes, IAB-Bibliothek, No. 367, ISBN 978-3-7639-4121-6, W. Bertelsmann Verlag (wbv), Bielefeld, <https://doi.org/10.3278/300969w>

This Version is available at:

<https://hdl.handle.net/10419/280220>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-sa/4.0/>

Institut für Arbeitsmarkt-
und Berufsforschung

Die Forschungseinrichtung der
Bundesagentur für Arbeit

IAB

IAB-Bibliothek

Die Buchreihe des Instituts für Arbeitsmarkt- und Berufsforschung

367

Cohort size and labour-market outcomes

Duncan Roth

Dissertationen

wbv

Institut für Arbeitsmarkt-
und Berufsforschung

Die Forschungseinrichtung der
Bundesagentur für Arbeit

IAB

IAB-Bibliothek

367

Die Buchreihe des Instituts für Arbeitsmarkt- und Berufsforschung

Cohort size and labour-market outcomes

Duncan Roth

Dissertationen



Bibliografische Information der Deutschen Nationalbibliothek

Die Deutsche Nationalbibliothek verzeichnet diese Publikation in der Deutschen Nationalbibliografie; detaillierte bibliografische Daten sind im Internet über <http://dnb.ddb.de> abrufbar.

Inaugural-Dissertation zur Erlangung der wirtschaftswissenschaftlichen Doktorwürde des Fachbereichs Wirtschaftswissenschaften der Philipps-Universität Marburg
eingereicht von Duncan Roth M. Sc. aus Darmstadt

Erstgutachter: Prof. Dr. Bernd Hayo
Zweitgutachter: Prof. Dr. Michael Kirk
Einreichungstermin: 10. Oktober 2016
Prüfungstermin: 20. Dezember 2016
Erscheinungsort: Marburg
Hochschulkennziffer: 1180

Dieses E-Book ist auf dem Grünen Weg Open Access erschienen. Es ist lizenziert unter der CC-BY-SA-Lizenz.



Herausgeber der Reihe IAB-Bibliothek: Institut für Arbeitsmarkt- und Berufsforschung der Bundesagentur für Arbeit (IAB), Regensburger Straße 100, 90478 Nürnberg, Telefon (09 11) 179-0
■ **Redaktion:** Martina Dorsch, Institut für Arbeitsmarkt- und Berufsforschung der Bundesagentur für Arbeit, Telefon (09 11) 179-32 06, E-Mail: martina.dorsch@iab.de ■ **Gesamtherstellung:** W. Bertelsmann Verlag, Bielefeld (wbv.de) ■ **Rechte:** Kein Teil dieses Werkes darf ohne vorherige Genehmigung des IAB in irgendeiner Form (unter Verwendung elektronischer Systeme oder als Ausdruck, Fotokopie oder Nutzung eines anderen Vervielfältigungsverfahrens) über den persönlichen Gebrauch hinaus verarbeitet oder verbreitet werden.

© 2018 Institut für Arbeitsmarkt- und Berufsforschung, Nürnberg/
W. Bertelsmann Verlag GmbH & Co. KG, Bielefeld

In der „IAB-Bibliothek“ werden umfangreiche Einzelarbeiten aus dem IAB oder im Auftrag des IAB oder der BA durchgeführte Untersuchungen veröffentlicht. Beiträge, die mit dem Namen des Verfassers gekennzeichnet sind, geben nicht unbedingt die Meinung des IAB bzw. der Bundesagentur für Arbeit wieder.

ISBN 978-3-7639-4120-9 (Print)
ISBN 978-3-7639-4121-6 (E-Book)
ISSN: 1865-4096
Best.-Nr. 300969

www.iabshop.de

www.iab.de

Inhalt

Acknowledgements	5
Einleitung – German Summary.....	7
Problem statement, structure and contribution of the dissertation....	17
John Moffat and Duncan Roth	
The cohort size–wage relationship in Europe	25
(published LABOUR: Review of Labour Economics and Industrial Relations, 2016, 30 (4): 415–32)	
Abstract	25
1 Introduction	25
2 Estimation	28
2.1 Data	28
2.2 Empirical model	30
2.3 Identification	33
3 Results	35
4 Conclusion	38
Acknowledgements	39
References	39
Appendix	42
Supplementary material	46
References	59
Alfred Garloff and Duncan Roth	
Regional population structure and young workers' wages.....	61
(forthcoming in U. Blien, K. Kourtit, P. Nijkamp and R. Stough (eds): Modelling Aging and Migration Effects on Spatial Labor Markets, Springer)	
Abstract	61
1 Introduction	61
2 Population structure and wages	64
3 Youth–population structure in Western Germany.....	66
4 Empirical analysis	69
4.1 Data	69
4.2 Sample and descriptive statistics	70
4.3 Empirical model and identification	73
5 Results	75
6 Conclusion	83
Acknowledgements	84
References	84
Appendix	88
Supplementary material	89
References	105

John Moffat and Duncan Roth	
Cohort size and youth labour-market outcomes: the role of measurement error.....	107
(forthcoming in Economics Bulletin)	
Abstract	107
1 Introduction.....	107
2 Literature review	109
3 Empirical analysis	111
3.1 Data	111
3.2 Variables and sample	113
3.3 Model	117
4 Results	119
5 Conclusion	123
Acknowledgements	124
References	124
Appendix	126
Supplementary material	134
References	164
Duncan Roth	
Cohort size and transitions into the labour market.....	165
Abstract	165
1 Introduction.....	165
2 Literature and hypotheses.....	167
3 Empirical analysis	170
3.1 Data	170
3.2 Sample and variables.....	171
3.3 Model	176
4 Results	178
4.1 Baseline results	178
4.2 Discussion of the hypotheses	182
4.3 Alternative explanations	183
4.4 Inclusion of individuals with zero search duration	185
5 Conclusion	186
Acknowledgements	188
References	188
Appendix	191
Supplementary material	192
References	203
Abstract	205
Kurzfassung	207

Acknowledgements

This book contains the doctoral thesis that I wrote at Philipps-Universität Marburg and the Institute for Employment Research. For his continued support, comments and suggestions I am particularly grateful to my supervisor, Prof Bernd Hayo. I would also like to thank Prof Michael Kirk for agreeing to supervise my work as well as for his feedback. Further, I would like to thank Prof Tim Friehe for acting as chair of the defence committee. I am also grateful to Stefan Fuchs for supporting me and giving me the opportunity to complete this thesis.

Furthermore, I would like to acknowledge the exchange that I had with a number of different people whose input contributed to my work: my co-authors John and Alfred, my colleagues at Marburg and at IAB as well as the participants of various workshops and conferences.

These acknowledgements would be incomplete without reference to my family and to Julia to whom this book is dedicated.

Duncan Roth

Marburg, February 2018

Einleitung – German Summary

Diese Arbeit setzt sich aus vier separaten Essays zusammen, die den Zusammenhang zwischen regionalen Bevölkerungsstrukturen und verschiedenen Arbeitsmarktergebnissen zum Thema haben.

The cohort size-wage relationship in Europe

Das erste Papier mit dem Titel *The cohort size-wage relationship in Europe* untersucht den Zusammenhang zwischen der Größe einer Gruppe, deren Mitglieder eine ähnliche Berufserfahrung (oder ein ähnliches Alter) und ein vergleichbares Ausbildungsniveau aufweisen, auf die Löhne, die von den Mitgliedern einer solchen „Kohorte“ realisiert werden. Basierend auf der Annahme, dass Personen innerhalb einer Kohorte substituierbar sind, dies über verschiedene Kohorten hinweg aber nur unvollständig möglich ist, lässt die ökonomische Theorie vermuten, dass Änderungen in der Größe einer Kohorte zunächst deren Grenzproduktivität beeinträchtigt. Auf Wettbewerbsmärkten sollte dies eine Anpassung in den kohortenspezifischen Löhnen verursachen. Im Fall abnehmender Grenzproduktivität lässt sich dieser Zusammenhang genauer spezifizieren: *Ceteris paribus*, sollte ein Anstieg in der Größe einer Kohorte dazu führen, dass die Grenzproduktivität innerhalb der Kohorte und dadurch auch die erzielten Löhne sinken. Theoretische Modelle legen darüber hinaus nahe, dass ein vergleichbarer Mechanismus auch im Fall unvollkommenen Wettbewerbs greift, wenn Löhne durch Verhandlungen zwischen Arbeitgeber- und Arbeitnehmervertretern festgesetzt werden.

In der bestehenden empirischen Literatur wird mehrheitlich ein negativer Lohn-effekt nachgewiesen. Darüber hinaus gibt es Hinweise, dass die Größe dieses Effekts mit dem Ausbildungsniveau der Kohorte ansteigt. Eine Schwierigkeit, den Lohn-effekt empirisch zu bestimmen, besteht darin, dass nicht davon ausgegangen werden kann, dass die Zugehörigkeit einer Person zu einer bestimmten Kohorte zufällig ist. Vielmehr ist in Betracht zu ziehen, dass Personen durch eigene Entscheidungen ihre Kohortenzugehörigkeit beeinflussen können. Im Fall einer durch Berufserfahrung (oder Alter) und Ausbildungsniveau bestimmten Kohorte kann dies einerseits dadurch geschehen, dass Personen in Regionen migrieren, die für die Höhe der von ihnen erzielten Löhne förderlich sind. Andererseits bestimmen Ausbildungsentscheidungen darüber, welcher Kohorte eine Person angehören wird. Beide Mechanismen verwandeln die Kohortengröße selbst in eine endogene Variable, sodass die Anwendung des Kleinsten-Quadrate-Schätzers möglicherweise verzerrte Ergebnisse liefert. Der Beitrag dieses Papiers besteht darin, eine Identifikationsstrategie zu verwenden, die in der Lage ist, beide Ursachen der Endogenität zu berücksichtigen, während

bestehende Untersuchungen zu diesem Thema sich darauf beschränken, die Auswirkungen der Ausbildungsentscheidung zu adressieren. Der Fokus auf Migration als einer Ursache der Endogenität wird auch dadurch gerechtfertigt, dass in diesem Papier regionale Einheiten als räumliche Grundlage für die Kohortenvariable herangezogen werden: Bevölkerungsstrukturen innerhalb solcher Einheiten sollten stärker von Binnenwanderungen betroffen sein als auf Ländern basierende Kohorten durch zwischenstaatliche Migration. Darüber hinaus erlauben kleinräumigere Einheiten eine bessere Annäherung an die Größe einer Kohorte innerhalb tatsächlicher Arbeitsmärkte, insofern diese auf sub-nationaler Ebene existieren, wie die Ergebnisse anderer Studien, die Pendlerströme zwischen Regionen berücksichtigen, nahelegen.

Die Grundlage für die empirische Untersuchung bilden verschiedene Wellen des Datensatzes *European Union Statistics on Income and Living Conditions* (EU-SILC), die zunächst miteinander kombiniert werden, sodass personenbezogene Beobachtungen aus 56 Regionen für den Zeitraum 2004–2010 vorliegen. Aus Gründen der Datenverfügbarkeit beschränkt sich die Analyse auf relativ junge Altersgruppen mit einer Berufserfahrung von bis zu 11 Jahren. Die abhängige Variable ist der durchschnittliche Stundenlohn, der aus Angaben zu dem Jahresarbeitseinkommen, der Dauer der Beschäftigung sowie der durchschnittlichen Zahl der geleisteten Stunden berechnet wird. Die Kohortenvariable misst die relative Größe einer Gruppe von Personen mit gleichem Ausbildungsniveau und einer vergleichbaren Berufserfahrung. Um mögliche Unterschiede in der Größe der Kohorteneffekte feststellen zu können, werden separate Schätzungen für jeden Bildungsgrad vorgenommen, wodurch implizit die Annahme getroffen wird, dass es getrennte Arbeitsmärkte für die verschiedenen Bildungsniveaus gibt. Um die Lohneffekte angesichts der oben geschilderten Endogenitätsprobleme konsistent schätzen zu können, wird ein zweistufiges Verfahren angewendet, das auf der Verwendung einer alters- und zeitversetzten Kohortenvariable als Instrument basiert. Zum Vergleich wird der Zusammenhang zwischen Kohortengröße und Löhnen auch mittels eines zweiten Instruments geschätzt, das bereits in der Literatur zu Lohneffekten Verwendung gefunden hat und Endogenität aufgrund von Ausbildungs-, aber nicht von Migrationsentscheidungen berücksichtigt.

Für die unterste Ausbildungsgruppe finden sich für beide Instrumente relativ kleine negative Lohneffekte, die allerdings statistisch nicht signifikant sind. Eine mögliche Erklärung für diesen Befund ist, dass innerhalb dieses Ausbildungsniveaus Personen trotz Unterschieden in der Berufserfahrung relativ gut miteinander substituierbar sind, sodass auch die Größe der eigenen Kohorte weniger relevant für die Erklärung der Löhne ist. Für das mittlere Ausbildungsniveau findet sich ein negativer Effekt, der statistisch signifikant ist, wenn das Instrument verwendet wird, welches auch die durch Migration hervorgerufene Endogenität adressiert.

Im Gegensatz dazu fällt der auf dem herkömmlichen Instrument basierende Effekt zwar auch negativ aus, ist aber deutlich kleiner und statistisch insignifikant. Dieses Ergebnis legt nahe, dass die bisherige Identifikationsstrategie die Höhe des Lohn-effekts unterschätzt, da der verzerrende Effekt, welcher von Migration in Regionen mit hohen Löhnen ausgeht, nicht berücksichtigt werden kann. Für die höchste Ausbildungsgruppe finden sich zwar ebenfalls negative Effekte, jedoch sind diese nicht statistisch signifikant. Dies könnte daran liegen, dass Arbeitsmärkte auf diesem Niveau stärker segmentiert sind, sodass die hier verwendete Kohortengröße womöglich kein gutes Maß für die Größe einer Gruppe darstellt, innerhalb derer Personen leicht substituierbar sind. Gleichzeitig sinkt auch die Zahl der zur Verfügung stehenden Beobachtungen und der Zusammenhang zwischen der Kohortenvariable und dem Instrument fällt schwächer aus – was auf eine höhere Mobilität der Hochausgebildeten zurückzuführen sein könnte –, wodurch die Identifikation bestehender Effekte erschwert wird.

Regional population structure and young workers' wages

Das nächste Papier befasst sich ebenfalls mit dem Zusammenhang zwischen Kohortengröße und Löhnen, setzt aber im Vergleich zum vorigen andere Schwerpunkte. Anstatt die Unterschiede in den Lohnwirkungen für verschiedene Bildungsgruppen zu identifizieren, werden hier zwei andere Aspekte in den Blick genommen: Erstens wird untersucht, wie sich Messfehler in der Kohortenvariable, die auf einer falschen räumlichen Abgrenzung beruhen, auf die Höhe des Lohneffekts auswirken, und zweitens werden die Mechanismen näher betrachtet, die für den negativen Zusammenhang zwischen Kohortengröße und Löhnen von Bedeutung sind.

Die Kohortenvariable soll das Angebot an Personen mit ähnlichen Eigenschaften abbilden, die einem Arbeitsmarkt zur Verfügung stehen. Wenn diese Größe auf der Grundlage von administrativen Einheiten – z. B. Bundesländer, Regierungsbezirke oder Kreise – berechnet wird, wie es in den meisten bestehenden Untersuchungen (auch in dem ersten Papier dieser Arbeit) der Fall ist, stellt die so bestimmte Kohortenvariable womöglich kein gutes Maß für das Arbeitsangebot auf einem tatsächlichen Arbeitsmarkt dar. Dieses Problem entsteht dadurch, dass administrative Einheiten für gewöhnlich nicht anhand ökonomischer Kriterien abgegrenzt sind, sondern vielmehr einen historischen Ursprung haben und daher auch keine Arbeitsmärkte abbilden. Beispielsweise ist es möglich, dass Personen zwar im Landkreis München wohnen, zum Arbeiten aber in die kreisfreie Stadt München pendeln (und umgekehrt). Eine kreisspezifische Kohortenvariable würde jedoch lediglich auf der Größe einer Altersgruppe innerhalb eines Kreises beruhen und könnte sich somit von der Größe der Kohorte auf dem entsprechenden Ar-

beitsmarkt unterscheiden. Das Vorliegen eines solchen Messfehlers kann die Höhe des geschätzten Lohneffekts beeinflussen (im klassischen Fall führen Messfehler dazu, dass die Koeffizienten zu null hin verzerrt werden).

Grundsätzlich stellt die Verwendung eines Instrumentvariablenschätzers eine Möglichkeit dar, um trotz des Vorhandenseins von Messfehlern die interessierenden Effekte konsistent zu schätzen. In diesem Papier wird jedoch argumentiert, dass die im Zusammenhang mit dem ersten Papier besprochene Identifikationsstrategie nicht geeignet ist, das Problem des Messfehlers zu lösen. Dies wäre nur unter der sehr starken Annahme der Fall, dass die Messfehler in der auf administrativen Einheiten beruhenden Kohortenvariable und dem entsprechenden Instrument – konditional auf die übrigen Kontrollvariablen – keine Korrelation aufweisen. Anstelle von administrativen Einheiten werden in diesem Papier sogenannte Arbeitsmarktregionen als Grundlage für die Berechnung der Kohortenvariable verwendet. Diese setzen sich aus einem oder mehreren Kreisen zusammen, die anhand der zwischen ihnen bestehenden Pendlerverflechtungen miteinander verbunden werden. Auf diese Weise nähern sich diese Einheiten tatsächlichen Arbeitsmärkten an, da die in diesem Gebiet arbeitende Bevölkerung auch weitestgehend dort wohnt und umgekehrt. Auf Arbeitsmarktregionen basierende Kohortenvariablen sollten daher die Größe einer Gruppe innerhalb eines Arbeitsmarkts besser abbilden und somit das Problem des Messfehlers reduzieren können.

Das zweite Ziel des Papiers besteht darin, genauere Aussagen über die Mechanismen zu treffen, die dem negativen Zusammenhang zwischen Löhnen und Kohortengröße zugrunde liegen, welcher – wie bereits im Kontext des ersten Papiers besprochen – theoretisch mit abnehmender Grenzproduktivität begründet wird. Konkret soll untersucht werden, welche Bedeutung der Selektion in bestimmte Wirtschaftszweige oder Berufe zukommt: Falls die Größe einer Kohorte einen Einfluss darauf hat, in welchen Wirtschaftszweigen oder Berufen eine Person Beschäftigung findet und sich diese systematisch in der Höhe der erzielten Entgelte unterscheiden, wäre ein Teil des negativen Lohneffekts auf einen Selektionsmechanismus zurückzuführen, bei dem Mitglieder größerer Kohorten eher in solchen Wirtschaftszweigen oder Berufen beschäftigt sind, in denen niedrigere Löhne gezahlt werden.

Im Gegensatz zum ersten Papier ist die empirische Analyse auf Deutschland beschränkt. Hinsichtlich der vergangenen und im Hinblick auf die für die Zukunft prognostizierten Entwicklungen in der Bevölkerungsstruktur bietet sich Deutschland jedoch als Untersuchungsgegenstand für die Auswirkungen von Veränderungen in Kohortengrößen besonders an. Die Datenbasis bildet die *Stichprobe der Integrierten Arbeitsmarktbiografien* (SIAB). Dieser Datensatz umfasst eine 2 %-Stichprobe der *Integrierten Erwerbsbiografien* (IEB), welche Individualdaten zu allen in Deutschland abhängig Beschäftigten, Arbeitsuchenden, Leistungsemp-

fängern und Teilnehmern an Programmen der aktiven Arbeitsmarktpolitik enthält. Für die empirische Untersuchung können etwas mehr als 100.000 Beobachtungen von Männern im Alter zwischen 15 und 24 Jahren im Zeitraum 1999–2010 genutzt werden, die sich auf 108 westdeutsche Arbeitsmarktregionen verteilen (da für die Untersuchung relevante Daten nicht zur Verfügung stehen, kann Ostdeutschland in dieser Analyse nicht berücksichtigt werden). Wie im vorigen Papier bildet auch hier der Lohn die abhängige Variable. Die Spezifikation der Kohortenvariable fällt an dieser Stelle indes vereinfacht aus, da einerseits von einer feineren Altersdifferenzierung abgesehen und stattdessen der Anteil der Altersgruppe 15–24 an der Bevölkerung im erwerbsfähigen Alter verwendet wird sowie andererseits auf eine Differenzierung nach Bildungsgruppen verzichtet wird. Stattdessen liegt der Fokus auf den eingangs beschriebenen Aspekten. Da die Kohortenvariable in diesem Papier nicht über eine Altersdimension verfügt, wird deren Effekt auf die Löhne allein aufgrund der Veränderung in der Kohortengröße innerhalb einer Region über die Zeit identifiziert. Um das bereits beschriebene Problem der Endogenität, das sich in diesem Fall auf migrationsbedingte Selektion beschränkt, zu adressieren, findet eine vergleichbare Identifikationsstrategie Verwendung wie im ersten Papier.

Die Ergebnisse zeigen, dass die Größe der Kohorte einen negativen Effekt auf die Löhne hat, die von deren Mitgliedern erzielt werden: *Ceteris paribus*, führt ein Anstieg um einen Prozentpunkt zu einem Lohnrückgang von etwa 3 Prozent. Um zu prüfen, wie sich die Verwendung von administrativen räumlichen Einheiten auf die Ergebnisse auswirkt, wird das Modell in einem zweiten Schritt mit einer auf Kreisen statt auf Arbeitsmarktregionen beruhenden Kohortenvariable geschätzt. In dieser Spezifikation fällt der Koeffizient der Kohortenvariable um zwischen 13 Prozent und 50 Prozent geringer aus. Eine Erklärung für die Unterschiede in den Ergebnissen bei der Verwendung von Kreisen statt Arbeitsmarktregionen könnte im oben beschriebenen Messfehler begründet liegen, mit dem die Kohortenvariable im ersten Fall behaftet ist. Für die Literatur ist dieser Befund von Bedeutung, da er nahelegt, dass die Nutzung administrativer regionaler Einheiten als Grundlage für die Bildung der Kohortenvariable zu einer Unterschätzung des Lohneffekts führt.

Um die Frage nach der Rolle von Selektion in Wirtschaftszweige oder Berufe beantworten zu können, werden Indikatorvariablen in das Modell aufgenommen, die die Zugehörigkeit einer Person zu einem bestimmten Wirtschaftszweig oder Beruf wiedergeben, sodass der Effekt von Kohortengröße nun konditional auf diese Informationen geschätzt wird. Wenn für den Wirtschaftszweig einer Person kontrolliert wird, sinkt die Größe des Effekts um zwischen 4 Prozent und 12 Prozent, während es im Fall der Berufe zu einem deutlich stärkeren Rückgang zwischen 30 Prozent und 40 Prozent kommt. Dieses Ergebnis legt nahe, dass ein Teil des negativen Lohneffekts dadurch erklärt werden kann, dass Personen in größeren

Kohorten eher in solchen Berufen oder Wirtschaftszweigen Beschäftigung finden, in denen niedrigere Löhne gezahlt werden.

Cohort size and youth labour-market outcomes: the role of measurement error

Dieses Papier befasst sich mit dem Zusammenhang zwischen Kohortengröße und dem Ausmaß von Arbeitslosigkeit bzw. Beschäftigung in dieser Kohorte. Wenn es infolge einer Änderung in der Größe einer Kohorte nicht oder nur teilweise zu einer Lohnanpassung kommt, ist es möglich, dass stattdessen Veränderungen in der kohortenspezifischen Arbeitslosigkeit bzw. Beschäftigung erfolgen. Im Gegensatz zum Lohneffekt finden sich in der vorliegenden Literatur jedoch verschiedene Hypothesen darüber, welches Vorzeichen dieser Zusammenhang hat.

Einerseits wird argumentiert, dass aufgrund verstärkter Konkurrenz Personen in größeren Kohorten einer höheren Wahrscheinlichkeit ausgesetzt sind, arbeitslos zu sein. Diese Hypothese wird durch einen großen Teil der empirischen Evidenz gestützt. Demgegenüber steht die Vermutung, dass eine relativ große Jugendbevölkerung die Arbeitslosigkeit allgemein sowie innerhalb jüngerer Altersgruppen senkt. Dieses Argument beruht auf der Annahme, dass Unternehmen in der Besetzung von Stellen geringere Kosten entstehen, wenn junge Altersgruppen relativ stark vertreten sind, da junge Personen eine höhere Wahrscheinlichkeit haben, nicht beschäftigt zu sein oder eine Beschäftigung zu haben, die nicht zu ihren Qualifikationen passt. Somit sollte in diesen Altersgruppen die Bereitschaft höher sein, eine neue Stelle aufzunehmen, was es Unternehmen wiederum leichter macht, Stellen zu besetzen. Da sich Veränderungen in der Bevölkerungsstruktur relativ gut prognostizieren lassen, schaffen Unternehmen in Erwartung eines Anstiegs in der Jugendbevölkerung neue Stellen, was wiederum die allgemeine und jugendspezifische Arbeitslosigkeit senkt. In der empirischen Literatur finden sich auch Ergebnisse, die mit dieser Art des Zusammenhangs kompatibel sind.

Dieses Papier liefert zunächst weitere empirische Erkenntnisse darüber, wie Arbeitslosigkeit und Beschäftigung innerhalb einer Kohorte von deren Größe abhängen. Hierzu wird der bereits beschriebene EU-SILC-Datensatz verwendet, der um weitere Jahre ergänzt wird. Jedoch steht in diesem Papier nicht der Effekt auf das Arbeitsmarktergebnis eines Individuums im Vordergrund, sondern es soll vielmehr untersucht werden, wie sich Änderungen in der Größe einer Altersgruppe auf die Anteile der Arbeitslosen und Beschäftigten in dieser Gruppe auswirken. Hierfür werden die in EU-SILC enthaltenen Individualdaten zunächst auf die Ebene einer Region-Jahr-Alter-Zelle aggregiert, woraus sich ein Datensatz von 49 Regionen und 5 Altersgruppen ergibt, die über den Zeitraum 2005–2012 be-

obachtet werden können. Wie in den anderen Papieren auch, wird die Analyse auf Männer beschränkt, um durch selektierte Arbeitsmarkteteiligung von Frauen hervorgerufene Probleme zu vermeiden; darüber hinaus wird eine vergleichbare Identifikationsstrategie angewendet. Die Ergebnisse zeigen, dass bei einer größeren Kohorte der Anteil der Arbeitslosen zurückgeht und der Anteil der Beschäftigten innerhalb der Kohorte steigt. Somit stützen diese Befunde die Hypothese, dass junge Personen von der Größe ihrer Kohorte profitieren. Einschränkend ist jedoch zu erwähnen, dass diese Ergebnisse keine Aussage über die Umstände der Beschäftigung zulassen: Auch wenn ein Anstieg der Kohortengröße zu einem höheren Beschäftigungsanteil führt, ist es möglich, dass diese Veränderung – wie die Ergebnisse der beiden vorigen Papiere nahelegen – mit einem Rückgang in der Höhe der Löhne einhergeht.

Der eigentliche Beitrag dieses Papiers besteht jedoch darin, zu zeigen, dass der geschätzte Effekt auf den Anteil der Arbeitslosen sowie der Beschäftigten stark davon abhängt, welche Altersgruppen in die Analyse aufgenommen werden. Im Speziellen dreht sich das Vorzeichen des Effekts um, wenn anstelle der Altersgruppen 25–29 die Gruppen 18–22 genutzt werden; für die dazwischenliegenden Gruppen bewegt sich das Ergebnis vom einen zum anderen Extrem. Um diesen Befund zu erklären, wird das Argument entwickelt, dass jüngere Altersgruppen weniger geeignet sind, um die Auswirkung von Kohortengröße auf Arbeitsmarktergebnisse zu untersuchen, da ein bedeutender Anteil dieser Gruppen dem Arbeitsmarkt nicht zur Verfügung steht und somit auch für die kohortenspezifischen Arbeitsmarktergebnisse nicht relevant sein sollte. Zum einen wird bei jüngeren Altersgruppen nicht der direkte Effekt von Kohortengröße auf Arbeitsmarktergebnisse gemessen, sondern vielmehr sind die geschätzten Effekte auf Arbeitslosigkeit und Beschäftigung konditional darauf zu interpretieren, dass sich eine Person zuvor dafür entschieden hat, in den Arbeitsmarkt einzutreten. Da die Partizipationsentscheidung allerdings auch von der Größe der Kohorte abhängig sein kann, nimmt der geschätzte Koeffizient beide Effekte auf. Zum anderen lässt sich die Veränderung im Vorzeichen des Effekts womöglich dadurch erklären, dass die Kohortenvariable mit einem Messfehler behaftet ist (allerdings anderer Art als im vorigen Papier). Diese Variable soll die Höhe des Angebots einer bestimmten Gruppe auf dem Arbeitsmarkt messen. Bei jüngeren Altersgruppen ist es allerdings sehr fraglich, ob eine altersspezifische Kohortenvariable ein gutes Maß für diese Größe darstellt, da ein beträchtlicher Anteil dieser Gruppe dem Arbeitsmarkt nicht zur Verfügung steht.

Weiterhin wird argumentiert, dass es sich in diesem Fall um einen nichtklassischen Messfehler handelt, bei dem die eigentlich interessierende Größe – das kohortenspezifische Arbeitsangebot – mit der Größe des Messfehlers korreliert ist. Ein Grund hierfür besteht darin, dass für eine bestimmte Region und einen bestimmten

Zeitpunkt jüngere Altersgruppen typischerweise kleiner sind als ältere und somit auch das Arbeitsangebot – gemessen an der Zahl der Person – geringer ausfallen sollte. Gleichzeitig sollte insbesondere der Anteil derer, die dem Arbeitsmarkt nicht zur Verfügung stehen, aufgrund verstärkter Teilnahme an Bildungsmaßnahmen höher ausfallen. Eine negative Korrelation zwischen der Höhe des kohortenspezifischen Arbeitsangebots und der Höhe des Messfehlers sollte sich dann einstellen, wenn der Unterschied im Anteil der Nichtteilnehmer die Unterschiede in der Größe der Altersgruppen überwiegt. Diese Beziehung wird auch durch die Instrumentierung nicht aufgelöst, da Kohorten, die in der Gegenwart relativ klein sind, auch zu einem früheren Zeitpunkt vergleichsweise klein gewesen sein sollten. Bei älteren Gruppen sollte diese Art des Messfehlers eine geringere Rolle spielen, da der Anteil der Arbeitsmarktteilnehmer deutlich höher ausfallen sollte. Für die Schätzung des Effekts von Kohortengröße auf Arbeitslosigkeit und Beschäftigung sind junge Altersgruppen daher weniger geeignet. Da bestehende Studien oftmals jüngere Altersgruppen in die empirische Analyse aufgenommen haben, ist dieses Ergebnis für die Literatur relevant, da es die Frage aufwirft, in welchem Maß die bisherigen Ergebnisse von Messfehlern in der Kohortenvariable beeinträchtigt sind.

Cohort size and transitions into the labour market

Das letzte Papier befasst sich mit dem Zusammenhang zwischen der Kohortengröße beim Eintritt in den Arbeitsmarkt und der Dauer bis zum Beginn der ersten Beschäftigung. Da die Auswirkungen auf die Suchdauer bisher noch nicht untersucht worden sind, leistet dieses Papier zum einen durch die Wahl einer neuen Ergebnisvariable einen Beitrag zur Literatur. Zum anderen unterscheidet es sich von den zuvor besprochenen Papieren dadurch, dass hier nicht der kontemporäre Zusammenhang zwischen der Größe einer Kohorte und einem bestimmten Arbeitsmarktergebnis betrachtet wird. Stattdessen geht es um die Auswirkung, die die Kohortengröße zu einem bestimmten Zeitpunkt – nämlich beim Eintritt in den Arbeitsmarkt – auf nachfolgende Entwicklungen, in diesem Fall die Suche nach Beschäftigung, hat. Aufgrund dieser Änderung im zeitlichen Kontext des untersuchten Zusammenhangs weist das Papier auch einen Bezug zu einer weiteren Literatur auf, in der der Einfluss von Konjunkturreffekten beim Arbeitsmarkteintritt auf zukünftige Arbeitsmarktergebnisse untersucht wird. Da in diesen Analysen andere Eintrittsbedingungen – z. B. die Größe der Eintrittskohorte – unberücksichtigt bleiben, können die Ergebnisse dieses Papiers auch für diese Literatur von Bedeutung sein.

Um Hypothesen zu bilden, wie sich die Größe der Eintrittskohorte auf die anschließende Dauer der Suche nach Beschäftigung auswirkt, wird auf die Literatur zum bereits im Kontext des vorigen Papiers beschriebenen Zusammenhang

zwischen Kohortengröße und Arbeitslosigkeit zurückgegriffen. Demnach wäre es zunächst möglich, dass in größeren Eintrittskohorten aufgrund der stärker ausgeprägten Konkurrenz auf dem Arbeitsmarkt länger gesucht werden muss, bevor eine Beschäftigung gefunden werden kann. Dieser Effekt könnte jedoch dadurch abgeschwächt (oder umgekehrt) werden, dass Personen, die den Arbeitsmarkt als Teil einer großen Kohorte betreten, Beschäftigungen aufnehmen, die unter ihrem Anforderungsprofil liegen. Schließlich besteht die Möglichkeit, dass es in größeren Eintrittskohorten zu kürzeren Suchdauern kommt, wenn Unternehmen angesichts eines gestiegenen Arbeitsangebots junger Altersgruppen Stellen schaffen.

Grundlage für die Untersuchung des beschriebenen Zusammenhangs bilden Daten zu Absolventen von Ausbildungsprogrammen. Dieser Fokus ist in mehrerer Hinsicht sinnvoll: Erstens ist mit den vorliegenden Daten eine Identifizierung des Orts und des Zeitpunkts des Ausbildungsabschlusses sowie des ersten nachfolgenden Beschäftigungsverhältnisses möglich (für andere Gruppen, z. B. die Hochschulabsolventen, liegen vergleichbare Angaben zum Studienabschluss nicht vor). Zweitens, beinhaltet diese Gruppe nicht nur Personen ähnlichen Alters, sondern auch einer vergleichbaren beruflichen Qualifikation. Im Gegensatz zu ausschließlich nach Alter abgegrenzten Kohorten sollte in diesem Fall, in dem die Eintrittskohorte auf dem Erwerb eines berufsqualifizierenden Abschlusses beruht, auch die Relevanz der Kohorte für den Arbeitsmarkt höher sein, was das im vorigen Papier beschriebene Problem des Messfehlers aufgrund fehlender Teilnahme am Arbeitsmarkt reduzieren sollte. Schließlich ist die Gruppe der Auszubildenden an sich relevant, da es sich hierbei um einen in Deutschland verbreiteten Weg handelt, mittels dessen junge Personen den Arbeitsmarkt betreten. Durch diese Einschränkung sind die Ergebnisse jedoch nicht zwangsläufig auf andere Gruppen, wie die der Hochschulabsolventen oder der Geringqualifizierten übertragbar, für die sich der untersuchte Zusammenhang womöglich anders dargestellt hätte.

In der empirischen Analyse werden zwei Datenquellen aus Deutschland verwendet, die bereits im Kontext des zweiten Papiers beschrieben worden sind: die *Integrierten Erwerbsbiografien* (IEB) sowie die *Stichprobe der Integrierten Arbeitsmarktbiografien* (SIAB). In einem ersten Schritt muss die zentrale erklärende Variable – die Größe der Eintrittskohorte – geschätzt werden, indem auf Grundlage der IEB die Zahl der Personen innerhalb eines bestimmten Zeitraums und in einer bestimmten Arbeitsmarktregion berechnet wird, die eine Reihe an Bedingungen erfüllen, sodass sie als Absolventen eines Ausbildungsprogramms gezählt werden können. Die der eigentlichen Regressionsanalyse zugrundeliegende Stichprobe wird hingegen aus SIAB-Daten gewonnen. Berücksichtigt werden männliche Personen, die zwischen Januar 1999 und Oktober 2012 im Alter von 19 bis 23 Jahren eine Ausbildung abgeschlossen haben. Um zu vermeiden, dass Ab-

solventen aus früheren Jahren systematisch längere Suchdauern aufweisen, werden unterschiedliche Analysen für verschiedene Zeiträume durchgeführt, über die alle Individuen in der Stichprobe ab dem Zeitpunkt des Ausbildungsabschlusses beobachtet werden (3 Monate, 6 Monate, 1 Jahr, 2 Jahre). Erfolgt innerhalb eines solchen Zeitraums ein Übergang in Beschäftigung, so wird er als solcher gezählt, wohingegen für Personen, deren Übergänge zu einem späteren Zeitpunkt erfolgen, die Information genutzt wird, dass es innerhalb des Beobachtungszeitraums nicht zu einer Beschäftigungsaufnahme gekommen ist. Da es sich bei der zu erklärenden Variable um eine Dauer handelt, werden für die empirische Untersuchung Methoden der Verweildaueranalyse und insbesondere das Cox-Modell genutzt.

Die Ergebnisse legen nahe, dass Absolventen, die als Teil einer größeren Kohorte in den Arbeitsmarkt eintreten, schneller eine Beschäftigung finden. Allerdings zeigt sich, dass dieser Effekt nur dann signifikant ist, wenn der dreimonatige Beobachtungszeitraum angewendet wird; bei längeren Zeiträumen ist der Effekt hingegen kleiner und statistisch insignifikant. Für den sechsmonatigen Beobachtungszeitraum stellen sich jedoch sehr ähnliche Ergebnisse ein, sobald nicht nur für die Größe der Kohorte beim eigenen Eintritt in den Arbeitsmarkt kontrolliert wird, sondern auch die Größe der nachfolgenden Eintrittskohorte berücksichtigt wird. Diese Ergebnisse liefern somit keine Evidenz für die erste Hypothese, dass Mitglieder größerer Eintrittskohorten aufgrund verstärkter Konkurrenz längere Suchdauern haben. Weitere Untersuchungen zeigen, dass die Größe der Kohorte keinen negativen Effekt auf die Höhe der Löhne hat, die im ersten Beschäftigungsverhältnis nach der Ausbildung erzielt werden, und auch nicht zu einer höheren Wahrscheinlichkeit führt, dass eine andere als eine sozialversicherungspflichtige Art der Beschäftigung – z. B. eine geringfügige Beschäftigung – aufgenommen wird. Diese Ergebnisse sprechen somit auch gegen die zweite Hypothese, dass sich kürzere Suchdauern bei größeren Eintrittskohorten durch eine Selektion in weniger anspruchsvolle Beschäftigungen erklären lassen.

Alternative Erklärungen für die empirischen Befunde – Selektion der Absolventen in Regionen mit kürzeren Suchdauern nach Beendigung der Ausbildung oder Unterschiede in der Zusammensetzung größerer Kohorten hinsichtlich der Produktivität ihrer Mitglieder – werden ebenfalls nicht durch die empirische Evidenz gestützt. Abschließend finden sich auch keine Belege dafür, dass die Ergebnisse auf die Tatsache zurückzuführen sind, dass das Cox-Modell Personen, die keine Suchdauer aufweisen, da sie direkt nach Beendigung der Ausbildung eine Beschäftigung finden, nicht berücksichtigen kann. Wenn die Ergebnisse auch keinen direkten Beleg für die dritte Hypothese darstellen, dass Unternehmen angesichts größerer Eintrittskohorten neue Stellen schaffen, so ist diese Erklärung doch mit dem Befund kompatibel, dass es Mitgliedern größerer Kohorten schneller gelingt, nach Beendigung der Ausbildung eine Beschäftigung zu finden.

Problem statement, structure and contribution of the dissertation

The aim of this thesis is to contribute to the understanding of how changes in cohort size affect various labour-market outcomes. It is therefore related to a large body of literature that has developed out of the desire to shed light on the implications of the large post-World-War-II birth cohorts entering the US labour markets from the late 1960s onwards (Freeman, 1979; Welch, 1979) and that has since continued to address the relationship between population structure and the labour market. In the part of this literature that is most relevant to my work the subject of interest is typically constituted by the effect that the size of an age group has on group-specific outcomes which is motivated by the assumption that members of different age groups are only imperfectly substitutable and as such compete for jobs mainly within their group. This assumption in turn reflects the view that differently aged individuals can be expected to differ with respect to the amount of work experience and human capital that they have acquired (Welch, 1979) and as long as human capital is a determinant of a worker's productivity on the job, there should be limits to the extent to which substitution across age groups is possible. In terms of economic models this assumption is reflected in workers of different age groups representing separate factors of production (Berger, 1983; Connelly, 1986; Card and Lemieux, 2001).

The central explanatory variable in this context is based on the concept of a cohort, which measures the size of a specific age group. The extant literature differs with respect to exactly how a cohort is defined, with the underlying age groups being either relatively broad, often representing the size of the youth population (Korenman and Neumark, 2000; Shimer, 2001; Biagi and Lucifora, 2008), or being based on single-year age groups (Wright, 1991; Mosca, 2009; Brunello, 2010). Other studies have employed specifications in which cohort size is delineated according to years of experience rather than age (Welch, 1979), with the former variable being argued to be more relevant to determining whether individuals are substitutable. Furthermore, the cohort that an individual belongs to may not only be determined by his age or experience, but also by his level of education (Welch, 1979; Wright, 1991; Mosca, 2009; Brunello, 2010). Such a specification allows for the effects of cohort size to differ between different levels of education but also imposes the assumption that differently educated individuals are active on separate labour markets.

The most commonly used outcome variables in this literature and the ones most relevant to this thesis are cohort-specific wages as well as employment and unemployment rates. In the case of a perfectly competitive labour market an

increase in cohort size should lead, *ceteris paribus*, to a fall in the wages earned in that age group if there is diminishing marginal productivity in production – an illustration of the effects of an outward shift in the labour-supply curve. This relation is shown formally by Brunello (2010), while Michaelis and Debus (2011) develop a model of imperfectly competitive labour markets in which wages are determined by bargaining between firms and monopoly unions. They show that in most cases an increase in the size of an age group will decrease the wages of that group. According to Stapleton and Young's (1988) diminishing-substitutability hypothesis the negative relationship between cohort size and wages should be more pronounced among the highly educated as the former are less easily substitutable across age groups. The majority of the available empirical research provides evidence for a negative wage effect of cohort size and often finds results to be in line with the diminishing-substitutability hypothesis (Welch, 1979; Wright, 1991; Brunello, 2010).

The possibility that wages might not fully adjust in response to changes in cohort size provides the possibility of a relationship between cohort size and cohort-specific employment or unemployment rates. Fertig and Schmidt (2004) argue that larger cohorts may have a higher degree of bargaining power which may help to prevent a downward wage adjustment, while a fixed number of jobs for a specific age group also constitutes a reason for changes in cohort size translating into (un-)employment adjustments (Korenman and Neumark, 2000). In contrast to the case of wage outcomes, there is no consensus on the sign of this relationship. A number of empirical analyses have yielded evidence that increases in cohort size lead to a larger group-specific (Korenman and Neumark, 2000; Biagi and Lucifora, 2008) or overall unemployment rate (Garloff et al., 2013) which would appear to suggest that there are negative labour-market consequences of belonging to a larger cohort. These findings, however, contrast with an argument proposed by Shimer (2001) – which he also supports with empirical evidence – that regions in which the share of young age groups is larger should experience lower youth and overall unemployment rates. This hypothesis rests on the assumption that an increase in the share of youths, who are often either unemployed or poorly matched and thus willing to take up or to switch jobs, makes it easier for firms to fill vacancies, so that an anticipated increase in the youth share is met by an expansion in the number of jobs offered. Skans (2005) also provides evidence that supports the hypothesis that the youth unemployment rate falls with the size of the youth cohort.

The above literature forms the basis for this thesis. The first three papers address issues which in my view represent shortcomings in the available research on cohort-size effects and provide empirical evidence to support this view. In

contrast, the fourth paper analyses the effect on an outcome variable that has so far not been the subject of research in this literature, and treats cohort size as a labour-market entry condition rather than a contemporaneous explanatory variable. The core of each paper is formed by an empirical analysis that assesses the effects of cohort size on individual-specific or group-specific outcomes. Moreover, each paper comes with supplementary material which further elaborates on arguments made in the corresponding paper and provides the results of various sensitivity analyses.

The topic of the first paper is the effect of cohort size on wages and how the former varies across educational groups. It argues that the identification strategy that has so far been used in studies on the wage effect is not suited to purge the endogeneity of the cohort-size variable that can arise because of selected migration into high-wage areas. While the limited amount of cross-national migration makes disregarding this possibility appear innocuous when the size of the cohort is measured at the country level, the former becomes much more of a concern at the regional level. Moreover, in light of what cohort size is supposed to measure – the supply of labour from a specified group within a labour market – it would appear more appropriate to base this variable on regions since they are likely to closer resemble the delineation of labour markets than countries. The results provide evidence – at least for the largest educational group – that the proposed identification strategy produces qualitatively different results – a negative significant effect as opposed to an insignificant one – compared to the previously employed identification strategy.

Identifying the effects of interest is complicated by the fact that the size of a cohort arguably cannot be treated as an exogenous variable: individuals are not randomly allocated to certain cohorts, but can influence which group they belong to at a given point in time through decisions pertaining to migration and investment in education. Since experimental data is not available, this paper – as well as the two subsequent ones – employs an instrumental-variables strategy in order to arrive at a consistent estimate of the cohort-size effect. This approach is not without problems of its own, though. Since two-stage least squares (2SLS) estimation is less efficient than ordinary least squares (OLS), the effect of interest is estimated less precisely. Moreover, its interpretation depends on the chosen instrument which in this case is given by the size of the cohort observed a certain number of years earlier when the members of the cohort were younger by the same amount of years. The estimated wage effect therefore stems from a change in (contemporaneous) cohort size that is caused by a change in its lagged value. This could be problematic if, for example, those who later on migrate represent a selected group of individuals. Finally, the instrument itself, while displaying a high degree of (partial) correlation with the endogenous cohort-size variable, might be

put into question since the problem of selected migration may simply be shifted from the individual to his parents.

The contribution of the second paper is twofold. First, it aims to produce insights into the mechanisms that are behind the negative wage effect of cohort size and finds that a substantial part of this effect is due to selection into lower-paying occupations and, to a lesser extent, industries. Second, it raises the question to what extent the cohort-size variables that are used in other studies contain measurement error. If this variable, as discussed above, is supposed to measure group-specific supply within an actual labour market, it is questionable whether the typically employed administrative units represent a reasonable basis as their delineations are not designed to produce entities within which a specified group of individuals competes for employment. Since (random) measurement error in an explanatory variable leads to attenuation bias, it is possible that the magnitude of the wage effect has been underestimated in previous studies (including the former). The paper proceeds by estimating two separate models in which the cohort-size variable is either derived from administrative units or from the functional labour-market regions derived by Eckey et al. (2006). The former model produces smaller cohort-size coefficients, thereby providing evidence that the choice of the underlying spatial entity is relevant in terms of the magnitude of the estimated effects.

The second paper also differs from the first with respect to the data it uses, which in this case come from register entries rather than from a survey, which may provide more reliable information about certain variables such as wages. Moreover, the data come from a single country, Germany, rather than from a sample of European countries – a feature which might be attractive in terms of reducing the potential of confounding influences. When data from different countries (or regions) is pooled in order to estimate a given model, the implicit assumption is made that the relationship is the same in each case, though the inclusion of appropriate fixed effects allows for country- or region-specific intercepts. However, differences in national labour-market institutions, for example, could lead to the relationship between cohort size and the outcome variable being structurally different between countries. Since the institutional framework can be expected to be more homogenous within a country, use of data from a single country arguably reduces this problem.

Estimating the effect of changes in cohort size on the (un-)employment rate within that cohort is the subject of the third paper. In light of the conflicting empirical evidence that has been produced by the extant literature, this paper provides new insights into this relationship. The main motivation for this analysis, however, is the hypothesis that cohort-size variables are subject to measurement

error when they contain very young age groups, which is often the case in the existing literature. Since a substantial share of individuals in these groups will not be available to the labour market – primarily, though not exclusively, due to participation in education – an age-specific cohort-size variable will provide only a poor measure of labour supply in that group, which in turn may affect size and sign of the estimated effects. The paper develops an argument of non-classical measurement error which the previously discussed identification strategy is unable to correct for. The results indeed show that the estimated cohort-size effects change drastically depending on the age range of the sample.

The final paper addresses the relationship between cohort size and the transition into the labour market by analysing the former's effect at the time of labour-market entry on the duration of search for employment. What sets this analysis apart from the other papers is not only that a new outcome variable is being analysed, but rather that cohort size is not treated as a contemporaneous explanatory variable. Instead the variable represents a condition under which entry into the labour market took place and which might affect subsequent outcomes. Given this setting, there are parallels between the subject of this paper and a recent literature that analyses the long-run effects of the state of the business cycle at the time of labour-market entry on future labour-market outcomes (Stevens, 2007; Kahn, 2010; Brunner and Kuhn, 2014; Cockx and Ghirelli, 2016).

In my view, the contribution of this thesis to the existing cohort-size literature has been to raise questions about the adequacy of the existing empirical methodology to identify the effects of interest and to provide evidence that these matters can have a substantial impact on the results. My understanding from reading the literature is that the cohort-size variable is supposed to quantify the supply of labour by a specified group whose members are reasonably similar so that they can be regarded as substitutable in production and who are active on the same labour market. If this reading is correct, questions about measurement are bound to arise and two have been addressed in this thesis: is it important to base the cohort-size variable on spatial units that approximate actual labour markets and how does the inclusion of very young age groups, substantial shares of which are often not available to the labour market, affect the results. Moreover, conceptualising cohort size as a labour-market entry condition, raises questions for future research that aim at assessing the long-run consequences of having entered the labour market as part of a large or small cohort.

References

- Berger, M.C. (1983) Changes in Labor Force Composition and Male Earnings: A Production Approach, *Journal of Human Resources*, **18**, 177–96.
- Biagi, F. and Lucifora, C. (2008) Demographic and education effects on unemployment in Europe, *Labour Economics*, **15**, 1076–101.
- Brunello, G. (2010) The effects of cohort size on European earnings, *Journal of Population Economics*, **23**, 273–90.
- Brunner, B. and Kuhn, A. (2014) The impact of labor market entry conditions on initial job assignments and wages, *Journal of Population Economics*, **27**, 705–38.
- Card, D. and Lemieux, T. (2001) Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis, *Quarterly Journal of Economics*, **116**, 705–46.
- Cockx, B. and Ghirelli, C. (2016) Scars of recessions in a rigid labor market, *Labour Economics*, **41**, 162–76.
- Connelly, R. (1986) A Framework for Analyzing the Impact of Cohort Size on Education and Labor Earning, *Journal of Human Resources*, **21**, 543–62.
- Eckey, H.-F., Kosfeld, R. and Türck, M. (2006) Abgrenzung deutscher Arbeitsmarktregionen, *Raumforschung und Raumordnung*, **64**, 299–309.
- Fertig, M. and Schmidt, C. (2004) Gerontocracy in Motion: European Cross-Country Evidence on the Labor Market Consequences of Population Ageing, in R. Wright (ed.) *Scotland's Demographic Challenge*, Scottish Economic Policy Network, Stirling, Glasgow.
- Freeman, R.B. (1979) The effect of demographic factors on age-earnings profiles, *Journal of Human Resources*, **14**, 289–318.
- Garloff, A., Pohl, C. and Schanne, N. (2013) Do small labor market entry cohorts reduce unemployment?, *Demographic Research*, **29**, 379–406.
- Kahn, L.B. (2010) The long-term labor market consequences of graduating from college in a bad economy, *Labour Economics*, **17**, 303–16.
- Korenman, S. and Neumark, D. (2000) Cohort Crowding and Youth Labor Markets: A Cross-National Analysis, in D.G. Blanchflower and R.B. Freeman (eds.) *Youth Employment and Joblessness in Advanced Countries*, University of Chicago Press, Chicago.
- Michaelis, J. and Debus, M. (2011) Wage and (un-)employment effects of an ageing workforce, *Journal of Population Economics*, **24**, 1493–511.
- Mosca, I. (2009) Population Ageing and the Labour Market, *Labour*, **23**, 371–95.
- Shimer, R. (2001) The Impact of Young Workers on the Aggregate Labor Market, *The Quarterly Journal of Economics*, **116**, 969–1007.

- Skans, O.N. (2005) Age effects in Swedish local labor markets, *Economics Letters*, **86**, 419–26.
- Stapleton, D.C. and Young, D.J. (1988) Educational Attainment and Cohort Size, *Journal of Labor Economics*, **6**, 330–61.
- Stevens, K. (2007) Adult Labour Market Outcomes: the Role of Economic Conditions at Entry into the Labour Market, available at http://www.iza.org/conference_files/SUMS2007/stevens_k3362.pdf
- Welch, F. (1979) Effects of Cohort Size on Earnings: The Baby Boom Babies' Financial Bust, *Journal of Political Economy*, **87**, S65–S97.
- Wright, R.E. (1991) Cohort size and earnings in Great Britain, *Journal of Population Economics*, **4**, 295–305.

John Moffat^a and Duncan Roth^b

The cohort size–wage relationship in Europe

Abstract

This paper estimates the impact of cohort size on wages using data on young males in European regions covering 2004–2010. The effect of cohort size on wages is identified through an instrumental-variables strategy which, in contrast to previous analyses of European data, addresses self-selection into geographical areas as well as into educational groups. The results suggest that cohort size has a significant negative effect on male wages for individuals with secondary education – the largest group – but not for individuals with less than secondary education or tertiary education. This effect is underestimated if self-selection into geographical areas is not addressed.

JEL classification: J10, J21, J31, R23

Keywords: Cohort size, wages, causal effect, instrumental variables, EU-SILC

1 Introduction

The demographic and educational composition of the European Union (EU) is changing. While the working-age population share is forecast to fall by 2030 (European Commission, 2014), among the population of working age, older groups will see a far smaller fall in population share than younger groups. At the same time, if current trends continue (Eurostat, 2015), the population of the EU will become better educated. In this paper, we provide evidence on the impact of changes in the experience and education profile of the labour force on the wages of men at the start of their career.

The analysis of the effects of cohort size – i.e. the relative size of a group of individuals sharing similar characteristics (such as gender, age/experience and/or education) – on labour market outcomes was initially driven by a desire to understand the economic consequences of the entry of large cohorts of young workers into the US labour market (known as the baby-boom cohorts) in the late 1960s. The literature has since been dominated by US research – a survey of which is provided

^a Durham University, Department of Economics and Finance.

^b Institute for Employment Research (IAB), Düsseldorf, Germany.

by Korenman and Neumark (2000). One question which has been the focus of much research is whether cohort size has a negative impact on wages. To address this question, the early literature proposed a production function with workers of different age or experience representing distinct factors of production (Freeman, 1979; Welch, 1979; Berger, 1983; Connelly, 1986; Stapleton and Young, 1988; and more recently Brunello, 2010). While the proposed models differ with respect to whether they allow for substitution across education, a common assumption is that within each educational group, workers of different age/experience are only imperfectly substitutable. Welch (1979) motivates the assumption of imperfect substitutability across experience levels by proposing a career-phase model in which differently experienced workers perform different tasks.

In a perfectly competitive economy in which factors of production are paid their marginal product, diminishing marginal productivity implies that an increase in the quantity of a specific production factor will reduce its returns. If workers with different levels of experience within a specific educational group are only imperfectly substitutable, an increase in the size of a specific experience–education group will affect mainly the wages of workers in this group. This is shown formally by Card and Lemieux (2001) and Brunello (2010). A large amount of North American empirical research (e.g. Freeman, 1979; Welch, 1979; Leveson et al., 1980; Alsalam, 1985; Berger, 1983; Berger, 1985; Dooley, 1986; Sapozhnikov and Triest 2007; Morin, 2015) has provided evidence in favour of the hypothesis that increases in the size of an experience–education group (i.e. cohort size) reduces wages.

However, Fertig and Schmidt (2004) point out that the effect of an increase in cohort size on wages is less clear in economies in which wages are rigid or the outcome of a bargaining process between employer associations and unions: if wages are rigid, changes in cohort size are likely to cause changes in experience–education–specific (un-)employment rates rather than wages; if they are the result of a bargaining process, a large cohort size has greater bargaining power which could mitigate the previously discussed negative effects of cohort size on wages. More formally, Michaelis and Debus (2011) specify a model in which output is produced by old and young workers and where age-specific wages and unemployment rates are determined by interaction between unions and firms in a right-to-manage framework. Their model suggests that changes in the size of age groups will usually lead to adjustment in age-specific wages, but when changes in the population structure also affect the weights that unions attach to the preferences of both groups, adjustment will take place through changes to age-specific unemployment. Once the framework of analysis is not restricted to a perfectly competitive set-up, the effects of changes in cohort size on wages are consequently *ex ante* uncertain.

Another question that has been addressed in the literature is whether the effect of cohort size on wages differs across educational groups. Stapleton and Young's (1988) 'diminishing-substitutability hypothesis' proposes that substitutability across experience decreases with a worker's level of education. Building on Welch's (1979) career-phase model, they argue that transition through the different career stages is more rapid for workers with less education as less training is required to perform the transition. Consequently, tasks of differently experienced workers are less differentiated for lower levels of education and workers are more easily substitutable across experience levels. In line with the 'diminishing-substitutability hypothesis', many studies (e.g. Welch, 1979; Leveson et al., 1980; Alsalam, 1985; Brunello, 2010) find that the effects of changes in cohort size are more pronounced for the highly educated.

Our focus is on individuals at the beginning of their careers. It is therefore important to note that some studies suggest that depressed earnings are only a temporary phenomenon (e.g. Welch, 1979) as workers in larger cohorts experience faster earnings growth, while others (e.g. Berger, 1985) suggest that cohort size has a permanently depressing effect on wages. By contrast, Berger (1989) finds that large cohorts have initially higher earnings but that, over time, their earnings fall below those of smaller cohorts. He argues that due to 'diminishing substitutability', individuals in large cohorts have less of an incentive to accumulate human capital than those in small cohorts. Larger cohorts therefore have higher wages than smaller cohorts when they are young but lower wages when they get older.

There is relatively little evidence on cohort-size effects on wages in Europe. Wright (1991), using UK data covering the period 1973–1982, finds that cohort size is negatively associated with wages for males with intermediate and higher qualifications with larger effects for the more educated group. However, these effects are only temporary, lasting for five years after assumed labour-market entry for the intermediate-qualifications group and 11 years for the high-qualifications group. Also for the UK, Nickell (1993) finds a negative effect of cohort size on the relative earnings of young men using time-series data covering 1961–1989. Mosca (2009), using Italian data for male workers from the European Community Household Panel (ECHP), obtains results that also support the negative relationship between cohort size and earnings. Opposing results are obtained in two papers that use Swedish data. Klevmarken (1993), using three waves (1984, 1986, 1988) of the Swedish Household Market and Nonmarket Activities (HUS) dataset, regresses average hourly male earnings by age group on a measure of age-specific relative cohort size and interactions with educational indicator variables and age and finds that none of the cohort size-related variables are significant. Dahlberg and Nahum (2003) use longitudinal data from various registers and find

that cohort size has a positive and significant effect on male wages which exists, to different extents, across gender and education groupings.¹ More recently, Brunello (2010) provides an analysis of the cohort size–earnings relationship using ECHP comprising data for the period 1995–2001 from 11 countries. Instrumental-variables (IV) estimation using age-specific cohort size as an instrument shows that cohort size depresses wages and does so to a larger extent for more educated individuals.

Interpreting the results of previous empirical studies is complicated by the potential endogeneity of the cohort-size variable. Most of the recent literature has acknowledged that cohort size is endogenous due to self-selection into educational groups through gaining qualifications. By contrast, self-selection into geographical areas through migration to high-wage areas remains unaddressed in cross-country European studies. Such an omission may be important due to the existence of free movement of individuals within the European Union (EU). One of the main contributions of this paper is therefore the use of an IV strategy that addresses this issue. Another contribution is the use of the Nomenclature of Territorial Units in Statistics (NUTS) region rather than the country as the spatial unit. This is advantageous as it provides greater variation in the measure of cohort size that facilitates the identification of the cohort-size effect. Moreover, labour market regions have been constructed empirically on the basis of observed commuting patterns for several European countries (e.g. Eckey et al., 2006). These entities are generally delineated at a sub-national level and, while a comparable system is not available for the whole of Europe, NUTS1 regions should provide a better approximation than countries to actual labour markets.

The next section provides a description of the dataset, the empirical specification and the identification strategy. The results are presented and discussed in the third section. The final section concludes.

2 Estimation

2.1 Data

The data are taken from various releases of the European Union Statistics on Income and Living Conditions (EU-SILC) survey which consists of cross-sectional and time-series data at the individual and household level for a large number of European

1 Dahlberg and Nahum (2003) use birth rates as a proxy for cohort size which means that their estimates are not direct estimates of the impact of cohort size on wages.

countries.² Different features of EU-SILC are described in Iacovou et al. (2012) and in Berger and Schaffner (2015). The primary purpose of this survey, which is the successor to ECHP, is to provide information on the distribution of income and social exclusion in Europe. However, EU-SILC also contains several variables related to labour-market outcomes, which in combination with the range of individual-level data makes it a suitable dataset for our purposes. Sampling weights are provided to account for the fact that the data does not constitute a random sample.

In contrast to ECHP, EU-SILC is a rolling panel, so individuals are not observed throughout their working life, but are followed for a maximum of 4 years in most countries. For a specific country and year, individual observations are grouped into sub-samples called rotational groups. In most countries, there are four rotational groups per year (there are 9 per year in France, 8 per year in Norway and 1 per year in Luxembourg). A typical longitudinal release covers four years, but will not contain data from all rotational groups. Instead, one rotational group is followed for four years, a second for three years and a third for two years (see Berger and Schaffner, 2015, for an illustration of this structure). Appending observations from different releases therefore allows an increase in the total sample size: starting from the 2011 release, which covers the years 2008–2011, we use the 2010 release to add observations from a new rotational group for the years 2007–2010 for each available country. We continue adding observations in this way back to the 2005 release, resulting in a final dataset that spans the years 2004–2011.³

This procedure makes it necessary to scale down the sampling weights since these are constructed on the basis of the number of rotational groups in a release (for each country-year combination). To assess the quality of the weights we estimate the size of the population in each region-year-age cell and compare these estimates with the corresponding values as reported by Eurostat. Where these two quantities are not identical, the weights of all observations within a cell are scaled so that they yield the true population size. The results reported in Section 3 are, however, robust to using the unadjusted weights.

A further issue is that the variables in EU-SILC refer to different periods. In the case of the labour-market variables, the number of hours worked refers to the time of the interview, while the income variables are based on the income-reference period, which is defined as the preceding calendar year for all countries except Ireland and the UK. To ensure that the variables refer to the same year, we replace the wage data in the sample by its leading value. This implies that

2 Data from the following longitudinal releases is used: 2005-1 from 15-09-2007, version 2006-2 from 01-03-2009, 2007-5 from 01-08-2011, 2008-4 from 01-03-2012, 2009-3 from 01-08-2012, 2010-3 from 01-08-2013 and 2011-1 from 01-08-2013.

3 Notice that not all countries provide observations for the whole period.

data from the year 2011 cannot be used and that only those individuals that are observed in adjacent years can be retained. In terms of countries our final sample includes observations from Austria, Belgium, Bulgaria, Cyprus, Czech Republic, Denmark, Estonia, France, Greece, Hungary, Italy, Latvia, Lithuania, Luxembourg, Malta, Norway, Poland, Romania, Slovakia, Spain and Sweden.⁴ For each of these countries EU-SILC provides information on an individual's residence at the NUTS1 level. This piece of information is crucial as it allows construction of the cohort-size variable at the regional level. The countries listed above provide us with a total of 56 NUTS1 regions.

2.2 Empirical model

The dependent variable of our model is given by the natural logarithm of the purchasing power parity (PPP)-adjusted hourly wage of individual i in experience group j , with educational qualification e , residing in region r at time t , w_{ijert} . This variable is constructed by first adjusting annual wage income for inflation using the GDP deflator (base year: 2010). This variable is then divided by the country-specific PPP-factor from the base year, as provided by Eurostat (see Friedrich, 2015). This quantity is then divided by the number of hours usually worked per week, which are multiplied by the number of weeks per year and the fraction of the year spent working as reported by the individual. To reduce the risk of measurement error due to changes in the number of hours worked over the year, we restrict our sample to those individuals that have been working either exclusively full-time or exclusively part-time during the income-reference period.

The main explanatory variable is the relative size of the experience cohort to which the individual belongs. This variable's specification follows from the assumptions made about the group with whom the individual is substitutable. First, we follow the literature (Card and Lemieux, 2001; Brunello, 2010) in assuming that substitutability is possible within but not across educational categories. The level of education in EU-SILC is given by the 1997 system of the International Standard Classification of Education (ISCED-97) which allows for cross-country comparisons of educational qualifications. This variable assigns a value from 0 (pre-primary education) to 5 (first stage of tertiary education) to every individual. Because of top-coding, individuals with ISCED 6 (second stage of tertiary education) cannot be identified separately but are subsumed into category 5. We

4 Observations from the following countries are excluded: Ireland and the UK (income-reference period is not the preceding calendar year as it is for other countries); Germany, the Netherlands and Portugal (no information on region of residence); Croatia (due to unavailability of data, the instrumental variable cannot be constructed); Slovenia (information on the degree of urbanisation missing); Finland and Iceland (year of birth as well as all age-related variables are not recorded precisely, presumably for disclosure reasons).

follow Brunello (2010) in combining individual categories into larger educational groupings: ISCED 0–2 includes all individuals with at most lower secondary education, ISCED 3–4 combines upper secondary and post-secondary, non-tertiary education and ISCED 5 contains individuals with completed tertiary education.

Second, we assume that individuals compete for jobs within regions rather than countries. This approach is preferable for two reasons. First, it allows the use of inter-regional variation in cohort size to identify the former's effect on wages. More substantively, we argue that labour markets are more likely to exist at a sub-national level because of limitations to mobility or because information about job opportunities decreases with distance from an individual's place of residence. Ideally, we would base cohort size on spatial entities which are delineated in a way that the working population residing in such an area would be exclusively employed there and vice-versa. But while such functional units have been designed for individual countries (see Eckey et al., 2006, for Germany), no comparable units have been defined for the European level. But the fact that functional labour markets tend to be found to be relatively small suggests that the use of NUTS1 regions as approximations of regional labour markets is preferable to the use of countries.

Finally, we choose to define cohort size in terms of labour-market experience rather than age. Within an educational grouping, years of work experience provides a measure of the human capital that individuals have had a chance to accumulate on the job. The use of experience thus provides a better measure of substitutability in the labour market than age and also ties in with Welch's (1979) proposed career-phase model in which workers with different levels of experience differ in terms of the tasks they can perform, making them only imperfectly substitutable. However, results comparable to those presented in Section 3 are obtained when an age-specific cohort-size variable is used.⁵

If individuals are not at all substitutable across experience groups, the appropriate cohort-size variable would be defined simply as the ratio of individuals of experience j with education e in region r at time t , N_{jert} , relative to the number of all individuals with education e in region r at time t , N_{ert} . But since it is likely that individuals are substitutable if they have similar but not necessarily the same level of experience, we follow Wright (1991) and Brunello (2010) in calculating the numerator of the cohort-size variable as a weighted average of the number of individuals with up to two years more or two years less work experience⁶:

5 The results of this and all subsequently mentioned robustness checks are available upon request.

6 Notice that the use of V-shaped weights implies that substitutability decreases with the difference in experience levels (see Wright, 1991, for a discussion). Comparable results to those presented in Section 3 are obtained when different specifications of the numerator are used.

$$CS_{jert} = \frac{(1/9)N_{j-2,ert} + (2/9)N_{j-1,ert} + (3/9)N_{jert} + (2/9)N_{j+1,ert} + (1/9)N_{j+2,ert}}{N_{ert}} \quad [1]$$

Because official statistics regarding the size of education–experience groups at a regional level are not available, these quantities are estimated from the EU-SILC dataset using the adjusted sampling weights. For each of the three educational groups, the sample from which cohort size is calculated includes males and females who are either employed or unemployed. Given that our focus is on individuals in the early stages of their career, a large share of inactive individuals are in the process of acquiring education and including those observations would, for example, mean including all individuals enrolled in tertiary education in the construction of cohort size for ISCED grouping 3–4, which in turn would lead to an artificial jump in cohort size once these individuals have completed education and entered the ISCED 5 grouping. The inactive are therefore excluded from the construction of the cohort-size variable. However, comparable results are obtained when cohort size is constructed from all individuals regardless of their economic status.

The sample from which the cohort-size variable is constructed is restricted to the working-age population (age groups 16–65) within each educational grouping. Cohort size is computed for up to 11, 9 and 5 years of experience for ISCED 0–2, ISCED 3–4 and ISCED 5, respectively. These upper limits are imposed for two reasons: first, our interest lies in individuals who are at an early stage of their career. Second and as discussed in Section 2.3, we want to ensure that the age groups which are used as instruments for the experience-based cohort-size variable do not contain individuals who are older than 15 in order to rule out issues of regional self-selection. Furthermore, the denominator includes individuals with up to 47, 45 and 41 years of experience in the case of ISCED 0–2, ISCED 3–4 and ISCED 5. These values are derived from assuming education-specific ages at entry into the labour market of 18, 20 and 24, respectively. In Section 2.3 we discuss how these assumptions fit the actual data from the regression sample.

Since there is an upper and a lower limit to experience, the construction of cohort size has to be adjusted at the corners of this range by reallocating the weights that would otherwise have been attached to the experience groups outside the specified range. At the lower limit, cohort-size for experience groups 0 and 1 are constructed as follows (with corresponding constructions at the education-specific upper limits):

$$CS_{0,ert} = \frac{(6/9)N_{0,ert} + (2/9)N_{1,ert} + (1/9)N_{2,ert}}{N_{ert}} \quad [1a]$$

$$CS_{1,ert} = \frac{(3/9)N_{0,ert} + (3/9)N_{1,ert} + (2/9)N_{2,ert} + (1/9)N_{3,ert}}{N_{ert}} \quad [1b]$$

In terms of control variables, x_{ijrt} , we include a constant, individual-level regressors (indicators for working part-time, being married, the degree of urbanisation of the place of residence, and occupational indicator variables), experience-related regressors (experience and squared experience), region-specific regressors (region dummies), time-specific regressors (year dummies) and region-by-time regressors (the regional unemployment rate). Further details on these variables as well as descriptive statistics are given in Tables A1 and A2 in the Appendix.

Our empirical model, which we estimate separately for each level of education (ISCED 0–2, ISCED 3–4, ISCED 5), is given by Equation 2 (throughout the remainder of the paper the e subscript is dropped):

$$\ln[w_{ijrt}] = \alpha CS_{jrt} + \beta x_{ijrt} + u_{ijrt} \quad [2]$$

We exclude female observations from the estimation of Equation 2 to avoid the issue of selected labour-market participation. To account for the sampling design weighted regressions are performed. Finally, as the main regressor in our model is defined at a higher level of aggregation than the dependent variable, standard errors are clustered at the level of the region-experience cell (see Moulton, 1990).

2.3 Identification

An obstacle to identifying the wage effect of cohort size using ordinary least squares (OLS) estimation is that individuals are not necessarily randomly allocated into cohorts. Rather membership of a specific cohort is potentially the result of individual self-selection into educational groups or regions. This would be the case if an individual's expectation about future wages affected the decision to acquire a specific level of education (thereby affecting education-specific cohort size) or if labour-market prospects induced migration into a different region (thereby affecting region-specific cohort size). Due to the comparatively low costs of moving between regions (as opposed to countries) the second type of self-selection is of particular concern although freedom of movement of labour within the EU implies that migration between countries may also be significant. OLS is likely to underestimate the depressing effect of cohort size, if individuals select into educational groups or regions that are characterised by higher wages. To identify the effect of cohort size on wage consistently we therefore employ IV estimation.

Recent contributions to the literature on cohort–size effects on wages either do not address the issue of endogeneity (Mosca, 2009) or acknowledge self-selection into educational groups, while implicitly disregarding self-selection through migration (Sapozhnikov and Triest, 2007; Brunello, 2010).⁷ The latter studies use contemporaneous age-specific cohort size as an instrument which is not differentiated by education. We argue that this approach suffers from the disadvantage of not addressing individual self-selection from migration. To assess this hypothesis we construct an instrumental variable (IV1) that corresponds to the one described above:

$$IV1_{gkt} = \frac{(1/9)N_{(g-2)rt} + (2/9)N_{(g-1)rt} + (3/9)N_{grt} + (2/9)N_{(g+1)rt} + (1/9)N_{(g+2)rt}}{N_{rt}} \quad [3]$$

Subscript g refers to age and the numerator is a weighted average of the number of individuals in a region that are two years younger, one year younger, the same age, one year older and two years older. Our preferred instrument (IV2) deals with both self-selection into educational groupings and self-selection into geographical areas. It is the relative size of the age group in the region that is fourteen years younger, fourteen years ago. Since the first year of sampling is the year 2004 and regional population data are available for most NUTS1 regions from the year 1990 onwards, fourteen years represents the longest feasible lag. Comparable instruments have been employed in the analysis of cohort-size and unemployment (Korenman and Neumark, 2000; Shimer, 2001; Garloff et al., 2013; Moffat and Roth, 2016).

$$IV2_{jrt} = \frac{(1/9)N_{g-2-14,rt-14} + (2/9)N_{g-1-14,rt-14} + (3/9)N_{g-14,rt-14} + (2/9)N_{g+1-14,rt-14} + (1/9)N_{g+2-14,rt-14}}{N_{rt-14}} \quad [4]$$

This variable is a natural predictor for our cohort-size variable as, in the absence of migration and natural changes in population, the individuals on which the instrument is based will be the same as those on which education-specific cohort-size is based, only that they are observed at different points in time. This association between the endogenous cohort-size variable and the instrument is supported by the first-stage test-statistics.

In addition to not varying across education, both of the above instruments are defined in terms of age rather than experience. This requires us to specify a link between an individual's age and years of experience. We do this with imputed

7 In contrast, Morin (2015) uses a natural experiment given by a reform to the educational system in a specific Canadian province to identify cohort-size effects.

age values which are defined as the sum of assumed entry age (18, 20 and 24 for educational groupings ISCED 0–2, ISCED 3–4 and ISCED 5) and number of years of experience. We compare actual and imputed age and find that the distribution of true age is centred on the imputed age values in the majority of cases. We prefer matching cohort size and the instrument using an imputed age rather than actual age as this ensures that individuals in the same experience cohort are assigned the same value of the instrument, thereby avoiding identification of cohort-size effects from within-cohort variation in the instrument. To avoid the inclusion of an age group where individuals may make their own decisions about where to reside, the age groups in the instrument are restricted to be no older than 15. This implies an upper age limit of 29 for those in the sample and we therefore exclude observations from the regression who are older than 29 (though raising the limit to 32 does not affect the results).

3 Results

Table 1 shows the coefficients of cohort size, experience and experience squared for each of the three education groups obtained by OLS, two-stage-least squares (2SLS) estimation using the instrument of Brunello (2010) and Wright (1991) – in the column headed IV1 – and 2SLS using an instrument based on lagged population sizes – in the column headed IV2. A full set of results can be found in Table A3 of the Appendix.

Each of the three specifications produces negative cohort-size coefficients for all ISCED groups and, with the exception of ISCED 5, the coefficients of model IV2 are more negative than those of either OLS or IV1. This finding is in line with the previous discussion that due to their inability to account for self-selection through migration into high-wage areas the identification strategies of the latter models will underestimate the negative wage effect of cohort size. However, the size of the standard errors of specifications IV1 and IV2 suggest that the difference between the two point estimates is not statistically significant.

In the case of ISCED 0–2 we find that none of the cohort-size coefficients is statistically significant. From a theoretical perspective (see Card and Lemieux, 2001; Brunello, 2010), this finding is compatible with individuals with different levels of experience in this educational category being easily substitutable for each other. Accordingly, the estimated effect of an increase in cohort-size of one standard deviation on an individual's wage is comparatively small at –3% for specification IV2. In contrast, we find that cohort size has a considerable and statistically significant effect on the wages of individuals with completed secondary and post-secondary, non-tertiary education (ISCED 3–4): based on specification IV2, an increase in

cohort size of one standard deviation decreases the wages of individuals in the affected cohort by 10%, *ceteris paribus*. The fact that the estimated effect is larger for ISCED 3–4 than for ISCED 0–2 is in line with Stapleton and Young’s (1988) ‘diminishing-substitutability hypothesis’ that differently aged/experienced workers are less easily substitutable at higher levels of education.

Table 1: Cohort size coefficients obtained from weighted regression (OLS and 2SLS)

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort Size	-1.42 (1.11)	-0.95 (3.22)	-2.54 (2.67)	-0.21 (0.94)	-3.91 (3.50)	-12.02** (4.70)	-1.87 (1.56)	-3.65 (7.44)	-1.94 (7.66)
Experience	0.08*** (0.01)	0.08*** (0.01)	0.08*** (0.01)	0.05*** (0.01)	0.06*** (0.01)	0.07*** (0.01)	0.10*** (0.03)	0.11** (0.05)	0.10** (0.05)
Experience ²	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.01* (0.00)	-0.01 (0.01)	-0.01 (0.01)
N(inds)	7,364	7,364	7,364	19,785	19,785	19,785	4,973	4,973	4,973
N(cells)	2,180	2,180	2,180	2,499	2,499	2,499	1,338	1,338	1,338
N(clusters)	596	596	596	553	553	553	319	319	319
F-stat		115.81***	124.82***		53.24***	44.63***		27.37***	21.43***
ME (std)	-1.71%	-1.14%	-3.05%	-0.17%	-3.28%	-10.07%**	-2.28%	-4.46%	-2.37%

Control variables from Table A1 are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region–experience level. N(inds): number of individual observations. N(cells): number of region–experience–year cells. N(clusters): number of clusters. F-stat refers to the first-stage F-statistic on the significance of the instrument in the first-stage regression of the endogenous cohort-size variable. ME(std) shows the percentage change in the hourly wage for a change in cohort size by one standard deviation.

The results for ISCED 5 do not display a similar pattern: though all coefficients are negative, the point estimates of specification IV2 are smaller than the corresponding results for ISCED 3–4 as well as the coefficients from IV1 in the same educational group. Moreover, none of the coefficients on cohort size are statistically significant. The only other study of which we are aware which has found larger negative cohort-size effects for those with secondary education than those with tertiary education is Dooley (1986) who obtained this result using Canadian data.

There are several potential reasons for this finding. From an empirical perspective, the comparatively small number of experience cells (5) reduces the variation from which the effect of cohort size can be identified (as evidenced by the number of region–year–experience or region–experience cells in the case of ISCED 5). The size of the standard errors in the IV estimations is also a consequence of the instrument’s decrease in predictive power with respect to cohort size (as evidenced by the comparatively small values of the F–statistic). From an economic

perspective, the smaller size of the coefficients may be explained by greater segmentation of the labour market at higher levels of education. In other words, individuals with higher levels of education operate in more heterogeneous labour markets and are therefore less substitutable with individuals with the same level of education, irrespective of their level of experience. An alternative explanation is that individuals, once they have attained tertiary education, are more affected by the mechanism discussed by Berger (1989) that leads individuals in larger cohorts to obtain less human capital and therefore relatively high wages when young. This would be the case if, as seems likely due to opportunities to pursue postgraduate education or receive advanced on-the-job training, individuals with ISCED 5 have more scope for differentiated levels of human capital than individuals with ISCED 0–2 or ISCED 3–4. If those with ISCED 5 in large cohorts choose not to take these opportunities, there then will be a weaker relationship between cohort size and wages within this group. Another possibility is that the effect of cohort size occurs more through unemployment than wages for individuals with tertiary education. However, it is unclear why the wages of those with tertiary education would be more rigid or more influenced by unions so we regard the previous explanations for our inability to find significant and negative effects for this group as more credible.

The first-stage F-statistics are above the rule-of-thumb value of 10 for each educational group, suggesting that there is no problem of weak instruments. The size of the test statistics decreases with the level of education which implies that lagged age-structures are a better predictor for education-specific cohort size of the less educated. A possible reason is that geographic mobility increases with the level of education. The experience variables show the standard positive but diminishing effect of experience on wages. If experience dummies are used, this pattern is also found and the estimated effects of cohort size are very similar to those reported above. The coefficients of the other control variables (reported in Table A3) are also in line with expectations. Specifically, higher regional unemployment is associated with lower wages while being married and living in a more urbanised environment have a positive effect on wages. The coefficients on the occupational dummies are also statistically significant and of the expected pattern.

To put our results for ISCED 3–4 into perspective, we compare them with those of Brunello (2010), who uses a dataset and empirical model that is comparable to the one used in this paper. One difference between the two analyses is that his data is aggregated at the level of the country-year-age cell, while this analysis is based on individual-level variables. However, estimating Equation 2 after averaging all variables over the observations within a region-year-experience cell and weighting the regression by the number of observations per cell (adjusted for

the sampling weights, see Angrist and Pischke, 2009), we obtain results that are very similar to those shown in Table 1. Another difference is Brunello's (2010) use of a log–log specification. When we adopt this approach, we find that an increase in cohort size by 1% is predicted to decrease the mean wage of individuals in that cohort by 0.098% (IV1). This result is comparable to the predicted change of –0.069%, as estimated by Brunello (2010) for those aged below 35 in educational group ISCED 3–4. However, employing our preferred instrument (IV2) yields a predicted decrease of –0.288%, four times the size of the effect found by Brunello (2010). Notwithstanding the possibility that the difference is due to differences in the sampling period and range of countries included in the analysis, this supports our contention that, as discussed previously, the contemporaneous age–cohort size is unable to deal with self–selection through migration and use of this instrument leads to an underestimation of the true cohort–size effect.

4 Conclusion

The aim of this paper has been the identification of the causal effect of cohort size on the wages of young men at the start of their career in Europe. Ex ante, the direction of this effect is unclear. If labour markets are perfectly competitive and differently aged workers are only imperfectly substitutable within each educational group, members of larger cohorts can be expected to receive lower wages as a result of their lower marginal productivity. However, in an environment of imperfect competition, increases in cohort size may have no or only a limited effect on wages if these are sufficiently rigid (in which case (un–)employment rates would be expected to change) or even a positive effect if larger cohorts are able to exert larger bargaining power. Identification of this effect is complicated by the fact that an individual's cohort is likely to be the result of self–selection into educational groups and self–selection into geographic areas. Unlike earlier papers that have looked at this question using cross–country European data, our approach addresses both types of self–selection by using the size of the population 14 years younger, 14 years ago as an instrument for cohort size. We also use regions rather than countries as the spatial unit since the former provides greater variation in the cohort–size variable and are also likely to provide a better approximation of actual labour markets than countries.

Our results show that cohort size represents a significant and negative determinant of wages for young males with secondary but not for those with less than secondary or tertiary education. This suggests that the projected fall in the share of young people in the labour force will put upward pressure on the wages of those with secondary education – the largest group in the labour force. The finding that those with lower levels of education do not experience

a negative effect is consistent with the 'diminishing-substitutability hypothesis'. We suggest that the failure to find a significant effect for those with tertiary education may be due to greater market segmentation among the more highly educated or greater scope for obtaining different levels of human capital after the completion of formal education. The effect of cohort size is not found to be statistically significant for any of the educational groupings if the IV strategy does not address the potential for self-selection through migration. This implies that the earlier work on the cohort size–wage relationship which did not address this source of endogeneity may have underestimated the true effect.

Acknowledgements

The authors would like to thank Eckhardt Bode, Bernd Hayo, Michael Kirk and Christian Traxler as well as the participants of the MACIE brownbag seminar, the MAGKS seminar, the 2012 Dutch Demographic Day, the 2013 Alpine Population Conference, the 2013 Congress of the European Association of Labour Economists (EALE), the 2013 Congress of the European Regional Science Association (ERSA) and the 8th European Workshop on Labour Markets and Demographic Change for valuable comments. This paper uses data from the European Union Statistics on Income and Living Conditions (EU-SILC). The results and conclusions are those of the authors and not those of Eurostat, the European Commission or any of the national statistical authorities whose data have been used.

References

- Alsalam, N. (1985) The Effect of Cohort Size on Earnings: an Examination of Substitution Relationships, Working Paper 21, Centre for Economic Research, University of Rochester, Rochester.
- Angrist, J. and Pischke, J.–S. (2009) *Mostly harmless econometrics: an empiricist's companion*, Princeton University Press, Princeton.
- Berger, M.C. (1983) Changes in Labor Force Composition and Male Earnings: A Production Approach, *Journal of Human Resources*, 18, 177–96.
- Berger, M.C. (1985) The Effect of Cohort Size on Earnings Growth: A Reexamination of the Evidence, *Journal of Political Economy*, 93, 561–73.
- Berger, M.C. (1989) Demographic Cycles, Cohort Size, an Earnings, *Demography*, 26, 311–21.
- Berger, M. and Schaffner, S. (2015) A Note on How to Realize the Full Potential of the EU-SILC Data, ZEW Discussion Paper 15-005, Centre for European Economic Research, Mannheim.

- Brunello, G. (2010) The effects of cohort size on European earnings, *Journal of Population Economics*, 23, 273–90.
- Card, D. and Lemieux, T. (2001) Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis, *Quarterly Journal of Economics*, 116, 705–46.
- Connelly, R. (1986) A Framework for Analyzing the Impact of Cohort Size on Education and Labor Earning, *Journal of Human Resources*, 21, 543–62.
- Dahlberg, S. and Nahum, R.–A. (2003) Cohort Effects on Earnings Profiles: Evidence from Sweden, Working Paper Series 2003–11, Department of Economics, University of Uppsala, Uppsala.
- Dooley, M. (1986) The Overeducated Canadian? Changes in the Relationship among Earnings, Education, and Age for Canadian Men: 1971–81, *The Canadian Journal of Economics/Revue canadienne d'Economie*, 19, 142–59.
- Eckey, H.–F., Kosfeld, R. and Türck, M. (2006) Abgrenzung deutscher Arbeitsmarktregionen, *Raumforschung und Raumordnung*, 64, 299–309.
- European Commission (2014) *The 2015 Ageing Report*, European Commission, Brussels.
- Eurostat (2015) Education and Training, 2015, available at <http://ec.europa.eu/eurostat/web/education-and-training/data/database>, (Accessed: 16/07/2015).
- Fertig, M. and Schmidt, C. (2004) Gerontocracy in Motion: European Cross-Country Evidence on the Labor Market Consequences of Population Ageing, in R. Wright (ed.) *Scotland's Demographic Challenge*, Scottish Economic Policy Network, Stirling, Glasgow.
- Freeman, R.B. (1979) The effect of demographic factors on age-earnings profiles, *Journal of Human Resources*, 14, 289–318.
- Friedrich, C. (2015) Does Financial Integration Increase Welfare? Evidence from International Household-Level Data, Bank of Canada Working Paper 4, Bank of Canada, Ottawa.
- Garloff, A., Pohl, C. and Schanne, N. (2013) Do small labor market entry cohorts reduce unemployment?, *Demographic Research*, 29, 379–406.
- Iacovou, M., Kaminska, O. and Levy, H. (2012) Using EU-SILC data for cross-national analysis: strengths, problems and recommendations, ISER Working Paper Series 3, Institute for Social and Economic Research, University of Essex, Colchester.
- Klevmarken, N. (1993) On ageing and earnings, in P. Johnson and K. Zimmermann (eds.) *Labour markets in an ageing Europe*, Cambridge University Press, Cambridge.

- Korenman, S. and Neumark, D. (2000) Cohort Crowding and Youth Labor Markets: A Cross-National Analysis, in D.G. Blanchflower and R.B. Freeman (eds.) *Youth Employment and Joblessness in Advanced Countries*, University of Chicago Press, Chicago.
- Leveson, I., Newitt, J., United States Department of Health Education and Welfare and Hudson Institute (1980) *Generational Crowding: Economic, Social and Demographic Effects of Changes in Relative Cohort Size*, Hudson Institute, Washington D.C..
- Michaelis, J. and Debus, M. (2011) Wage and (un-)employment effects of an ageing workforce, *Journal of Population Economics*, 24, 1493–511.
- Moffat, J. and Roth, D. (2016) Cohort size and youth labour-market outcomes: the role of measurement error, IAB Discussion Paper No. 37/2016, Institute for Employment Research, Nuremberg.
- Morin, L.-P. (2015) Cohort size and youth earnings: Evidence from a quasi-experiment, *Labour Economics*, 32, 99–111.
- Mosca, I. (2009) Population Ageing and the Labour Market, *Labour*, 23, 371–95.
- Moulton, B.R. (1990) An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units, *Review of Economics and Statistics*, 72, 334–38.
- Nickell, S. (1993) Cohort Size Effects on the Wages of Young Men in Britain, 1961–1989, *British Journal of Industrial Relations*, 31, 459–69.
- Sapozhnikov, M. and Triest, R.K. (2007) Population Aging, Labor Demand, and the Structure of Wages, Research Department Working Paper 07-8, Federal Reserve Bank of Boston, Boston.
- Shimer, R. (2001) The Impact of Young Workers on the Aggregate Labor Market, *The Quarterly Journal of Economics*, 116, 969–1007.
- Stapleton, D.C. and Young, D.J. (1988) Educational Attainment and Cohort Size, *Journal of Labor Economics*, 6, 330–61.
- Welch, F. (1979) Effects of Cohort Size on Earnings: The Baby Boom Babies' Financial Bust, *Journal of Political Economy*, 87, S65–S97.
- Wright, R.E. (1991) Cohort size and earnings in Great Britain, *Journal of Population Economics*, 4, 295–305.

Appendix

Table A1: Variable definitions

Variable	Definition	Source
Wage	Hourly wage in Euros, adjusted by a purchasing-power-parity factor (see the appendix of Friedrich, 2015 for details)	EU-SILC
Cohort size	See Eq. 1 and related discussion	EU-SILC
Married	Dummy variable coded one if the individual is married	EU-SILC
Part-time	Dummy variable coded one if the individual defines himself as part-time	EU-SILC
Self-employed	Dummy variable coded one if the individual defines himself as self-employed	EU-SILC
Occupation	Dummy variables for each of the following occupational groupings defined according to the International Standard Classification of Occupations, (ISCO) – 88: 1. Legislators 2. Senior officials and managers 3. Professionals 4. Technicians and associate professionals 5. Clerks, Service workers and shop and market sales workers 6. Skilled agricultural and fishery workers 7. Craft and related trades workers 8. Plant and machine operators and assemblers 9. Elementary occupations The omitted category is those in the armed forces	EU-SILC
Urbanisation	Dummy variables for residence in the following: 1. An 'intermediate' area – an area with a population density of more than 100 inhabitants per square kilometre (km) and either a population of at least 50,000 inhabitants or adjacent to a 'densely populated' area of at least 500 inhabitants per square km and a population of least 50,000 inhabitants 2. A 'thinly populated' area – an area with fewer than 100 inhabitants per square km and a population of less than 50,000 inhabitants The omitted category is 'densely populated' – an area with a population density of more than 500 inhabitants per square km and a population of at least 50,000 inhabitants.	EU-SILC
Experience	Years of experience	EU-SILC
Unemployment	Regional unemployment rate	Eurostat
Region Dummies	Dummy variables for residence in particular region (see footnote 6 for a list of countries included in the analysis)	EU-SILC
Year Dummies	Dummy variables for 2005, 2006, 2007, 2008, 2009 or 2010	EC-SILC

Table A2: Descriptive statistics

	ISCED 0–2		ISCED 3–4		ISCED 5	
	Mean	Standard deviation	Mean	Standard deviation	Mean	Standard deviation
Wage	7.31	3.89	7.11	4.54	10.01	6.20
Cohort size	0.03	0.01	0.03	0.01	0.04	0.01
Cohort size (instrument 1)	0.02	0.00	0.02	0.00	0.02	0.00
Cohort size (instrument 2)	0.02	0.00	0.02	0.00	0.02	0.00
Experience	5.10	3.16	4.41	2.58	2.70	1.52
Unemployment	8.44	4.31	7.49	3.40	7.89	3.58
Married	0.14	0.34	0.16	0.37	0.13	0.33
Part-time	0.08	0.27	0.05	0.21	0.07	0.25
<i>Occupation dummies</i>						
1. Armed forces	0.01	0.11	0.03	0.16	0.01	0.07
2. Legislators	0.00	0.06	0.01	0.10	0.05	0.21
3. Senior officials and managers	0.01	0.09	0.02	0.13	0.40	0.49
4. Professionals	0.03	0.18	0.13	0.33	0.23	0.42
5. Technicians and associate professionals	0.03	0.17	0.07	0.25	0.10	0.30
6. Clerks, service workers and shop and market sales workers	0.10	0.30	0.14	0.35	0.06	0.25
7. Skilled agricultural and fishery workers	0.03	0.16	0.02	0.14	0.01	0.09
8. Craft and related trade workers	0.40	0.49	0.32	0.47	0.07	0.25
9. Plant and machine operators and assemblers	0.17	0.37	0.17	0.38	0.04	0.21
10. Elementary occupations	0.22	0.42	0.10	0.29	0.03	0.17
<i>Urbanisation dummies</i>						
1. Densely populated	0.37	0.48	0.41	0.49	0.60	0.49
2. Intermediately populated	0.27	0.44	0.22	0.42	0.18	0.38
3. Thinly populated	0.37	0.48	0.37	0.48	0.22	0.41
Observations		7,364		19,785		4,973

Table A3: OLS and 2SLS regression results

	ISCED 0–2		ISCED 3–4		ISCED 5		IV2	IV1	IV2	IV1	IV2
	OLS	IV1	IV2	OLS	IV1	IV2					
Cohort size	-1.42 (1.11)	-0.95 (3.22)	-2.54 (2.67)	-0.21 (0.94)	-3.91 (3.50)	-12.02** (4.70)	-1.87 (1.56)	-3.65 (7.44)	-1.94 (7.66)		
Experience	0.08*** (0.01)	0.08*** (0.01)	0.08*** (0.01)	0.05*** (0.01)	0.06*** (0.01)	0.07*** (0.01)	0.10*** (0.03)	0.11** (0.05)	0.10** (0.05)		
Experience ²	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.01* (0.00)	-0.01 (0.01)	-0.01 (0.01)		
Unemployment	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.02*** (0.00)	-0.02*** (0.00)	-0.02*** (0.00)	-0.00 (0.01)	-0.01 (0.01)	-0.00 (0.01)		
Married	0.07*** (0.02)	0.08*** (0.02)	0.07*** (0.02)	0.07*** (0.02)	0.06*** (0.02)	0.05*** (0.02)	0.12*** (0.03)	0.12*** (0.03)	0.12*** (0.03)		
Part-time	0.01 (0.05)	0.01 (0.05)	0.01 (0.05)	0.03 (0.04)	0.03 (0.03)	0.02 (0.03)	0.02 (0.07)	0.03 (0.07)	0.02 (0.07)		
<i>Occupation</i>											
Legislators	-0.05 (0.14)	-0.05 (0.14)	-0.05 (0.14)	0.02 (0.06)	0.01 (0.06)	0.01 (0.06)	0.22 (0.20)	0.21 (0.20)	0.22 (0.20)		
Senior officials	-0.12 (0.11)	-0.12 (0.11)	-0.12 (0.11)	-0.01 (0.07)	-0.00 (0.06)	-0.00 (0.06)	0.19 (0.20)	0.19 (0.19)	0.19 (0.19)		
Professionals	-0.24*** (0.06)	-0.24*** (0.06)	-0.24*** (0.06)	-0.04 (0.05)	-0.04 (0.05)	-0.04 (0.05)	-0.01 (0.20)	-0.01 (0.19)	-0.01 (0.19)		
Technicians	-0.23*** (0.05)	-0.23*** (0.05)	-0.23*** (0.05)	-0.10** (0.05)	-0.11** (0.05)	-0.11** (0.05)	-0.02 (0.20)	-0.02 (0.19)	-0.02 (0.19)		
Clerks	-0.34*** (0.06)	-0.34*** (0.05)	-0.34*** (0.06)	-0.22*** (0.05)	-0.22*** (0.05)	-0.23*** (0.05)	-0.19 (0.20)	-0.19 (0.20)	-0.19 (0.20)		

	ISCED 0–2		ISCED 3–4			ISCED 5			
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Skilled agricultural and fishery workers	-0.31*** (0.06)	-0.31*** (0.06)	-0.32*** (0.06)	-0.28*** (0.06)	-0.28*** (0.06)	-0.28*** (0.06)	-0.23 (0.22)	-0.24 (0.22)	-0.23 (0.22)
Craft and related trade workers	-0.26*** (0.04)	-0.26*** (0.04)	-0.26*** (0.04)	-0.19*** (0.05)	-0.19*** (0.05)	-0.20*** (0.05)	-0.06 (0.19)	-0.05 (0.19)	-0.06 (0.19)
Plant and machine operators	-0.13*** (0.04)	-0.13*** (0.04)	-0.13*** (0.04)	-0.14*** (0.05)	-0.14*** (0.05)	-0.14*** (0.05)	-0.06 (0.21)	-0.06 (0.20)	-0.06 (0.20)
Elementary occupations	-0.30*** (0.05)	-0.30*** (0.05)	-0.30*** (0.05)	-0.24*** (0.05)	-0.24*** (0.05)	-0.24*** (0.05)	-0.17 (0.21)	-0.18 (0.21)	-0.17 (0.21)
<i>Urbanisation</i>									
Intermediately	-0.01 (0.03)	-0.01 (0.03)	-0.01 (0.03)	-0.01 (0.01)	-0.01 (0.01)	-0.01 (0.01)	-0.03 (0.03)	-0.03 (0.03)	-0.03 (0.03)
Thinly	-0.05** (0.02)	-0.05** (0.02)	-0.05* (0.02)	-0.06*** (0.01)	-0.06*** (0.01)	-0.06*** (0.01)	-0.03 (0.03)	-0.03 (0.03)	-0.03 (0.03)
Constant	2.00*** (0.12)	1.97*** (0.19)	2.05*** (0.16)	2.42*** (0.08)	2.54*** (0.12)	2.78*** (0.16)	2.28*** (0.24)	2.33*** (0.31)	2.28*** (0.31)
Dummies									
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations									
Individuals	7,364	7,364	7,364	19,785	19,785	19,785	4,973	4,973	4,973
Cells	2,180	2,180	2,180	2,499	2,499	2,499	1,388	1,388	1,388
Clusters	596	596	596	553	553	553	319	319	319
R ²	0.41	0.41	0.41	0.58	0.58	0.57	0.42	0.42	0.42
F-stat		115.81***	124.82***		53.24***	44.63***		27.37***	21.43***

***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region-experience level. Observations refer either to the number of individuals, the number of region-year-experience cells or the number of clusters (region-experience cells) in the sample. F-stat refers to the first-stage F-statistic on the significance of the instrument in the first-stage regression of the endogenous cohort-size variable.

Supplementary material

The first part of this section provides further information on how the endogenous cohort-size variable, which is defined in terms of experience, is matched with the age-based instrumental variable. In the second part, we present the results of a variety of sensitivity analyses as evidence for the robustness of our findings.

S1 Experience, age and imputed age

In our empirical model, cohort size is constructed on the basis of experience, whereas the variable which the former is instrumented with is age-specific. This requires us to specify a relation between age and experience. We do this by imputing an age variable that is constant for all individuals within an experience group. This variable is defined as the sum of years of experience and the education-specific age at which entry into the labour market is assumed to take place. Specifically, for each of the three educational groupings imputed age is defined as follows:

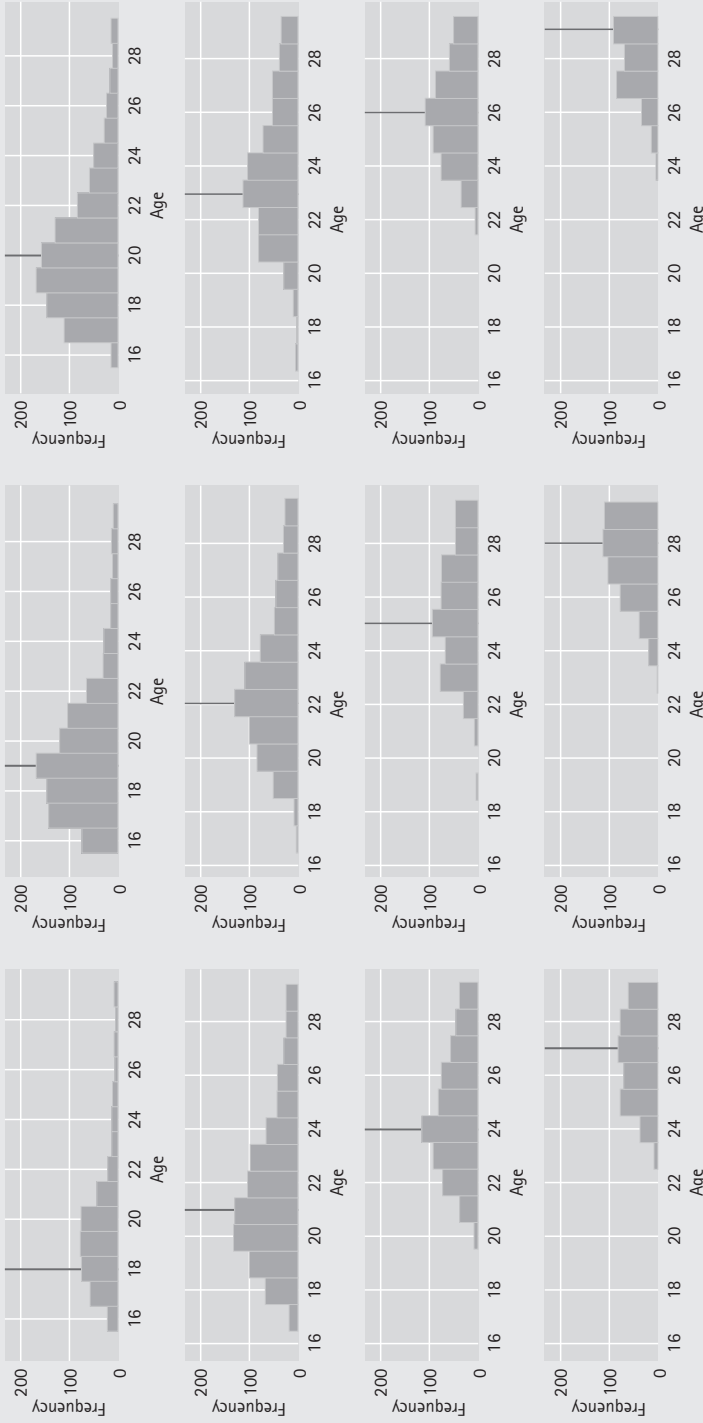
ISCED 0–2: $\text{age_imputed} = 18 + \text{years of experience}$ (range: 18–29)

ISCED 3–4: $\text{age_imputed} = 20 + \text{years of experience}$ (range: 20–29)

ISCED 5: $\text{age_imputed} = 24 + \text{years of experience}$ (range: 24–29)

The distribution of actual age for given values of imputed age is shown below for each educational group (Figures S1–S3). The graphs illustrate that actual age is indeed centred on the corresponding value of imputed age in the vast majority of cases. For higher values imputed age does not appear to lie in the centre of the actual age distribution. This, however, results from individuals whose actual age is above 29 being excluded from the sample.

Figure S1: Age and imputed age (ISCED 0-2)



Notes: The dark line indicates the value of imputed age from 18 (top-left) to 29 (bottom-right)
 Source: EU-SILC (authors' calculations).




Figure S3: Age and imputed age (ISCED 5)




Notes: The dark line indicates the value of imputed age from 24 (top-left) to 29 (bottom-right)
 Source: EU-SILC (authors' calculations).

S2 Sensitivity analysis

This section starts by illustrating the effect on the 2SLS cohort-size coefficients (corresponding to specification IV2) of excluding individual regions, years or experience groups from the sample. As can be seen from Figures S4–S8, the results are typically very close to the coefficient of the full model and always lie within the former's 95% confidence interval.

Figure S4: Excluding individual regions from the sample (ISCED 0–2)



Source: EU-SILC (authors' calculations).

Figure S5: Excluding individual regions from the sample (ISCED 3–4)

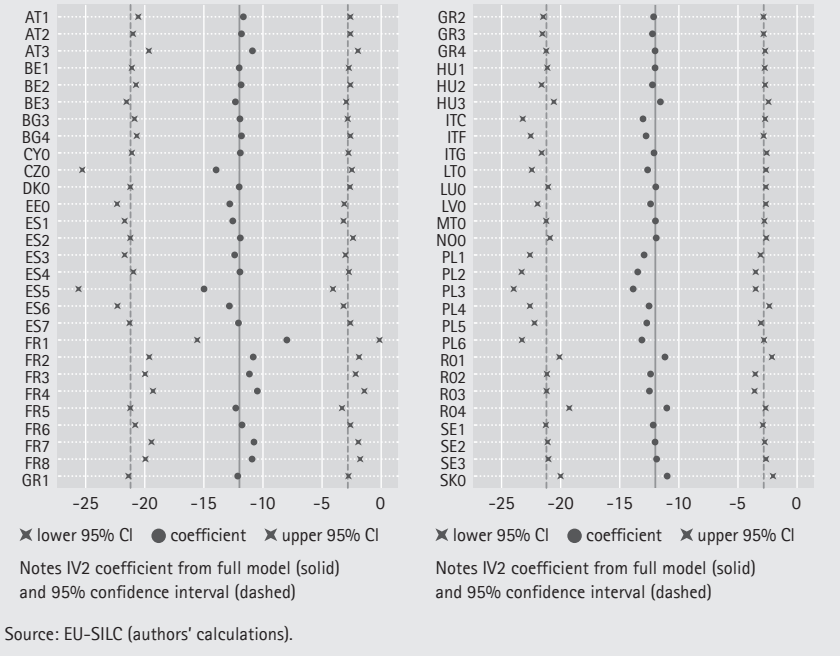


Figure S6: Excluding individual regions from the sample (ISCED 5)




Figure S7: Excluding individual years from the sample

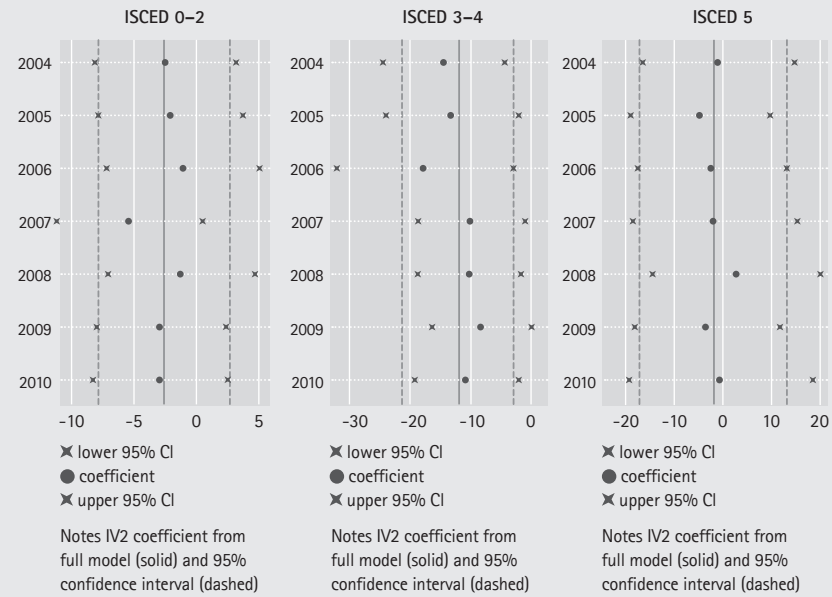
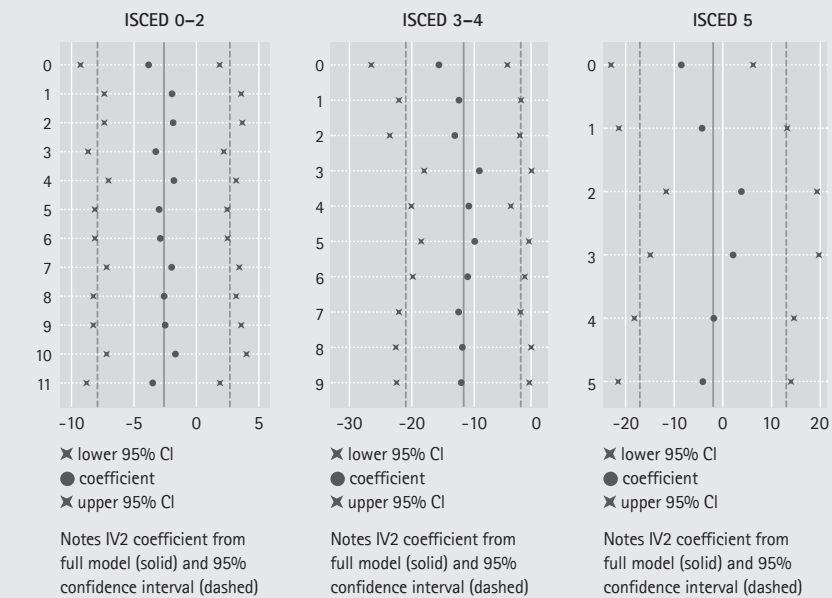


Figure S8: Excluding individual experience groups from the sample



Next, we show the results of a number of changes in the specification of the model. Table S1 shows the results that are obtained when a double-log specification is estimated: cohort size continues to have a significant negative effect for educational group ISCED 3–4 with an elasticity of approximately -0.3.

Table S1: Double-log specification

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size (log)	-0.03 (0.03)	-0.01 (0.08)	-0.04 (0.07)	-0.01 (0.03)	-0.10 (0.09)	-0.29** (0.12)	-0.07 (0.06)	-0.10 (0.23)	-0.05 (0.25)
N(inds)	7,364	7,364	7,364	19,785	19,785	19,785	4,973	4,973	4,973
N(cells)	2,180	2,180	2,180	2,499	2,499	2,499	1,338	1,338	1,338
N(clusters)	596	596	596	553	553	553	319	319	319
R ²	0.41	0.41	0.41	0.58	0.58	0.58	0.42	0.42	0.42

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region-experience level. N(inds): number of individual observations. N(cells): number of region-experience-year cells. N(clusters) indicates the number of clusters.

The use of experience dummies rather than its first two polynomials yields very similar results to those shown in the paper (cf. Table S2).

Table S2: Experience dummies

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-1.46 (1.12)	-0.95 (2.97)	-2.42 (2.49)	-0.30 (0.96)	-4.33 (3.47)	-12.45*** (4.60)	-1.91 (1.54)	-3.68 (7.43)	-1.91 (7.71)
ME (std)	-1.75%	-1.14%	-2.91%	-0.25%	-3.63%	-10.43%***	-2.33%	-4.49%	-2.33%
N(inds)	7,364	7,364	7,364	19,785	19,785	19,785	4,973	4,973	4,973
N(cells)	2,180	2,180	2,180	2,499	2,499	2,499	1,338	1,338	1,338
N(clusters)	596	596	596	553	553	553	319	319	319
R ²	0.42	0.42	0.42	0.58	0.58	0.57	0.42	0.42	0.42

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region-experience level. N(inds): number of individual observations. N(cells): number of region-experience-year cells. N(clusters) indicates the number of clusters. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

In the paper the individual sampling weights are adjusted so as to ensure that the estimated size of a region-year-age cell coincides with the values provided by Eurostat. Table S3 shows that when this is not done and the initial weights are used instead, cohort size retains its negative effect for educational group ISCED 3–4, but the coefficient as well as its marginal effect decrease in magnitude.

Table S3: Unadjusted weights

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-1.43 (1.22)	-0.78 (3.98)	-2.28 (2.95)	-0.39 (0.92)	-3.00 (3.30)	-8.64** (3.91)	-1.07 (1.50)	1.33 (6.91)	2.96 (6.88)
ME (std)	-1.63%	-0.89%	-2.60%	-0.31%	-2.40%	-6.88%**	-1.23%	1.53%	3.41%
N(inds)	7,364	7,364	7,364	19,785	19,785	19,785	4,973	4,973	4,973
N(cells)	2,180	2,180	2,180	2,499	2,499	2,499	1,338	1,338	1,338
N(clusters)	596	596	596	553	553	553	319	319	319
R ²	0.40	0.40	0.40	0.58	0.58	0.57	0.40	0.40	0.40

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region–experience level. N(inds): number of individual observations. N(cells): number of region–experience–year cells. N(clusters) indicates the number of clusters. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

In order to provide a better measure of the supply of individuals in an experience group to the labour market the cohort-size variable that is used in the paper takes into consideration only those individuals that are employed or unemployed. If instead all individuals are included in the construction of the cohort-size variable irrespective of their labour-market status, very similar results are obtained as shown in Table S4.

Table S4: Cohort-size variable based on all observations in a region–year–experience cell

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-0.18 (1.47)	-0.95 (3.05)	-2.33 (2.45)	-1.16 (0.92)	-4.41 (3.90)	-12.23*** (4.66)	-0.50 (1.47)	-3.24 (6.62)	-1.69 (6.67)
ME (std)	-0.20%	-1.05%	-2.56%	-1.02%	-3.85%	-10.70%***	-0.57%	-3.67%	-1.92%
N(inds)	7,364	7,364	7,364	19,785	19,785	19,785	4,973	4,973	4,973
N(cells)	2,180	2,180	2,180	2,499	2,499	2,499	1,338	1,338	1,338
N(clusters)	596	596	596	553	553	553	319	319	319
R ²	0.41	0.41	0.41	0.58	0.58	0.58	0.42	0.42	0.42

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region–experience level. N(inds): number of individual observations. N(cells): number of region–experience–year cells. N(clusters) indicates the number of clusters. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

The specific form of the cohort-size variable has already been used in the literature (Wright, 1991; Brunello, 2010) and reflects the idea that individuals are substitutable with those who are slightly more or less experienced than themselves, while the degree of substitutability is assumed to decrease with the difference in

experience. However, this assumption as well as the inclusion of individuals that differ by at most two years of experience can be argued to be arbitrary. We further assess this issue by specifying alternative cohort-size variables which differ in terms of the number of adjacent experience groups as well as in terms of the use of weights. Specifically, we compute a weighted sum across three experience groups (Equation S1) as well as sums of one, three and five experience groups (Equations S2–S4). In the latter case individuals in the included experience groups are assumed to be perfectly substitutable.

$$CS_{-2}_{jert} = \frac{(1/4)N_{j-1,ert} + (2/4)N_{jert} + (1/4)N_{j+1,ert}}{N_{ert}} \quad [S1]$$

$$CS_{-3}_{jert} = \frac{N_{jert}}{N_{ert}} \quad [S2]$$

$$CS_{-4}_{jert} = \frac{N_{j-1,ert} + N_{jert} + N_{j+1,ert}}{N_{ert}} \quad [S3]$$

$$CS_{-5}_{jert} = \frac{N_{j-2,ert} + N_{j-1,ert} + N_{jert} + N_{j+1,ert} + N_{j+2,ert}}{N_{ert}} \quad [S4]$$

Tables S5–S8 contain the corresponding regression results. In each case the coefficient remains negative and significant for ISCED 3–4, but since the means and standard deviations of these variables differ from those of the initial cohort-size variable, the magnitude of the effect can be better compared by looking at the marginal effects rather than at the coefficients. This effect is larger for the three-year weighted average and particularly when only the own experience group is used. The results of the latter specification especially should be treated with caution: the size of region-year-experience-education groups is estimated from EU-SILC data and its accuracy clearly depends on the sampling. In small cells minor changes in the number of observations can have a profound effect on the estimated cohort-size variable. Comparable results to those presented in the paper are found for the three-year sum and a smaller effect for the five-year sum.

Table S5: Cohort-size variable (weighted sum of three experience groups)

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-2.08** (0.85)	-0.61 (3.06)	-1.92 (2.43)	-0.18 (0.83)	-4.94 (4.25)	-15.97** (6.81)	-1.22 (1.28)	-3.04 (6.42)	-1.82 (6.66)
ME (std)	-2.84%**	-0.83%	-2.62%	-0.17%	-4.61%	-14.91%**	-1.69%	-4.21%	-2.53%
N(inds)	7,364	7,364	7,364	19,785	19,785	19,785	4,973	4,973	4,973
N(cells)	2,180	2,180	2,180	2,499	2,499	2,499	1,338	1,338	1,338
N(clusters)	596	596	596	553	553	553	319	319	319
R ²	0.42	0.41	0.42	0.58	0.58	0.56	0.42	0.42	0.42

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region-experience level. N(inds): number of individual observations. N(cells): number of region-experience-year cells. N(clusters) indicates the number of clusters. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

Table S6: Cohort-size variable (unweighted sum of one experience group)

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-1.36** (0.53)	-0.23 (2.66)	-1.36 (2.12)	0.47 (0.62)	-7.68 (5.99)	-22.85* (13.48)	-0.26 (0.73)	-1.06 (4.26)	-0.68 (4.51)
ME (std)	-2.76%**	-0.48%	-2.77%	0.59%	-9.60%	-28.55%*	-0.50%	-2.06%	-1.33%
N(inds)	7,364	7,364	7,364	19,785	19,785	19,785	4,973	4,973	4,973
N(cells)	2,180	2,180	2,180	2,499	2,499	2,499	1,338	1,338	1,338
N(clusters)	596	596	596	553	553	553	319	319	319
R ²	0.42	0.41	0.42	0.58	0.57	0.45	0.42	0.42	0.42

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region-experience level. N(inds): number of individual observations. N(cells): number of region-experience-year cells. N(clusters) indicates the number of clusters. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

Table S7: Cohort-size variable (unweighted sum of three experience groups)

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-0.38 (0.32)	-0.24 (1.02)	-0.70 (0.84)	-0.22 (0.28)	-1.41 (1.30)	-4.73* (1.94)	-0.57 (0.46)	-1.31 (2.36)	-0.80 (2.49)
ME (std)	-1.52%	-0.95%	-2.75%	-0.59%	-3.79%	-12.69%*	-2.23%	-5.16%	-3.14%
N(inds)	7,364	7,364	7,364	19,785	19,785	19,785	4,973	4,973	4,973
N(cells)	2,180	2,180	2,180	2,499	2,499	2,499	1,338	1,338	1,338
N(clusters)	596	596	596	553	553	553	319	319	319
R ²	0.41	0.41	0.41	0.58	0.58	0.57	0.42	0.42	0.42

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region-experience level. N(inds): number of individual observations. N(cells): number of region-experience-year cells. N(clusters) indicates the number of clusters. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

Table S8: Cohort-size variable (unweighted sum of five experience groups)

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-0.16 (0.25)	-0.26 (0.69)	-0.63 (0.58)	-0.04 (0.19)	-0.67 (0.63)	-2.04** (0.78)	-0.42 (0.31)	-0.79 (1.56)	-0.37 (1.58)
ME (std)	-0.89%	-1.42%	-3.47%	-0.16%	-2.70%	-8.21%**	-2.46%	-4.59%	-2.15%
N(inds)	7,364	7,364	7,364	19,785	19,785	19,785	4,973	4,973	4,973
N(cells)	2,180	2,180	2,180	2,499	2,499	2,499	1,338	1,338	1,338
N(clusters)	596	596	596	553	553	553	319	319	319
R ²	0.41	0.41	0.41	0.58	0.58	0.58	0.42	0.42	0.42

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region–experience level. N(inds): number of individual observations. N(cells): number of region–experience–year cells. N(clusters) indicates the number of clusters. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

The effect of using an age-specific rather than an experience-based cohort-size variable is displayed in Table S9. The marginal effect of the cohort-size variable is smaller in size and less significant for ISCED 3–4.

Table S9: Age-specific cohort-size variable

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-0.76 (1.38)	-1.93 (4.80)	-2.99 (3.79)	-0.60 (1.92)	-1.03 (3.81)	-7.77* (4.36)	3.70* (1.91)	-6.72 (11.57)	-1.54 (10.75)
ME (std)	-0.67%	-1.72%	-2.67%	-0.41%	-0.71%	-5.36%*	4.80%*	-8.72%	-2.01%
N(inds)	7,364	7,364	7,364	19,785	19,785	19,785	4,973	4,973	4,973
N(cells)	2,372	2,372	2,372	2,855	2,855	2,855	1,652	1,652	1,652
N(clusters)	671	671	671	668	668	668	457	457	457
R ²	0.42	0.42	0.42	0.59	0.59	0.59	0.42	0.41	0.41

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region–experience level. N(inds): number of individual observations. N(cells): number of region–age–year cells. N(clusters) indicates the number of clusters. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

We also estimate the model after averaging individual variables within region-year–experience cells. As discussed in Angrist and Pischke (2009) we weight the regression by the implied number of observations per cell and obtain results that are very similar to those in the paper (cf. Table S10).

Table S10: Averaging individual data within region–year–experience cells

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-1.20 (1.10)	-0.60 (3.37)	-2.49 (2.70)	-0.36 (1.01)	-4.95 (3.74)	-12.74*** (4.49)	-2.09 (1.63)	-3.91 (6.97)	-2.04 (7.46)
ME (std)	-1.45%	-0.72%	-3.00%	-0.30%	-4.14%	-10.67%***	-2.56%	-4.77%	-2.48%
N(cells)	2,180	2,180	2,180	2,499	2,499	2,499	1,338	1,338	1,338
R ²	0.59	0.59	0.59	0.85	0.85	0.83	0.61	0.61	0.61

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are used. N(cells): number of region–experience–year cells. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

In addition, to making adjustments to the specification of the model and its variables, we assess the effect of changes in the sample. In Table S11 we present the effect of extending the age range up to the year 32: while the coefficient remains negative and significant for ISCED 3–4, its marginal effect decreases in magnitude.

Table S11: Increasing the upper age limit to 32

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-1.21 (1.04)	0.02 (3.00)	-2.61 (2.58)	0.18 (0.84)	-2.01 (3.08)	-9.64** (4.28)	-1.55 (1.68)	-5.01 (6.97)	-2.26 (7.14)
ME (std)	-1.41%	0.03%	-3.03%	0.15%	-1.70%	-8.14%**	-1.86%	-6.04%	-2.73%
N(inds)	8,450	8,450	8,450	21,992	21,992	21,992	5,994	5,994	5,994
N(cells)	2,359	2,359	2,359	2,557	2,557	2,557	1,414	1,414	1,414
N(clusters)	624	624	624	553	553	553	323	323	323
R ²	0.43	0.43	0.43	0.59	0.59	0.59	0.43	0.42	0.43

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region–experience level. N(inds): number of individual observations. N(cells): number of region–experience–year cells. N(clusters) indicates the number of clusters. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

To ensure that the results are not driven by wage outliers, we remove observations from the top and bottom 0.5% and 1% of the wage distribution (Tables S12 and S13). In both cases the results are comparable to those of the paper.

Table S12: Excluding top and bottom 0.5% of the wage distribution

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-1.10 (1.11)	0.54 (2.97)	-1.43 (2.65)	0.10 (0.92)	-3.17 (3.23)	-11.38*** (4.43)	-1.60 (1.41)	-2.87 (6.62)	-1.59 (6.68)
ME (std)	-1.31%	0.64%	-1.72%	0.09%	-2.66%	-9.54%***	-1.95%	-3.50%	-1.94%
N(inds)	7,292	7,292	7,292	19,589	19,589	19,589	4,925	4,925	4,925
N(cells)	2,166	2,166	2,166	2,494	2,494	2,494	1,332	1,332	1,332
N(clusters)	594	594	594	552	552	552	319	319	319
R ²	0.44	0.44	0.44	0.62	0.62	0.61	0.47	0.47	0.47

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region–experience level. N(inds): number of individual observations. N(cells): number of region–experience–year cells. N(clusters) indicates the number of clusters. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

Table S13: Excluding top and bottom 1% of the wage distribution

	ISCED 0–2			ISCED 3–4			ISCED 5		
	OLS	IV1	IV2	OLS	IV1	IV2	OLS	IV1	IV2
Cohort size	-0.95 (1.04)	0.15 (2.79)	-1.55 (2.54)	0.01 (0.94)	-3.14 (3.27)	-10.63** (4.31)	-1.63 (1.40)	-2.48 (6.64)	-1.23 (6.61)
ME (std)	-1.14%	0.18%	-1.85%	0.01%	-2.64%	-8.92%**	-1.98%	-3.02%	-1.50%
N(inds)	7,219	7,219	7,219	19,391	19,391	19,391	4,875	4,875	4,875
N(cells)	2,157	2,157	2,157	2,493	2,493	2,493	1,329	1,329	1,329
N(clusters)	594	594	594	552	552	552	319	319	319
R ²	0.45	0.45	0.45	0.62	0.62	0.61	0.48	0.48	0.48

Control variables are included. ***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors are clustered at the region–experience level. N(inds): number of individual observations. N(cells): number of region–experience–year cells. N(clusters) indicates the number of clusters. ME(std) shows the percentage change in the hourly wage given a change in cohort size by one standard deviation.

References

- Angrist, J. and Pischke, J.-S. (2009) *Mostly harmless econometrics: an empiricist's companion*, Princeton University Press, Princeton.
- Brunello, G. (2010) The effects of cohort size on European earnings, *Journal of Population Economics*, 23, 273–90.
- Wright, R.E. (1991) Cohort size and earnings in Great Britain, *Journal of Population Economics*, 4, 295–305.

Alfred Garloff^a and Duncan Roth^b

Regional population structure and young workers' wages

Abstract

This paper estimates the effect that changes in the size of the youth population have on the wages of young workers. Assuming that differently aged workers are only imperfectly substitutable, economic theory predicts that individuals in larger age groups earn lower wages. We test this hypothesis for a sample of young, male, full-time employees in Western Germany during the period 1999–2010. In contrast to other studies, functional rather than administrative spatial entities are used as they provide a more accurate measure of the youth population in an actual labour market. Based on instrumental variables estimation, we show that an increase in the youth share by one percentage point is predicted to decrease a young worker's wages by 3%. Our results also suggest that a substantial part of this effect is due to members of larger age groups being more likely to be employed in lower-paying occupations.

JEL classification: J21, J31, R23

Keywords: Population structure, wages, youth share, labour-market regions, instrumental variables, occupational selection

1 Introduction

Germany is in the middle of a demographic transition. The size of its population was on the decline between 2003 – when positive net immigration started falling short of the natural population decrease – and 2010 and is projected to continue shrinking over the coming decades, falling by 11% between 2010 and 2040 (Statistisches Bundesamt, 2009).¹ However, this transition also has a second dimension: during the second half of the twentieth century fertility rates declined

a Institute for Employment Research (IAB), Frankfurt am Main, Germany.

b Institute for Employment Research (IAB), Düsseldorf, Germany.

1 To ensure comparability with the empirical analysis of this paper, the reported numbers refer to Western Germany (excluding West Berlin). With the availability of the 2011 census, the basis for estimating population variables has changed. As the population measures in this paper are based on pre-census data, we also use the population projections that are derived from this data rather than the recently released projections that make use of the 2011 census.

permanently and eventually fell below replacement level. Coupled with increases in life expectancy, these processes are having a substantial effect on the age structure of Germany's population as evidenced by the ongoing increases in the size of older age groups at the expense of younger ones.

Between 1990 and 2010 the ratio of the working-age to the total population fell by over three percentage points, a downward trend that is expected to be exacerbated by the entry into retirement of the large post-World War II birth cohorts. Moreover, demographic change has affected the age composition of the working-age population: while the share of individuals aged 15–24 in the working-age population increased between 2000 and 2010, this development is expected to reverse in the near future with the youth share projected to fall by 2.5 percentage points between 2010 and 2025. The implications of these changes – the combination of a shrinking and ageing population – for the future standard of living constitutes a widely discussed area of research (see Börsch-Supan, 2013). In this context, the question of how labour productivity will be affected by the changes in the population-age structure will be of prime importance (see Bloom and Sousa-Poza, 2013). Likewise, the sustainability of health care and public pension systems in light of demographic pressure has received considerable attention (see Arnds and Bonin, 2002; Jimeno et al., 2008).

The objective of this paper is to empirically analyse the impact of changes in the size of the youth population within regional labour markets on the wages of young workers. In the light of the projected population developments, this type of analysis is relevant as it provides a basis for evaluating how demographic processes can be expected to affect the wages of future cohorts of young workers. Given its focus, this paper belongs to a larger body of literature that analyses the effects of changes in the age structure on labour-market outcomes. In addition to wage adjustments, a considerable amount of research has addressed the impact on age-specific (un-) employment (Zimmermann, 1991; Shimer, 2001; Skans, 2005; Biagi and Lucafora, 2008; Ochsén, 2009; Garloff et al., 2013; Moffat and Roth, 2016b) and educational attainment (Connelly, 1986; Stapleton and Young, 1988; Fertig et al., 2009).

While wage differences and wage trends between different cohorts in Western Germany are documented in Fitzenberger (1999), his analysis does not focus on the consequences of changes in the age structure, which is the concern of this paper. In a world with a single type of labour input, an increase in the size of the labour force will lead to an outward shift of the labour supply curve. If the labour market works in a way that the wage rate adjusts so as to equate the demand for and the supply of labour and diminishing marginal productivity implies a downward-sloping labour demand curve, the effect of an increase in the labour force will be a lower equilibrium wage rate. If instead labour inputs are not homogenous but rather only

imperfectly substitutable across age groups, the effects of a change in age-specific labour supply will – depending on the degree of substitutability – be concentrated on the members of that age group. Within such a framework, an increase in the share of young individuals should be accompanied by a decrease in their wages.

Our contribution is threefold. First, our assessment of the relationship between the youth share and young workers' wages in Western Germany addresses the lack of recent empirical evidence on this topic. Second, we use functional entities in order to identify the size of the youth population within an actual labour market rather than within an administrative unit as is done by earlier studies, which reduces the potential for measurement error in this variable. Third, we assess the channels through which changes in the size of the youth population affect young workers' wages by controlling for industrial and occupational up- or downgrading. Gertler and Trigari (2009) argue that individuals have a better chance of moving into higher-paying industries, firms or jobs during boom periods than during recessions. We propose that a similar argument can be made with respect to age-group size, as increased competition may lead individuals to take up positions in lower-paying industries or occupations than they would have done, had they been part of a smaller age group. In order to distinguish between the direct and the selection-related, indirect effect of belonging to a larger age group, we compare the estimated wage effect of the youth share from models that exclude or include detailed information about an individual's industrial and occupational affiliation.

In our model the effect that the regional youth share has on the wages of young workers is identified solely through the within-variation of this variable. However, as the relative size of the youth population within a labour market is potentially endogenous due to migration into high-wage areas, an instrumental variables (IV) identification strategy is employed: within a given region the instrument is defined as the share of individuals that are fifteen years younger and that are observed fifteen years earlier than the age group of the endogenous regressor. We find that the youth share has a statistically significant negative effect on the wages of young workers. Specifically, an increase by one percentage point is predicted to decrease wages by 3% in our baseline model. When using a district-based measure of the youth-share variable, the estimated coefficients are smaller by between 13% and 48%. Finally, we find that controlling for an individual's industry and, particularly, occupation reduces the estimated wage decrease from 3% to 2%, which suggests that a substantial part of the negative effect of age-group size is the result of individuals in larger age groups being more likely to be employed in lower-paying occupations. According to these results, future generations of young workers can expect to benefit from demographic developments. Specifically, a decrease in the youth share by 2.5 percentage points,

as projected to occur between 2010 and 2025, would be predicted to lead to an increase in young workers' wages of about 5%, *ceteris paribus*.

The remainder of the paper is structured as follows. Section 2 addresses the relationship between age structure and wage outcomes and reviews the relevant theoretical and empirical literature. Section 3 provides descriptive statistics on the youth population in Germany. The empirical analysis is the topic of Section 4, while Section 5 discusses the regression results. Section 6 presents the conclusion.

2 Population structure and wages

Differently aged workers are not perfectly substitutable. Age can be expected to be correlated with a worker's set of skills, which in turn affects his suitability for different tasks. First, age is a good predictor for work experience, and, *ceteris paribus*, more experienced workers will usually have more firm-specific, occupation-specific, industry-specific or general human capital. If this type of knowledge is relevant for on-the-job performance, differently aged workers can be expected to be only imperfectly substitutable. Indeed, Welch's (1979) career-phase model can be interpreted as an example of a model in which imperfect substitutability arises from differences in firm-specific human capital. Second, jobs vary with respect to the tasks that they contain and therefore also concerning the abilities that workers are required to have in order to perform these tasks. Older workers may be less easily substitutable for younger workers in occupations requiring physical or certain types of cognitive skills (Mazzonna and Perracchi, 2012). As a consequence of imperfect substitutability a change in the relative size of an age group will mainly affect the labour market outcomes of the members of that group.

As a starting point to analysing the effects of a change in the size of a specific age group on the wages of its members, it is useful to assume a production function with differently aged workers as distinct factors of production (see Card and Lemieux, 2001; Fitzenberger and Kohn, 2006). In the benchmark case of a perfectly competitive labour market, in which each factor of production is paid the monetary value of his marginal product, a change in the supply of a specific production factor will cause the wage to adjust in a way that the market is again cleared. In the case of each factor of production exhibiting diminishing marginal productivity, an increase in the size of an age group will reduce the wages paid to its members. Labour markets, however, do not necessarily clear. The existence of minimum or efficiency wages as well as collective wage bargaining are possible sources that can prevent the wage rate from fully adjusting in response to a change in labour supply, while the coexistence of unemployment and vacancies provides evidence against the existence of a market-clearing equilibrium as predicted by the benchmark model

of a competitive labour market. Existing theoretical models, however, suggest that even in the absence of clearing labour markets, changes in the relative supply of an age group will have an effect on age-specific wages (Michaelis and Debus, 2011).

The extant empirical literature, though differing with respect to the time periods and countries (or regions) under study, the model specification and identification strategy, provides evidence that increases in the size of an age group are associated with depressed wage outcomes for the members of that group.² Early studies using US data estimate a negative relationship between the relative size of an age group and the average wages that are earned by individuals within that group for different levels of educational qualification (Welch, 1979; Berger, 1985). Alternatively, Freeman (1979) finds a negative effect of the young-to-old population ratio on the average wages of young workers relative to those of old workers. The existence of a negative effect of age-group size is also supported by evidence from Sapozhnikov and Triest (2007). Most recently, Morin (2015) exploits an exogenous shock to the supply of high-school graduates in Canada due to a reform of the secondary schooling system and finds negative cohort-size effects on wages. Empirical evidence from Europe is scarcer but also supports the hypothesis that wages earned in larger age groups are depressed compared to those of smaller age groups (see Wright, 1991, for the UK and Brunello, 2010, as well as Moffat and Roth, 2016a, for a sample of European countries).

A drawback with respect to identifying the effect of interest is that the size of an age group within a given spatial unit is arguably endogenous due to self-selection of individuals into high-wage areas. Korenman and Neumark (2000) proposed birth rates as an instrument, while other authors have since used the lagged relative size of age groups as exogenous predictors (Skans, 2005; Garloff et al., 2013; Moffat and Roth, 2016a and 2016b). Whereas cross-country migration might be deemed too small to influence the size of nationally defined age groups, endogeneity resulting from self-selection through migration becomes a larger concern when the spatial units that are used to construct the measure of population structure are defined at a sub-national level.

While many empirical studies in this field of research have used measures of population structure at the national level, it appears questionable whether a country indeed constitutes the appropriate delineation of a labour market. If individuals are restricted in their mobility or if awareness of job openings in other regions decreases with distance, a nationally defined youth-share variable groups together young individuals that are not active in the same labour market and that are hence not

2 Notable exceptions can be found in the migration literature where many studies conclude that natives' wages are not negatively affected by age-specific immigration (Ottaviano and Peri, 2012). A possible explanation for this finding is that migrants are complements rather than substitutes for native labour.

substitutable for one another. Such a variable would be subject to measurement error if labour markets existed at a sub-national level and the size of the youth population varied across them. And while more recent studies have made use of administrative units at a sub-national level, so-constructed youth-share variables may still be measured with error as administrative units are generally not delineated in a way as to coincide with actual labour markets, meaning that they would not necessarily capture the relative supply of young labour that is relevant for the determination of a young worker's wages. To address this issue, we employ the functional labour-market regions that are defined by Eckey et al. (2006). These regions consist of one or more districts (*Kreise*) and are constructed on the basis of observed commuting flows with a typical labour-market region combining an economic centre with the surrounding *Umland* from which people commute to work in the centre. They approximate self-contained local labour markets in as far as they aim to maximise the overlap between the population living and working within such a region. Functional units therefore provide a better measure of the size of the youth population in an actual labour market than administrative units. The self-contained nature of these units also reduces the need to consider the youth population in surrounding labour markets as a factor determining the wages of young workers in a given region.

It should be noted that changes in the age structure of the population do not necessarily imply changes in age-specific labour supply as participation rates as well as the number of hours worked could in principle adjust in a way as to completely counteract changes in age-group size. However, such a reaction seems unlikely as empirical evidence suggests that male labour supply is inelastic – at the extensive and the intensive margin – to changes in the wage rate (Blundell and MaCurdy, 1999). More specifically, Garloff et al. (2013) show that a counteracting development in participation rates has not taken place in Germany in response to changes in the age structure at the national level in recent years.

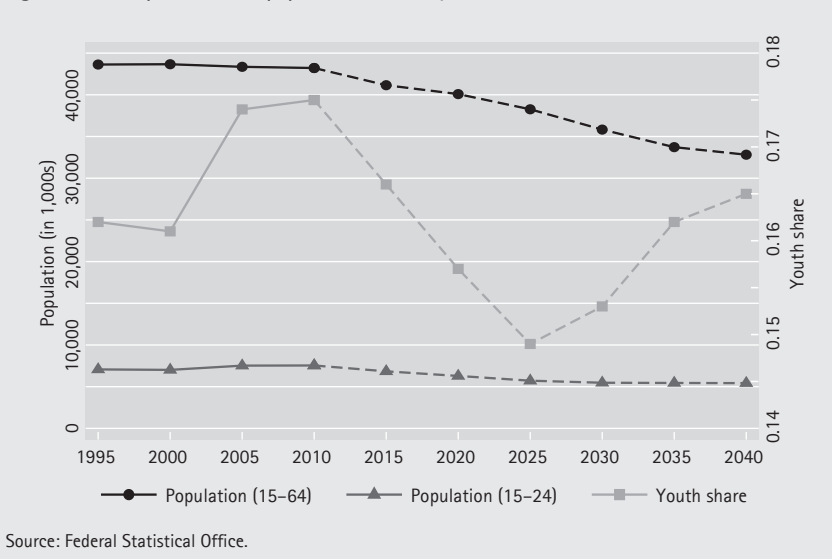
3 Youth-population structure in Western Germany

This section provides information about the development of the working-age (15–64) and the youth population (15–24) in Western Germany at the national level and at the level of the labour-market region. Figure 1 shows the absolute size of both populations at five-year intervals between 1995 and 2040. While the actual values are shown up to the year 2010³, subsequent developments represent projections based on the variant *Untergrenze der mittleren Bevölkerung*, which

3 Data comes from the Federal Statistical Office and has been obtained through the following link: [https://www.genesis.destatis.de/genesis/online/link/tabellen/12411*](https://www.genesis.destatis.de/genesis/online/link/tabellen/12411)

assumes an annual net immigration of 100,000 individuals and a fertility rate of 1.4 and which represents the lower bound of corridor within which population development is expected to take place (Statistisches Bundesamt, 2010).⁴

Figure 1: Development of the population and the youth share at the national level



Except for a small increase between 1995 and 2000, the working-age population has been shrinking steadily and is projected to continue decreasing in size over the coming decades. By 2040 it will have fallen by almost 25% compared to its 2010 value, which reflects the effect of the large post-World War II birth cohorts reaching retirement age. In contrast, the number of young individuals grew by half a million between the years 2000 and 2010⁵, but this development is expected to reverse in the near future with the size of the age group 15–24 projected to fall continuously until 2040. Reflecting changes in these two populations' relative rate of growth, the youth share, i.e. the size of the population aged 15–24 relative to the working-age population, displays a cyclical development: from 2000 to 2010 the share of young individuals expanded by approximately one percentage point (equivalently, 7%). However, as the youth population is expected to decrease at a faster rate than the working-age population, its share is projected to fall by

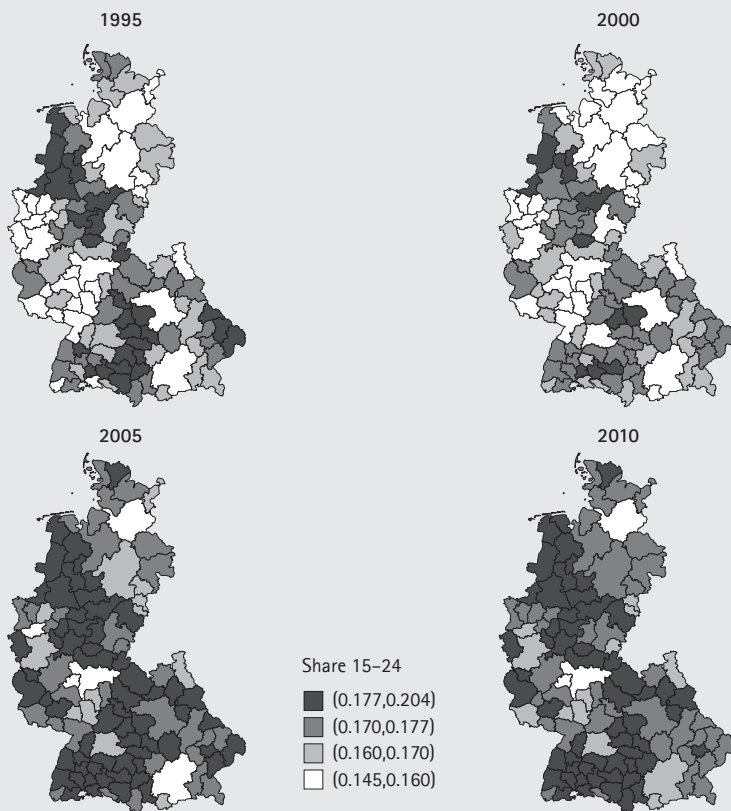
4 The upper bound of this corridor (*Obergrenze der mittleren Bevölkerung*) differs by assuming that annual net immigration will increase steadily to 200,000 in the year 2020 before plateauing at that level. Despite this difference the projection for the youth share is very similar (the largest difference between both projections amounts to 0.25 percentage points in the year 2040).

5 These age groups are the children of the large post-World War II birth cohorts. This increase therefore reflects the large size of the parental generation.

2.5 percentage points (equivalently, 15%) between 2010 and 2025. At the national level, the increase in the youth share during most of the sample period therefore contrasts with its projected development in the immediate future, which implies that changing demographics may contribute positively towards the development of young workers' wages in the coming years.

Figure 2 illustrates the existing regional heterogeneity in the share of individuals aged between 15 and 24 in the working-age population by reporting the value of this variable for the West-German regional labour markets. The extent of cross-sectional variation in the youth-share variable is revealed for the year 1995 in the top left map, in which the labour-market regions are grouped into quartiles based on the size of the youth share. Compared to a value of about 16% at the national level, the regional youth share varies between 14% and 21%.

Figure 2: Variation in the youth share (15–24) at the regional level



Source: Federal Statistical Office (population data) and Federal Institute for Research on Building, Urban Affairs and Spatial Development (geodata).

The other maps show the cross-sectional variation in the youth-share variable for the years 2000, 2005 and 2010, respectively. Moreover, they reveal the within-region variation in this variable, i.e. its development over time (to allow for a comparison of the different years, the same intervals are chosen as for the year 1995). Reflecting the drop in the national youth share in the year 2000, the share has also generally fallen at the regional level as illustrated by a number of regions that were in the fourth or third quartile in 1995 now being in the third or second quartile, respectively. Likewise, an increasing number of regions are registered in higher quartiles in the years 2005 and 2010, reflecting the increase in the youth share at the national level.

4 Empirical analysis

The different steps of empirically analysing the relationship between the youth share and young workers' wages are the subject of this section: the relevant datasets are introduced in the first part, which is followed by a description of how the sample is constructed and how the model's main variables are defined. The final part discusses the empirical model and the identification strategy.

4.1 Data

Three data sources are used for the empirical analysis. The first source is population data for Germany on the regional level according to age groups which is used to construct the relative size of the youth population within a regional labour market. The information reported by the statistical offices refers to the end of the year (31 December). There is no information beyond age and sex in these data. Particularly, there is no information available on the educational composition. Corrections have to be made to account for changes in the delineations of municipalities and districts which results in a dataset that is spatially consistent over time back until 1978. However, the available age brackets differ for the time before and after 1985. Second, we use statistics from the Federal Employment Agency (FEA) to gather information on employment numbers and rates as well as unemployment rates. Employment numbers and rates can be obtained at the level of the labour-market regions starting in 1987 for employment at place of work and from 1999 for employment at place of residence. The data is available by single-age cohorts, sex and education and refers to the middle of the year (30 June), since those values are typically close to yearly averages. In order to better compare the results from a model using a youth-share variable based on an individual's place of residence with those derived from an individual's place of employment, the year 1999 is chosen as the start of the sample period.

The final source is the *Stichprobe der Integrierten Erwerbsbiografien* (SIAB), a large micro-dataset from the Institute for Employment Research (IAB), that includes information on a 2% random sample of all individuals in Germany that were employed, unemployed or participating in measures of active labour-market policy between 1975 and 2010 (civil servants and the self-employed are excluded). For employed individuals in the dataset we have information on their employment relationship on a daily basis. Moreover, it contains a wealth of additional information that we use in part as control variables. The data further contains information about an employee's place of residence and place of employment, though the former only becomes available in 1999. A detailed description of the dataset can be found in vom Berge et al. (2013).

4.2 Sample and descriptive statistics

The observations contained in SIAB refer to spells of an individual (e.g. an employment spell) with given start and end dates as well as characteristics of the spell (e.g. the average daily wage earned during this period). We use the setting-up routines by Eberle et al. (2013) to transform the structure of the data so that it contains data from a single spell per individual and year. In doing so, we choose 15 June as the annual reference date, which means that only those spells are retained that include the reference date in a given year. As employers are required to report the wages of their employees once a year and this is typically done on 31 December, the longest spells run from 1 January to 31 December in a given year. Using 15 June as the reference date implies that spells starting and ending before (or after) 15 June within a given year are not being considered. This specific reference date is chosen because June values of employment figures are usually close to annual averages, while the middle of the month is used to avoid any end-of-calendar-month effects. However, the results are robust to using 31 December as the reference date.⁶

The sample covers the period 1999–2010 and consists of regularly employed (*sozialversicherungspflichtig Beschäftigte*) males who are between 15 and 24 years old. Individuals in vocational training are excluded because the mechanisms determining their remuneration are considered to be different from the rest of the labour market. As there is no information about the number of hours worked in the data, the sample is further restricted to full-time employees. While 95% of the observations have one full-time job, some observations hold other jobs in addition to being full-time employed, e.g. 3% of observations are also in minor

⁶ The results of this and all other robustness checks can be found in the Supplementary Material to this paper.

employment (*geringfügige Beschäftigung*). In such a case only information about the first full-time job is retained.⁷ We do not restrict employment spells to have a minimum duration. However, the results are robust to keeping only observations with employment spells of at least 90 days in the sample.

The model's dependent variable is an individual's inflation-adjusted daily wage including social security contributions and taxes.⁸ The reported wage is censored at the value of the corresponding year's upper social security threshold; but given that our sample is restricted to individuals aged between 15 and 24 only a small fraction of observations will have wages above the threshold, and since imputation procedures (see Gartner, 2005) suggest that in such a case the true wage values are close to the censoring value, we use the censored wage for these observations. At the other end of the spectrum, we also observe unrealistically low daily wages. To remove these observations we truncate the wage distribution at twice the value of the minor-employment threshold (*Geringfügigkeitsgrenze*) – an approach that has also been taken by other authors working with the same data source (e.g. Gürtzgen, 2016). This implies that observations with wages of less than 650 Euro per month (21.26 Euro per day or, alternatively, 2.57 Euro per hour, assuming an eight-hour working day) between 1999 and 2002 or less than 800 Euro per month (26.28 Euro per day or 3.29 Euro per hour) between 2003 and 2010 are dropped.⁹

The main variable of interest is the youth share, which measures the number of individuals aged between 15 and 24 relative to the number of working-age individuals (ages 15–64) within a regional labour market as defined by Eckey et al. (2006).¹⁰ Due to limitations pertaining to the availability of population data preceding re-unification, our empirical analysis is restricted to the 108 labour-market regions (313 districts) of Western Germany. This restriction is unfortunate: the demographic processes that have seen the youth share in Eastern Germany fall

-
- 7 For individuals holding more than one job at the same time it would in principle be possible to use total earnings from all jobs rather than just the wage earned in one job as the relevant dependent variable. We abstain from doing so as our focus is on how the supply of young workers affects the wages earned in a particular job. Similar results to those shown in Table 1 are obtained when observations with more than a full-time job are removed from the sample.
- 8 Inflation-adjustment is done using the consumer price index (base year: 2010). The data comes from the Federal Statistical Office and has been obtained through the following link: https://www.destatis.de/DE/ZahlenFakten/GesamtwirtschaftUmwelt/Preise/Verbraucherpreisindizes/Tabellen/_/VerbraucherpreiseKategorien.html?cms_gtp=145110_slot%253D2&thttps=1
- 9 If individuals with wages below the specified thresholds are not excluded from the analysis, the youth-share coefficients are smaller in size and less significant. Compared to the sample used in the empirical analysis of this paper, individuals below the threshold are more likely to have a lower secondary education without apprenticeship training (56% compared to 19%) and are employed in firms with on average a smaller number of employees (446 compared to 970), whereas the average size of the youth share is similar. In addition to measurement error in the wages, the decrease in the effect of the youth share might also be due to the wages of this group being less responsive to changes in the supply of young workers, possibly because they are downward-rigid due to institutional constraints (e.g. sector-specific minimum wages).
- 10 Similar results are obtained when we use an employment-based youth-share variable that is defined as the number of employed youths aged 15–24 relative to the workforce.

from 19% in 2004 to 14% in 2012 (Fuchs and Weyh, 2014) certainly warrant an analysis of the corresponding wage effects.

Using a sub-national variable allows us to identify the effect of the youth share on young workers' wages while also controlling for macroeconomic shocks at the national level in a flexible way. As discussed in Section 2, the main advantage of employing labour-market regions as opposed to administrative units is that they provide a more accurate measure of the size of the youth population in an actual labour market, thereby reducing the potential for measurement error. For comparative purposes, however, we also estimate a model using a youth-share measure that is based on districts, which represent administrative units at the third level of the Nomenclature of Territorial Units for Statistics (NUTS 3). Furthermore, we are able to define two versions of the youth-share variable that refer to either the relative size of the youth population within the labour-market region (or district) that an individual works in or within which he resides. Owing to the way in which labour-market regions are designed, the fraction of observations for which the region of residence and the region of employment are identical stands at 85%, whereas the value is considerably smaller in the case of districts (66%).

A range of control variables are included in the model. At the individual level, SIAB contains information on age and labour-market experience as well as on an employee's level of education and his nationality. At the firm level, we use the size of the establishment and, in an extension to the baseline model, we also include two-digit indicators for an individual's occupation and industry which allows us to address the issue of industrial and occupational up- and downgrading (see Gertler and Trigari, 2009). In order to control for local macroeconomic effects we use the region-specific (district-specific) unemployment rate and, as the corresponding youth-unemployment rate is not available, the share of unemployed young individuals in the population. Descriptive statistics of the variables included in the baseline model are shown in Table A1 in the Appendix.

The average log real daily wage earnings are equal to 4.28 (approximately 72.24 Euro). The share of individuals aged between 15 and 24 in the working-age population is 17%. Since only employed individuals are included in the sample and individuals in vocational training are not considered, over 95% of observations are 20 years or older and a similar share has acquired up to four years of work experience. In terms of educational qualification the sample is rather homogenous as more than nine out of ten observations have lower secondary education and about three quarters of the cases also have a completed apprenticeship. The average firm size is slightly below 1,000 employees, while the regional unemployment rate has a mean value of about 8%, which is slightly higher than the share of unemployed youths in the population.

4.3 Empirical model and identification

In order to estimate the relationship between the wages of young workers and their relative supply, we specify an enhanced Mincer equation (Mincer, 1958) and regress the natural logarithm of an individual's inflation-adjusted daily wage earnings w_{irt} on the youth share y_{rt} and a set of control variables x_{irt} as formulated in Equation 1.¹¹ The indexes i , r and t denote individuals, spatial units and years, respectively. As described in the previous sub-section, separate models are estimated in which the spatial unit refers either to an individual's place of residence or to the place of employment. The variables δ_r and μ_t represent dummies for the spatial unit an individual resides or is employed in and for the sample year, respectively. Due to the inclusion of the region dummies it is only the within-region variation from which the coefficient of the youth share is identified. The error term ε_{irt} captures stochastic shocks as well as the effects of all other variables that are not explicitly controlled for¹²:

$$\log(w_{irt}) = \alpha + \beta y_{rt} + x_{irt}'\gamma + \delta_r + \mu_t + \varepsilon_{irt} \quad [1]$$

Consistent estimation of the effect that the youth share has on the wages of young workers by pooled ordinary least squares (OLS) requires that the regressor y_{rt} be conditionally uncorrelated with the error term. We argue that this requirement is unlikely to hold because individuals are able to self-select into regions where they can expect to earn higher wages, *ceteris paribus*, thereby turning the youth share into an endogenous variable. This endogeneity can be viewed as being the result of either omitted variables or reverse causality. The underlying mechanism is shown in the following set of equations:

$$\log(w_{irt}) = \alpha^a + \beta^a y_{rt} + x_{irt}'\gamma^a + \delta_r^a + \mu_t^a + \psi_{rt}'\chi^a + \varepsilon_{irt}^a \quad [2a]$$

$$y_{rt} = \alpha^b + \beta^b w_{rt} + x_{rt}^b\gamma^b + \delta_r^b + \mu_t^b + \psi_{rt}'\chi^b + \varepsilon_{rt}^b \quad [2b]$$

First, there might be unobserved regional characteristics (e.g. regional industrial structure, regional labour-market conditions), ψ_{rt} , that jointly determine a young

11 The specification of Equation 1 can also be interpreted as a special case of the model provided by Card and Lemieux (2001) in as far as our analysis also assumes imperfect substitutability across age groups but considers only the age group 15–24 in the empirical analysis.

12 We abstain from estimating a model that includes fixed effects at the individual level. Since 44% of observations come from individuals that are included in the sample only once, estimation of such a model suffers from an insufficient degree of within-variation. Notice that for consistent estimation of the youth share's marginal effect, a fixed effects approach would only be required in the presence of unobserved, time-invariant heterogeneity at the individual level that is correlated with the youth share.

individual's wages (Equation 2a) as well as his decision to reside (work) in a specific region and hence the size of the youth share (Equation 2b). Assuming that individuals are likely to select into regions with characteristics that are favourable to their earnings ($\chi^a > 0$ and $\chi^b > 0$), pooled OLS estimates of the coefficient β in Equation 1 will be on average less negative than its true value (or even positive). The use of regional dummy variables, which capture unobserved time-invariant regional heterogeneity, and regional unemployment variables should help us to control for these characteristics. In an extension, we also fit a model with fixed effects for state-year combinations to further control for unobserved heterogeneity.

Second, even in the absence of omitted regional characteristics, endogeneity may yet arise from reverse causality. In Equation 2b the youth share is modelled as a function of the mean daily log-earnings of young workers in a given region, w_{rt} . As this variable is a linear function of the variable $\log(w_{irt})$, it follows that the youth share is correlated with the error term of Equation 1.¹³ If the size of the youth share depends positively on the mean earnings of young workers in that region ($\beta^b > 0$), the correlation between ε_{irt} and y_{rt} in Equation 1 will be positive. Under these assumptions and assuming further that the youth share has a negative effect on individual earnings ($\beta < 0$), pooled OLS estimates of the corresponding coefficient would be expected to be less negative compared to the true value (or even positive). What is therefore required to identify the true relationship between individual wages and the youth share is a source of exogenous variation in the latter variable.

To consistently estimate the causal effect that changes in the youth share have on the earnings of young workers we employ an IV strategy. Our instrument is the variable that has also been used by Skans (2005), Garloff et al. (2013) and Moffat and Roth (2016a, 2016b). This variable is defined as the relative size of the group of individuals who are 15 years younger than the age group on which the youth-share variable is based and who are observed 15 years earlier, i.e. we instrument the current share of those aged 15–24 (relative to the age group 15–64) with the share of those aged 0–9 (relative to the age group 0–49) 15 years earlier. The strength of the instrument derives from the fact that in the absence of migration and natural population changes the instrument and the youth-share variable would be based on the same group of individuals and both variables would actually be identical. We argue that migration and natural changes do not purge the association between the instrument and the endogenous regressor, meaning that if an age group in a given spatial unit was comparatively large (relative to the size of the same age group in other spatial units and years), the group of individuals in the same region

13 Specifically, $w_{rt} = \frac{1}{N} \sum_{i=1}^N \log(w_{irt})$.

who are 15 years older will still be relatively large in the present. This argument is supported by the results of the first-stage statistics, which show the instrument to have a high degree of explanatory power.

The identifying assumption is that individuals in the age group 0–9 do not choose where to reside based on the anticipation of their earnings 15 years in the future. If this condition is satisfied, the causal effect of the relative supply of young individuals in a given spatial unit on young workers' earnings can be identified by using the two-stage least squares (2SLS) estimator with the time-lagged and age-lagged population variable as an instrument. An argument that can be brought forward against the validity of the instrument is that the relative size of the age group 0–9 will depend on the locational choices of their parents. If parents, and thus their children, self-selected into high-wage areas and their wages were correlated with the wages of their children fifteen years in the future, this would lead the proposed identification strategy to fail. Notice, however, that if the parental generation's choice of location and the correlation between their own and their children's wages are due to time-invariant factors, these will be accounted for by the region dummies of Equation 1.

Another source of endogeneity due to omitted variables is the fact that we only observe daily but not hourly wages. If the number of hours worked varies systematically with the youth share, pooled OLS estimation will again produce inconsistent results, but as long as the supply of hours is uncorrelated with the proposed instrument, 2SLS estimation will be consistent. Finally, a feature of the model in Equation 1 is that the explanatory variable of interest, y_{rt} , is defined at a higher level of aggregation than the dependent variable, which also varies across individuals.¹⁴ To account for this feature we cluster at the level of the spatial unit in order to avoid biased standard errors (see Moulton, 1990).

5 Results

Table 1 shows the results of estimating the baseline specification of Equation 1 (i.e. excluding indicators for an individual's industry and occupation). In the first two columns labour-market regions refer to an individual's place of residence, while the results for the place of employment are shown in the third and the fourth column. In both cases, the model is estimated by OLS as well as by 2SLS.

In line with the prediction that, *ceteris paribus*, members of larger age groups earn lower wages, the youth-share variables draw negative and statistically

¹⁴ Comparable results are obtained when all variables are averaged across the individuals in a region-year cell and the regression is weighted by the number of observations per cell (see Angrist and Pischke, 2009).

significant coefficients when estimated by 2SLS. Measured at an individual's place of residence, a decrease in the youth share by one percentage point is predicted to increase a young worker's wages by 3.2%. The corresponding Figure for the place of employment is slightly smaller at 2.9%. The fact that these effects are similarly sized is not surprising given the way in which functional labour markets are constructed (see Section 2) and the large share of observations for which the region of residence and the region of employment are identical (see Section 4.2). We also calculate the marginal effect of a change in the youth share by one standard deviation, which is reported at the bottom of the table. The estimated effects are -4.1% at the place of residence and -3.7% at the place of employment. In terms of magnitude these changes are comparable to the average return to an additional year of experience during the first four years of a worker's career.

The first-stage coefficients of the instrument are positive and highly significant in both specifications. There is, however, no one-to-one relationship between the current and the lagged value of the youth share, which suggests that within the 15 years over which the instrument is lagged the size of the youth share is affected by natural population changes and migration; instead an increase in the instrument by 1 percentage point is associated with an increase in the current youth-size variable by 0.46 percentage points. The first-stage F-statistics, which measure the significance of the excluded instruments, are considerably larger than the rule-of-thumb value 10 (Staiger and Stock, 1997), and the instrument's explanatory power is further evidenced by the value of Shea's partial R^2 . Identification does therefore not appear to be hampered by the presence of weak instruments.¹⁵

The OLS point estimates, while still negative, are considerably smaller than their 2SLS counterparts and their values lie outside the formers' 90% confidence interval. This finding is in line with the discussion of Section 4.3: if the value of the youth share is influenced by individuals migrating into economically attractive regions, OLS estimation can be expected to produce coefficients that are on average less negative than the true value of the youth share's marginal effect. The coefficients of the control variables display a large degree of similarity across the four different specifications of Table 1. Wages are predicted to increase at a decreasing rate in age and experience – the latter being suggestive of the widely documented concave experience-earnings profile (Polachek, 2008).¹⁶ Higher levels of schooling and professional qualification are associated with higher earnings.

15 A number of studies find that the magnitude of cohort-size effects differs across educational groupings (e.g. Brunello, 2010). The results of Table 1 are not affected by excluding either those observations with tertiary education or all observations with either tertiary or upper secondary education.

16 We have also estimated Equation 1 using mutually exclusive sets of age and experience dummies. Changing the specification in this way has no effect on the estimated youth-share coefficients.

Table 1: Baseline model

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.46 (0.63)*	-3.22 (0.97)***	-0.85 (0.73)	-2.89 (1.22)*
Age	0.21 (0.02)***	0.21 (0.02)***	0.21 (0.02)***	0.21 (0.02)***
Age ²	-0.00 (0.00)***	-0.00 (0.00)***	-0.00 (0.00)***	-0.00 (0.00)***
Experience	0.05 (0.00)***	0.05 (0.00)***	0.05 (0.00)***	0.05 (0.01)***
Experience ²	-0.00 (0.00)***	-0.00 (0.00)***	-0.00 (0.00)***	-0.00 (0.00)***
<i>Education</i>				
Lower secondary (with apprenticeship)	0.22 (0.01)***	0.22 (0.01)***	0.22 (0.01)***	0.22 (0.01)***
Upper secondary (without apprenticeship)	0.05 (0.01)***	0.05 (0.01)**	0.04 (0.02)**	0.04 (0.02)**
Upper secondary (with apprenticeship)	0.30 (0.01)***	0.30 (0.01)***	0.30 (0.01)***	0.30 (0.01)***
Tertiary (University of Applied Sciences)	0.31 (0.02)***	0.31 (0.02)***	0.31 (0.02)***	0.31 (0.02)***
Tertiary (University)	0.46 (0.03)***	0.46 (0.03)***	0.46 (0.03)***	0.46 (0.03)***
<i>Nationality</i>				
Turkey	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)
Switzerland/Austria	-0.02 (0.04)	-0.01 (0.04)	-0.02 (0.04)	-0.02 (0.04)
Western Europe	-0.03 (0.03)	-0.03 (0.03)	-0.03 (0.03)	-0.03 (0.03)
Northern Europe	-0.07 (0.08)	-0.07 (0.08)	-0.07 (0.07)	-0.07 (0.07)
Central Europe	-0.03 (0.02) [†]	-0.03 (0.02) [†]	-0.03 (0.02)*	-0.03 (0.02)*
Eastern Europe	-0.08 (0.03)**	-0.08 (0.03)**	-0.08 (0.03)***	-0.08 (0.03)***
South-East Europe	-0.02 (0.01) [†]	-0.02 (0.01)*	-0.02 (0.01) [†]	-0.02 (0.01) [†]
Southern Europe	-0.07 (0.01)***	-0.07 (0.01)***	-0.07 (0.01)***	-0.07 (0.01)***
Africa	-0.11 (0.02)***	-0.11 (0.02)***	-0.11 (0.02)***	-0.11 (0.02)***
Asia	-0.12 (0.02)***	-0.12 (0.02)***	-0.12 (0.02)***	-0.12 (0.02)***
America/Oceania	0.08 (0.05)	0.08 (0.05)	0.08 (0.06)	0.08 (0.06)
Firm size (in 1,000s)	0.20 (0.02)***	0.20 (0.02)***	0.21 (0.03)***	0.21 (0.03)***
Unemployment rate	0.00 (0.00)	0.00 (0.00)	-0.00 (0.00)	0.00 (0.00)
Youth unemployment rate	0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)	-0.00 (0.00)
Constant	1.78 (0.28)***	2.06 (0.29)***	1.69 (0.29)***	2.02 (0.32)***
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market region	Yes	Yes	Yes	Yes
<i>First-stage regression</i>				
Instrument	-	0.46 (0.00)***	-	0.46 (0.00)***
<i>First-stage test statistics</i>				
F-statistic	-	136.60***	-	131.80***
Shea's partial R ²	-	0.32	-	0.32
<i>Observations</i>				
Individuals	107,351	107,352	107,351	107,351
Labour-market region-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.24	0.24	0.24	0.24
ME (stdev)	-1.84%*	-4.05%***	-1.09%	-3.67%*
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/*/ [†] indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

Nationals from Eastern European, South-East and Southern European countries are predicted to earn significantly less than Germans, while the largest difference is found for Africans and Asians with earnings lower by more than 10%. Individuals who are employed in firms with larger workforces are found to have higher

earnings, which is in line with evidence by Lehmer and Möller (2010). Finally, the effects of the unemployment rate and the share of young unemployed individuals are small. The youth-unemployment variable draws a negative coefficient in the 2SLS estimations, but in contrast to findings by Baltagi and Blien (1998) its effect is not statistically significant.

Related studies have used administrative units at the sub-national level as the basis for constructing population variables. As discussed in Section 2, the drawback of such an approach is that these units do not necessarily represent actual labour markets and that, consequently, the size of age groups within a given labour market is potentially measured with error (see the Supplementary Material for a discussion). We assess the effect of using administrative rather than functional units by estimating Equation 1 on the basis of a district-specific youth-share variable. The results are shown in Table 2.

Table 2: District-based youth share variable

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.31 (0.45)***	-2.79 (0.81)***	-0.12 (0.42)	-1.50 (0.92)
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
District	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage regression</i>				
Instrument	-	0.44 (0.00)***	-	0.43 (0.00)***
<i>First-stage test statistics</i>				
F-statistic	-	300.92***	-	181.95***
Shea's partial R ²	-	0.27	-	0.22
<i>Observations</i>				
Individuals	107,351	107,351	107,351	107,351
District-year cells	3,756	3,756	3,756	3,756
Districts (clusters)	313	313	313	313
R ²	0.25	0.25	0.26	0.26
ME(stdev)	-1.80%***	-3.84%***	-0.17%	-2.18%
Cluster-robust standard errors in parentheses (clustered at the district level). ***/**/* indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

The 2SLS coefficients of the youth-share variables remain negative and larger in absolute value than the corresponding OLS estimates, but only the specification referring to an individual's place of residence produces statistically significant results. However, compared to the results of Table 1, using districts rather than labour-market regions leads to an underestimation of the youth share's negative

effect: the point estimates referring to the place of residence are smaller by 13%, while the size of the coefficient for the place of employment drops by almost 50%.¹⁷ The increased discrepancy between the youth-share coefficients of these specifications reflects the fact that individuals are more likely to live and work in different districts than is the case for labour-market regions.

It can be shown that the negative and significant youth-share coefficients of Table 2 are driven by those districts in which individuals in the sample are more likely to live than to work. At the same time, the average absolute difference between the district-based youth-share variable and its value at the corresponding labour-market region is smaller for these districts, which suggests that measurement error in the size of the youth-share variable is less pronounced. The fact that these districts account for a larger fraction of observations in the place-of-residence specification suggests that size and significance of the youth-share coefficients will be less affected in that specification. It turns out that individuals in the sample are more likely to work in cities and to live in rural areas. The rationale behind the above argument could therefore be that the youth share within a city-district provides only an inaccurate measure of the size of the youth population that is relevant for the determination of wage outcomes as cities will also draw workers from surrounding districts.

As discussed in Section 1, the size of an individual's age group could have an effect on the conditions of his employment. Specifically, if young workers in larger age groups are more likely to be in positions in lower-paying occupations or industries, the estimated wage effect of the youth share in Table 1 would be confounded by these types of selection effects. In particular, the negative effect would be overestimated. To address this issue, we successively add indicator variables to the model of Equation 1 which are derived from two-digit codes referring to an individual's industry and occupation. The results are shown in Table 3 for the place-of-residence specification and in Table 4 for the place of employment.

In both cases we find that adding industry and, especially, occupation indicators has a sizeable impact on the estimated youth-share effects compared to the baseline specification: the inclusion of industry indicators decreases the size of the 2SLS coefficients by 12% (place of residence) and 4% (place of employment) compared to the results of the baseline model, while the reduction resulting from adding occupation indicators is considerably larger at 40% and 30%, respectively. Similar results are obtained when both sets of indicator variables are used. Moreover, it can be seen that when dummies for industry or occupation are added the difference in

¹⁷ Due to the higher variance of the district-based youth-share variable the proportional changes in the marginal effects for a change of one standard deviation are less pronounced.

the size of the 2SLS coefficients between the place of residence and the place of employment decrease in magnitude. This supports the argument that once labour-market regions are used as the spatial entities from which the youth-share variable is constructed both types of places produce similar results.

Table 3: Industry and occupation indicators (place of residence)

<i>Dependent variable: log real daily earnings</i>	Baseline	+industry	+occupation	+industry +occupation
Youth share (2SLS)	-3.22 (0.97)***	-2.81 (0.92)***	-1.88 (0.91)*	-1.89 (0.86)*
Youth share (OLS)	-1.46 (0.63)*	-1.36 (0.57)*	-0.90 (0.60)	-0.97 (0.53) [†]
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market region	Yes	Yes	Yes	Yes
Industry	No	Yes	No	Yes
Occupation	No	No	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage regression</i>				
Instrument	0.46 (0.00)***	0.46 (0.00)***	0.46 (0.00)***	0.46 (0.00)***
<i>First-stage test statistics</i>				
F-statistic	136.60***	137.17***	137.32***	137.70***
Shea's partial R ²	0.32	0.32	0.32	0.32
<i>Observations</i>				
Individuals	107,351	107,351	107,351	107,351
Labour-market region-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ² (2SLS)	0.24	0.46	0.40	0.51
R ² (OLS)	0.24	0.46	0.40	0.51
ME(stdev, 2SLS)	-4.05%***	-3.54%***	-2.36%*	-2.39%*
ME(stdev, OLS)	-1.84%*	-1.71%*	-1.14%	-1.22% [†]
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/*/† indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

While in the baseline model an increase in the size of the youth-share variable by one percentage point was predicted to decrease an individual's wages by about 3%, ceteris paribus, the size of this effect is reduced once an individual's industrial and, in particular, occupational affiliation are controlled for. This finding suggests that the estimated youth-share coefficients of the baseline specification were indeed confounded by the positive association between young workers being in larger age groups and being employed in lower-paying industries and occupations. We conclude that in addition to the direct negative effect of the size of the youth share, there is an indirect effect driven by selection into specific industries and occupations. A possible explanation for this finding is that, ceteris paribus, a larger supply of young individuals increases competition for higher-quality jobs, forcing some individuals to take up employment in lower-paying occupations.

Table 4: Industry and occupation indicators (place of employment)

<i>Dependent variable: log real daily earnings</i>	Baseline	+industry	+occupation	+industry +occupation
Youth share (2SLS)	-2.89 (1.22)*	-2.77 (1.06)**	-1.96 (1.08) ⁺	-2.11 (1.01)*
Youth share (OLS)	-0.85 (0.73)	-0.89 (0.62)	-0.52 (0.66)	-0.62 (0.57)
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market region	Yes	Yes	Yes	Yes
Industry	No	Yes	No	Yes
Occupation	No	No	Yes	Yes
Control variables	Yes	Yes	Yes	Yes
<i>First-stage regression</i>				
Instrument	0.46 (0.00)***	0.46 (0.00)***	0.46 (0.00)***	0.46 (0.00)***
<i>First-stage test statistics</i>				
F-statistic	131.80***	132.31***	132.33***	132.67***
Shea's partial R ²	0.32	0.32	0.32	0.32
<i>Observations</i>				
Individuals	107,351	107,351	107,351	107,351
Labour-market region-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ² (2SLS)	0.24	0.46	0.41	0.51
R ² (OLS)	0.24	0.46	0.41	0.51
ME(stdev, 2SLS)	-3.67%*	-3.52%**	-2.49% ⁺	-2.68%*
ME(stdev, OLS)	-1.09%	-1.13%	-0.66%	-0.79%
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/*/+ indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

This interpretation is in line with recent results pertaining to the wage effects of labour market conditions. Kahn (2010) and Brunner and Kuhn (2014) find that adverse labour market conditions (measured by the unemployment rate at the time of labour-market entry) depress wages and increase the probability of employment in lower-quality occupations. Morin (2015) studies the wage effects of the increase in labour supply due to the double cohort of high-school graduates in Ontario and provides evidence that part of the negative wage effect is due to selection into lower-paying occupations. Alternatively, higher-quality jobs may require a specific type of qualification. If the supply of training positions does not adjust to the supply of young individuals, the number of individuals barred from entering higher-paying occupations will increase in larger age groups. The effect of age-group size on selection into industries and occupations certainly warrants further research.

The Supplementary Material contains the corresponding output tables for the case in which districts provide the basis for the construction of the youth-share variable. These show that region-specific and district-specific variables continue to produce different results once industry and occupation dummies have been added and also illustrate that the difference between the results of the place-of-residence and the place-of-employment specifications are more pronounced at the district level.

To assess to what extent the results of Table 1 merely reflect unobserved heterogeneity at the federal state-year level, we add dummy variables for the interaction between federal states and years to the model of Equation 1. Doing so allows us to control for annual shocks that affect states differently and that are relevant for the determination of individual wages, e.g. the effects of macroeconomic shocks may vary between states due to differences in industrial structure. The results displayed in Table 5 suggest that, at least for the place-of-residence specification, the estimated effects of the youth-share variable in the baseline specification are not driven by unobserved heterogeneity at the state-year level.

Table 5: State-by-year interactions

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.01 (0.63)	-2.58 (1.27)*	-0.46 (0.74)	-2.35 (1.56)
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market region	Yes	Yes	Yes	Yes
Federal state-by-year	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage regression</i>				
Instrument	-	0.43 (0.00)***	-	0.43 (0.00)***
<i>First-stage test statistics</i>				
F-statistic	-	85.39***	-	84.74***
Shea's partial R ²	-	0.26	-	0.26
<i>Observations</i>				
Individuals	107,351	107,351	107,351	107,351
Labour-market region-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.24	0.24	0.24	0.24
ME(stdev)	-1.28%	-3.25%*	-0.58%	-2.99%
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/*/+ indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

The 2SLS point estimates fall by approximately 20% as part of the explained variation in the earnings variable is now picked up by the additional dummies. The standard errors increase presumably because parts of the variation in the youth-share variable are now explained by the additional dummy variables, which results in less precise estimates. For a similar reason there is a drop in the values of the first-stage F-statistic and the partial R² of the excluded instrument: the explanatory power of the instrument is reduced as a consequence of including the state-by-year dummies in the first-stage equation.

6 Conclusion

This paper empirically analyses how changes in the size of the youth population affect the wages of young workers. Under the assumption that differently aged individuals are only imperfectly substitutable because of differences in firm-specific, occupation-specific, industry-specific or general human capital, economic theory predicts that an increase in the size of an age group reduces the earnings of the members of that group. This hypothesis is tested using a sample of young male employees from Western Germany. The demographic forces that are currently changing the age-structure of the German population illustrate the relevance of this analysis. Specifically, the share of young individuals is projected to fall by 2.5 percentage points (equivalently, by 15%) at the national level over the period 2010–2025 following a period of an increasing youth share.

Besides providing an analysis of this relationship using recent data from administrative records, this paper makes two additional contributions. First, functional labour-market regions rather than administrative units are used as the spatial entities within which the size of the youth population is measured. These units provide a better measure of the number of young individuals in an actual labour market than administrative units, which are usually not delineated according to economic criteria, and hence of the supply of young labour that is relevant for the determination of young workers' wages. Use of a youth-share variable based on labour-market regions therefore reduces the potential for measurement error in this variable. Second, we address the channels through which an increase in the supply of young individuals affects their wages by controlling for industrial and occupational upgrading, i.e. for the possibility that changes in the size of the youth population affect the chances of finding employment in higher-paying industries or occupations.

The empirical analysis employs an IV approach in order to account for the possibility that the youth share is endogenous due to young individuals migrating into high-wage areas. In line with the hypothesis that increases in age-group size reduce the wages of the members of that group, the 2SLS coefficients are negative and significant: an increase in the youth share by one percentage point is predicted to decrease young workers' wages by 3%. Consistent with the argument that migration into high-wage regions induces endogeneity, the corresponding OLS estimates are less negative. Estimating our model using a youth-share variable that is based on districts rather than labour-market regions reduces the size of the 2SLS coefficients by either 13% (place of residence) or 48% (place of employment), which is consistent with the hypothesis that the use of administrative units induces measurement error in the youth-share variable. Finally, adding indicators for an individual's occupation and industry reduces the size of the youth-share

coefficients from -3% to -2%. We interpret this result as providing evidence for the hypothesis that belonging to a larger age group increases the likelihood of being employed in lower-paying occupations or industries.

What are the implications of these findings for the wages of the coming cohorts of young workers in light of Western Germany's changing demographics? As the youth share is projected to decrease over the coming years, demographic processes appear to be favourable to the development of the wages that young workers can expect in the future. But as the development of population structures is likely to differ between regions, regional variation in the extent to which young workers stand to benefit is to be expected. Finally, it should be borne in mind that these results come from a specific sample consisting of young, male, full-time employees with a few years of work experience and, predominantly, lower secondary education. Whether the relationship between the youth share and young workers' wages is similar for other groups, such as females or the highly educated, remains a topic for future research.

Acknowledgements

The authors would like to thank Stefan Fuchs, Bernd Hayo, John Moffat and Norbert Schanne for their advice and are grateful for comments from the participants of IAB's regional research network meeting in Aalen, the 54th Conference of the European Regional Science Association (ERSA), the 5th ifo Workshop Arbeitsmarkt und Sozialpolitik, the joint IAB Regional Science Academy workshop in Amsterdam and the 28th Conference of the European Association of Labour Economists (EALE). The Federal Institute for Research on Building, Urban Affairs and Spatial Development kindly provided the shapefile of the German labour-market regions. Our thanks also go to Annie Roth for proof-reading an earlier version of this paper.

References

- Angrist, J. and Pischke, J.-S. (2009) *Mostly harmless econometrics: an empiricist's companion*, Princeton University Press, Princeton.
- Arnds, P. and Bonin, H. (2002) Arbeitsmarkteffekte und finanzpolitische Folgen der demographischen Alterung in Deutschland, IZA Discussion Paper No. 667, Institute of Labor Economics, Bonn.
- Baltagi, B.H. and Blien, U. (1998) The German wage curve: evidence from the IAB employment sample, *Economics Letters*, 61, 135–42.

- Berger, M.C. (1985) The Effect of Cohort Size on Earnings Growth: A Reexamination of the Evidence, *Journal of Political Economy*, 93, 561–73.
- Biagi, F. and Lucifora, C. (2008) Demographic and education effects on unemployment in Europe, *Labour Economics*, 15, 1076–101.
- Bloom, D.E. and Sousa-Poza, A. (2013) Aging and productivity: Introduction, *Labour Economics*, 22, 1–4.
- Blundell, R. and MaCurdy, T. (1999) Labour Supply: A Review of Alternative Approaches, in O. Ashenfelter and D. Card (eds.) *Handbook of Labor Economics*, Vol. 3, North-Holland, Amsterdam.
- Börsch-Supan, A. (2013) Myths, Scientific Evidence and Economic Policy in an Aging World, *Journal of the Economics of Ageing*, 1–2, 3–15.
- Brunello, G. (2010) The effects of cohort size on European earnings, *Journal of Population Economics*, 23, 273–90.
- Brunner, B. and Kuhn, A. (2014) The impact of labor market entry conditions on initial job assignments and wages, *Journal of Population Economics*, 27, 705–38.
- Card, D. and Lemieux, T. (2001) Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis, *Quarterly Journal of Economics*, 116, 705–46.
- Connelly, R. (1986) A Framework for Analyzing the Impact of Cohort Size on Education and Labor Earning, *Journal of Human Resources*, 21, 543–62.
- Eberle, J., Schmucker, A. and Seth, S. (2013) Programmierbeispiele zur Datenaufbereitung der Stichprobe der Integrierten Arbeitsmarktbiografien (SIAB) in Stata. Generierung von Querschnittsdaten und biografischen Daten, FDZ Methodenreport No. 04/2013, Institute for Employment Research, Nuremberg.
- Eckey, H.-F., Kosfeld, R. and Türck, M. (2006) Abgrenzung deutscher Arbeitsmarktregionen, *Raumforschung und Raumordnung*, 64, 299–309.
- Fertig, M., Schmidt, C. and Sinning, M. (2009) The Impact of Demographic Change on Human Capital Accumulation, *Labour Economics*, 16, 659–68.
- Fitzenberger, B. (1999) *Wages and employment across skill groups: An analysis for West Germany*, Physica-Verlag: Heidelberg.
- Fitzenberger, B. and Kohn, K. (2006) Skill Wage Premia, Employment, and Cohort Effects: Are Workers in Germany All of the Same Type?, IZA Discussion Paper No. 2185, Institute of Labor Economics, Bonn.
- Freeman, R.B. (1979) The effect of demographic factors on age-earnings profiles, *Journal of Human Resources*, 14, 289–318.
- Fuchs, M. and Weyh, A. (2014) Demography and unemployment in East Germany. How close are the ties?, IAB Discussion Paper No. 26/2014, Institute for Employment Research, Nuremberg.

- Garloff, A., Pohl, C. and Schanne, N. (2013) Do small labor market entry cohorts reduce unemployment?, *Demographic Research*, 29, 379–406.
- Gartner, H. (2005) The imputation of wages above the contribution limit with the German IAB employment sample, FDZ Methodenreport No. 02/2005, Institute for Employment Research, Nuremberg.
- Gertler, M. and Trigari, A. (2009) Unemployment Fluctuations with Staggered Nash Wage Bargaining, *Journal of Political Economy*, 117, 38–86.
- Gürtzgen, N. (2016) Estimating the Wage Premium of Collective Wage Contracts – Evidence from Longitudinal Linked Employer–Employee Data, *Industrial Relations*, 55, 294–322.
- Jimeno, J.F., Rojas, J.A. and Puente, S. (2008) Modelling the impact of aging on social security expenditures, *Economic Modelling*, 25, 201–24.
- Kahn, L.B. (2010) The long-term labor market consequences of graduating from college in a bad economy, *Labour Economics*, 17, 303–16.
- Korenman, S. and Neumark, D. (2000) Cohort Crowding and Youth Labor Markets: A Cross-National Analysis, in D.G. Blanchflower and R.B. Freeman (eds.) *Youth Employment and Joblessness in Advanced Countries*, University of Chicago Press, Chicago.
- Lehmer, F. and Möller, J. (2010) Interrelations between the urban wage premium and firm-size wage differentials: a microdata cohort analysis for Germany, *The Annals of Regional Science*, 45, 31–53.
- Mazzonna, F. and Peracchi, F. (2012) Ageing, cognitive abilities and retirement, *European Economic Review*, 56, 691–710.
- Michaelis, J. and Debus, M. (2011) Wage and (un-)employment effects of an ageing workforce, *Journal of Population Economics*, 24, 1493–511.
- Mincer, J. (1958) Investment in Human Capital and the Personal Income Distribution, *Journal of Political Economy*, 66, 281–302.
- Moffat, J. and Roth, D. (2016a) The Cohort Size-Wage Relationship in Europe, *Labour*, forthcoming.
- Moffat, J. and Roth, D. (2016b) Cohort size and youth labour-market outcomes: the role of measurement error, IAB Discussion Paper No. 37/2016, Institute for Employment Research, Nuremberg.
- Morin, L.-P. (2015) Cohort size and youth earnings: Evidence from a quasi-experiment, *Labour Economics*, 32, 99–111.
- Moulton, B.R. (1990) An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units, *Review of Economics and Statistics*, 72, 334–38.
- Ochsen, C. (2009) Regional labor markets and aging in Germany, Thünen-Series of Applied Economic Theory Working Paper No. 102, Department of Economics, Faculty of Economics and Social Sciences, University of Rostock, Rostock.

- Ottaviano, G. and Peri, G. (2012) Rethinking the effect of immigration on wages, *Journal of the European Economic Association*, 10, 152–97.
- Polachek, S.W. (2008) Earnings Over the Life Cycle: The Mincer Earnings Function and Its Applications, *Foundations and Trends in Microeconomics*, 4, 165–272.
- Sapozhnikov, M. and Triest, R.K. (2007) Population Aging, Labor Demand, and the Structure of Wages, Research Department Working Paper 07–8, Federal Reserve Bank of Boston, Boston.
- Shimer, R. (2001) The Impact of Young Workers on the Aggregate Labor Market, *The Quarterly Journal of Economics*, 116, 969–1007.
- Skans, O.N. (2005) Age effects in Swedish local labor markets, *Economics Letters*, 86, 419–26.
- Staiger, D. and Stock, J.H. (1997) Instrumental Variables Regression with Weak Instruments, *Econometrica*, 65, 557–86.
- Stapleton, D.C. and Young, D.J. (1988) Educational Attainment and Cohort Size, *Journal of Labor Economics*, 6, 330–61.
- Statistisches Bundesamt (2009) *Bevölkerung Deutschlands bis 2060: 12. koordinierte Bevölkerungsvorausberechnung*, Statistisches Bundesamt, Wiesbaden.
- Statistisches Bundesamt (2010) *Bevölkerung und Erwerbstätigkeit: Bevölkerung in den Bundesländern, dem früheren Bundesgebiet und den neuen Ländern bis 2060: Ergebnisse der 12. koordinierten Bevölkerungsvorausberechnung*, Statistisches Bundesamt, Wiesbaden.
- vom Berge, P., König, M. and Seth, S. (2013) Sample of Integrated Labour Market Biographies (SIAB) 1975 – 2010, FDZ Datenreport No. 01/2013, Institute for Employment Research, Nuremberg.
- Welch, F. (1979) Effects of Cohort Size on Earnings: The Baby Boom Babies' Financial Bust, *Journal of Political Economy*, 87, S65–S97.
- Wright, R.E. (1991) Cohort size and earnings in Great Britain, *Journal of Population Economics*, 4, 295–305.
- Zimmermann, K. (1991) Ageing and the Labour Market. Age structure, cohort size and unemployment, *Journal of Population Economics*, 4, 177–200.

Appendix

Table A1: Descriptive statistics

Variable	Mean	Standard deviation	Minimum	Maximum
Log daily earnings	4.28	0.31	3.18	6.21
<i>Youth share</i>				
<i>Labour market region</i>				
Population-based (place of residence)	0.17	0.01	0.14	0.21
Population-based (place of employment)	0.17	0.01	0.14	0.21
<i>District</i>				
Population-based (place of residence)	0.17	0.01	0.13	0.23
Population-based (place of employment)	0.17	0.01	0.13	0.23
<i>Instrument</i>				
<i>Labour market region</i>				
Population-based (place of residence)	0.16	0.02	0.12	0.20
Population-based (place of employment)	0.16	0.02	0.12	0.20
<i>District</i>				
Population-based (place of residence)	0.16	0.02	0.10	0.21
Population-based (place of employment)	0.16	0.02	0.10	0.21
<i>Age</i>	22.33	1.46	15	24
15	0.00	0.01	0	1
16	0.00	0.01	0	1
17	0.00	0.03	0	1
18	0.00	0.07	0	1
19	0.03	0.17	0	1
20	0.09	0.29	0	1
21	0.16	0.37	0	1
22	0.20	0.40	0	1
23	0.24	0.43	0	1
24	0.28	0.45	0	1
<i>Experience</i>	2.04	1.48	0	10
0	0.14	0.35	0	1
1	0.27	0.44	0	1
2	0.24	0.43	0	1
3	0.18	0.38	0	1
4	0.11	0.31	0	1
5	0.04	0.20	0	1
6	0.01	0.11	0	1
7	0.00	0.06	0	1
8	0.00	0.04	0	1
9	0.00	0.02	0	1
10	0.00	0.01	0	1
<i>Education</i>				
Lower secondary (without apprenticeship)*	0.19	0.39	0	1
Lower secondary (with apprenticeship)	0.76	0.42	0	1
Upper secondary (without apprenticeship)	0.01	0.12	0	1
Upper secondary (with apprenticeship)	0.03	0.17	0	1
Tertiary (University of Applied Sciences)	0.01	0.08	0	1
Tertiary (University)	0.00	0.05	0	1

Variable	Mean	Standard deviation	Minimum	Maximum
<i>Nationality</i>				
Germany*	0.90	0.30	0	1
Turkey	0.04	0.20	0	1
Switzerland/Austria	0.00	0.03	0	1
Western Europe	0.00	0.04	0	1
Northern Europe	0.00	0.02	0	1
Central Europe	0.01	0.09	0	1
Eastern Europe	0.00	0.04	0	1
South-East Europe	0.02	0.14	0	1
Southern Europe	0.02	0.13	0	1
Africa	0.00	0.05	0	1
Asia	0.01	0.08	0	1
America/Oceania	0.00	0.02	0	1
Firm size	970.47	3,897.56	1	42,626
<i>Unemployment rate</i>				
Labour-market region (place of residence)	8.16	2.53	2.60	18.04
Labour-market region (place of employment)	8.13	2.51	2.60	18.04
District (place of residence)	8.01	2.94	1.90	25.59
District (place of employment)	8.28	3.01	1.90	25.59
<i>Youth unemployment share</i>				
Labour-market region (place of residence)	7.26	2.65	1.90	24.38
Labour-market region (place of employment)	7.23	2.63	1.90	24.38
District (place of residence)	7.20	2.95	1.70	24.92
District (place of employment)	7.36	2.99	1.70	24.92
Observations	107,351			
* Base category in the regression analysis				

Supplementary Material

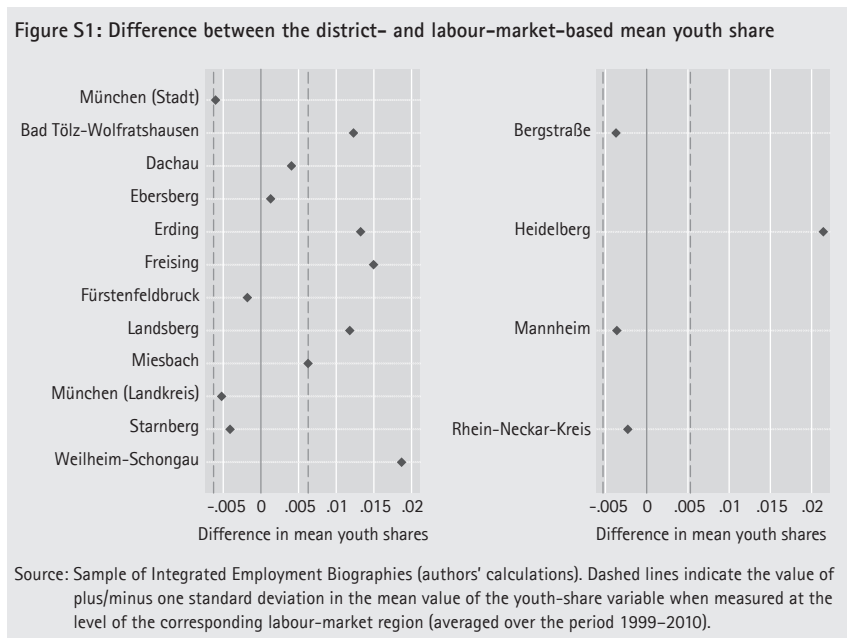
The supplementary material serves two purposes: first, it provides a discussion of the implications for the estimated coefficients when the youth share is constructed from districts rather than from labour-market regions. Since the former are administrative units, which typically do not correspond with local labour market, a district-based youth-share variable is potentially subject to measurement error in that it groups together individuals that are not active within the same labour market. It is, moreover, argued, that the identification strategy employed in the paper may not lead to a consistent estimate of the youth-share coefficient when the former variable is district-specific. The second part addresses the robustness of the results of the empirical analysis and presents various sensitivity analyses.

S1 Measurement error

The paper uses the functional labour-market regions defined by Eckey et al. (2006) as the spatial units from which the empirical model's main variable, the share of the population aged 15–24 relative to the population aged 15–64, is constructed.

As these entities are designed to approximate regional labour markets, their use provides a measure of the potential supply of young workers within an actual labour market. In contrast, administrative units, such as districts, are not delineated accordingly and therefore only provide an incorrect measure of the size of the youth share in the corresponding labour market. The aim of this section is to provide evidence for the existence of measurement error in a district-based youth-share variable and to discuss the implications for the estimation of a model that uses district-specific variables.

Figure S1 illustrates the potential for measurement error in a district-based youth-share variable through use of two exemplary labour-market regions. The labour-market region of Munich (left panel) consists of twelve districts and combines the city of Munich with the surrounding periphery, whereas the region Mannheim-Heidelberg (right panel) is an example of two cities sharing a joint labour market. The graphs show the difference between the average value of the youth-share variable at the level of the labour-market region and at the level of the districts it contains (values are averaged over the period 1999–2010). For some districts this difference can be substantial, exceeding the standard deviation of the average youth share at the level of the labour-market region (as indicated by the dashed lines). In these cases a district-based youth-share variable would appear to provide only an inaccurate measure of the size of the youth population within the corresponding labour-market region.



Under the assumption that the labour-market regions of Eckey et al. (2006) indeed represent the labour markets that individuals are active on, the relationship between an individual's (log) wage and the size of the youth share in his labour-market region can be specified according to the following model (which corresponds to Equation 1 in the paper):

$$\log(w_{irt}) = \alpha + \beta y_{rt} + x_{irt}' \gamma + \delta_r + \mu_t + \varepsilon_{irt} \quad [S1]$$

We propose that a district-specific version of the youth-share variable, y_{rt}^{dis} , provides an incorrect measure of the youth share within the labour market that an individual belongs to which we specify in form of a district-level dummy variable, ψ_{irt}^{dis} , and an additive random measurement error ξ_{rt}^{dis} . The variable ψ_{irt}^{dis} allows for the possibility of the youth share in a specific district being permanently smaller or larger than its value in the corresponding labour market region:

$$y_{rt}^{dis} = y_{rt} + \psi_r^{dis} + \xi_{rt}^{dis} \quad [S2]$$

Substituting Equation S2 into the model of Equation S1 shows that a district-specific youth-share variable is correlated with the composite error term ε_{irt}^{dis} , which contains the measurement error component ξ_{rt}^{dis} .

$$\log(w_{irt}) = \alpha + \beta y_{rt}^{dis} + x_{irt}' \gamma + \eta_r^{dis} + \mu_t + \varepsilon_{irt}^{dis} \quad [S3]$$

$$\eta_r^{dis} = \delta_r - \beta \psi_r^{dis} \quad [S4]$$

$$\varepsilon_{irt}^{dis} = \varepsilon_{irt} - \beta \xi_{rt}^{dis} \quad [S5]$$

As labour-market regions are comprised of one or more districts, it can be shown that the youth share of a given labour-market region k is equal to the weighted sum of the youth shares in the districts l ($l = 1, \dots, L$) that are contained in region k , where the weights are given by the fraction of the population aged 15–64 in region k , $N_{15-64, kt}$, that can be ascribed to district l , $N_{15-64, lt}^{dis}$ (consequently, the weights add up to unity):

$$y_{kt} = \sum_{l=1}^L \omega_{lt}^{dis} y_{lt}^{dis} \quad [S6]$$

$$\omega_{lt}^{dis} = \frac{N_{15-64, lt}^{dis}}{N_{15-64, kt}} \quad [S7]$$

The relationship between the region-based and the district-based youth-share variables shown in Equation S2 implies that the measurement errors of those districts within a given labour-market region are linearly dependent:

$$\xi_{1t}^{dis} = -\psi_1^{dis} - \sum_{l=2}^L \frac{\omega_{lt}^{dis}}{\omega_{1t}^{dis}} \psi_l^{dis} - \sum_{l=2}^L \frac{\omega_{lt}^{dis}}{\omega_{1t}^{dis}} \xi_{lt}^{dis} \quad [S8]$$

Consequently, use of a district-based youth-share variable not only induces endogeneity due to measurement error, but also leads to the error terms of observations from different districts being correlated if they belong to the same labour-market region.

In principle, IV estimation can be used to obtain consistent estimates in the presence of measurement error (Hausman, 2001) if the instrument is uncorrelated with the composite error term of Equation S5. In this paper the instrument is defined as the ratio of the number of individuals up to the age of 9 and the number of individuals up to the age of 49 observed 15 years earlier. As in the case of the youth-share variable the instrument can be constructed from either districts or labour-market regions and, analogously to Equation S2, it is possible to interpret the district-based version of the instrument as an incorrect measure of the regional variable:

$$z_{r,t-15}^{dis} = z_{r,t-15} + \varphi_r^{dis} + v_{r,t-15}^{dis} \quad [S9]$$

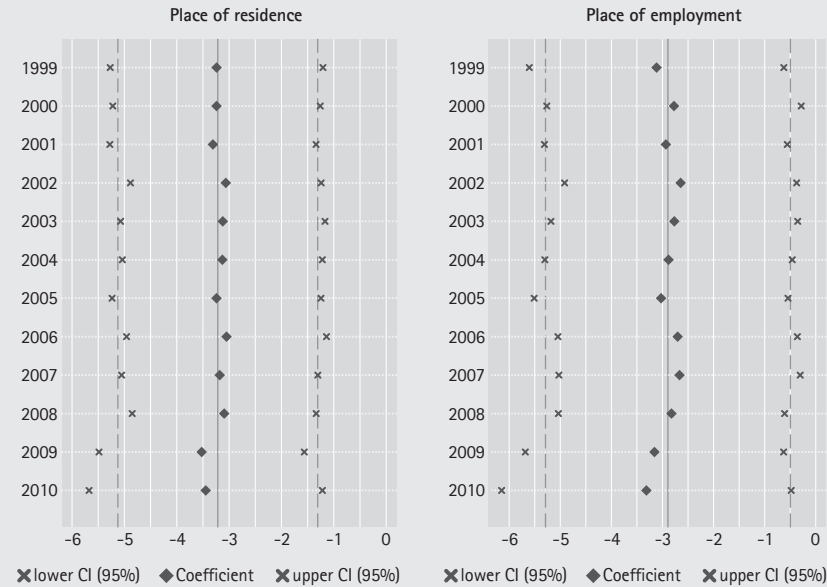
Consistent estimation of a model with a district-based youth-share variable in combination with a district-specific instrument requires that the current and the lagged measurement errors are uncorrelated. This condition would not be satisfied if the extent of measurement error exhibited persistence over time, e.g. if a large difference between the district-based and the region-based instrument was associated with a large difference in the district-specific and region-specific youth-share variable. Under these circumstances, application of the identification strategy of the paper to a model with a district-based youth-share variable would not yield a consistent estimate of the former's effect on an individual's wages.

S2 Sensitivity analysis

In this section we perform various sensitivity analyses in order to assess the robustness of the paper's findings. First, we show that the 2SLS coefficients of the baseline model are not driven by individual regions or years. Figure S2 presents the youth-share coefficients that are obtained when observations from a single year are excluded from the sample: regardless whether an individual's region of

residence or region of employment is used, the estimated coefficients are always very close to those of the full model and always lie within the formers' 95% confidence interval.

Figure S2: 2SLS baseline coefficients after excluding individual years



Source: Sample of Integrated Employment Biographies (authors' calculations). The youth-share coefficients are derived from the baseline model; the solid line represents the youth-share coefficients from the full model, the dashed lines the corresponding 95% confidence interval.

Due to the relatively large number of regions (108), illustrating the effects of dropping individual regions would be unwieldy. Instead of showing individual coefficients, their histogram is depicted in Figure S3. As can be seen, the distribution of the coefficients is centred on the coefficient of the full model and displays a spread which is small compared to the 95% confidence interval.

The following tables contain the coefficient of the youth-share variable as well as the former's marginal effect for a change of one standard deviation for a number of sensitivity analyses in which either the sample (Tables S1–S10) or the empirical specification are modified. In Tables S11 and S12 the question is addressed why a change from a region-specific to a district-specific youth-share variable leads to a larger decrease in the size of the coefficient when the variable is measured at an individual's place of employment. Finally, Tables S12 and S13 report the youth-share coefficients from a district-specific variable when indicators for an individual's industry and/or occupation are added.

Figure S3: Histogram of 2SLS baseline coefficients after excluding individual regions



Table S1: Exclusion of observations with more than a full-time job

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.44 (0.68)*	-3.29 (1.01)***	-0.81 (0.77)	-2.90 (1.22)*
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market regions	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	-	0.46 (0.00)***	-	0.46 (0.00)***
<i>First-stage statistics</i>				
F-statistic	-	135.19***	-	131.05***
Shea's partial R ²	-	0.32	-	0.32
<i>Observations</i>				
Individuals	102,387	102,387	102,387	102,387
Labour-market regions-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.24	0.24	0.24	0.24
ME(stdev)	-1.81%*	-4.14%***	-1.03%	-3.68%*
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/+ indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

As discussed in the paper, the analysis is restricted to those individuals who are subject to social security contributions. Tables S1 and S2 show the results from further homogenising the sample by either excluding those individuals with more than a full-time job (Table S1) or by dropping individuals whose employment spells contain less than 90 days (Table S2). In the first case the marginal effects are very close to those of the paper's baseline specification, while they become slightly smaller in the second case.

Table S2: Exclusion of observations with employment spells of less than 90 days

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.21 (0.62) [†]	-2.75 (0.97) ^{***}	-0.69 (0.71)	-2.53 (1.21) [*]
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market regions	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	-	0.46 (0.00) ^{***}	-	0.46 (0.00) ^{***}
<i>First-stage statistics</i>				
F-statistic	-	137.34 ^{***}	-	132.74 ^{***}
Shea's partial R ²	-	0.32	-	0.32
<i>Observations</i>				
Individuals	103,652	103,652	103,652	103,652
Labour-market regions-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.23	0.23	0.23	0.23
ME(stdev)	-1.52% [*]	-3.47% ^{***}	-0.88%	-3.22% [*]
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ^{***} / ^{**} / [*] / [†] indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

Since the effect of the size of the youth share may vary between differently educated individuals, the sample is homogenised by first excluding those with tertiary education (Table S3) and then those with tertiary or upper secondary education (Table S4). In neither case is the size of the marginal effects substantially changed.

Table S3: Exclusion of observations with tertiary education

<i>Dependent variable: log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.47 (0.62)*	-3.22 (0.96)***	-0.91 (0.73)	-2.97 (1.22)*
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market regions	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	-	0.46 (0.00)***	-	0.46 (0.00)***
<i>First-stage statistics</i>				
F-statistic	-	137.01***	-	132.41***
Shea's partial R ²	-	0.32	-	0.32
<i>Observations</i>				
Individuals	106,422	106,422	106,422	106,422
Labour-market regions-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.24	0.24	0.24	0.24
ME(stdev)	-1.86%*	-4.05%***	-1.16%	-3.77%*

Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/*/† indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.

Table S4: Exclusion of observations with tertiary or upper secondary education

<i>Dependent variable: log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.39 (0.64)*	-3.29 (0.95)***	-0.84 (0.72)	-3.04 (1.21)*
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market regions	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	-	0.46 (0.00)***	-	0.46 (0.00)***
<i>First-stage statistics</i>				
F-statistic	-	140.36***	-	135.72***
Shea's partial R ²	-	0.33	-	0.32
<i>Observations</i>				
Individuals	101,820	101,820	101,820	101,820
Labour-market regions-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.24	0.24	0.24	0.24
ME(stdev)	-1.75%*	-4.14%***	-1.06%	-3.85%*

Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/*/† indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.

The SIAB dataset contains observations with unrealistically low average daily wages. In order to get a handle on this issue, observations with daily wages below twice the value of the minor-employment threshold were excluded from the sample. Table S5 shows that when these observations are included, the size of the coefficients decrease in size and they become less significant.

Table S5: No truncation of the wage distribution

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-0.99 (0.78)	-2.51 (1.09)*	-0.33 (0.85)	-2.14 (1.26) [†]
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market regions	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	-	0.46 (0.00)***	-	0.46 (0.00)***
<i>First-stage statistics</i>				
F-statistic	-	136.17***	-	131.61***
Shea's partial R ²	-	0.32	-	0.32
<i>Observations</i>				
Individuals	110,651	110,651	110,651	110,651
Labour-market regions-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.24	0.24	0.24	0.24
ME(stdev)	-1.25%	-3.16%*	-0.41%	-2.72% [†]
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/† indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

The paper's empirical analysis is restricted to those individuals who are employed on 30 June of a given year and while June values are usually representative of average annual employment levels, the selection of a specific date is essentially arbitrary. Table S6 shows the results when the reference date is set to 31 December. Doing so produces comparable results in terms of sign and significance but the size of the marginal effects increases in magnitude.

Table S6: Alternative reference date (31 December)

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.36 (0.65)*	-3.80 (0.96)***	-0.75 (0.74)	-3.26 (1.16)***
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market regions	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	-	0.46 (0.00)***	-	0.46 (0.00)***
<i>First-stage statistics</i>				
F-statistic	-	135.93***	-	130.83***
Shea's partial R ²	-	0.32	-	0.32
<i>Observations</i>				
Individuals	113,748	113,748	113,748	113,748
Labour-market regions-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.30	0.30	0.30	0.30
ME(stdev)	-1.71%*	-4.79%***	-0.95%	-4.14%***
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/*/+ indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

Table S7: Employment-based youth share variable

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-0.45 (0.37)	-3.06 (1.08)***	-0.31 (0.41)	-2.79 (1.28)*
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market regions	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	-	0.48 (0.01)***	-	0.47 (0.01)***
<i>First-stage statistics</i>				
F-statistic	-	56.91***	-	56.01***
Shea's partial R ²	-	0.18	-	0.18
<i>Observations</i>				
Individuals	107,351	107,351	107,351	107,351
Labour-market regions-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.24	0.24	0.24	0.24
ME(stdev)	-0.69%	-4.72%***	-0.48%	-4.30%*
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/*/+ indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

The youth-share variable is meant to measure the potential supply of young individuals to the labour market. In the paper this variable is constructed from the size of the population in the corresponding age group. However, it is likely that parts of this group are not available to the labour market and as such a population-based variable might provide an inaccurate measure of age-specific labour supply. When a youth-share variable is used instead that is defined as the number of employees aged between 15 and 24 relative to the number of employees between 15 and 64, similarly sized coefficients are obtained, but since the standard deviation of the employment-based youth-share variable is larger than in the case of the population-based variable the marginal effects increase in size (Table S7).

The paper estimates the effect of the youth-share on individual wages. Alternatively, it is possible to first average individual-level variables within a region-year cell and to then regress the average daily wage within such a cell on the youth share and weighting the regression by the number of observations in a region-year cell (see Angrist and Pischke, 2009). As can be seen from Table S8, the aggregate-level analysis yields comparable results, though the marginal effects are slightly larger.

Table S8: Aggregated model (variables averaged at the level of the region-year cell)

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.65 (0.68)*	-3.71 (1.14)***	-0.89 (0.77)	-3.37 (1.43)*
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market regions	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	-	0.45 (0.04)***	-	0.44 (0.04)***
<i>First-stage statistics</i>				
F-statistic	-	134.84***	-	124.29***
Shea's partial R ²	-	0.30	-	0.30
<i>Observations</i>				
Labour-market regions-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.82	0.81	0.82	0.81
ME(stdev)	-2.08%*	-4.68%***	-1.14%	-4.28%*
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/*/t indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

In the empirical specification the effects of age and experience on an individual's wage are approximated through the inclusion of these variables' first two polynomials. However, very similar results are obtained if mutually exclusive dummy variables are used instead (Table S9).

Table S9: Age and experience dummies

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.47 (0.63)*	-3.27 (0.98)***	-0.86 (0.73)	-2.96 (1.23)*
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market regions	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes
Experience	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	-	0.46 (0.00)***	-	0.46 (0.00)***
<i>First-stage statistics</i>				
F-statistic	-	136.64***	-	131.81***
Shea's partial R ²	-	0.32	-	0.32
<i>Observations</i>				
Individuals	107,351	107,351	107,351	107,351
Labour-market regions-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.29	0.29	0.29	0.29
ME(stdev)	-1.85%*	-4.12%***	-1.10%	-3.75%*
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/*/+ indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

Table S10 contains the results of estimating a double-log specification. The coefficients of the youth-share variable continue to be negative and significant.

Comparing the results of Tables 1 and 2 in the paper shows that when a district-specific youth-share variable is used rather than one based on labour-market regions, the decrease in the size of the coefficient is considerably stronger for the place of employment than the place of residence. In the following, all districts are ordered according to the difference between the number of observations in the sample that live and that work in a district. The model of Equation 1 is then estimated separately for the districts from the top half (i.e. for which the difference is largest) and for those from the bottom half (i.e. for which the difference is smaller) of this ordering. The results are shown in Tables S11 and S12, respectively. As was already discussed in the paper, negative and significant effects are only found for the set of districts from the top half of the

ordering. The row *Fraction of full sample* shows that for the place-of-residence specification the majority of observations (55%) are from such districts. In the place-of-employment specification the corresponding Figure stands at only 44%.

Table S10: Double-log specification

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-0.24 (0.11)*	-0.52 (0.15)***	-0.13 (0.13)	-0.46 (0.19)*
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour-market regions	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	-	0.42 (0.00)***	-	0.42 (0.00)***
<i>First-stage statistics</i>				
F-statistic	-	129.62***	-	124.52***
Shea's partial R ²	-	0.35	-	0.35
<i>Observations</i>				
Individuals	107,351	107,351	107,351	107,351
Labour-market regions-year cells	1,296	1,296	1,296	1,296
Labour-market regions (clusters)	108	108	108	108
R ²	0.24	0.24	0.24	0.24
Cluster-robust standard errors in parentheses (clustered at the level of the labour-market region). ***/**/*/+ indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression.				

Assuming that the districts from the top half represent those which individuals are more likely to live in than to work in, the larger decrease in the size of the youth-share coefficient (i.e. a larger degree of attenuation) in the place-of-employment specification may be explained by the fact that the degree of measurement error is more pronounced in regions that people are more likely to live in and that this type of district is over-represented in the place-of-employment specification. To support this argument, the bottom rows of Tables S11 and S12 show the mean difference between the youth-share variable at the level of the labour-market region and of the district in the corresponding sample. A comparison of these differences between Table S11 and Table S12 shows that regardless of whether the place-of-residence (0.51 as opposed to 0.42) or the place-of-employment specification (0.55 as opposed to 0.40) is used, the extent of measurement error is larger for those districts that individuals are more likely to work in than to live in.

Table S11: Districts in which individuals are more likely to live

<i>Dependent variable: log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.78 (0.62)***	-3.63 (0.98)***	-1.07 (0.67)	-2.21 (1.12)*
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Districts	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	-	0.46 (0.03)***	-	0.45 (0.03)***
<i>First-stage statistics</i>				
F-statistic	-	253.16***	-	230.56***
Shea's partial R ²	-	0.45	-	0.43
<i>Observations</i>				
Individuals	58,705	58,705	47,146	47,146
Fraction of full sample	54.69%	54.69%	43.92%	43.92%
District-year cells	1,884	1,884	1,884	1,884
Districts (clusters)	157	157	157	157
R ²	0.23	0.23	0.23	0.23
ME(stdev)	-2.12%***	-4.33%***	-1.27%	-2.63%*
Mean difference	0.42		0.40	
Cluster-robust standard errors in parentheses (clustered at the district level). ***/**/+ indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation. Mean difference gives the average absolute difference between the district-level youth share and the value at the level of the corresponding labour-market region (multiplied by 100).				

Finally, Tables S13 and S14 provide the analogues to Tables 3 and 4 but use a youth-share variable that is constructed from districts rather than labour-market regions. First, the results continue to be considerably larger in magnitude for the place of residence than the place of employment when industry and occupation dummies are added; for the place of employment, none of the coefficients are statistically significant. Second, the youth-share coefficients of the district-specific model remain smaller than the ones from the region-specific model. In the case of the place of residence the district-specific coefficients are smaller by between 27% (industry dummies) and 10% (occupation dummies). In contrast, the differences in size are much more pronounced at the place of employment where the inclusion of dummies for an individual's industrial or occupational affiliation further reduces the magnitude of the youth-share coefficients relative to those from the labour-market specification. This finding illustrates that the distinction between place of employment and place of residence is of particular importance for the estimated size and significance of the effects at the district level.

Table S12: Districts in which individuals are more likely to work

<i>Dependent variable:</i> <i>log real daily earnings</i>	Place of residence		Place of employment	
	OLS	2SLS	OLS	2SLS
Youth share	-1.05 (0.61) [†]	-1.75 (1.39)	0.11 (0.52)	-0.87 (1.43)
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Districts	Yes	Yes	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage statistics</i>				
Instrument	–	0.43 (0.05) ^{***}	–	0.42 (0.05) ^{***}
<i>First-stage statistics</i>				
F-statistic	–	85.01 ^{***}	–	60.10 ^{***}
Shea's partial R ²	–	0.17	–	0.15
<i>Observations</i>				
Individuals	48,646	48,646	60,205	60,205
Fraction of full sample	45.31%	45.31%	56.08%	56.08%
District-year cells	1,872	1,872	1,872	1,872
Districts (clusters)	156	156	156	156
R ²	0.27	0.27	0.29	0.29
ME(stdev)	-1.63% [†]	-2.70%	0.18%	-1.40%
Mean difference	0.51		0.55	
Cluster-robust standard errors in parentheses (clustered at the district level). ***/**/† indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation. Mean difference gives the average absolute difference between the district-level youth share and the value at the level of the corresponding labour-market region (multiplied by 100).				

Table S13: Industry and occupation indicators (place of residence)

<i>Dependent variable: log real daily earnings</i>	Baseline	+industry	+occupation	+industry +occupation
Youth share (2SLS)	-2.79 (0.81)***	-2.05 (0.69)***	-1.68 (0.75)**	-1.40 (0.66)**
Youth share (OLS)	-1.31 (0.45)***	-1.10 (0.35)***	-0.93 (0.44)**	-0.86 (0.34)**
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour market region	Yes	Yes	Yes	Yes
Industry	No	Yes	No	Yes
Occupation	No	No	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage regression</i>				
Instrument	0.44 (0.00)***	0.44 (0.00)***	0.44 (0.00)***	0.44 (0.00)***
<i>First-stage test statistics</i>				
F-statistic	300.92***	301.78***	301.96***	302.71***
Shea's partial R ²	0.27	0.27	0.27	0.27
<i>Observations</i>				
Individuals	107,351	107,351	107,351	107,351
District-year cells	3,756	3,756	3,756	3,756
Districts (clusters)	313	313	313	313
R ² (2SLS)	0.25	0.46	0.41	0.51
R ² (OLS)	0.25	0.46	0.41	0.51
ME(stdev, 2SLS)	-3.84%***	-2.82%***	-2.31%**	-1.92%**
ME(stdev, OLS)	-1.80%***	-1.51%***	-1.28%**	-1.19%**
Cluster-robust standard errors in parentheses (clustered at the district level). ***/**/+ indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

Table S14: Industry and occupation indicators (place of residence)

<i>Dependent variable: log real daily earnings</i>	Baseline	+industry	+occupation	+industry +occupation
Youth share (2SLS)	-1.50 (0.92)	-1.19 (0.76)	-0.87 (0.79)	-0.83 (0.69)
Youth share (OLS)	-0.12 (0.42)	-0.23 (0.34)	0.04 (0.40)	-0.05 (0.32)
<i>Dummies</i>				
Year	Yes	Yes	Yes	Yes
Labour market region	Yes	Yes	Yes	Yes
Industry	No	Yes	No	Yes
Occupation	No	No	Yes	Yes
<i>Control variables</i>	Yes	Yes	Yes	Yes
<i>First-stage regression</i>				
Instrument	0.43 (0.00)***	0.43 (0.00)***	0.43 (0.00)***	0.43 (0.00)***
<i>First-stage test statistics</i>				
F-statistic	181.95***	182.17***	182.12***	182.29***
Shea's partial R ²	0.22	0.22	0.22	0.22
<i>Observations</i>				
Individuals	107,351	107,351	107,351	107,351
District-year cells	3,756	3,756	3,756	3,756
Districts (clusters)	313	313	313	313
R ² (2SLS)	0.26	0.47	0.42	0.52
R ² (OLS)	0.26	0.47	0.42	0.52
ME(stdev, 2SLS)	-2.18%	-1.74%	-1.27%	-1.21%
ME(stdev, OLS)	-0.17%	-0.33%	0.05%	-0.08%
Cluster-robust standard errors in parentheses (clustered at the district level). ***/**/*+ indicate significance at the 0.005/0.01/0.05/0.10 level, respectively. Instrument shows the coefficient of the instrument in the first-stage regression. ME(stdev) gives the percentage change in daily earnings given an increase in the youth share by one standard deviation.				

References

- Angrist, J. and Pischke, J.-S. (2009) *Mostly harmless econometrics: an empiricist's companion*, Princeton University Press, Princeton.
- Eckey, H.-F., Kosfeld, R. and Türck, M. (2006) Abgrenzung deutscher Arbeitsmarktregionen, *Raumforschung und Raumordnung*, **64**, 299–309.
- Hausman, J. (2001) Mismeasured Variables in Econometric Analysis: Problems from the Right and Problems from the Left, *Journal of Economic Perspectives*, **15**, 57–67.

John Moffat^a and Duncan Roth^b

Cohort size and youth labour-market outcomes: the role of measurement error

Abstract

Using data from 49 European regions covering 2005–2012, this paper finds that the estimated effect of cohort size on employment and unemployment outcomes is very sensitive to the age range of the sample. We argue that this is because the identification strategy commonly used in this literature is unable to eliminate the bias caused by measurement error in the cohort-size variable. The latter arises because large shares of the young choose to acquire education and consequently the size of an age group provides a poor measure of age-specific labour supply. In our view older age groups provide a more suitable sample to test the implications of cohort crowding since the former will have largely entered the labour market. Using a sample aged 25–29, which has relatively low rates of participation in education, we find robust evidence that an increase in cohort size increases employment and reduces unemployment.

JEL classification: J10, J21, R23

Keywords: Cohort size, cohort crowding, unemployment, employment, measurement error, EU-SILC

1 Introduction

The effect of the size of the youth population upon its labour-market prospects is of critical importance, particularly in light of demographic trends which will cause the youth share of the population to fall in most countries in coming decades (United Nations, 2015). The cohort-crowding hypothesis suggests that this will be beneficial for young individuals (Easterlin, 1961; Welch, 1979). By contrast, the model of Shimer (2001) implies that smaller youth cohorts will have a detrimental impact as firms create fewer jobs in areas with smaller youth shares. While the bulk of the empirical literature has focused on earnings and generally found negative effects of cohort size (e.g. Welch, 1979; Wright, 1991; Brunello, 2010;

^a Durham University, Department of Economics and Finance.

^b Institute for Employment Research (IAB), Düsseldorf, Germany.

Moffat and Roth, 2016; Garloff and Roth, 2016), the effect on unemployment and employment has received less attention and the empirical evidence is so far mixed (Korenman and Neumark, 2000; Shimer, 2001; Skans, 2005; Foote, 2007; Biagi and Lucifora, 2008; Garloff et al., 2013).

In this paper, we propose that the standard identification strategy that has been used in the cohort-size literature does not allow for consistent estimation of the effect of cohort crowding for young age groups. There are two reasons for this, both of which are based on the observation that, due to high rates of participation in education, the relative size of an age group represents a poor measure of age-specific labour supply among the young, the latter being the relevant variable for age-specific employment and unemployment outcomes. First, since the proportion of young people that choose to defer entry to the labour market in order to acquire education may be influenced by cohort size (Fertig et al., 2009), this complicates the interpretation of estimated effects of cohort size since they reflect effects on participation and, conditional on participation, on (un-)employment. More importantly, the use of the number of individuals in an age group as the basis for the cohort-size variable creates measurement error that the standard instrumental variables (IV) approach to estimating the effects of cohort-size is unable to overcome.

We assess this argument by estimating the effect of cohort size on employment and unemployment shares using data from the longitudinal European Union Statistics on Income and Living Conditions (EU-SILC) survey which provides us with data on 49 regions for the period 2005–2012. Our results show that the estimated cohort-size effects are very sensitive to the chosen age range of the sample. Our preferred results come from a sample of individuals aged 25–29 since most of that group has entered the labour market and therefore the decision to participate in the labour market as well as the degree of measurement error are less of a concern. Among this group, we find, in contradiction of the cohort-crowding hypothesis, a negative effect of cohort size on the unemployment share. These results are robust to a variety of changes in the sample and in the empirical specification. This finding is relevant because it casts doubt on the conclusions from previous studies, which have defined the youth population as individuals aged 15/16–24, regarding the relationship between the size of the youth population and its members' employment and unemployment outcomes.

Section 2 reviews the extant theoretical and empirical literature on the relationship between population structure and labour market outcomes. Section 3 discusses the dataset and empirical model. The results are presented in Section 4 and Section 5 concludes.

2 Literature review

Competing theoretical predictions and conflicting empirical evidence exist regarding the question of how changes in the size of an age group affect its (un-) employment prospects. The cohort-crowding hypothesis is based on the assumption that differently aged workers are only imperfectly substitutable due to differences in human capital (Welch, 1979) so that changes in the size of an age group have implications predominantly for members of that age group (see Moffat and Roth, 2016, for a more detailed discussion). In perfectly competitive labour markets, changes in age-group size would only be reflected in changes to age-specific wages. If labour markets are imperfectly competitive, however, wages need not be fully flexible and an increase in the size of an age group may lead to an increase in the unemployment rate of that group (a theoretical model of this relationship in imperfectly competitive markets is provided by Michaelis and Debus, 2011).

In line with the cohort-crowding hypothesis, Korenman and Neumark (2000) provide empirical evidence that large youth cohorts (measured as the ratio of individuals aged 15–24 to individuals aged 25–54) increase the youth unemployment rate. Their findings are robust to a number of specifications, including the use of lagged birth rates as an instrument for the potentially endogenous youth-share variable. Moreover, the use of cross-national variation in their dataset of Organisation for Economic Cooperation and Development countries allows the authors to separately identify the effects of changes in youth-cohort size from the effects of other macroeconomic developments and as such provides an improvement on earlier studies that relied solely on time-series variation (e.g. Zimmermann, 1991; Schmidt, 1993).

Rather different results are obtained by Shimer (2001). Using data on a panel of US states for the period 1970–1996, he finds that increases in the youth share – measured as the ratio of those aged 16–24 to those aged 16–64 – are associated with decreases in the state-level unemployment rate. This is surprising for two reasons: first, since the overall unemployment rate is the sum of age-specific unemployment rates weighted by the share of the respective age group in the labour force and the youth unemployment rate generally exceeds that of older individuals, the direct effect of an increase in the youth share should be to increase the overall rate. Second, according to the cohort-crowding hypothesis the indirect effect of an increase in the youth share should be to increase the youth unemployment rate, thereby reinforcing the direct effect. Shimer's (2001) empirical results, however, not only show a negative effect on the overall unemployment rate, but also that the youth share reduces the unemployment rate of youths as well as other age groups.

Shimer (2001) provides a theoretical foundation to his empirical findings in the form of a search and matching model with on-the-job search. Changes in the size of the youth population tend to be predictable, as evidenced by the explanatory power of lagged birth rates for the size of the current youth share. Moreover, young individuals are more often either without a job or less well matched than older individuals and are therefore, on average, more willing to take up or switch jobs. This makes it easier for firms to make a productive match with workers in markets with a large number of potential employees. They therefore react to an expected change in the youth share by creating vacancies, to the benefit of all age groups.

Aiming to explain the substantial differences between his own and Korenman and Neumark's (2000) empirical findings, Shimer (2001) points out that the former ignored the possibility of changes in the youth share having an effect on the unemployment rate of other age groups. Specifically, Korenman and Neumark's (2000) model includes the adult unemployment rate, alongside the youth share, as a regressor in the model of the youth unemployment rate. According to Shimer (2001), if changes in the youth share affect the unemployment rates of both age groups, the former's coefficient will be biased upwards and he is able to show this using his own dataset. However, applying his empirical model to the data of Korenman and Neumark (2000) produces inconclusive results, which casts doubt on the applicability of his theoretical model to other countries and time periods.

The small number of studies that have since looked at the relationship between age structures and unemployment outcomes have yielded mixed results. Using data on Swedish labour markets for the years 1985–1999, Skans (2005) finds no evidence for an effect of the relative size of the group aged 16–24 on the total unemployment rate, but his results are otherwise in line with Shimer (2001) since they show that the youth unemployment rate falls when the size of young age groups increases. In contrast, Foote (2007) shows that when the time dimension of Shimer's (2001) dataset is extended to 2005 the negative effect of the youth share on the overall unemployment rate decreases considerably and becomes insignificant in most specifications. The empirical evidence of Biagi and Lucifora (2008) also contradicts the findings of Shimer (2001): their analysis of a dataset of European countries spanning the late 1970s to the early 2000s suggests that the share of individuals aged 15–24 has a positive effect on the unemployment rate of the young and is not statistically significant for the unemployment outcomes of prime-age individuals. Finally, Garloff et al. (2013), using data on West German labour-market regions for the years 1993–2008, find that increases in the share of individuals aged 15–24 years are associated with increases in the overall unemployment rate.

In light of the conflicting results produced by previous studies this analysis provides new evidence on the relationship between age-group size and age-specific unemployment outcomes. Our dataset is a longitudinal sample of European regions covering 2005–2012 which provides us with more heterogeneity to separate the effects of cohort size from other influences than has generally been available in the literature. However, the paper's main contribution is to consider the effect of the definition of the youth population on the estimates obtained. The previous literature has used the share of individuals aged either 15–24 or 16–24 as a definition of the youth share. Since a high proportion of this group will be in education and therefore potentially unavailable to the labour market, this will, as discussed in the introduction and in more detail below, have important implications for both the interpretation and econometric identification of the cohort-size effect.

3 Empirical analysis

3.1 Data

The major part of the dataset that is used in the empirical analysis is constructed by combining different longitudinal EU-SILC releases.¹ Appending data from different releases not only allows the extension of the sample period beyond the four years provided by a single longitudinal release, but also increases the number of observations within a given year. In order to match observations from different releases that refer to the same individual, a unique personal identifier is constructed.² This is then used to verify that there are very few individuals with inconsistencies in age and sex over time³ (see Moffat and Roth, 2016, for further details on the process of appending the different datasets and Berger and Schaffner, 2015, for general information about EU-SILC).

Individuals in EU-SILC are not randomly sampled and weights are therefore provided so that unbiased population estimates may be calculated. We use these to construct two new weighting variables: the first of these variables corrects

1 The longitudinal releases are: 2013 (version 1 from 01-08-2015), 2012 (version 3 from 01-08-2015), 2011 (version 4 from 01-03-2015), 2010 (version 5 from 01-08-2014), 2009 (version 4 from 01-03-2013), 2008 (version 4 from 01-03-2012), 2007 (version 5 from 01-08-2011), 2006 (version 2 from 01-03-2009) and 2005 (version 1 from 15-09-07).

2 This identifier is defined as a combination of an observation's identification number (which is not unique across countries), his country of residence and the rotational group to which he belongs.

3 In total, there are 36 individuals (182 observations) with inconsistencies. All of these individuals are from France, Luxembourg or Norway (i.e. countries in which individuals can be followed for more than 4 years). For these individuals, the inconsistent observations are dropped. If there are only two observations per individual, both are dropped.

the initial weights for the number of rotational groups within a country-year combination that change as a result of appending data from different releases (see Moffat and Roth, 2016). The second weighting variable also re-scales the weights so that the size of the estimated population within a region-year-age-sex cell is identical to the statistics reported by Eurostat.⁴

The so-constructed dataset contains 2.76 million observations on just over 1 million individuals and covers the years 2004–2013. In addition to the country that an individual resides in, EU-SILC provides information about the region of residence at the first level of the Nomenclature of Territorial Units in Statistics (NUTS). Availability of this information allows us to construct the relevant variables at the regional rather than at the national level, which is attractive because estimates of functional labour markets have tended to show them to be defined at the sub-national level (see Moffat and Roth, 2016).

Rather than focussing on outcomes at the individual level, the empirical analysis in this paper is concerned with estimating the effect of age-specific cohort size on unemployment and employment outcomes at the level of the corresponding age group. For this reason, the dataset is aggregated to the level of region-year-age cells. The resulting dataset is further supplemented by variables taken from Eurostat's publicly available database⁵: the level of regional GDP and the size of relevant age groups between 1991 and 1998 which are used as instruments in the empirical analysis.⁶

Due to data limitations, observations from the following countries are dropped: Germany, the Netherlands and Portugal (information on NUTS1 regions is not provided); Croatia (lagged population data for the construction of the instrument is not available); Finland, Iceland and Slovenia (age-related variables are randomly perturbed to prevent disclosure); Ireland and the United Kingdom (the age variable is measured at a different time of year for these countries, see footnote 4). Moreover, we exclude observations from Bulgaria, Cyprus, Malta, Norway and Romania because the necessary variables are not available throughout the whole sample period. This leaves a panel of 49 NUTS1 regions from the following countries for which age groups can be observed from 2005–2012 (number of regions per country in parentheses): Austria (3), Belgium (3), Czech Republic (1),

4 Note that while the Eurostat statistics refer to 1 January of a given year, use of the variable *age at the end of the income reference period* ensures that the population sizes estimated from EU-SILC data refer to 31 December of the preceding year.

5 The data can be obtained through the following link: http://epp.eurostat.ec.europa.eu/portal/page/portal/statistics/search_database

6 Due to a change in delineation lagged population data is not available before the year 2003 for the two regions ITH (Northeast Italy) and ITI (Central Italy). Since these changes are minor compared to the total size of the regions we instead use lagged age-group size based on the predecessor regions ITD and ITE, which we obtain from the homepage of the Italian Statistical Office (www.istat.it).

Denmark (1), Estonia (1), Greece (4), Spain (7), France (8), Hungary (3), Italy (5), Lithuania (1), Luxemburg (1), Latvia (1), Poland (6), Sweden (3), Slovakia (1).

3.2 Variables and sample

This section serves several purposes: first, it defines the main variables of the empirical model; second, it discusses the age range of the sample; finally, an illustration is provided of the variation in the cohort-size variable that is used for identification.

The analysis separately estimates the effect of changes in cohort size on the share of individuals in age group j , region r and year t that are unemployed ($unemp_{jrt}$) and employed (emp_{jrt}). As discussed in the previous section, these shares are derived from individual-level data. Specifically, the weighted sum of male individuals who report to be (un-)employed in a given region-year-age group is calculated and divided by the total male population in that cell. Female observations are excluded in order to avoid the results being affected by selected labour-market participation. As these variables are standardised on the population rather than the labour force, the outcome variables differ from the unemployment and the employment rate. An advantage of this specification is that any effects that changes in cohort size, if measured without error, might have on participation rates could be ignored in the interpretation of the results.

Figure A1 in the Appendix shows the development of the dependent variables $unemp_{jrt}$ and emp_{jrt} as well as of a similarly defined variable that shows the share of individuals reporting to be in education in a given age group ($educ_{jrt}$). These variables are plotted for the age group 18–29 in selected regions and years to illustrate the variation in age-specific labour-market outcomes across Europe. While there are differences in the slope of the profiles, a common feature of all region-year combinations is that the employment share tends to increase and the share of individuals in education decreases with age. In contrast, there is no obvious trend in the unemployment share. In order to understand the implications of the high share of young individuals in education, the empirical model is firstly estimated for overlapping five-year age groups (beginning with individuals aged 18–22 and ending with individuals aged 25–29). The reason for adopting this strategy is that for younger age groups the coefficients will capture the effect of cohort size on labour market participation and, conditional on participation, the effect on (un-)employment. If the decision to participate in the labour market is also affected by cohort size, the estimated effects on employment and unemployment would be confounded by the effect of cohort size on participation. Moreover, the existence of measurement error in the cohort-size variable among

young age groups, as described further in Section 3.3, may also lead to biased estimates. We therefore focus on individuals aged 25–29 since the estimates for this group will be less susceptible to these problems since, as shown in Figure A1, the share of individuals in education has decreased substantially by that age.

Means and standard deviations of the three dependent variables are shown in the first two columns of Table 1 for the age range 25–29. On average 78% of individuals in a region-year-age group cell are employed compared to 13% that are unemployed. The three remaining columns provide an insight into whether these variables tend to vary most across regions, years or age groups. This is done by regressing each of the dependent variables on a set of dummy variables for two of the aforementioned dimensions and then comparing the adjusted R^2 . Dummies for years and age groups explain only 14% of the variation in the employment share but this value increases considerably once region dummies are included, which suggests that most of the variation in this variable exists between regions. While the explanatory power of the dummy variables is generally lower, the between-region variation also appears to be largest for the unemployment share.

Table 1: Descriptive statistics (employment and unemployment share)

	Mean	Standard deviation	Adjusted R^2 (year, age)	Adjusted R^2 (region, age)	Adjusted R^2 (region, year)
Emp_{jrt}	0.777	0.156	0.136	0.459	0.394
$Unemp_{jrt}$	0.126	0.109	0.063	0.281	0.333

Means and standard deviations are weighted by the weight-adjusted number of individuals per region-year-age group cell.
Adjusted R^2 is derived from a regression of the dependent variables on dummies for the indicated variables; the regression is weighted by the weight-adjusted number of individuals per region-year-age group cell.

The main explanatory variable measures age-specific cohort size which refers to the number of individuals in age group j , region r and year t , N_{jrt} , relative to the size of the population aged between 16 and 65, $N_{16-65, rt}$. While most studies instead use a measure of the youth share, e.g. the relative size of the age group 16–24, we choose a specification that also varies across age to better capture the assumption of imperfect substitutability across age groups which has been posited in theoretical models (Card and Lemieux, 2001). Since it seems overly restrictive to assume that individuals only compete with individuals of the same age, we adopt another specification that has been previously used in this literature (Wright, 1991; Brunello, 2010).⁷ This defines the cohort-size variable as a weighted sum

7 We show in the Supplementary Material that alternative specifications of the cohort-size variable, including unweighted sums across three and five age groups, yield comparable results to those shown in Table 3.

that takes into account the size of the age groups that are up to two years older or younger than the reference group as shown in Equation 1:

$$CS_{jrt} = \frac{(1/9)N_{j-2,rt} + (2/9)N_{j-1,rt} + (3/9)N_{jrt} + (2/9)N_{j+1,rt} + (1/9)N_{j+2,rt}}{N_{16-65,rt}} \quad [1]$$

These quantities are estimated from the EU-SILC dataset by computing the weighted sum of male and female observations in the corresponding region-year-age cells. As they are not available to the labour market, individuals reporting to be either in the military or disabled or unfit to work are omitted but individuals reporting that they are in education are included (the implications of this are discussed in Section 3.3).

The size of an age group in a given region and year is not necessarily exogenous because individuals might react to contemporaneous economic shocks by migrating into regions that offer better economic prospects. If such self-selection takes place, cohort-size would be endogenous to the share of individuals that are (un-)employed and estimation by ordinary least squares (OLS) would yield an inconsistent estimate of the cohort-size effect. We address this issue by employing an IV strategy in which the cohort size of the age group that is fourteen years younger than the reference group as observed fourteen years earlier serves as an instrument. Identification strategies based on time-lagged and age-lagged instruments or, as a special case of the former, birth rates are common in this literature (Korenman and Neumark, 2000; Shimer, 2001; Skans, 2005; Garloff et al., 2013; Moffat and Roth, 2016).⁸ Instruments of this type are appealing because a cohort that was relatively large (small) in the past is likely to remain large (small) in the present despite migration and natural population changes⁹:

$$CS_{Ins_{jrt}} = \frac{(1/9)N_{j-16,r,t-14} + (2/9)N_{j-15,r,t-14} + (3/9)N_{j-14,r,t-14} + (2/9)N_{j-13,r,t-14} + (1/9)N_{j-12,r,t-14}}{N_{2-51,r,t-14}} \quad [2]$$

8 If cohort-size effects are heterogeneous across age, region and/or time, 2SLS estimates a local average treatment effect (LATE) (Imbens and Angrist, 1994). This estimate is the weighted average of the region-year-age cell-specific effects of cohort size with the largest weights attached to cells for which the relationship between the instrument and cohort-size is strongest (Angrist and Imbens, 1995). Since the strength of the relationship between the instrument and cohort-size will be mainly determined by net migration, greater weight will be attached to cells with low levels of net migration. If immigrants are less attractive to employers as a result of having less country-specific human capital (Kim and Park, 2013) than individuals that lived in the region fourteen years ago, this suggests that the LATE will be more positive (more negative) in the employment (unemployment) model than the average treatment effect (ATE). 2SLS estimates may then be larger than OLS estimates of the cohort-size effects if this effect outweighs that of self-selection bias, which would tend to cause OLS to overestimate the positive (negative) effect on employment (unemployment).

9 Further information on the instrument can be found in Moffat and Roth (2016), while the validity of time- and age-lagged instruments is discussed in Garloff and Roth (2016).

Table 2 contains descriptive statistics on the cohort-size variable and its instrument. On average, the five-year weighted sum of an age group in the range 25–29 accounts for about 2% of the population aged between 16 and 65, while the value is slightly smaller in the case of the instrument. For both variables, the larger part of the variation exists between regions.

Table 2: Descriptive statistics (cohort-size variable and instrument)

	Mean	Standard deviation	Adjusted R ² (year, age)	Adjusted R ² (region, age)	Adjusted R ² (region, year)
CS_{jrt}	0.021	0.003	0.073	0.749	0.778
CS_Ins_{jrt}	0.020	0.003	0.080	0.780	0.826

Means and standard deviations are weighted by the weight-adjusted number of individuals per region-year-age group cell.
Adjusted R² is derived from a regression of the dependent variables on dummies for the indicated variables; the regression is weighted by the weight-adjusted number of individuals per region-year-age group cell.

Figures A2 and A3 plot the dependent variables and the cohort-size variable (depicted as the fitted value from a weighted regression on the instrument) across time and age groups, respectively, for the same set of regions as in Figure A1 and thereby illustrate the variation from which cohort-size effects can be identified. Variation over time for given combinations of regions and age groups can be seen in Figure A2; the chosen regions are representative of the larger parts of Europe to which they belong: in Western and Northern Europe (represented by regions BE2 and SE1), the cohort-size profiles are rather flat. In contrast, in region ES5 there is a clear decrease in cohort size over time which affects all age groups – similar profiles can be found in the remaining regions of Spain as well as in Greece and Italy. Finally, different types of profiles can be found in Eastern Europe: on the one hand, the decreasing trend in cohort size in region HU1 resembles the developments in Southern Europe, while on the other hand age groups have increased in size in the Baltic country Latvia. Figure A3 suggests that variation across age groups is less pronounced: older age groups tend to be larger in ES5 and HU1, but the differences become smaller in later years. The profiles in the remaining regions are comparatively flat. At the same time both figures also illustrate the variation in cohort size across regions for given years and age groups. For example, the share of older age groups is larger in regions ES5 and HU1 in earlier years, whereas younger cohorts are relatively big in LVO at the end of the sample period. While the regression analysis in Section 4 makes use of variation across each of these dimensions, in the Appendix we show results that are obtained from a single source of variation.

3.3 Model

According to the theory outlined in the literature review, age-specific labour market outcomes are determined by the supply of age-specific labour. Therefore the effect of cohort size on the outcome variables is modelled as shown in Equation 3 where $share_{jrt}$ represents either the unemployment or employment share, CS_{jrt}^* represents measurement error-free cohort size (i.e. the size of the age cohort that is available to the labour market), x_{jrt} represents a vector of control variables and ε_{jrt} is an error term:

$$share_{jrt} = \alpha + \beta CS_{jrt}^* + x_{jrt}' \gamma + \varepsilon_{jrt} \quad [3]$$

In addition to the problem of regional self-selection that is addressed by IV estimation, there is also a problem of measurement error. This has so far not been addressed in this literature. It arises because of the inclusion of individuals, many of whom will be in education, that are unavailable to the labour market in the cohort-size variable. Moreover, datasets usually do not allow distinguishing individuals that are committed to long-term educational programmes and therefore unavailable to the labour market from individuals in education that would enter the labour market if an attractive opportunity arose (Jones and Riddell, 2006; Moffat and Yoo, 2015). The existence of the latter group means that the alternative approach of excluding those in education from the cohort-size variable would not provide a solution to the measurement-error problem.¹⁰ Formally, the relationship between the observable age-specific cohort-size variable CS_{jrt} and the unobservable measurement error-free variable can be represented as follows:

$$CS_{jrt} = CS_{jrt}^* + u_{jrt} \quad [4]$$

In Equation (4), u_{jrt} is the part of observed cohort size that is not available to the labour market (i.e. the measurement error). Rearranging and substituting Equation (4) into Equation (3) gives:

$$share_{jrt} = \alpha + \beta CS_{jrt} + x_{jrt}' \gamma + \varepsilon_{jrt} - \beta u_{jrt} \quad [5]$$

¹⁰ In the Supplementary Material we provide the regression results from a model in which the numerator of the cohort-size variable is constructed from individuals reporting to be employed or unemployed. For the age group 25–29 the obtained results are very similar to those reported in Table 3. Using younger age groups produces a pattern of cohort-size coefficients which is close to the one in Figure 1 which suggests that exclusion of those reporting to be in education does not remove the problem of measurement error.

If the measurement error is 'classical', there is no correlation between the error-free measure of cohort size and the measurement error and this leads to attenuation of the estimated effect of cohort size. However, empirical evidence suggests that members of large cohorts are less likely to acquire education (Fertig et al., 2009), which suggests the existence of a correlation between the size of an age group CS_{jrt} and u_{jrt} . Arguably, the number of individuals who are available to the labour market is larger in larger age groups and therefore the correlation between the degree of measurement error and the observable cohort size also carries over to the latent variable CS_{jrt}^* , which measures the size of an age group that is available to the labour market. In this 'non-classical' case, it is not possible to state a priori the direction of bias since it will be dependent on the relative variances of CS_{jrt}^* and u_{jrt} , the size of the covariance of CS_{jrt}^* and u_{jrt} and the partial correlations between the measurement error and the dummy variables in the model (Bound et al., 2001).

A second reason for the existence of non-classical measurement error is given by the current demographic processes, as a result of which younger age groups tend to be smaller than older ones in a given region and year (support for this hypothesis is provided in the Supplementary Material). Moreover, given the assumption that the share of non-participants is larger in younger age groups – for which the substantially larger education shares in younger age groups provide some evidence – it is possible for the latent cohort-size variable and the degree of measurement error to be negatively correlated across age groups. This will be the case as long as the ratio of the non-participation share in younger and older groups exceeds the ratio of the size of older and younger groups (details on this argument are provided in the Supplementary Material).

While two-stage least squares (2SLS) estimation is one approach to tackling measurement error (Hausman, 2001), the instrument which is standard in the literature does not purge the correlation with u_{jrt} . The instrument is based on the size of the same cohort observed at an earlier point in time and since an age group that is relatively large in the present can be expected to have also been relatively large in the past, the instrument would also be correlated with the degree of measurement error. As a result, 2SLS will not provide a consistent estimate of the cohort-size effect.

For the sample of individuals aged 25–29, the empirical analysis is based on 1,959 region-year-age cells¹¹. Two specifications of Equation 5 are estimated for each of the outcome variables. Analogously to the use of control variables in Shimer (2001), in the baseline specification vector x_{jrt} only contains a constant

11 In principle, 5 age groups (25–29) are observed in 49 regions for 8 years (2005–2012), but since there are no observations for age group 26 in region FR1 and year 2010 in the sample, the total number of observations is reduced by one.

and three sets of dummy variables for each of the three dimensions of the cohort-size variable: regions, years and age groups. In the second specification a set of control variables is added to the model (definitions and summary statistics are given in Table A1 in the Appendix). One part of these variables is assumed to affect the (un-)employment probability at the individual level and has therefore been aggregated in order to control for compositional differences between region-year-age cells. They include the share of individuals in such cells that a) belong to different educational groups according to the International Standard Classification of Education (ISCED), b) are married and c) reside in areas that differ with respect to their degree of urbanisation. Moreover, we add the level of regional GDP. While the use of year dummies accounts for shocks that are common to all region-age cells, this variable is useful in order to control for the region-specific economic environment in a given year. The inclusion of regional GDP therefore helps to avoid the estimated cohort-size effects being confounded by regional economic shocks.

4 Results

Figure 1 shows the estimated coefficients and confidence intervals on the cohort-size variable using overlapping samples of differently aged individuals when the dependent variable is the unemployment and employment share, respectively. For both outcome variables, the effect of cohort size varies substantially across age groups. When the dependent variable is the unemployment share, the effects are positive and statistically significant for individuals aged 18–22 but are negative and statistically significant for older groups. The effect appears to converge to between -10 and -20 for the older groups. The shift in sign and magnitude of the coefficients coincides with a decrease in the share of individuals reporting to be in education (see Figure A1 in the Appendix). In the employment model, cohort-size effects are significant and negative for individuals aged 18–22 but positive and significant for older age groups, converging to a value of approximately 25.

The results for the younger age groups appear to be supportive of the cohort-crowding hypothesis. However, our view is that the estimated effects for younger age groups cannot be regarded as a direct test of this hypothesis since they capture both the effect of cohort size on labour-market participation and the effect on (un-)employment. For example, the finding that cohort size reduces the employment share of individuals aged 18–22 may indicate either that large cohorts lead young individuals to acquire education and thereby defer entry to the labour market or that young individuals in the labour market are disadvantaged by belonging to a large age group. In addition to this problem of interpretation, the change in the coefficients may be driven by measurement error in the cohort-size

variable. As discussed above, this variable is supposed to measure the availability of similarly aged individuals on the labour market, but in light of the large share of young individuals in education, some of whom will be committed to long-term programmes, it is less suitable as a measure of labour-market availability in younger than in older groups.

Figure 1: Cohort-size coefficients for different age groups



Source: EU-SILC (authors' calculations). Coefficients are obtained from weighted 2SLS estimation of a model containing dummy variables for regions, years and age groups. Robust standard errors are used.

In order to mitigate this problem, the remainder of this section focuses on individuals aged 25–29. As can be seen from Figure A1, the share of individuals in education is considerably smaller for those age groups. In this age range, the cohort-size variable should therefore present a better measure of the degree of labour-market crowding, while any confounding effects resulting from the preceding decision to enter the labour market or to acquire further education will be less relevant. Table 3 contains OLS and 2SLS estimation results for each of the two specifications discussed in Section 3.3 using a sample of individuals aged 25–29 (full results including the coefficients of the control variables can be found in Tables A2 and A3 in the Appendix and the results of the first-stage regressions are shown in Table A4).

Table 3: OLS and 2SLS regression results

Panel A: Unemployment share				
	OLS	2SLS	OLS	2SLS
Cohort size	-10.32*** (1.70)	-17.30*** (2.10)	-7.98*** (1.73)	-15.06*** (2.05)
<i>Dummies</i>				
Region	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes
<i>Observations</i>				
Region-year-age cells	1,959	1,959	1,959	1,959
R ²	0.38	0.37	0.41	0.40
F-stat	–	1,540.67***	–	1,642.59***
ME(std)	-0.03***	-0.05***	-0.02***	-0.05***
Panel B: Employment share				
	OLS	2SLS	OLS	2SLS
Cohort size	14.39*** (2.03)	24.32*** (2.64)	11.91*** (2.02)	22.07*** (2.52)
<i>Dummies</i>				
Region	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes
<i>Observations</i>				
Region-year-age cells	1,959	1,959	1,959	1,959
R ²	0.53	0.52	0.56	0.55
F-stat	–	1,540.67***	–	1,642.59***
ME(std)	0.04***	0.08***	0.04***	0.07***
***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.				

The first two columns of panel A show that in the baseline model an increase in cohort size is predicted to decrease the share of individuals in the corresponding age group that are unemployed. OLS and 2SLS estimates have the same sign and are statistically significant at the 1% level. The finding that the latter are larger (in absolute terms) was also obtained by Shimer (2001) in some specifications and is consistent with the argument (see footnote 8) that cohort-size effects are heterogeneous across region-year-age cells and that immigrants are less attractive to employers than individuals that have lived in the region for 14 years. The third

and fourth columns show that when the set of control variables, described in Section 3.3, are added to the model, the cohort-size coefficients decrease somewhat in magnitude. To give a better impression of the size of the coefficients, marginal effects for changes in cohort size of one standard deviation are shown at the bottom of panel A. Such an increase is predicted to reduce the share of unemployed in an age group by 5 percentage points, which is a sizeable effect given that the average unemployment share is 13% (see Table 1). Finally, the size of the F-statistics suggests that the excluded instrument has predictive power for the endogenous cohort-size variable with values considerably larger than the threshold value of 10 (Staiger and Stock, 1997). The results for the employment model are shown in panel B. The cohort-size variable is found to have a statistically significant and positive effect on the employment share. Adding control variables slightly reduces the size of the coefficients. For 2SLS estimation, an increase in cohort size by one standard deviation is predicted to increase the employment share by between 7 and 8 percentage points. In light of an average employment share of 77% this change is comparatively smaller than the corresponding effect on the unemployment share.

As discussed in Section 3.2, the above results use variation across regions, years and age groups. Table A5 shows cohort-size coefficients that are obtained when the identifying variation is restricted to a single source. This is accomplished by adding dummy variables for interactions between regions and age groups (identification is based on variation over time), between years and age groups (variation across regions only) or between regions and years (variation across age groups only). Except for an increase in the marginal effect of cohort-size on the unemployment share when only variation over time is used, the key results are not materially affected in the first two cases. By contrast, the cohort-size variable is not statistically significant in the unemployment model when region-year dummies are included. This is unsurprising since there is relatively little variation in cohort size across age within the sample. The results of various sensitivity analyses are available in the Supplementary Material.

The signs of the estimated coefficients suggest that members of large cohorts do not fare worse in terms of unemployment and employment outcomes. As such the results of this paper contradict the cohort-crowding hypothesis that increases in the size of an age group lead to increased unemployment within that group. Our findings rather provide evidence in support of Shimer (2001) that young individuals benefit from being part of large cohorts. However, even though increases in cohort size are found to increase the share of employed individuals in the corresponding age group, these results do not provide any evidence regarding the type and conditions of employment. Indeed results by Moffat and Roth (2016) that are also based on EU-SILC data show that individuals with completed secondary education

command lower wages when they are part of a larger cohort. Similarly, using German microdata Garloff and Roth (2016) find that an increase in the share of youths in the population reduces young workers' wages; moreover, their analysis provides evidence that belonging to a larger youth cohort increases the likelihood of being employed in occupations and industries that pay lower wages.

5 Conclusion

A prominent research question of the cohort-size literature concerns the effect that the size of an age group has on its members' employment and unemployment outcomes. Based on the assumption of imperfect substitutability of differently aged workers, these outcomes should be determined by the size of an age group that is available to the labour market. As this quantity is typically not observable, the common approach has been to use the size of an age group as a proxy for age-specific labour supply instead. However, this ignores the fact that among the young the size of an age group will only be a poor measure of the size of the group that is available to the labour market because of the large share of individuals who participate in education.

This gives rise to two problems. First, for young age groups the estimated effect of cohort size on (un-)employment will be confounded by the former's effect on the decision to participate in the labour market in the first place. Second, using the size of an age group induces a problem of measurement error that the standard IV approach is unable to solve. For these reasons, the standard identification strategy is unsuited to produce informative insights into the effects of cohort crowding for young age groups regardless of whether an age-specific cohort-size variable is used that also varies across age or, as in other papers, a youth-share variable is employed.

To illustrate this, we estimate the effect of cohort size on age-specific employment and unemployment outcomes using data comprising information on 49 regions covering the period 2005–2012. In a first step we show that the estimated effects of cohort size are indeed highly sensitive to the chosen age range. In particular, we find that the sign of the coefficient changes as successively younger age groups are used. In a second step we apply these models to the age group 25–29 for which the above-mentioned problems should be less of a concern because participation rates in education are considerably lower. The results of this analysis suggest that an increase in cohort size reduces the unemployment share in an age group and increases the employment share, which is consistent with the mechanism between the youth share and (un-)employment outcomes that is described in Shimer (2001).

Acknowledgements

The authors would like to thank Georgios Marios Chrysanthou, Bernd Hayo and Michael Kirk as well as the participants of IAB's regional research network meeting, the 63rd Annual Meeting of the French Economic Association (AFSE), the 7th Summer Conference in Regional Science, the 4th World Meeting of the Society of Labor Economists (SOLE) and the European Association of Labour Economists (EALE) and the 2014 Congress of the European Regional Science Association (ERSA) for valuable comments. This paper uses data from the European Union Statistics on Income and Living Conditions (EU-SILC). The results and conclusions are those of the authors and not those of Eurostat, the European Commission or any of the national statistical authorities whose data have been used.

References

- Angrist, J. and Imbens, G. (1995) Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity, *Journal of the American Statistical Association*, **90**, 431–42.
- Berger, M. and Schaffner, S. (2015) A Note on How to Realize the Full Potential of the EU-SILC Data, ZEW Discussion Paper 15-005, Centre for European Economic Research, Mannheim.
- Biagi, F. and Lucifora, C. (2008) Demographic and education effects on unemployment in Europe, *Labour Economics*, **15**, 1076–101.
- Bound, J., Brown, C. and Mathiowetz, N. (2001) Measurement Error in Survey Data, in J. Heckman and E. Leamer (eds.) *Handbook of Econometrics*, North-Holland, Amsterdam.
- Brunello, G. (2010) The effects of cohort size on European earnings, *Journal of Population Economics*, **23**, 273–90.
- Card, D. and Lemieux, T. (2001) Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis, *Quarterly Journal of Economics*, **116**, 705–46.
- Easterlin, R. (1961) The American baby boom in historical perspective, *American Economic Review*, **51**, 869–911.
- Fertig, M., Schmidt, C. and Sinning, M. (2009) The Impact of Demographic Change on Human Capital Accumulation, *Labour Economics*, **16**, 659–68.
- Foote, C. (2007) Space and time in macroeconomic panel data: young workers and state-level unemployment revisited, Research Department Working Paper No. 07–10, Federal Reserve Bank of Boston, Boston.
- Garloff, A., Pohl, C. and Schanne, N. (2013) Do small labor market entry cohorts reduce unemployment?, *Demographic Research*, **29**, 379–406.

- Garloff, A. and Roth, D. (2016) Regional population structure and young workers' wages, IAB Discussion Paper No. 06/2016, Institute for Employment Research, Nuremberg.
- Hausman, J. (2001) Mismeasured Variables in Econometric Analysis: Problems from the Right and Problems from the Left, *Journal of Economic Perspectives*, 15, 57–67.
- Imbens, G. and Angrist, J. (1994) Identification and Estimation of Local Average Treatment Effects, *Econometrica*, 62, 467–75.
- Jones, S. and Riddell, C. (2006) Unemployment and Nonemployment: Heterogeneities in Labor Market States, *Review of Economics and Statistics*, 88, 314–23.
- Kim, J. and Park, J. (2013) Foreign Direct Investment and Country-Specific Human Capital, *Economic Inquiry*, 51, 198–210.
- Korenman, S. and Neumark, D. (2000) Cohort Crowding and Youth Labor Markets: A Cross-National Analysis, in D.G. Blanchflower and R.B. Freeman (eds.) *Youth Employment and Joblessness in Advanced Countries*, University of Chicago Press, Chicago.
- Michaelis, J. and Debus, M. (2011) Wage and (un-)employment effects of an ageing workforce, *Journal of Population Economics*, 24, 1493–511.
- Moffat, J. and Roth, D. (2016) The Cohort Size-Wage Relationship in Europe, *Labour*, forthcoming.
- Moffat, J. and Yoo, H.I. (2015) Who are the unemployed? Evidence from the United Kingdom, *Economics Letters*, 132, 61–4.
- Schmidt, C. (1993) Ageing and unemployment, in P. Johnson and K. Zimmermann (eds.) *Labour markets in an ageing Europe*, Cambridge University Press, Cambridge.
- Shimer, R. (2001) The Impact of Young Workers on the Aggregate Labor Market, *The Quarterly Journal of Economics*, 116, 969–1007.
- Skans, O.N. (2005) Age effects in Swedish local labor markets, *Economics Letters*, 86, 419–26.
- Staiger, D. and Stock, J.H. (1997) Instrumental Variables Regression with Weak Instruments, *Econometrica*, 65, 557–86.
- United Nations (2015) World Population Prospects: The 2015 Revision, Key Findings and Advance Tables, Working Paper No. ESA/P/WP.241, Department of Economics and Social Affairs, Population Division, New York.
- Welch, F. (1979) Effects of Cohort Size on Earnings: The Baby Boom Babies' Financial Bust, *Journal of Political Economy*, 87, S65–S97.
- Wright, R.E. (1991) Cohort size and earnings in Great Britain, *Journal of Population Economics*, 4, 295–305.
- Zimmermann, K. (1991) Ageing and the Labour Market. Age structure, cohort size and unemployment, *Journal of Population Economics*, 4, 177–200.

Appendix

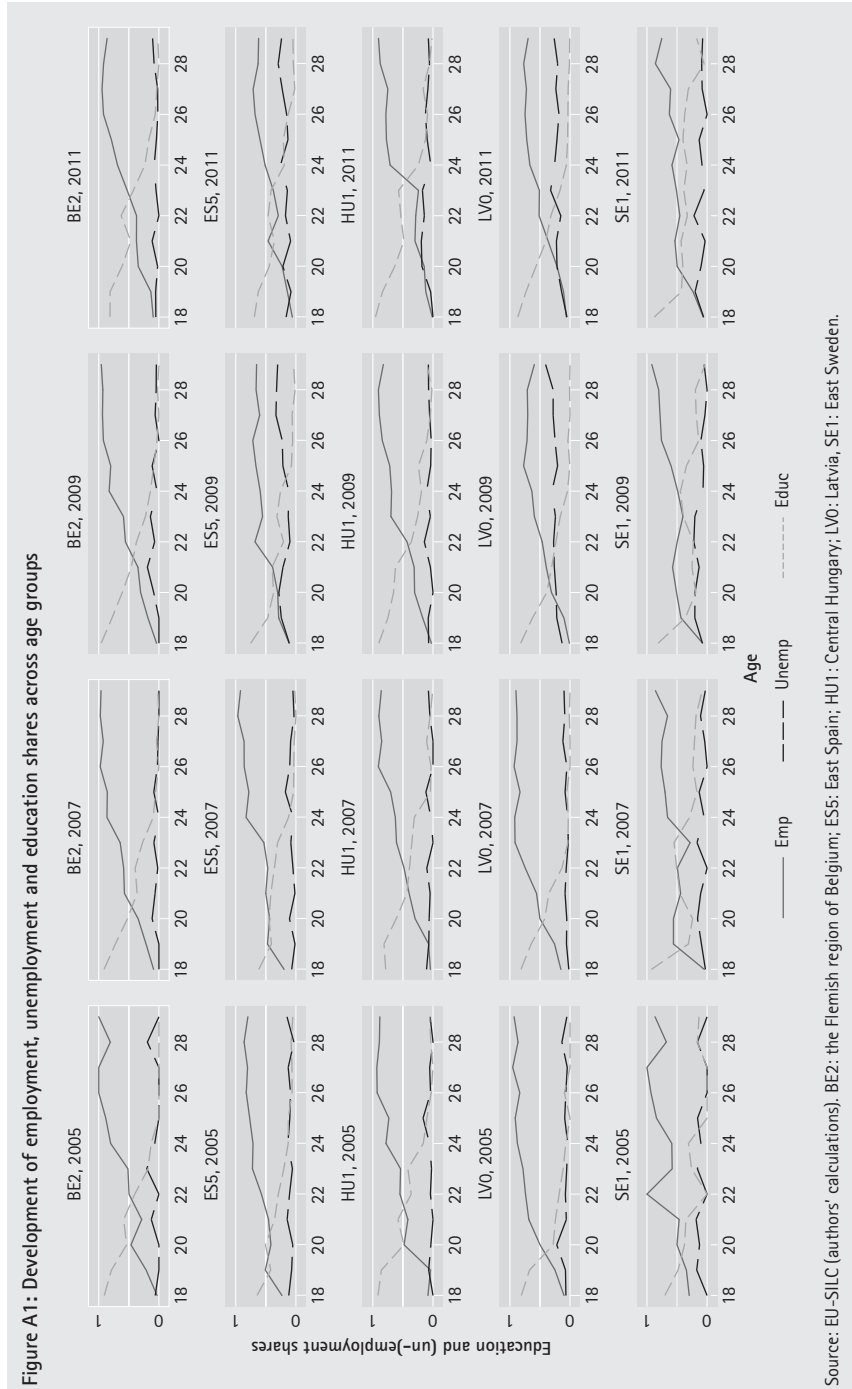



Figure A2: Development of employment and unemployment shares and of fitted cohort-size variable over time



Source: EU-SILC (authors' calculations). BE2: the Flemish region of Belgium; ESS: East Spain; HU1: Central Hungary; LVO: Latvia; SE1: East Sweden.

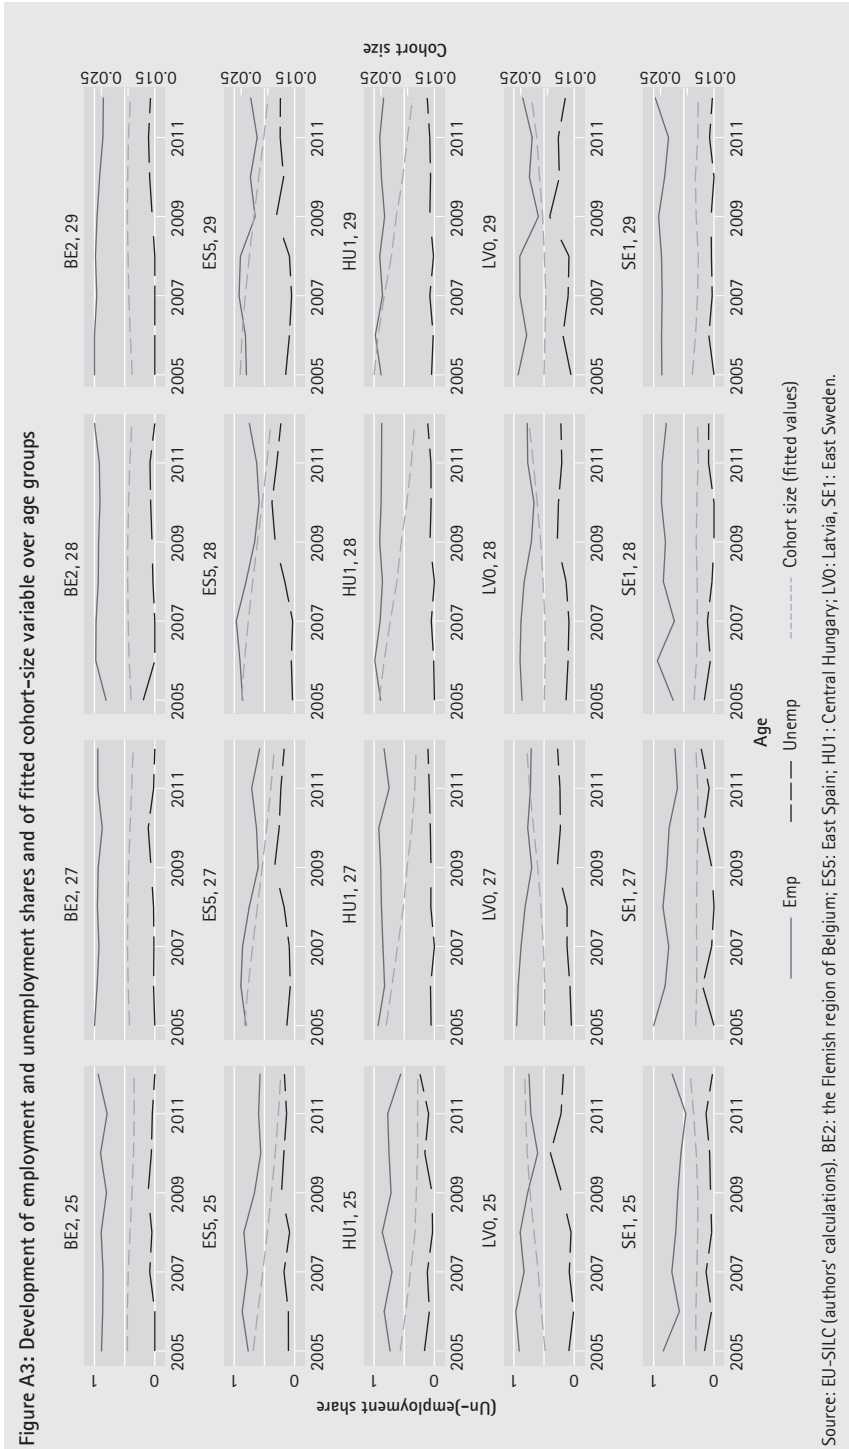


Table A1: Definitions and descriptive statistics of control variables

Name	Definition	Source	Mean	Standard deviation
ISCED_0	Share of individuals in region-year-age cell with pre-primary education	EU-SILC	0.006	0.027
ISCED_1	Share of individuals in region-year-age cell with primary education	EU-SILC	0.040	0.060
ISCED_2	Share of individuals in region-year-age cell with lower secondary education	EU-SILC	0.136	0.133
ISCED_3	Share of individuals in region-year-age cell with upper secondary education	EU-SILC	0.479	0.187
ISCED_4	Share of individuals in region-year-age cell with post-secondary, non-tertiary education	EU-SILC	0.035	0.052
ISCED_5	Share of individuals in region-year-age cell with tertiary education (also includes category ISCED_6, i.e. individuals with second stage of tertiary education)	EU-SILC	0.304	0.168
Married	Share of individuals in region-year-age cell that are married	EU-SILC	0.195	0.153
Urban_1	Share of individuals in region-year-age cell living in densely populated areas (an area with a population density of more than 500 inhabitants per square kilometre (km) and a population of at least 50,000 inhabitants)	EU-SILC	0.461	0.216
Urban_2	Share of individuals in region-year-age cell living in intermediately populated areas (an area with a population density of more than 100 inhabitants per square km and either a population of at least 50,000 inhabitants or adjacent to a 'densely populated' area)	EU-SILC	0.248	0.170
Urban_3	Share of individuals in region-year-age cell living in thinly populated areas (an area with fewer than 100 inhabitants per square km and a population of less than 50,000 inhabitants)	EU-SILC	0.291	0.222
GDP	Gross domestic product at the NUTS1 level (in billion Euros, adjusted for purchasing-power-parity)	Eurostat	188.391	127.737

Means and standard deviations are weighted by the weight-adjusted number of individuals per region-year-age group cell.

Table A2: Full OLS and 2SLS regression results (Unemployment share)

<i>Unemployment share</i>	OLS	2SLS	OLS	2SLS
Cohort size	-10.32*** (1.70)	-17.30*** (2.10)	-7.98*** (1.73)	-15.06*** (2.05)
<i>Dummies</i>				
Region	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes
<i>Control variables</i>				
ISCED_1	-	-	0.07 (0.14)	0.08 (0.14)
ISCED_2	-	-	0.06 (0.13)	0.06 (0.13)
ISCED_3	-	-	-0.01 (0.12)	-0.01 (0.12)
ISCED_4	-	-	-0.11 (0.13)	-0.09 (0.13)
ISCED_5	-	-	-0.08 (0.13)	-0.08 (0.13)
Married	-	-	-0.10*** (0.02)	-0.09*** (0.02)
Urban_2	-	-	-0.03 (0.03)	-0.03 (0.03)
Urban_3	-	-	-0.03 (0.03)	-0.03 (0.03)
GDP	-	-	-0.00*** (0.00)	-0.00*** (0.00)
<i>Observations</i>				
Region-year-age cells	1,959	1,959	1,959	1,959
R ²	0.38	0.37	0.41	0.40
F-stat	-	1,540.67***	-	1,642.59***
ME(std)	-0.03***	-0.05***	-0.02***	-0.05***
<p>***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.</p>				

Table A3: Full OLS and 2SLS regression results (Employment share)

<i>Employment share</i>	OLS	2SLS	OLS	2SLS
Cohort size	14.39*** (2.03)	24.32*** (2.64)	11.91*** (2.02)	22.07*** (2.52)
<i>Dummies</i>				
Region	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes
<i>Control variables</i>				
ISCED_1	–	–	0.47** (0.20)	0.46** (0.20)
ISCED_2	–	–	0.51** (0.20)	0.51** (0.20)
ISCED_3	–	–	0.57*** (0.19)	0.58*** (0.19)
ISCED_4	–	–	0.66*** (0.20)	0.64*** (0.20)
ISCED_5	–	–	0.62*** (0.19)	0.62*** (0.19)
Married	–	–	0.09*** (0.03)	0.09** (0.03)
Urban_2	–	–	0.06* (0.03)	0.06* (0.03)
Urban_3	–	–	0.07* (0.04)	0.07* (0.04)
GDP	–	–	0.00*** (0.00)	0.00*** (0.00)
<i>Observations</i>				
Region-year-age cells	1,959	1,959	1,959	1,959
R ²	0.53	0.52	0.56	0.55
F-stat	–	1,540.67***	–	1,642.59***
ME(std)	0.04***	0.08***	0.04***	0.07***
<p>***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.</p>				

Table A4: First-stage regression results

	Unemployment share		Employment share	
Instrument	0.93*** (0.02)	0.93*** (0.02)	0.93*** (0.02)	0.93*** (0.02)
<i>Dummies</i>				
Region	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	Yes	No	Yes
<i>Observations</i>				
Region-year-age cells	1,959	1,959	1,959	1,959
R ²	0.92	0.92	0.92	0.92
F-stat	1,540.67***	1,642.59***	1,540.67***	1,642.59***
<p>***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. <i>F-stat</i> represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables.</p>				

Table A5: OLS and 2SLS results

<i>Panel A:</i>						
<i>Unemployment share</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-12.84*** (1.91)	-23.01*** (2.37)	-10.76*** (1.67)	-17.35*** (2.07)	-2.16 (2.15)	-1.46 (2.64)
<i>Dummies</i>						
Region	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	Yes	Yes	No	No	No	No
Year-by-age	No	No	Yes	Yes	No	No
Region-by-year	No	No	No	No	Yes	Yes
<i>Control variables</i>	No	No	No	No	No	No
<i>Observations</i>						
Region-year-age cells	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.43	0.41	0.39	0.38	0.57	0.57
F-stat	–	1,140.11***	–	1,582.57***	–	674.72***
ME(std)	-0.04***	-0.07***	-0.03***	-0.05***	-0.01	-0.00
<i>Panel B:</i>						
<i>Employment share</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	13.88*** (2.24)	26.55*** (2.89)	14.41*** (2.02)	24.40*** (2.59)	7.24*** (2.73)	11.15*** (3.37)
<i>Dummies</i>						
Region	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	Yes	Yes	No	No	No	No
Year-by-age	No	No	Yes	Yes	No	No
Region-by-year	No	No	No	No	Yes	Yes
<i>Control variables</i>	No	No	No	No	No	No
<i>Observations</i>						
Region-year-age cells	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.58	0.57	0.54	0.53	0.66	0.66
F-stat	–	1,140.11***	–	1,582.57	–	674.72***
ME(std)	0.04***	0.08***	0.04***	0.08***	0.02***	0.03***

***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.

Supplementary material

S1 Selection of the age range and measurement error

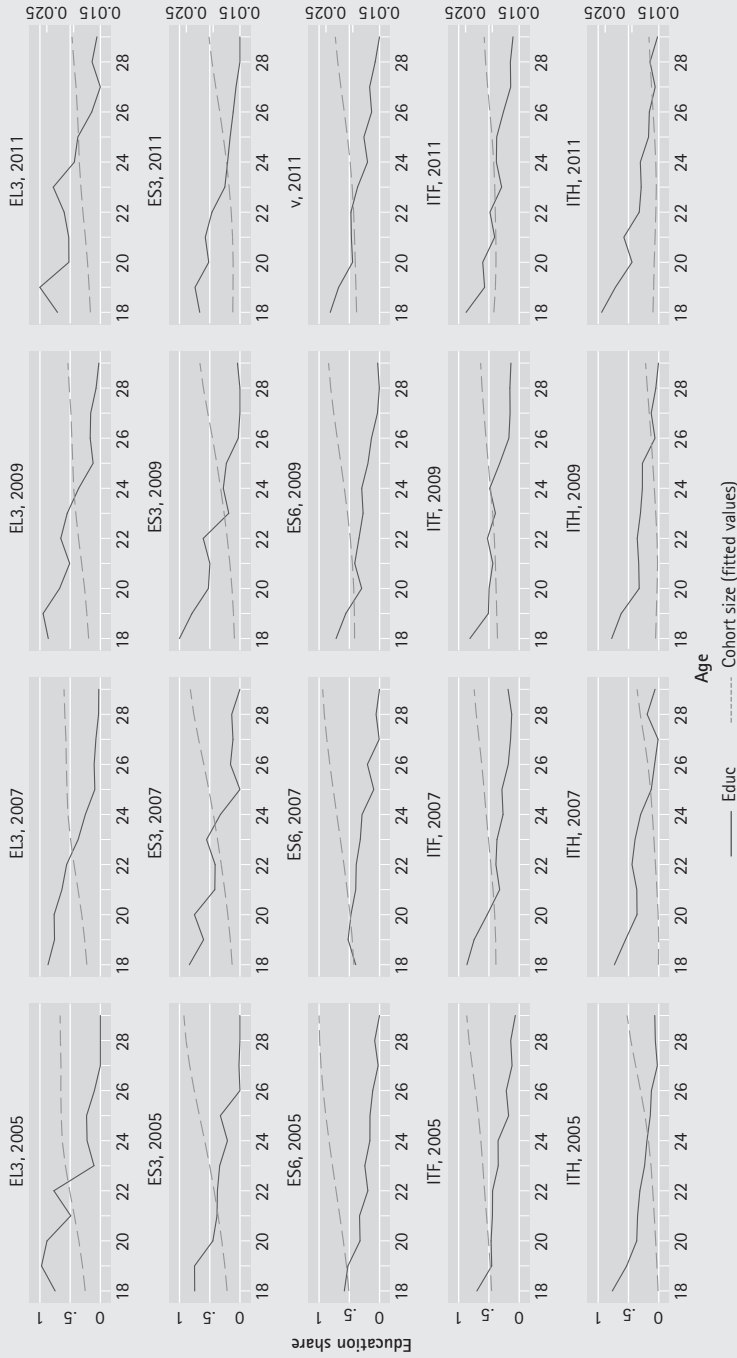
The paper's main finding is that the estimated effect of cohort size on the (un-)employment share is sensitive to the selected age range of the sample (see Figure 1 in the paper). We propose two explanations for the observed pattern of the coefficients and in both cases the core of the argument is that for young age groups the cohort-size variable can be a poor measure of the age-specific supply of labour: first, a population-based cohort-size variable will include a substantial number of individuals that are not on the labour market, primarily because they are acquiring education; second, given the large share of non-participants among young age groups the estimated effect of cohort-size on the (un-)employment share will be confounded by the former's effect on the decision to participate in the labour market. In the following, we provide further detail on the former point.

Figures S1 and S2 plot the share of individuals reporting to be in education against age for different region-year combinations.¹² As can be seen, the education share can be close to 100% at age 18 and usually is in excess of 50% at age 20, whereas the share is considerably smaller in the age range 25–29, which is used in the empirical analysis of this paper.¹³ This observation provides support for the hypothesis that the share of individuals that are included in a population-based cohort-size variable but that are not on the labour market can be substantial, especially among young age groups. However, it is important to note that simply excluding those individuals that report to be in education from the construction of the cohort-size variable does not necessarily lead to a better measure of age-specific labour supply. First, a part of the group of individuals reporting to be in education may be enticed to enter the labour market depending on the conditions of employment and as such should be treated as being available to the labour market, whereas participants in lengthy degree programmes are less likely to do so (these groups cannot be separated in the data); second, switching between periods of participation and non-participation is more likely to occur among young individuals compared to older age groups whose members tend to be more established in the labour market.

12 The regions are EL3 (Attica), ES3 (Madrid), ES6 (Andalusia), ITF (Southern Italy) and ITH (Northeast Italy), CZ0 (Czech Republic), DKO (Denmark), FR1 (Île de France), LT0 (Lithuania) and PL1 (Central Poland).

13 The main exception is Denmark where the education share takes longer to decrease and can be large at later ages (e.g. age 26 in the year 2011). However, we are able to show in Figures S3 and S4 that the exclusion of Denmark from the sample has virtually no effect on the size of the coefficient in the unemployment and the employment model, respectively, while allowing the sample to start at age 26 instead of 25 also yields comparable coefficients in both models (see Figure S6).

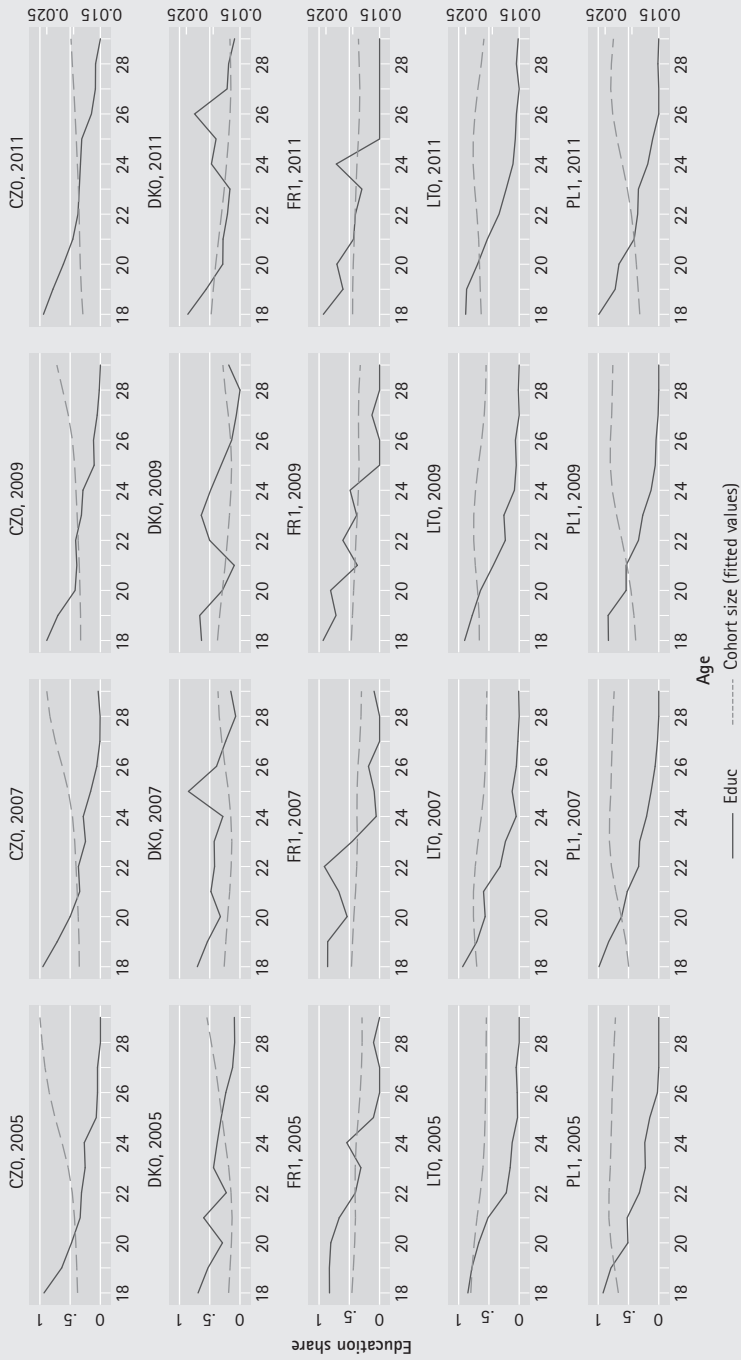
Figure S1: Development of education share and fitted cohort-size variable (set 1)



Source: EU-SILC (authors' calculations).

The case of non-classical measurement error arises when the degree of measurement error is correlated with the measurement error-free cohort-size variable which in this case is given by the size of an age group that is also available to the labour market. One reason why such a correlation might arise is that the degree of measurement error is larger in younger age groups (as shown in Figures S1 and S2, the share of individuals in education is considerably higher among younger age groups), while at the same time younger age groups tend to be smaller than older ones (for given regions and years). This implies a (negative) correlation between observed cohort size and the degree of measurement error.

Figure S2: Development of education share and fitted cohort-size variable (set 2)



Source: EU-SILC (authors' calculations).

This hypothesis is supported by Figures S1 and S2 which also show the development of the fitted value of the cohort-size variable (obtained from a regression on the instrument). From Figure S1 it can be seen that in the Southern European regions of Spain (ES3, ES6), Italy (ITF, ITH) and Greece (EL3) younger cohorts are indeed smaller than older ones, especially in earlier years. Figure S2 illustrates that similar patterns can be found in the Czech Republic (CZO) and Central Poland (PL1). In contrast, younger cohorts are larger than older ones in Lithuania (LTO). The profiles of most Western European regions tend to be flat, as exemplified by the Île de France (FR1); an exception is given by Denmark (DK0) where older age groups also tend to belong to larger cohorts than younger ones.

It is argued in the paper that under certain conditions there will be a negative correlation between the latent cohort-size variable, which measures age-specific labour supply, and the degree of measurement error. According to Equation 4 in the paper the observed cohort-size variable can be expressed as the sum of age-specific labour supply and measurement error:

$$CS_{jrt} = CS_{jrt}^* + u_{jrt} \quad [S1]$$

This condition can be re-written in form of the size of the age-group j in region r at time t , N_{jrt} , the number of individuals in that age group that are available to the labour market, N_{jrt}^* , those that are not available, N_{jrt}^{out} , and the overall population, N_{rt} :

$$\frac{N_{jrt}}{N_{rt}} = \frac{N_{jrt}^*}{N_{rt}} + \frac{N_{jrt}^{out}}{N_{rt}} \quad [S2]$$

The degree of measurement error can be expressed in terms of the share of non-participants, N_{jrt}^{out} , in an age group, N_{jrt} :

$$\alpha_{jrt} = \frac{N_{jrt}^{out}}{N_{jrt}} \quad [S3]$$

Since for a given region and year, the denominators are identical for different age groups, it is sufficient to focus on the numerators. There will be a negative correlation between the latent cohort-size variable and the degree of measurement error across age groups, if the number of participants, N_{jrt}^* , increases in older age groups while the number of non-participants, N_{jrt}^{out} , becomes smaller. This can be formalised in terms of two age groups k and l ($k < l$):

$$N_{krt}^* < N_{lrt}^* \quad [S4]$$

$$N_{krt}^{out} > N_{lrt}^{out} \quad [S5]$$

Substituting Equation S3 into S5 and re-formulating yields the condition that the ratio of the non-participation shares in younger and older age groups exceeds the ratio of the size of the older and the younger age group (since older age groups are typically larger than younger ones, the condition in Equation S4 will hold if condition S5 is satisfied):

$$\frac{\alpha_{krt}}{\alpha_{lrt}} > \frac{N_{lrt}}{N_{krt}} \quad [S6]$$

If the size of the education share is used as a proxy for the degree of measurement error, Figures S1 and S2 suggest that the above condition is not unreasonable since the difference in cohort size between age groups often appears less pronounced than the difference between education shares.


S2 Robustness of the empirical results

This section addresses the robustness of the estimated cohort-size coefficients to a variety of changes in the empirical model and in the underlying sample.

S2.1 Robustness to the exclusion of individual regions, year and age groups

This part starts by assessing the sensitivity of the results to dropping individual regions, years and age groups. The cohort-size coefficients and their 95% confidence interval that are estimated from the reduced sample using an analogue of the specification that includes region, year and age dummies are shown in Figures S3 to S6. For better comparability these figures also contain the cohort-size coefficient and confidence interval from the full sample.

Figure S3: Exclusion of single regions (Unemployment share)

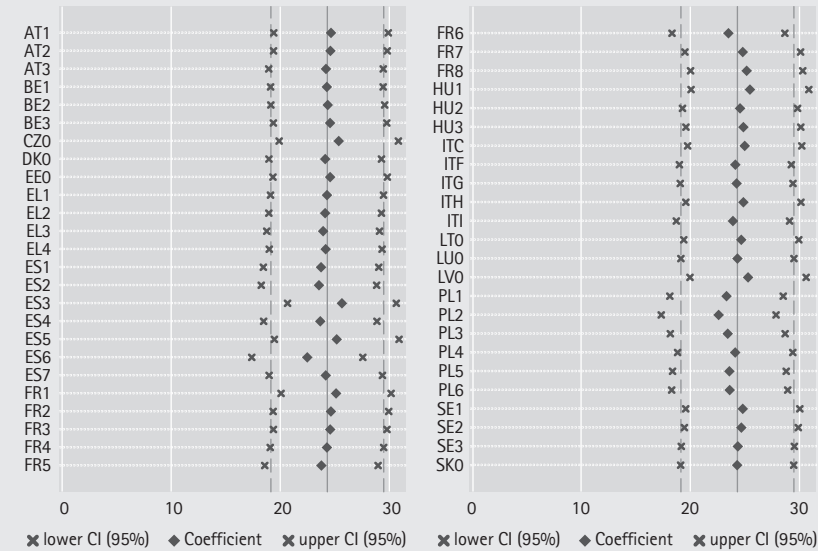


Source: EU-SILC (authors' calculations). Cohort-size coefficients are estimated as described in Section 3; the estimated model also includes region, year and age dummies; the solid line represents the cohort-size coefficients from the full model, the dashed lines the corresponding 95% confidence interval.

As illustrated by Figure S3, for most regions it is the case that their exclusion does not have a large effect on the cohort-size coefficient as can be seen by the former's closeness to the solid line. Some regions, however, do affect the size of the coefficient if they are excluded: in the unemployment model dropping the Czech Republic (CZO) or Latvia (LVO) increases the magnitude of the coefficient, while exclusion of the Spanish region Andalusia (ES6) or the Polish regions PL1–PL3 leads to a decrease. The resulting estimates do, however, remain well within the 95% confidence interval of the full sample's cohort-size coefficient (given by the dashed lines). Those regions that, when excluded, decrease or increase the magnitude of the cohort-size coefficient tend to have the same effect in the employment model, while there are also some additional regions that now have a larger effect on the size of the coefficient (ES3, ES5, FR8), as shown in Figure S4. As with the unemployment share, the estimates always lie within the confidence interval of the full sample's coefficient. An increase in the magnitude of the cohort-size coefficient implies that in the specific sub-sample labour-market shares are more responsive to changes in cohort size: the decreasing effect on the unemployment share as well as the increasing effect on the employment share both become larger – and vice-versa for a decrease in the magnitude of the coefficient. However, when interpreting the

change in the coefficients it should be borne in mind that omission of a certain region (or year or age group) will also have an effect on the distribution of the cohort-size variable in the sample. The effect of an increase (decrease) in the coefficient's magnitude can be mitigated if the change in the underlying sample reduces (increases) the standard deviation of the cohort-size variable.

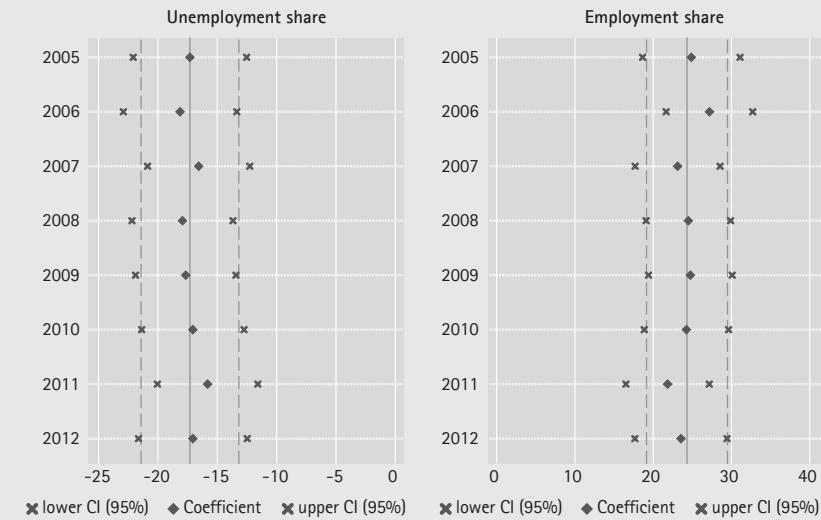
Figure S4: Exclusion of single regions (Employment share)



Source:EU-SILC (authors' calculations). Cohort-size coefficients are estimated as described in Section 3; the estimated model also includes region, year and age dummies; the solid line represents the cohort-size coefficients from the full model, the dashed lines the corresponding 95% confidence interval.

Figure S5 provides an overview of the effect that the exclusion of individual years has on the estimated cohort-size coefficients. While there are changes in the estimates in some cases, the former always remain within the confidence interval of the full sample's coefficients. Comparing the unemployment and the employment model, the coefficients appear to change in a symmetric manner, e.g. omission of the year 2006 increases the magnitude of the coefficients in both model, while dropping observations from the year 2011 leads to a decrease in size.

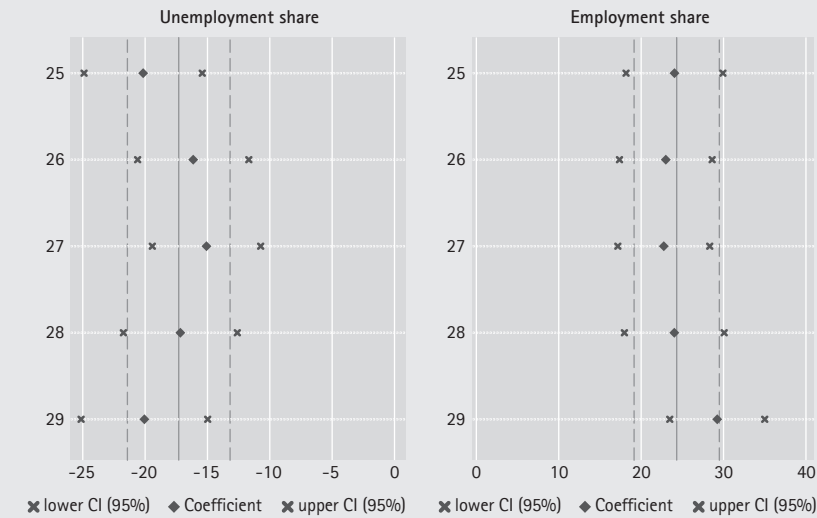
Figure S5: Exclusion of single years



Source: EU-SILC (authors' calculations). Cohort-size coefficients are estimated as described in Section 3; the estimated model also includes region, year and age dummies; the solid line represents the cohort-size coefficients from the full model, the dashed lines the corresponding 95% confidence interval.

The effects of excluding individual age groups from the sample are illustrated in Figure S6. The largest change in the coefficient can be observed when the age group 29 is dropped, in which case the magnitude of the coefficient increases in both models and almost moves outside of the full sample's confidence interval in the employment model. The responsiveness of labour-market shares to changes in cohort size therefore appears less pronounced for this age group. Unfortunately, the unavailability of lagged population data prevents the inclusion of older age groups in the sample and thus the possibility to check whether a further decrease in the strength of the relationship between cohort size and labour-market shares could be found at older ages. Such a development would be in line with the underlying mechanism that is proposed by Shimer (2001): firms create vacancies in areas where the share of young individuals is large because the former are usually not well matched to their jobs and a large pool of such individuals makes it easier for firms to find good matches for these vacancies. However, if the degree to which individuals are matched to their job increases with age, larger older age groups would not necessarily induce the same reaction on the firms' side because members of those age groups would not be as easily enticed to engage in on-the-job search as younger individuals, thereby reducing the incentive to firms to create vacancies. In addition, dropping age 25 also increases the magnitude of the coefficient in the unemployment model but has no sizeable effect in the employment model.

Figure S6: Exclusion of single age groups



Source: EU-SILC (authors' calculations). Cohort-size coefficients are estimated as described in Section 3; the estimated model also includes region, year and age dummies; the solid line represents the cohort-size coefficients from the full model, the dashed lines the corresponding 95% confidence interval.

To further assess to what extent the estimated cohort-size effects vary between different groups of regions, we estimate Equation 3 separately for regions from three parts of Europe: Southern Europe (16 regions from Greece, Italy and Spain), Eastern Europe (14 regions from the Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland and Slovakia) and a combination of Northern and Western Europe (19 regions from Austria, Belgium, Denmark, France, Luxembourg and Sweden). Tables S1–S3 show the results of the baseline model as well as the coefficients from the model containing region-by-age dummies (with and without control variables).

Estimating separate models for each of the three regions reduces the degrees of freedom compared to the pooled sample, which is reflected in higher standard errors. Moreover, the explanatory power of the instrument appears to be lower as evidenced by a reduction in the first-stage F-statistics. Nevertheless, in many specifications the 2SLS coefficients remain negative and significant in the unemployment model and positive and significant in the employment model when the Southern European regions are used. All of the coefficients have the expected sign and are significant at the 1% level for the sample of Eastern European regions. While there are no significant effects for the remaining regions of Northern and Western Europe, this need not imply that the relationship between cohort size and labour-market outcomes is structurally different in this part of Europe, but may rather be a reflection of the limited variation in the cohort-size variable as could already be seen in Figures A2 and A3.

Table S1: OLS and 2SLS regression results (Southern European regions)

<i>Panel A: Unemployment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-5.54 (3.58)	-11.16** (5.13)	-3.51 (3.68)	-6.63 (5.36)	-8.93** (3.86)	-16.87*** (5.44)	-6.56* (3.95)	-12.23** (5.59)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	640	640	640	640	640	640	640	640
R ²	0.53	0.52	0.55	0.55	0.57	0.57	0.59	0.59
F-stat	-	338.11***	-	323.28***	-	340.42***	-	300.82***
ME(std)	-0.02	-0.04**	-0.01	-0.02	-0.03**	-0.06***	-0.02*	-0.04**
<i>Panel B: Employment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	7.55* (4.05)	8.97 (6.58)	6.07 (4.11)	5.92 (7.27)	11.76** (4.61)	14.73** (6.78)	10.05** (4.69)	11.23 (7.24)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	640	640	640	640	640	640	640	640
R ²	0.65	0.65	0.66	0.66	0.69	0.69	0.70	0.70
F-stat	-	338.11***	-	323.28***	-	340.42***	-	300.82***
ME(std)	0.03*	0.03	0.02	0.02	0.04**	0.05**	0.03**	0.04
***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.								

Table S2: OLS and 2SLS regression results (Eastern European regions)

<i>Panel A: Unemployment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-9.21*** (1.87)	-8.94*** (2.06)	-5.78*** (2.00)	-5.35*** (2.20)	-9.90*** (2.18)	-10.56*** (2.30)	-5.96*** (2.29)	-6.33*** (2.47)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	560	560	560	560	560	560	560	560
R ²	0.32	0.32	0.40	0.40	0.36	0.36	0.44	0.44
F-stat	–	843.57***	–	899.35***	–	705.96***	–	732.99***
ME(std)	-0.02***	-0.02***	-0.01***	-0.01***	-0.02***	-0.02***	-0.01***	-0.01***
<i>Panel B: Employment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	15.91*** (2.33)	20.28*** (2.63)	11.99*** (2.49)	16.06*** (2.80)	14.74*** (2.64)	19.99*** (2.94)	9.87*** (2.84)	14.71*** (3.33)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	560	560	560	560	560	560	560	560
R ²	0.41	0.41	0.48	0.48	0.46	0.45	0.53	0.52
F-stat	–	843.57***	–	899.35***	–	705.96***	–	732.99***
ME(std)	0.03***	0.04***	0.02***	0.03***	0.03***	0.04***	0.02***	0.03***

***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.

Table S3: OLS and 2SLS regression results (Northern and Western European regions)

<i>Panel A: Unemployment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-0.23 (3.93)	5.73 (7.88)	1.74 (4.07)	5.08 (7.51)	0.44 (4.22)	1.18 (7.51)	2.38 (4.46)	1.18 (6.99)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	759	759	759	759	759	759	759	759
R ²	0.13	0.12	0.17	0.17	0.19	0.19	0.22	0.22
F-stat	–	83.20***	–	94.78***	–	67.78***	–	77.30***
ME(std)	-0.00	0.01	0.00	0.01	0.00	0.00	0.00	0.00
<i>Panel B: Employment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-1.55 (5.45)	-0.23 (11.00)	-2.74 (5.31)	4.75 (10.36)	-6.66 (5.39)	-3.28 (9.58)	-7.75 (5.36)	1.20 (9.07)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	759	759	759	759	759	759	759	759
R ²	0.24	0.24	0.30	0.29	0.32	0.32	0.37	0.37
F-stat	–	83.20***	–	94.78***	–	67.78***	–	77.30***
ME(std)	-0.00	-0.00	-0.01	0.01	-0.01	-0.01	-0.02	0.00
<p>***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.</p>								

S2.2 Robustness to changes in the model specification and the sample

This part assesses the robustness of the estimated relationship between cohort-size and the unemployment and the employment shares to a variety of changes in the specification of the empirical model or the underlying sample.

In Table S4 we first show that the paper's results also hold when instead of aggregating the dependent variable to the level of the region-year-age group the underlying microdata is used (Angrist and Pischke, 2009). In this case the dependent variable is defined as a binary variable that indicates whether an individual i in age group j , region r and year t is unemployed ($unemp_{ijrt}$) or employed (emp_{ijrt}). In light of the strong assumptions that have to be made to ensure consistency in a binary dependent variable model with endogenous regressors (Cameron and Trivedi, 2009) and since the focus of the analysis is on estimating marginal effects rather than on making predictions, a linear probability model is used to which we apply the same IV estimation strategy that is outlined in Section 3. As the cohort-size variable is defined at a higher level of aggregation than the dependent variable, which now may also vary across individuals in the same region-year-age group, standard errors are clustered at the level of the region-age group cell (Moulton, 1990). Observations are weighted by the individual-level weights which have been provided as part of the EU-SILC data and which have then been calibrated so that the estimated size of a region-year-age-sex cell matches the population size as reported by Eurostat (see Section 2). The size of the standard errors increases compared to the aggregate-level analysis but all coefficients remain statistically significant at the 1% level.

Table S4: OLS and 2SLS regression results (individual-level analysis)

<i>Panel A: Unemployment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-10.32*** (1.97)	-17.30*** (2.55)	-8.38*** (1.95)	-15.50*** (2.44)	-12.84*** (2.35)	-23.01*** (3.13)	-10.17*** (2.34)	-20.47*** (3.00)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations</i>								
Individual-level	64,387	64,387	64,387	64,387	64,387	64,387	64,387	64,387
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
Region-age	243	243	243	243	243	243	243	243
R ²	0.04	0.04	0.06	0.06	0.05	0.04	0.07	0.07
F-stat	–	1,352.68***	–	1,568.95***	–	1,024.29***	–	1,385.52***
ME (std)	-0.03***	-0.05***	-0.03***	-0.05***	-0.04***	-0.07***	-0.03***	-0.06***
<i>Panel B: Employment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	14.39*** (2.29)	24.32*** (3.05)	12.20*** (2.32)	22.06*** (3.02)	13.88*** (2.61)	26.55*** (3.71)	10.73*** (2.59)	23.15*** (3.63)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations</i>								
Individual-level	64,387	64,387	64,387	64,387	64,387	64,387	64,387	64,387
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
Region-age	243	243	243	243	243	243	243	243
R ²	0.07	0.07	0.10	0.10	0.08	0.08	0.10	0.10
F-stat	–	1,352.68***	–	1,568.95***	–	1,024.29	–	1,385.52
ME(std)	0.04***	0.08***	0.04***	0.07***	0.04***	0.08***	0.03***	0.07***
***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors that are clustered at the level of the region-age group cell are shown in parentheses. The regression is weighted using calibrated individual-level weights. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.								

In this paper a specific form of the cohort-size variable is used which, first, includes age groups that are up to two years younger and older and, second, assigns lower weights to age groups that are further away from the reference group. This specification is chosen to incorporate the assumption that members of an age group also compete with individuals that are slightly younger and older, but that substitutability decreases with the age difference. However, Wright (1991) already notes that this specific formulation is arbitrary. We therefore show that the results are robust to using a weighted cohort-size variable that only includes age groups that are up to one year younger or older (Equation S7), the relative size of the own-age group which does not consider any other age groups (Equation S8) as well as a three-year sum (Equation S9) and a five-year sum (Equation S10) in which each group receives an equal weight. Tables S5 to S8 show that the cohort-size coefficients retain their sign and significance. Since the distribution of these variables differ, it is useful to look at the marginal effects of a change in the corresponding cohort-size variable by one standard deviation instead of the cohort-size coefficients in order to compare the magnitude of the effects across the different specifications.

$$CS_{jrt} = \frac{(1/4)N_{j-1,rt} + (1/2)N_{jrt} + (1/4)N_{j+1,rt}}{N_{16-64,rt}} \quad [S7]$$

$$CS_{jrt} = \frac{N_{jrt}}{N_{16-64,rt}} \quad [S8]$$

$$CS_{jrt} = \frac{N_{j-1,rt} + N_{jrt} + N_{j+1,rt}}{N_{16-64,rt}} \quad [S9]$$

$$CS_{jrt} = \frac{N_{j-1,rt} + N_{jrt} + N_{j+1,rt}}{N_{16-64,rt}} \quad [S10]$$

Table S5: OLS and 2SLS regression results (3-year weighted cohort-size variable)

<i>Panel A: Unemployment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-8.58*** (1.54)	-16.30*** (2.00)	-6.72*** (1.57)	-14.20*** (1.94)	-10.55*** (1.71)	-21.97*** (2.29)	-8.04*** (1.76)	-18.98*** (2.16)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.38	0.36	0.41	0.40	0.42	0.40	0.45	0.43
F-stat	–	1,216.01***	–	1,242.94***	–	885.58***	–	972.80***
ME(std)	-0.03***	-0.05***	-0.02***	-0.05***	-0.03***	-0.07***	-0.03***	-0.06***
<i>Panel B: Employment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	11.27*** (1.91)	23.02*** (2.51)	9.47*** (1.87)	21.04*** (2.39)	10.66*** (2.06)	25.23*** (2.78)	8.30*** (2.06)	22.40*** (2.59)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.53	0.51	0.56	0.54	0.57	0.56	0.60	0.59
F-stat	–	1,216.01***	–	1,242.94***	–	885.58***	–	972.80***
ME(std)	0.04***	0.07***	0.03***	0.07***	0.03***	0.08***	0.03***	0.07***
***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.								

Table S6: OLS and 2SLS regression results (own-age cohort-size variable)

<i>Panel A: Unemployment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-3.07*** (0.95)	-14.72*** (1.93)	-1.99** (0.93)	-12.96*** (1.88)	-3.41*** (1.02)	-20.02*** (2.32)	-2.11*** (1.02)	-17.62*** (2.18)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.37	0.29	0.40	0.33	0.41	0.26	0.45	0.32
F-stat	–	326.62***	–	337.13***	–	226.31***	–	249.71***
ME(std)	-0.01***	-0.06***	-0.01**	-0.05***	-0.01***	-0.08***	-0.01***	-0.07***
<i>Panel B: Employment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	3.01** (1.23)	20.84*** (2.55)	1.96* (1.19)	19.29*** (2.41)	2.45* (1.26)	22.92*** (2.84)	1.21 (1.23)	20.73*** (2.64)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.52	0.42	0.55	0.46	0.57	0.46	0.60	0.50
F-stat	–	326.62***	–	337.13***	–	226.31***	–	249.71***
ME(std)	0.01**	0.08***	0.01*	0.08***	0.01*	0.09***	0.00	0.08***
***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.								

Table S7: OLS and 2SLS regression results (3-year non-weighted cohort-size variable)

<i>Panel A:</i>								
<i>Unemployment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-2.99*** (0.54)	-5.56*** (0.68)	-2.46*** (0.55)	-4.83*** (0.66)	-3.69*** (0.61)	-7.45*** (0.78)	-2.96*** (0.62)	-6.39*** (0.73)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.38	0.37	0.41	0.40	0.42	0.40	0.46	0.44
F-stat	–	1,356.30***	–	1,410.33***	–	1,029.92***	–	1,118.76***
ME(std)	-0.03***	-0.05***	-0.02***	-0.05***	-0.04***	-0.07***	-0.03***	-0.06***
<i>Panel B:</i>								
<i>Employment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	4.30*** (0.67)	7.85*** (0.84)	3.77*** (0.66)	7.14*** (0.81)	4.12*** (0.74)	8.57*** (0.94)	3.43*** (0.72)	7.55*** (0.87)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.53	0.52	0.56	0.55	0.58	0.56	0.61	0.60
F-stat	–	1,356.30***	–	1,410.33***	–	1,029.92***	–	1,118.76***
ME(std)	0.04***	0.08***	0.04***	0.07***	0.04***	0.08***	0.03***	0.07***

***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.

Table S8: OLS and 2SLS regression results (5-year non-weighted cohort-size variable)

<i>Panel A: Unemployment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-2.03*** (0.34)	-3.60*** (0.44)	-1.54*** (0.35)	-3.13*** (0.43)	-2.50*** (0.39)	-4.72*** (0.49)	-1.85*** (0.40)	-4.08*** (0.47)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.38	0.37	0.41	0.40	0.42	0.41	0.45	0.44
F-stat	–	1,483.63***	–	1,605.32***	–	1,096.43***	–	1,202.69***
ME(std)	-0.03***	-0.06***	-0.02***	-0.05***	-0.04***	-0.07***	-0.03***	-0.06***
<i>Panel B: Employment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	2.99*** (0.41)	5.04*** (0.55)	2.43*** (0.41)	4.55*** (0.53)	2.86*** (0.46)	5.47*** (0.60)	2.14*** (0.46)	4.80*** (0.56)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.53	0.52	0.56	0.55	0.58	0.57	0.61	0.60
F-stat	–	1,483.63***	–	1,605.32***	–	1,096.43***	–	1,202.69***
ME(std)	0.05***	0.08***	0.04***	0.07***	0.04***	0.08***	0.03***	0.07***

***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.

In the paper, individuals reporting to be in education are not excluded from the construction of the cohort-size variable. This is done because a part of these individuals may be willing to join the labour market if an attractive opportunity became available, while others are unlikely to do so because they are enrolled in long-term degree programmes. Crucially, distinguishing between these groups is not possible and consequently both approaches – including or excluding individuals in education – lead to measurement error in the cohort-size variable. However, Table S9 shows the results when a cohort-size variable is constructed

in which the numerator is derived only from individuals who are either employed or unemployed (to ensure a better comparison with the results in the paper, the construction of the denominator is left unchanged). In the unemployment model the 2SLS coefficients and corresponding marginal effects are similar in size to those reported in Table 3, while there is a decrease in the magnitude of the OLS estimates. In the employment model, there is a pronounced increase in the magnitude of the OLS coefficients and marginal effects, while the 2SLS effects are only slightly larger.

Table S9: OLS and 2SLS regression results (cohort-size variable from employed and unemployed individuals)

<i>Panel A:</i>								
<i>Unemployment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-6.15*** (1.63)	-18.13*** (2.30)	-3.72** (1.65)	-15.91*** (2.26)	-7.77*** (1.83)	-23.84*** (2.57)	-4.44** (1.86)	-20.73*** (2.44)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.37	0.34	0.40	0.38	0.41	0.38	0.45	0.41
F-stat	–	892.61***	–	920.50***	–	740.37***	–	857.18***
ME(std)	-0.02***	-0.06***	-0.01**	-0.05***	-0.03***	-0.08***	-0.02**	-0.07***
<i>Panel B:</i>								
<i>Employment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	21.50*** (1.99)	25.47*** (2.64)	18.93*** (2.04)	23.31*** (2.59)	21.48*** (2.38)	27.51*** (2.88)	18.04*** (2.46)	24.43*** (2.74)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.56	0.55	0.58	0.58	0.60	0.59	0.62	0.62
F-stat	–	892.61***	–	920.50***	–	740.37***	–	857.18***
ME(std)	0.07***	0.09***	0.06***	0.08***	0.07***	0.09***	0.06***	0.08***
***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.								

Next we report aggregate-level results which are not derived from weights that have been modified so that the weighted sum of observations per region-year-age-sex cell matches the corresponding population values reported by Eurostat. Instead this analysis is based on the weights provided as part of the EU-SILC dataset which have only been modified to take account of the change in the number of rotational groups per year by appending different longitudinal releases (see Section 3.1 and Moffat and Roth, 2016). Table S10 shows that using calibrated weights does not affect sign and significance of the cohort-size coefficients, though it increases the size of the marginal effects.

Table S10: OLS and 2SLS regression results (non-calibrated weights)

<i>Panel A: Unemployment</i>		OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size		-4.02*** (1.23)	-23.84*** (3.20)	-2.67** (1.19)	-20.87*** (3.09)	-4.19*** (1.39)	-28.36*** (3.45)	-2.62** (1.32)	-24.32*** (3.30)
<i>Dummies</i>									
Region		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age		No	No	No	No	Yes	Yes	Yes	Yes
Control variables		No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>									
Region-year-age		1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²		0.40	0.30	0.44	0.36	0.45	0.31	0.48	0.38
F-stat		-	242.66***	-	245.20***	-	215.47***	-	209.95***
ME(std)		-0.02***	-0.09***	-0.01**	-0.08***	-0.02***	-0.11***	-0.01**	-0.09***
<i>Panel B: Employment</i>		OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size		7.70*** (1.50)	35.48*** (3.95)	6.13*** (1.50)	32.08*** (3.82)	6.99*** (1.67)	34.52*** (4.09)	5.19*** (1.67)	30.08*** (3.91)
<i>Dummies</i>									
Region		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age		Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age		No	No	No	No	Yes	Yes	Yes	Yes
Control variables		No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>									
Region-year-age		1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²		0.57	0.47	0.60	0.51	0.62	0.53	0.64	0.57
F-stat		-	242.66***	-	245.20***	-	215.47***	-	209.95***
ME(std)		0.03***	0.13***	0.02***	0.12***	0.03***	0.13***	0.02***	0.11***

***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.

Table S11 shows the results when standard errors are estimated that are clustered at the level of the region-age cell, as is done in the individual-level analysis, instead of standard errors that are merely robust against heteroscedasticity. Despite the increase in the size of the standard errors, the cohort-size coefficients remain significant at the 1% level.

Table S11: OLS and 2SLS regression results (clustered standard errors)

<i>Panel A:</i>								
<i>Unemployment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-10.32*** (2.00)	-17.30*** (2.55)	-7.98*** (1.98)	-15.06*** (2.41)	-12.84*** (2.52)	-23.01*** (3.13)	-9.70*** (2.54)	-19.88*** (2.95)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
Region-age	243	243	243	243	243	243	243	243
R ²	0.38	0.37	0.41	0.40	0.43	0.41	0.45	0.44
F-stat	–	1,312.45***	–	1,561.99**	–	895.96***	–	1,226.07***
ME(std)	-0.03***	-0.05***	-0.02***	-0.05***	-0.04***	-0.07***	-0.03***	-0.06***
<i>Panel B:</i>								
<i>Employment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	14.39*** (2.32)	24.32*** (3.05)	11.91*** (2.39)	22.07*** (2.93)	13.88*** (2.79)	26.55*** (3.71)	10.63*** (2.82)	23.43*** (3.56)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
Region-age	243	243	243	243	243	243	243	243
R ²	0.53	0.52	0.56	0.55	0.58	0.57	0.61	0.60
F-stat	–	1,312.45***	–	1,561.99***	–	895.96***	–	1,226.07***
ME(std)	0.04***	0.08***	0.04***	0.07***	0.04***	0.08***	0.03***	0.07***
***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors that are clustered at the level of the region-age group cell are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.								

Table S12: OLS and 2SLS regression results (data aggregated from unemployed and employed individuals only)

<i>Panel A: Unemployment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-12.40*** (1.84)	-21.12*** (2.28)	-9.51*** (1.86)	-18.23*** (2.21)	-14.37*** (2.08)	-26.00*** (2.58)	-10.72*** (2.12)	-22.39*** (2.47)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.43	0.42	0.46	0.45	0.47	0.46	0.50	0.49
F-stat	–	1,499.69***	–	1,619.54***	–	1,104.85***	–	1,216.07***
ME(std)	-0.04***	-0.07***	-0.03***	-0.06***	-0.04***	-0.08***	-0.03***	-0.07***
<i>Panel B: Employment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	12.40*** (1.84)	21.12*** (2.28)	9.51*** (1.86)	18.23*** (2.21)	14.37*** (2.08)	26.00*** (2.58)	10.72*** (2.12)	22.39*** (2.47)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,959	1,959	1,959	1,959	1,959	1,959	1,959	1,959
R ²	0.43	0.42	0.46	0.45	0.47	0.46	0.50	0.49
F-stat	–	1,499.69***	–	1,619.54***	–	1,104.85***	–	1,216.07***
ME(std)	0.04***	0.07***	0.03***	0.06***	0.04***	0.08***	0.03***	0.07***

***/**/* indicate significance at the 1%/5%/10% level, respectively. Standard errors that are clustered at the level of the region-age group cell are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.

In the paper, the empirical analysis is conducted for the age range 25–29 in order to avoid the estimated effects of cohort size on the unemployment and the employment share being confounded by the decision to enter the labour market or to acquire education. If indeed only a small share of individuals participates in education in this age range, we would expect to obtain similar results when the empirical analysis is restricted to the regression sample of individuals who are either employed or unemployed (notice that this restriction does not affect the construction of the cohort-size variable, which is population-based and

therefore independent of the distribution of individuals across different labour-market states). The results in Table S12 show that if this restriction is imposed, the marginal effects for a change of one standard deviation increase slightly in the unemployment model (notice that since there are only two labour-market states in the sample, the coefficients in the employment model have the same magnitude but opposite sign compared to those in the unemployment model).

As can be seen from Figures A2 and A3 in the Appendix there is fluctuation in the share of individuals in a particular labour-market state across age groups and over time for a given region. While these fluctuations may reflect 'true' variation in the dependent variables, it is also possible that they are the result of labour-market shares being derived from small cell sizes: if the number of observations per region-year-age cell is small, the estimated shares may no longer be representative of the actual distribution of labour-market status in the population. While measurement error in the dependent variable generally reduces estimation precision, estimates may also be biased if the fluctuations vary systematically with the cohort-size variable. In order to assess the sensitivity of the results we drop cells containing less than 3, less than 5 and less than 10 observations. As shown in Tables S13 to S15, the resulting coefficients and marginal effects are close to those reported in the paper.

Table S13: OLS and 2SLS regression results (cells with less than three observations are excluded)

<i>Panel A: Unemployment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-10.25*** (1.70)	-17.32*** (2.10)	-7.91*** (1.73)	-15.08*** (2.05)	-12.75*** (1.91)	-23.04*** (2.38)	-9.61*** (1.96)	-19.90*** (2.25)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,953	1,953	1,953	1,953	1,953	1,953	1,953	1,953
R ²	0.38	0.37	0.41	0.40	0.43	0.41	0.45	0.44
F-stat	–	1,538.49***	–	1,639.52***	–	1,138.38***	–	1,267.89***
ME(std)	-0.03***	-0.05***	-0.02***	-0.05***	-0.04***	-0.07***	-0.03***	-0.06***
<i>Panel B: Employment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	14.19*** (2.03)	24.15*** (2.64)	11.66*** (2.02)	21.88*** (2.52)	13.62*** (2.24)	26.39*** (2.89)	10.33*** (2.25)	23.26*** (2.70)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,953	1,953	1,953	1,953	1,953	1,953	1,953	1,953
R ²	0.53	0.53	0.56	0.55	0.58	0.57	0.61	0.60
F-stat	–	1,538.49***	–	1,639.52***	–	1,138.38***	–	1,267.89***
ME(std)	0.04***	0.08***	0.04***	0.07***	0.04***	0.08***	0.03***	0.07***

***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.

Table S14: OLS and 2SLS regression results (cells with less than five observations are excluded)

<i>Panel A:</i>								
<i>Unemployment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-10.55*** (1.71)	-17.58*** (2.11)	-8.15*** (1.74)	-15.27*** (2.05)	-13.43*** (1.91)	-23.56*** (2.39)	-10.29*** (1.96)	-20.45*** (2.26)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,937	1,937	1,937	1,937	1,937	1,937	1,937	1,937
R ²	0.38	0.38	0.42	0.41	0.43	0.42	0.46	0.45
F-stat	–	1,531.02***	–	1,637.04***	–	1,134.46***	–	1,272.82***
ME(std)	-0.03***	-0.05***	-0.03***	-0.05***	-0.04***	-0.07***	-0.03***	-0.06***
<i>Panel B:</i>								
<i>Employment</i>	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	14.48*** (2.03)	24.43*** (2.64)	11.86*** (2.02)	22.03*** (2.51)	14.36*** (2.22)	27.02*** (2.90)	11.03*** (2.23)	23.81*** (2.69)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,937	1,937	1,937	1,937	1,937	1,937	1,937	1,937
R ²	0.54	0.53	0.57	0.56	0.58	0.58	0.61	0.61
F-stat	–	1,531.02***	–	1,637.04***	–	1,134.46***	–	1,272.82***
ME(std)	0.05***	0.08***	0.04***	0.07***	0.04***	0.08***	0.03***	0.07***

***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.

Table S15: OLS and 2SLS regression results (cells with less than ten observations are excluded)

<i>Panel A: Unemployment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-10.94*** (1.70)	-17.26*** (2.09)	-8.40*** (1.72)	-14.82*** (2.03)	-14.09*** (1.92)	-23.23*** (2.36)	-10.77*** (1.98)	-19.94*** (2.25)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,836	1,836	1,836	1,836	1,836	1,836	1,836	1,836
R ²	0.41	0.40	0.44	0.43	0.46	0.45	0.49	0.48
F-stat	–	1,512.88***	–	1,600.87***	–	1,106.95***	–	1,224.71***
ME(std)	-0.03***	-0.05***	-0.03***	-0.05***	-0.04***	-0.07***	-0.03***	-0.06***
<i>Panel B: Employment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	15.70*** (1.99)	25.50*** (2.48)	12.87*** (1.97)	22.85*** (2.35)	15.40*** (2.21)	27.82*** (2.78)	11.76*** (2.24)	24.21*** (2.61)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
Control variables	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	1,836	1,836	1,836	1,836	1,836	1,836	1,836	1,836
R ²	0.57	0.56	0.60	0.59	0.62	0.61	0.65	0.64
F-stat	–	1,512.88***	–	1,600.87***	–	1,106.95***	–	1,224.71***
ME(std)	0.05***	0.08***	0.04***	0.07***	0.05***	0.09***	0.04***	0.08***

***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.

The empirical analysis of the paper is based on a balanced panel of regions which can be observed throughout the whole sample period 2005–2012. Table S16 shows the cohort-size coefficients which are obtained if the following regions are not excluded from the analysis: 2 regions from Bulgaria (2006–2012), 1 region from Cyprus (2007–2012; due to unavailability of the instrumental variable, age group 25 can only be included from 2009 onwards), 1 region from Malta (2006–2012), 1 region from Norway (2008–2012) and 4 regions from Romania (2007–2012). The marginal effects are slightly smaller than those shown in the paper, but retain their sign and significance at the 1% level.

Table S16: OLS and 2SLS regression results (inclusion of all available regions)

<i>Panel A: Unemployment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-9.47*** (1.59)	-16.32*** (2.03)	-7.33*** (1.62)	-13.99*** (1.97)	-11.54*** (1.78)	-21.39*** (2.26)	-8.77*** (1.82)	-18.39*** (2.14)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	2,236	2,236	2,236	2,236	2,236	2,236	2,236	2,236
R ²	0.38	0.38	0.41	0.41	0.43	0.42	0.46	0.44
F-stat	–	1,531.47***	–	1,656.76***	–	1,124.69***	–	1,249.18***
ME(std)	-0.03***	-0.05***	-0.02***	-0.04***	-0.03***	-0.06***	-0.03***	-0.06***
<i>Panel B: Employment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	13.27*** (1.91)	22.91*** (2.55)	10.96*** (1.90)	20.42*** (2.43)	12.57*** (2.09)	24.71*** (2.76)	9.61*** (2.10)	21.39*** (2.58)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	2,236	2,236	2,236	2,236	2,236	2,236	2,236	2,236
R ²	0.53	0.52	0.56	0.55	0.57	0.57	0.60	0.59
F-stat	–	1,531.47***	–	1,656.76***	–	1,124.69***	–	1,249.18***
ME(std)	0.04***	0.07***	0.03***	0.06***	0.04***	0.07***	0.03***	0.06***
***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.								

Since the EU-SILC dataset also contains observations from some regions for the years 2004 and 2013, the sample period can in principle be extended by another two years, though this is only possible for the regions from the following set of countries: Austria (3 regions, 2004, 2013), Belgium (3 regions, 2004), Bulgaria (2 regions, 2013), Cyprus (1 region, 2013), Czech Republic (1 region, 2013), Denmark (1 region, 2004, 2013), Estonia (1 region, 2004, 2013), Greece (4 regions, 2004), Spain (7 regions, 2004, 2013), France (8 regions, 2004, 2013), Hungary (3 regions, 2013), Italy (3 regions, 2004, 2013), Italy (2 regions, 2013), Lithuania

(1 region, 2013), Luxembourg (1 region, 2004, 2013), Latvia (1 region, 2013), Malta (1 region, 2013), Poland (6 regions, 2013) and Slovakia (1 region, 2013). As can be seen from Table S17, the cohort-size coefficients retain their sign and continue to be significant at the 1% level. In the unemployment model the size of the marginal effects is slightly smaller than in Table S16, whereas the size of the marginal effects in the employment model stays the same.

Table S17: OLS and 2SLS regression results (inclusion of all available regions and years)

<i>Panel A: Unemployment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	-8.31*** (1.54)	-13.70*** (2.07)	-6.73*** (1.53)	-12.20*** (1.97)	-9.66*** (1.68)	16.86*** (2.26)	-7.67*** (1.69)	-14.82*** (2.16)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	2,606	2,606	2,606	2,606	2,606	2,606	2,606	2,606
R ²	0.38	0.37	0.40	0.40	0.42	0.41	0.44	0.43
F-stat	–	2,104.78***	–	2,309.70***	–	1,674.20***	–	1,899.38***
ME(std)	-0.03***	-0.04***	-0.02***	-0.04***	-0.03***	-0.05***	-0.02***	-0.05***
<i>Panel B: Employment</i>								
	OLS	2SLS	OLS	2SLS	OLS	2SLS	OLS	2SLS
Cohort size	12.78*** (1.76)	21.60*** (2.36)	10.99*** (1.75)	19.63*** (2.26)	12.21*** (1.90)	22.34*** (2.52)	10.10*** (1.91)	19.67*** (2.41)
<i>Dummies</i>								
Region	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Region-by-age	No	No	No	No	Yes	Yes	Yes	Yes
<i>Control variables</i>	No	No	Yes	Yes	No	No	Yes	Yes
<i>Observations (cells)</i>								
Region-year-age	2,606	2,606	2,606	2,606	2,606	2,606	2,606	2,606
R ²	0.54	0.53	0.56	0.55	0.57	0.57	0.60	0.59
F-stat	–	2,104.78***	–	2,309.70***	–	1,674.20***	–	1,899.38***
ME(std)	0.04***	0.07***	0.03***	0.06***	0.04***	0.07***	0.03***	0.06***

***/**/* indicate significance at the 1%/5%/10% level, respectively. Robust standard errors are shown in parentheses. The regression is weighted by the estimated number of male observations in a region-year-age cell. F-stat represents the first-stage F-statistic from a regression of the endogenous cohort-size variable on the instrument and control variables. ME(std) shows the change in the dependent variable if the cohort-size variable increases by one standard deviation.

References

- Angrist, J. and Pischke, J.-S. (2009) *Mostly harmless econometrics: an empiricist's companion*, Princeton University Press, Princeton.
- Cameron, A.C. and Trivedi, P.K. (2009) *Microeconometrics Using STATA*, STATA Press, College Station.
- Moffat, J. and Roth, D. (2016) The Cohort Size-Wage Relationship in Europe, *Labour*, forthcoming.
- Moulton, B.R. (1990) An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units, *Review of Economics and Statistics*, 72, 334–38.
- Shimer, R. (2001) The Impact of Young Workers on the Aggregate Labor Market, *The Quarterly Journal of Economics*, 116, 969–1007.
- Wright, R.E. (1991) Cohort size and earnings in Great Britain, *Journal of Population Economics*, 4, 295–305.

Duncan Roth^a

Cohort size and transitions into the labour market

Abstract

This paper estimates the effect that the size of an individual's labour-market entry cohort has on the subsequent duration of search for employment. Survival-analysis methods are applied to empirically assess this relationship using a sample of apprenticeship graduates who entered the German labour market between 1999 and 2012. The results suggest that apprentices from larger graduation cohorts take less time to find employment, but this effect appears to be significant only for a period of up to six months after graduation. These results therefore do not support the cohort-crowding hypothesis that members of larger cohorts face depressed labour-market outcomes. Moreover, there is no evidence that shorter search durations are the result of graduates being pushed into lower-quality employment. The finding that graduating as part of a larger cohort leads to shorter search durations is in line with those parts of the cohort-size literature that find larger youth cohorts being associated with lower unemployment rates. A possible explanation is that firms react to an anticipated increase in the number of graduates by creating jobs.

JEL classification: J21, J64, R23

Keywords: Survival analysis, entry conditions, cohort size, apprentices, search duration

1 Introduction

The extant cohort-size literature has predominantly focussed on how the size of a specifically defined age group affects the wage (Mosca, 2009; Brunello, 2010; Morin, 2015; Garloff and Roth, 2016; Moffat and Roth, 2016a) and (un-) employment outcomes (Korenman and Neumark, 2000; Shimer, 2001; Skans, 2005; Biagi and Lucifora, 2008; Garloff et al., 2013; Moffat and Roth, 2016b) of that group. In contrast, the question how cohort-size shapes an individual's transition into the labour market and subsequent career has so far been left largely unaddressed, although the demographic processes which are projected to lead to reductions in population size and changes in age structures throughout

^a Institute for Employment Research (IAB), Düsseldorf, Germany.

Europe (European Commission, 2014) and in Germany in particular (Statistisches Bundesamt, 2015) would appear to provide a motivation to better understand this relationship.

This paper addresses this question by estimating the effect that an increase in the size of the cohort of graduates from Germany's apprenticeship system has on the duration that apprentices spend searching for employment following graduation. Specifying a cohort-size variable in terms of the group as part of which an individual enters the labour market sets this study apart from the majority of the above-mentioned literature in which cohort size typically refers to the contemporaneous size of an age group. As such, this paper is also related to a recent literature on the effects of the state of the local labour market at the time of entry – usually, based on a measure for the business cycle – on subsequent labour-market outcomes (Stevens, 2007; Kahn, 2010; Brunner and Kuhn, 2014; Cockx and Ghirelli, 2016) since the size of the graduation cohort within a local labour market also represents a feature of the conditions under which labour-market entry takes place.

The use of apprenticeship graduates in this paper as opposed to population-based age groups, which is common in the extant cohort-size literature, also provides a better measure of a group that is relevant to the labour market and therefore allows a better assessment of the consequences of labour-market crowding. It is typically assumed that individuals within a cohort are substitutable for each other, but that there is imperfect substitution across cohorts. This assumption is more likely to hold among apprenticeship graduates since the majority of the former are not only of a similar age but also share a comparable level of qualification which makes it more likely that they will be competing on the same labour market than two individuals that only belong to the same age group. Constructing cohort size from apprenticeship graduates who have completed their training and are therefore ready to enter the labour market should, moreover, reduce the problem of measurement error in this variable. This problem arises when cohorts that are based on young age groups are included, as large parts of the former are likely to be unavailable to the labour market (see Moffat and Roth, 2016b).

From a theoretical perspective the sign of the effect that the size of an individual's graduation cohort has on his subsequent search duration is *ex ante* unclear. The results of the empirical analysis suggest that belonging to a larger graduation cohort is predicted to reduce search duration. Specifically, the effect of a rise in the size of the entry cohort by one standard deviation is predicted to increase the hazard rate of finding employment by approximately 8%, which is comparable in magnitude to the effect of a corresponding increase in the unemployment at the time of entry. This effect, however, is only significant

within a relatively short period following graduation. The empirical analysis therefore does not provide any evidence that members of larger entry cohorts face longer search durations. Moreover, the results do not suggest that shorter search durations come at the price of taking up employment in lower-quality jobs. Alternative explanations for the pattern of the regression results relating to selected migration after graduation or changes in the productivity composition in larger cohorts are also not supported by the data. While offering no direct evidence for the mechanisms suggested by parts of the literature that find that larger youth cohorts reduce the youth unemployment rate, these results are nevertheless in line with the hypothesis that an increase in the size of an entry cohort induces an expansion in labour demand.

The remainder of the paper is structured as follows: Section 2 provides an overview of the extant literature; the empirical analysis is the subject of Section 3, while Section 4 contains the results; Section 5 concludes.

2 Literature and hypotheses

The subject of this paper is related to a large body of literature that analyses the impact that the size of a cohort has on the labour-market outcomes of its members. In this literature the term cohort usually refers to a group of individuals that fall into a specified age range, though in some cases cohorts are also differentiated with respect to educational attainment. The main motivation for defining cohorts in this way is the assumption that differently aged individuals are only imperfectly substitutable for each other and can be thought of as distinct factors of production (Card and Lemieux, 2001). The reason for this assumption is that older individuals tend to have more years of work experience, which in turn makes it more likely that they have acquired more human capital of various types (general, industry-specific, occupation-specific and job-specific). As long as a worker's productivity is related to the amount of human capital he has acquired, it follows that differently aged individuals should only be imperfectly substitutable (a more detailed discussion can be found in Garloff and Roth, 2016, and Moffat and Roth, 2016a).

Most research within this literature has so far concentrated on the effect of cohort size on wages as well as on employment and unemployment outcomes. In the case of wages the benchmark model of a perfectly competitive labour-market predicts that if there is diminishing marginal productivity of labour an increase in cohort size reduces the wages earned by its members (Brunello, 2010), while Michaelis and Debus (2011) show that a similar result holds in the case of an imperfectly competitive labour market in which wages are set by monopoly

unions. Findings by Garloff and Roth (2016) suggest that a considerable part of the negative effect can be ascribed to members of larger age groups being more likely to find employment in lower-paying occupations and industries. A substantial body of empirical research from different countries and time periods provides evidence in support of the hypothesis that increases in cohort size reduce the wages of its members (Freeman, 1979; Welch, 1979; Berger, 1983; Dooley, 1986; Wright, 1991; Mosca, 2009; Brunello, 2010; Morin, 2015; Garloff and Roth, 2016; Moffat and Roth, 2016a). However, if age-specific wages are rigid or the number of jobs for members of an age group are limited, changes in cohort size might rather affect age-specific employment or unemployment. The empirical literature provides conflicting evidence on this issue with some studies finding that larger youth cohorts lead to depressed employment and unemployment outcomes (Korenman and Neumark, 2000; Biagi and Lucifora, 2008; Garloff et al., 2013), while others provide evidence of a positive effect (Shimer, 2001; Skans, 2005; Moffat and Roth, 2016b).

One feature of the cohort-size literature is that the former's impact is typically analysed for contemporaneous outcomes. While this paper also utilises the concept of a cohort as a group of individuals with similar characteristics, it differs by defining a cohort-size variable that refers to a specific point in time – the time of entry into the labour market – and estimates its effect on the subsequent duration of search for employment. In light of this set-up, the paper is also relevant to a recent literature analysing the effects that the conditions prevailing at the time of an individual's entrance into the labour market have on subsequent labour-market outcomes. In this literature these conditions refer to the state of the economy when an individual enters the labour market which is typically measured by the local or national unemployment rate, the most commonly used outcome variables being an individual's subsequent wages or earnings, though some studies also consider the effect on annual hours worked or the employment rate. Initially, entering the labour market during an economic downturn has the effect of increasing the probability of being unemployed, while individuals may also be pushed into lower-paying jobs. This initial effect can become persistent if these jobs offer fewer opportunities to acquire productivity-enhancing human capital and if individuals fail to transfer to a higher-quality job at a later stage. Evidence for the hypothesis that labour-market entry during an economic downturn can lead to lasting depressed labour-market outcomes is provided by a number of studies (Stevens, 2007; Kahn, 2010; Brunner and Kuhn, 2014; Cockx and Ghirelli, 2016).

However, as the literature on cohort-size effects suggests, the state of the economy does not necessarily constitute the only factor that is relevant to an individual's labour-market outcomes and the supply of similarly aged and qualified

individuals may also represent an important entry condition. So far, evidence on the effects of cohort size at the time of labour-market entry is scarce – Morin (2015) analyses changes in the size of the Canadian school graduation cohorts on subsequent wage outcomes and the quality of employment – and the former's relationship with search duration, which is the subject of this paper, has so far not been studied. The effect that an increase in the size of the entry cohort has on the amount of time that its members have to search before finding employment is *ex ante* unclear. However, the cohort-size literature and especially the mechanisms underlying the relationship with (un-)employment outcomes provide a basis from which to derive hypotheses. The cohort-crowding argument states that in the absence of a full and immediate adjustment in cohort-specific wages, an increase in cohort size leads to depressed employment and unemployment outcomes due to increased competition. Within such a framework members of larger entry cohorts can be expected to have longer search durations. However, the relationship between entry-cohort size and search duration would become indeterminate if members of larger cohorts avoided prolonged search durations by (temporarily) moving into lower-quality jobs. In such a scenario the effect of increased competition may be fully or partially countered depending on how many individuals would be prepared to select into such jobs and how quickly this would happen.

Finally, a possible rationale for members of larger cohorts having shorter search durations is provided by Shimer (2001) who finds that an increase in the size of the youth cohort reduces the unemployment rate of that age group (as well as of other groups). In his model the primary difference between younger and older individuals is that the former are more likely to be either unemployed or employed but poorly matched and therefore more prepared to either take up or switch jobs. An increase in the size of the youth cohort therefore leads to a larger supply of individuals that can be recruited by firms. The central assumption is the existence of a trading externality: that there is a higher probability of employers and job searchers realising a match if the number of trading partners is large. Given this assumption, firms are predicted to react to an increase in the size of the youth cohort by creating vacancies because the larger number of unemployed or poorly matched young individuals increases the probability of making a match. However since new matches can also be poor matches – in which case an individual would continue searching for other job opportunities – firms have an incentive to continue creating vacancies with the result that the overall unemployment rate and the unemployment rate of the young decreases. Within this framework it is conceivable that members of larger entry cohorts have shorter search durations. In order to assess the validity of the above hypotheses, the relationship between entry-cohort size and search duration is analysed empirically based on a sample

of graduates from Germany's apprenticeship system who enter the labour market between 1999 and 2012.

In addition to analysing the effect of changes in cohort size on an outcome that has so far not been considered, this paper is also able to deal with two sources of measurement error which are usually not addressed in the cohort-size literature. First, cohorts are supposed to measure the amount of individuals with similar characteristics that are active on the same labour market. Usually, administrative units at different levels of aggregation are used as the spatial basis from which to construct the cohort-size variable. These units do not necessarily provide good measures of actual labour markets because they are typically not delineated according to economic criteria. As a result, a cohort-size variable derived from administrative units is subject to measurement error because it is likely to group together individuals that are not active on the same labour market. This paper addresses this concern by employing the labour-market regions defined by Kosfeld and Werner (2012), which combine one or more administrative units based on the degree of commuting between these units. By creating as large an overlap as possible between the resident and the working population, these functional entities approximate actual labour markets.

Second, cohort-size variables are usually derived from the size of different age groups. Concerning the fact that members of a cohort are supposed to be available to the labour market, this approach can be problematic if considerable parts of an age group are non-participants and as such do not influence the labour-market outcomes of their age group. This is a particular concern for young age groups as their members are often engaged in education and are therefore not available to the labour market. Moffat and Roth (2016b) show that the inclusion of young age groups in the analysis of the relationship between cohort size and (un-)employment outcomes has considerable implications for size and sign of the cohort-size coefficient. This problem should be less of a concern in this study as apprenticeship graduates should be more likely to be available to the labour market.

3 Empirical analysis

3.1 Data

The empirical analysis of this paper utilises two different data sources. To construct the model's main explanatory variable – the number of graduates from an apprenticeship programme – the Integrated Employment Biographies (IEB) are used. This dataset contains information on all individuals who belong to one of

the following groups: employees subject to social security contributions, marginal employees, individuals receiving unemployment benefits, individuals registered as seeking employment and participants in the Federal Employment Agency's (FEA) measures of labour-market policy (groups that are not covered are civil servants and the self-employed). For each individual the dataset consists of different records that correspond to episodes in one of the above-mentioned states with specified start and end dates. Moreover, each episode is supplemented with two different sorts of information: first, characteristics of the individual are provided which refer to the beginning of the episode (among others, these characteristics include sex, nationality, year of birth, place of residence and level of education); second, details are provided that describe the state an individual is in (in the case of an employment episode, information would be available on the average daily wage during the episode, the occupation and industry of employment, place of employment as well as on the type of employment).¹

Participation in apprenticeship programmes constitutes a separate type of employment (employment subject to social security contributions and marginal employment constitute other major categories) and as such it is possible to determine whether an individual is participating in such a training programme at any given point in time. Because a change in the type of employment – e.g. when an individual completes an apprenticeship and takes up another form of employment – entails that a new episode is defined, it is further possible to identify when participation in an apprenticeship programme has ended. Based on this information, the number of individuals graduating from such a training programme in a given month, year and region can be estimated (Section 3.2. provides further details on the conditions that are imposed for an individual to be regarded as having completed training). Due to its size working directly with IEB records can be cumbersome and therefore the regression analysis of this paper uses a 2% sample, the so-called Sample of Integrated Employment Biographies (SIAB).

3.2 Sample and variables

The sample consists of male individuals aged between 19 and 23 who have completed an apprenticeship. Construction of the sample from the SIAB dataset proceeds as follows: first, those individuals without any episode as an apprentice are removed. For the remaining individuals it is then decided whether the information

¹ Variables differ in the extent to which they are provided. An individual's level of education is an example of a variable for which information can often be missing. Moreover, changes in classifications, e.g. in the coding of occupations, can cause problems in constructing a consistent coding scheme over longer periods of time.

on the registered apprenticeship episodes also warrants the assumption that training was completed. This is done by imposing two criteria: first, the combined duration of apprenticeship episodes has to be at least 730 days. While completion of training can often require more than two years, it is the case that individuals with a higher secondary education degree are able to complete an apprenticeship faster than those without a comparable schooling certificate. The rationale for setting a comparatively low threshold is thus to avoid excluding those who have completed secondary school. On the other hand, the risk of including individuals in the sample who have not completed training appears limited since they have already been participating in training for at least two years and dropping out of such schemes can be expected to typically happen earlier. Second, it is required that any gaps between two apprenticeship episodes are no longer than 100 days. A possible reason for such breaks is that training also includes a coursework component which does not take place within the training company. No additional restrictions are imposed; in particular, changes in the training company, in the occupation or industry during the apprenticeship are disregarded because parts of the training should be sufficiently general so as to be transferable to a different company, occupation or industry.

In order to avoid any confounding effects of selected female labour-market participation, the sample is restricted to men. Moreover, the age range of the sample is homogenised to include only those between the age of 19 and 23 because the majority of graduates complete their training within this age range.² Applying this procedure yields a sample of 52,234 individuals³ who have graduated between January 1999 and October 2012 and for whom transition into employment can be observed.

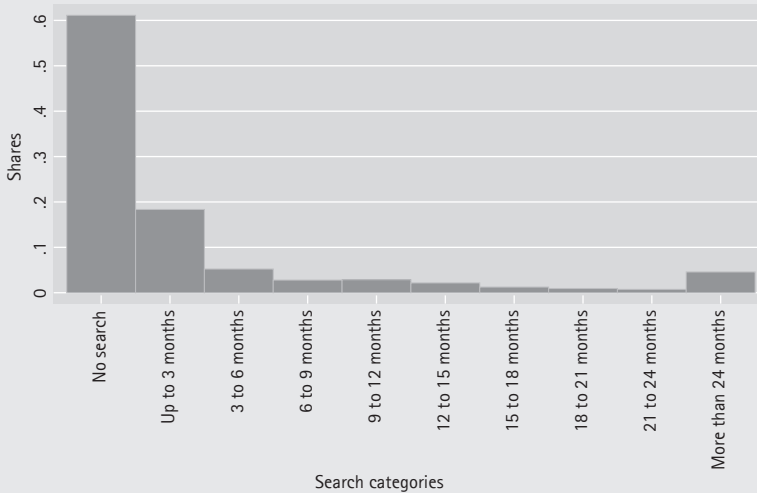
The model's dependent variable, *search*, is defined as the number of days it takes an individual to find employment after graduating from an apprenticeship programme. Figure 1 shows the distribution of this variable for the sample of individuals described above. The distribution is highly skewed as the majority (61%) falls into the category *No search*, which means that the employment episode of these individuals starts the day after graduation from apprenticeship training. Approximately 80% of graduates are able to find employment within 3 months after graduation, with this figure increasing to over 85% after 6 months.

2 SIAB only includes an individual's year of birth. Age at the time of graduation is defined as the difference between the year of graduation and the year of birth. Some individuals who are registered as being 25 upon graduation will therefore actually be between 24 and 26. Out of all male observations in SIAB with a completed apprenticeship for which all control variables are available 79% fall into the age range 19–23. As shown in the Supplementary Material, comparable results are obtained if this restriction is not imposed.

3 The maximum number of observations that can be used in the empirical analysis decreases to 46,408 due to missing values for the covariates.

Despite the large number of individuals who find employment directly upon graduation, this group is excluded from the empirical analysis. This is primarily due to technical reasons for the empirical model of Section 3.3 requires strictly positive durations. Moreover, zero and strictly positive durations may not be the outcomes of the same process. Instead firms may first decide whether to offer an apprentice a position after graduation from the training programme, with this decision being based on the performance of apprentices during training as well as on the economic condition of the firm. Apprentices are then free to either accept or decline the offer. If no match between training firm and apprentice is reached, individuals enter the labour market and search for employment. The empirical analysis therefore models search duration conditional on an apprentice not having been directly employed by his training firm (or having found employment immediately at a different firm).

Figure 1: Distribution of search durations



Source: Sample of Integrated Employment Biographies (author's calculations).


The obvious drawback of this approach is that those individuals who are employed directly might not constitute a random sample of graduates. In contrast, it is more likely that firms employ those apprentices which they believe to be especially productive. These individuals might possess characteristics which are not directly observable but which are relevant for on-the-job performance. If these characteristics also increased employability at other firms, graduates who are directly employed would be expected to experience shorter search periods in the counter-factual case of not being directly employed by their training firm.

Table 1 assesses this hypothesis by comparing average values of a number of characteristics between those apprentices who are employed directly and those who experience a strictly positive search duration. The first three variables refer to characteristics of the apprenticeship episode and while the difference in average duration of training and the share of Germans is statistically significant at the 0.01 level, the absolute difference in the variables is very small compared to the mean values of both groups. There is no statistically significant difference in average age at graduation. In contrast, there are sizeable and significant differences between characteristics of the employment spell that follows graduation: average daily earnings are about 20 Euro smaller for individuals who are not directly employed and the share of individuals working in part time is higher by about 15 percentage points. These latter findings suggest that both groups differ with respect to characteristics that are relevant for labour-market performance. Ideally, one would like to explicitly model this selection, but doing so would require an exogenous piece of information that would explain whether an individual is employed directly or experiences a positive search duration. In the absence of a suitable instrument, Section 4.4 provides an alternative way of including individuals with a zero search duration; the results of this analysis suggest that their inclusion reduces the magnitude of the cohort-size effect but does not affect its sign.

Table 1: Comparison of individuals with no and strictly positive search duration

Variable	Observations	Group 1 (<i>search = 0</i>)	Group 2 (<i>search > 0</i>)	Mean difference
<i>Apprenticeship episode</i>				
Duration of training	52,234	1,095.00	1,086.84	-8.16***
Age at graduation	52,234	21.16	21.17	0.01
German	52,226	0.97	0.95	-0.01***
<i>Employment episode</i>				
Average daily earnings	52,234	61.78	42.32	-19.46***
Part-time share	52,196	0.00	0.16	0.15***
Values derived from a regression on a group indicator as well as dummies for period and region of graduation. Robust standard errors are used. ***/**/* signifies significance at the .01/0.05/0.1 level. Differences in the number of observations are due to missing values of the corresponding variables.				

Figure 2: Monthly number of graduates (1999–2012)



Source: Integrated Employment Biographies (author's calculations).

The main explanatory variable, *cohort*, measures the regional supply of apprenticeship graduates and is based on the number of individuals that complete training within a given 6-month period and thus become available to the labour-market. Figure 2 shows the monthly number of graduates for the years 1999–2012. The annual distribution displays two peaks – one in January and another in June and July – which suggests that the bulk of apprentices complete training at two distinct points in time each year. To better reflect this pattern, the size of graduation cohorts is not computed for the whole year, but separately for two periods that cover six months each and that are centred on the peaks: November–April and May–October.

It is assumed that the duration of search for employment is influenced by the conditions of the labour market that an individual enters after graduation. The variable *cohort* measures the characteristic that is most relevant to this analysis: the degree of labour-market crowding among recently graduated apprentices. In order to avoid measurement error, graduation cohorts are constructed at the level of the 141 labour-market regions that are defined by Kosfeld and Werner (2012). As discussed in Section 2, these entities approximate self-contained units in which the employed population is exclusively recruited from the resident population. Since administrative units are typically not delineated according to economic criteria, they cannot be relied upon to provide an accurate measure of the size of a graduation cohort within an actual labour market. This argument is supported by findings of Garloff and Roth (2016) that the effect of cohort size on wages appears to be biased downwards when measured at the district level as compared to the level of labour-market regions.⁴ Finally, to ensure comparability of the size of the graduation cohort across different labour-market regions, this quantity is standardised by total employment in the region.⁵

Additional control variables are given by dummy variables for an individual's age at graduation, for whether an individual is of German nationality, for the occupation of the apprenticeship and industry of the training firm⁶ as well as the labour-market-specific unemployment rate. Summary statistics of these variables are given in the Appendix.

3.3 Model

To evaluate empirically the effect that the size of the graduation period has on an individual's search duration, the following Cox model is specified where subscripts i , r and p refer to the individual, the region and the period of graduation:

4 Labour-market regions refer to the individual's place of employment at the time of graduation. More than 80% of individuals in the sample live and work in the same region.

5 For the first period (November–April) employment numbers refer to 31 March of the year, while it is 31 October for the second period (May–October).

6 Occupation indicators are derived from the coding scheme *Klassifikation der Berufe 2010*, while industry indicators are based on the *Klassifikation der Wirtschaftszweige 1993*. Details are provided in the Appendix.

$$h_{irp}(t) = h_0(t)e^{(\gamma \text{cohort}_{rp} + \delta' x_{irp})} \quad [1]$$

Instead of formulating a relationship between the search duration and covariates, this model is specified in terms of the hazard rate $h_{irp}(t)$, which can be interpreted as the instantaneous probability that an individual realises a transition from search into employment. The term $h_0(t)$ represents the baseline hazard, i.e. the hypothetical hazard rate of an individual for whom all covariates are equal to zero. The Cox model belongs to the class of proportional hazard models meaning that changes in covariates shift the hazard rate up or down relative to the baseline hazard. The variable cohort_{rp} captures the size of the graduation cohort in region r and period p relative to the number of employed individuals in that region. Sign and significance of the coefficient γ therefore provide the basis for assessing the effect that the size of the entry cohort has on the duration of job search. The vector x_{irp} contains the above-mentioned set of control variables as well as dummy variables for period and region of graduation. The coefficients of the model are derived by maximum partial likelihood estimation (MPLE).⁷ To account for the difference in the level of aggregation of the dependent variable and the cohort variable, standard errors are clustered at the level of the labour-market region.

Four different specifications of this model are estimated which differ with respect to the specified period of time during which transitions into employment are observed. An inherent asymmetry in the data is given by the fact that individuals that complete their apprenticeship training earlier can be observed for a longer period of time (up to 31 December 2014) and as such can also accumulate longer search durations. To ensure comparability between graduates from different periods, four common periods of observation following graduation are defined: 3 months, 6 months, 1 year and 2 years. Individuals that find employment after the end of the common observation period are treated as not having realised a transition (i.e. they are right-censored) and their search durations are set equal to the corresponding common period of observation.⁸

7 The term partial refers to the fact that in contrast to fully parametric models, information on the search durations themselves is not used in the estimation. Instead, the relationship between the hazard rate and the covariates is derived solely from the ordering of the search durations.

8 The empirical model is based on two pieces of information: an indicator for whether transition into employment took place and the number of days an individual survived before transition. In the case of a 3-month period of observation an individual who found employment after six months would be recorded as not having experienced transition and his duration of search would be set to three months. Right-censored observations are not dropped from the regression. While they are treated as not having experienced transition, they are included in the 'risk set', i.e. the set of observations that are at risk of realising a transition into employment at each of the recorded transition times. The share of right-censored observations is 53% (3-month period), 39% (6-month), 25% (1-year) and 11% (2-year), respectively.

To consistently estimate the effect that the size of the graduation cohort has on the duration of search, it has to be assumed that individuals did not systematically select a region in which to undertake the apprenticeship training on the basis of their expectations regarding the probability of finding employment upon graduation. While the absence of regional selection appears unlikely in the context of other studies on cohort-size effects (see Moffat and Roth, 2016a), it is argued that this possibility is less of a concern in this case. First, since individuals are typically young when they start training, the region in which an apprenticeship is being undertaken will usually be determined by the region they live in at that time. Second, it appears unlikely that reliable expectations can be formed about the economic conditions prevailing in a region at the time of graduation. Moreover, if self-selection occurs into regions that constantly provide better employment opportunities for apprentices (and hence shorter search durations), this effect would be captured by the region dummies.

4 Results

4.1 Baseline results

Table 2 contains the coefficients from estimating the model of Equation 1 for each of the four common observation periods (3-month, 6-month, 1-year and 2-year periods).

In each case the estimated coefficient for the size of the entry cohort is positive though it is significant only if transitions into employment are counted as such if they take place during the first three months following graduation. For this observation period an increase in the size of the entry cohort significantly increases the hazard rate of finding employment. This means that individuals who complete their apprenticeship training as part of a larger cohort have shorter search durations. As presented in further detail in the Supplementary Material, the finding that belonging to a larger entry cohort is associated with shorter search durations is robust to a number of changes in the sample as well as in the empirical model.

To assess the size of the estimated effects, hazard ratios are computed, which show the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation. This value, which is shown in the bottom row of Table 2, is given by the exponentiated product of the cohort coefficient and the corresponding standard deviation (see Table A1). For the 3-month period such a change is predicted to increase the hazard rate by about 8%. Performing a similar computation for the regional unemployment rate at the

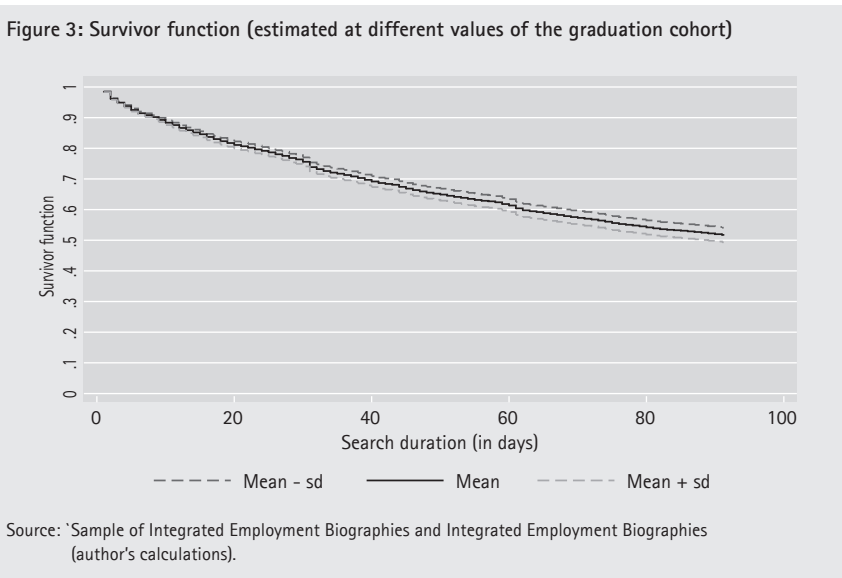
time of graduation shows that the effects of both variables are of similar size (though opposite sign) as the hazard rate is predicted to fall by approximately 10% if the unemployment rate increased by one standard deviation.

Table 2: Regression results

	3 months	6 months	1 year	2 years
Cohort	26.04** (11.20)	15.16 (9.33)	11.81 (9.32)	5.79 (8.57)
Unemployment rate	-1.99* (1.09)	-0.95 (0.94)	-1.50* (0.84)	-1.31* (0.77)
German	0.08 (0.06)	-0.00 (0.05)	-0.07 (0.05)	-0.07* (0.04)
<i>Age</i>				
20	-0.04 (0.04)	-0.04 (0.03)	-0.05* (0.03)	-0.06 (0.03)*
21	-0.06* (0.03)	-0.05* (0.03)	-0.06** (0.03)	-0.05 (0.03)
22	-0.13*** (0.04)	-0.11*** (0.04)	-0.12*** (0.03)	-0.09*** (0.03)
23	-0.06 (0.05)	-0.04 (0.04)	-0.03 (0.04)	-0.03 (0.04)
<i>Dummies</i>				
Occupation	Yes	Yes	Yes	Yes
Industry	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes
Log pseudo-likelihood	-81,452.15	-103,503.43	-126,212.95	-145,330.79
Observations	18,133	18,133	18,133	18,133
Clusters	141	141	141	141
ME(std)	1.07**	1.04	1.03	1.02
All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The Breslow method is used to handle tied observations. ME(std) shows the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation.				

An alternative way to illustrate the size of the estimated effect is by means of the survivor function, which shows how the share of individuals that have not yet found employment changes with the duration of search. Figure 3 plots the survivor function for the 3-month period and for different values of the entry cohort: the solid line corresponds to the case in which all explanatory variables are equal to zero (for the cohort-size variable and the unemployment rate this implies that they are equal to their mean), while the dashed line above (below) shows the survivor function when the entry cohort is smaller (larger) by one

standard deviation. Naturally, the survivor function is decreasing as the share of graduates finding employment increases with time; at the end of the observation period, between 50% and 55% of graduates have taken up employment. In line with the finding that larger entry cohorts increase the hazard rate of finding a job, Figure 3 shows that the share of survivors is generally smaller for larger cohorts. After 90 days the survivor function takes on a value of 52% when all variables are equal to zero, with the corresponding value equal to 54% (50%) when the size of the entry cohort is smaller (larger) by one standard deviation. A change in the relative number of apprenticeship graduates by one standard deviation is therefore predicted to change the share of individuals who are still searching after 90 days by 2 percentage points or, equivalently, by approximately 4%.



Another feature of the results presented in Table 2 is that the size of the cohort coefficient decreases as the period of observation is extended and the estimation results become increasingly affected by graduates with longer search durations (while they are always included in the sample, censored observations only contribute to the estimation by belonging to the set of individuals that are at risk of transition into employment). On the one hand this finding could be seen as evidence that the effect that the size of the entry cohort has on subsequent search durations is not persistent. On the other hand, it is conceivable that for those individuals that require more time to find employment current labour-market conditions matter in addition to the conditions prevailing at the time of graduation. To assess this hypothesis, the model of Equation 1 is supplemented

with measures of cohort size and the unemployment rate that refer to later points in time: 6 months after graduation in the case of the 6-month observation period, 6 and 12 months for the 1-year period as well as 6, 12, 18 and 24 months for the 2-year period.

Table 3 shows that once these measures for the current labour-market conditions are added, the cohort coefficient in the 6-month period is almost identical in terms of size and significance to the corresponding effect that is measured when only transitions occurring within three months after graduation are treated as such. This suggests that in this case the positive effect of the size of the entry condition on the hazard rate of finding employment continues to exist once current conditions are controlled for. Similar results are, however, not obtained for the two remaining observation periods as, first, the effect of the entry cohort decreases in magnitude and, second, all of the cohort coefficients are individually insignificant, though in the case of the 2-year period they remain jointly significant at the 5% level.

Table 3: Regression results (when current labour-market conditions are controlled for)

	3 months	6 months	1 year	2 years
Cohort	26.04** (11.20)	26.78** (12.77)	2.29 (19.63)	-3.94 (18.97)
Cohort (+6 months)	–	-1.00 (12.08)	-7.54 (11.70)	1.37 (19.51)
Cohort (+12 months)	–	–	23.30 (20.48)	19.77 (22.49)
Cohort (+18 months)	–	–	–	-3.42 (19.70)
Cohort (+24 months)	–	–	–	8.94 (22.36)
Control variables	Yes	Yes	Yes	Yes
Log pseudo-likelihood	-81,452.15	-103,498.96	-126,203.29	-138,947.97
Observations	18,133	18,133	18,133	17,466
Clusters	141	141	141	141
ME(std)	1.07**	1.07**	1.01	0.99
Variables refer to the time of graduation from apprenticeship training unless indicated otherwise. The set of control variables also includes current values of the unemployment rate. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The Breslow method is used to handle tied observations. ME(std) shows the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation.				

4.2 Discussion of the hypotheses

The results of Table 2 provide no support for the cohort-crowding hypothesis that members of larger entry cohorts have longer search durations; on the contrary, the empirical evidence suggests that graduating as part of a larger group reduces the time required to find a job. A possible explanation for this relationship, as discussed in Section 2, is that in the face of increased competition graduates from larger cohorts choose to take up lower-quality jobs. If entering the labour market as part of a larger group indeed pushes apprentices into jobs that do not match their qualifications, characteristics of the first employment spell should differ between graduates from large and from small cohorts. This hypothesis is assessed by means of two outcome variables: the natural logarithm of the average daily wage earned in the first employment spell and an indicator for whether this spell refers to regular employment subject to social-security contributions.⁹ These variables are regressed on the size of the entry cohort as well as on the set of control variables used in the estimation of Equation 1.

Table 4: Cohort-size effects on wages and regular employment status

	All	3 months	6 months	1 year	2 years
<i>Ln(average daily wage)</i>					
Cohort	2.57 (6.19)	11.61 (9.42)	7.19 (8.15)	7.29 (7.40)	1.40 (6.69)
Control variables	Yes	Yes	Yes	Yes	Yes
Observations	17,995	8,567	10,985	13,586	15,938
Clusters	141	141	141	141	141
<i>Indicator for regular employment</i>					
Cohort	2.91 (3.04)	5.60 (4.17)	3.87 (4.09)	3.42 (3.55)	0.98 (2.96)
Control variables	Yes	Yes	Yes	Yes	Yes
Observations	18,133	8,605	11,056	13,685	16,057
Clusters	141	141	141	141	141
Estimation by ordinary least squares (OLS). All variables refer to the time of graduation from apprenticeship training. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region.					

Table 4 contains the estimated cohort-size coefficients for the full set of observations as well as separately for those individuals that fall into each of

9 The smaller number of observations in the top panel is due to some individuals being assigned wages of a value zero. For approximately 76% of observations with a strictly positive search duration the first employment spell is of the regular type, with 18% being registered as working in marginal employment (*Geringfügige Beschäftigung*) and 5% having started a new apprenticeship.

the four periods of observation. The top panel reports the results pertaining to average daily wages, while the effects on the probability of being in non-regular employment are recorded in the bottom half. The results do not support the hypothesis that graduates from larger cohorts are pushed into lower-quality jobs. If anything, the findings suggest that the size of the graduation cohort is positively associated with the wages earned in the first job as well as with the probability of being in regular employment, though none of the estimated coefficients is statistically different from zero.

The positive impact of cohort size on the hazard rate of finding employment and the lack of evidence in support of the hypothesis that graduating as part of a large group drives apprentices into lower-quality jobs leaves the possibility that firms react to changes in the number of apprenticeship completers by creating jobs, though this effect appears to be restricted to a relatively short period after graduation. This explanation would appear to challenge findings by Garloff et al. (2013) whose empirical analysis for West German labour-market regions shows that larger cohorts increase the overall unemployment rate.¹⁰ If firms creating jobs in expectation of large entry cohorts is indeed the explanation for the finding that members of larger cohorts have shorter search durations, the empirical evidence presented in Tables 2 and 3 suggests that these beneficial effects are limited to a period of about six months following graduation.

4.3 Alternative explanations

Two alternative explanations for the findings should be considered which address the role of regional selection following graduation and changes in the composition of the group of graduates. First, the positive relationship between the hazard rate and the size of the graduation cohort could be spurious if it is driven by apprentices that graduate as part of a large group choose to search for employment in regions where search durations are shorter. In the sample of graduates with strictly positive search durations approximately 31% of individuals register their first employment spell in a different region to the one in which they have graduated. If belonging to a large entry cohort induces some individuals to search for employment elsewhere, the size of the graduation cohort and the probability of finding employment in a different region should be positively related.

Table 5 shows the results from regressing a binary dependent variable that takes the value 1 if the region of an individual's first employment spell is not the

¹⁰ Their use of the overall unemployment rate as the dependent variable and the focus on the share of young individuals aged between 15 and 24 rather than on the number of graduates from an apprenticeship programme, however, limit the comparability to this paper.

same as the one in which he graduated on the set of explanatory variables that are used in the estimation of Equation 1. The cohort-size coefficients, however, provide no evidence for the hypothesis that selecting into another region after graduation is the reason for the results of Table 2 as none of the estimated effects are significantly different from zero.¹¹

Table 5: Cohort-size effects on the probability of finding employment in a different region

<i>Indicator for employment in a different region</i>	All	3 months	6 months	1 year	2 years
Cohort	1.56 (3.08)	-1.73 (4.88)	0.07 (3.87)	0.60 (3.26)	1.70 (3.13)
Control variables	Yes	Yes	Yes	Yes	Yes
Observations	18,133	8,605	11,056	13,685	16,057
Clusters	141	141	141	141	141

Estimation by ordinary least squares (OLS). All variables refer to the time of graduation from apprenticeship training. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region.

Second, the composition in terms of productivity may differ between small and large graduation cohorts. If the number of graduates that are employed directly is fixed, some highly productive individuals will have to engage in job search if they belong to a larger cohort. In such a scenario the fact that search durations are shorter in larger cohorts might be the result of a change in the productivity composition of the cohort towards more individuals with a higher level of productivity. More productive graduates are likely to find employment and to require less time to do so, which might explain the positive cohort-size coefficients, especially shortly after graduation.

This hypothesis is assessed by estimating the effect of the size of the entry cohort on the probability of having a strictly positive search duration. If the above argument is correct, belonging to a larger cohort should be associated with a higher probability of having to search for employment. Table 6 shows the results from regressing a binary indicator for whether an individual has to search on the same set of explanatory variables as used in Equation 1.¹² Compared to Table 2 the number of observations increases because those individuals with a zero search duration are now also included. The coefficient of the entry cohort is significant

11 Similar conclusions can be drawn from estimating a logit model instead of a linear probability model. The results are available upon request.

12 Similar conclusions can be drawn from estimating a logit model instead of a linear probability model. The results are available upon request.

only at the 10% level and suggests that belonging to a larger group of graduates reduces the probability of having a positive search duration. This effect, however, appears to be small with an increase in the size of the graduation cohort by one standard deviation being predicted to increase the probability of search by one percentage point compared to a mean value of the dependent variable of 0.39. The hypothesis that the results of Table 2 reflect a change in the productivity composition of the group of individuals in larger graduation cohorts that have to engage in search is therefore not supported by the data.

Table 6: Cohort-size effects on the probability of having a strictly positive search duration

<i>Indicator for having a strictly positive search duration</i>	All
Cohort	-4.01 (2.31)
Control variables	Yes
Observations	46,408
Clusters	141
<small>Estimation by ordinary least squares (OLS). All variables refer to the time of graduation from apprenticeship training. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The smaller number of observations in the top panel is due to some individuals being assigned wages of a value zero.</small>	

4.4 Inclusion of individuals with zero search duration

The use of survival models prevents the inclusion of individuals that are employed upon graduation and therefore have a zero search duration. As discussed in Section 3.2, omitting this set of observations potentially raises a problem of sample selection if the two groups of individuals differ in terms of unobserved characteristics which in turn may have an effect on their employability. In order to assess the impact of this selection on the estimated effect of cohort size, the search-duration variable is adjusted by adding 1 to each value (and adjusting the censoring variables accordingly). Doing so allows the inclusion of those individuals for whom search duration is actually zero in the estimation of the Cox model. The results of this analysis are presented in Table 7.

Table 7: Regression results (when individuals with zero search durations are included)

	3 months	6 months	1 year	2 years
Cohort	11.58*** (4.46)	7.82* (4.33)	7.04 (4.92)	4.35 (4.77)
Unemployment rate	-2.46*** (0.38)	-2.12*** (0.38)	-2.24*** (0.41)	-2.10*** (0.40)
Control variables	Yes	Yes	Yes	Yes
Log pseudo-likelihood	-384,802.52	-406,547.02	-429,408.37	-448,434.23
Observations	46,408	46,408	46,408	46,408
Clusters	141	141	141	141
ME(std)	1.03***	1.02*	1.02	1.01

All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The Breslow method is used to handle tied observations. ME(std) shows the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation.

In terms of their pattern the estimated coefficients are comparable to the results of Table 2: the coefficients are positive and decrease in size as the observation period becomes longer; moreover, the effects are significant for the 3-month period, but also for the 6-month period. The main difference is that the coefficients are smaller, suggesting that once those individuals are included that are employed directly upon graduation, the strength of the relationship between the size of an individual's graduation cohort and the duration of her search for employment is reduced. A possible explanation for the weaker relationship between cohort size and search duration is that a number of graduates will always be employed directly regardless of the size of their graduation cohort. This explanation is in line with the results of Table 6, which show that the probability of having to search (i.e. of not becoming employed directly) is only marginally affected by the number of apprentices completing training.¹³

5 Conclusion

How the size of the cohort that an individual belongs to affects his contemporaneous labour-market outcomes constitutes a widely analysed field of research, with particular attention being paid to the effects on wages as well as employment and unemployment. In contrast, how cohort size measured at a specific point in an

¹³ An alternative way of including observations with zero search durations is to estimate count-data models. The results of these models, which are available upon request, also provide evidence that members of larger graduation cohorts have shorter search durations.

individual's career affects future outcomes has so far not attracted a large amount of attention, while there has recently been a substantial amount of research on the effect of the state of the business cycle at the time of labour-market entry on an individual's subsequent wages and employment opportunities. The contribution of this paper to the cohort-size literature is to analyse the effect on the amount of time an individual spends searching for employment after entering the labour market, which represents an outcome that has so far not been addressed. Moreover, in doing so, this paper conceptualises the size of the cohort as a factor affecting the conditions under which an individual's entry to the labour market takes place rather than as a contemporaneous explanatory variable.

From a theoretical perspective the relationship between the size of the entry cohort and an individual's subsequent duration of search can take various forms. Longer durations would be expected if increased competition makes it harder for members of larger cohorts to find employment – a relationship that would be in line with the standard cohort-crowding hypothesis. Individuals, however, may counteract this effect if they are willing to downgrade by taking up employment in a lower-quality job. Finally, if large cohorts indeed lead to lower unemployment rates, as has been argued by parts of the cohort-size literature, a negative impact on search durations is also conceivable. As such the results of the analysis may not only shed light on the relationship between cohort size at the point of labour-market entry and the subsequent duration of search, but may also provide insights into the former's effect on employment and unemployment outcomes. Since economic theory does not provide a clear indication on the nature of the relationship, the above hypotheses are assessed by means of an empirical analysis. The sample is based on register data and consists of graduates from Germany's apprenticeship programme who completed their training between January 1999 and October 2012.

The results of the empirical analysis suggest that the hazard rate of finding employment increases with the size of the cohort as part of which an individual graduates and enters the labour market. While this effect appears to apply only to individuals that find employment within a relatively short period of three months following graduation, once contemporaneous economic conditions are controlled for this effect is also found for individuals who take up a job within six months of graduating. Overall, the empirical analysis provides no evidence to suggest that members of larger entry cohorts suffer depressed labour-market outcomes in terms of longer search durations. Further analyses show that shorter search durations among members of larger graduation cohorts are not associated with employment in lower-quality jobs as there is no empirical evidence for a negative effect of cohort size on wages or on the probability of finding regular employment. The possibility that the observed effects are driven by either selection into regions

with better employment opportunities (and hence shorter search durations) after graduation or changes in the productivity composition of those graduates that have to search for employment is also not supported by the data. Finally, the fact that those apprentices who find employment directly upon graduation cannot be included in the baseline results does not appear to materially affect the conclusions regarding the nature of the relationship between cohort size and search duration.

A possible explanation for the positive effect of the size of the entry cohort on the hazard rate of finding employment is that firms anticipate such changes in the supply of young workers and react by creating jobs which in turn causes a shorter duration of search. Such an interpretation would be compatible with the view that larger cohorts also lead to lower unemployment rates, which potentially challenges existing evidence for (Western) Germany that larger cohorts are associated with higher unemployment rates. In light of the demographic processes that are projected to lead to a lower share of young age groups in the population and of a rising preference for tertiary education, future cohorts of apprenticeship graduates may be expected to decrease in size, which, at least according to this analysis, would suggest that search durations might become longer in the future.

Acknowledgements

I would like to thank Bernd Hayo as well as the participants of IAB's regional research network meeting, the 2016 Congress of the European Regional Science Association (ERSA), the XXXI National Conference of the Italian Association of Labour Economists (AIEL), the 63rd Annual North American Meetings of the Regional Science Association International (NARSC).

References

- Berger, M.C. (1983) Changes in Labor Force Composition and Male Earnings: A Production Approach, *Journal of Human Resources*, **18**, 177–96.
- Biagi, F. and Lucifora, C. (2008) Demographic and education effects on unemployment in Europe, *Labour Economics*, **15**, 1076–101.
- Brunello, G. (2010) The effects of cohort size on European earnings, *Journal of Population Economics*, **23**, 273–90.
- Brunner, B. and Kuhn, A. (2014) The impact of labor market entry conditions on initial job assignments and wages, *Journal of Population Economics*, **27**, 705–38.
- Card, D. and Lemieux, T. (2001) Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis, *Quarterly Journal of Economics*, **116**, 705–46.

- Cockx, B. and Ghirelli, C. (2016) Scars of recessions in a rigid labor market, *Labour Economics*, 41, 162–76.
- Dooley, M. (1986) The Overeducated Canadian? Changes in the Relationship among Earnings, Education, and Age for Canadian Men: 1971–81, *The Canadian Journal of Economics/Revue canadienne d'Economique*, 19, 142–59.
- European Commission (2014) *The 2015 Ageing Report*, European Commission, Brussels.
- Freeman, R.B. (1979) The effect of demographic factors on age-earnings profiles, *Journal of Human Resources*, 14, 289–318.
- Garloff, A., Pohl, C. and Schanne, N. (2013) Do small labor market entry cohorts reduce unemployment?, *Demographic Research*, 29, 379–406.
- Garloff, A. and Roth, D. (2016) Regional population structure and young workers' wages, IAB Discussion Paper No. 06/2016, Institute for Employment Research, Nuremberg.
- Kahn, L.B. (2010) The long-term labor market consequences of graduating from college in a bad economy, *Labour Economics*, 17, 303–16.
- Korenman, S. and Neumark, D. (2000) Cohort Crowding and Youth Labor Markets: A Cross-National Analysis, in D.G. Blanchflower and R.B. Freeman (eds.) *Youth Employment and Joblessness in Advanced Countries*, University of Chicago Press, Chicago.
- Kosfeld, R. and Werner, A. (2012) Deutsche Arbeitsmarktregionen – Neuabgrenzung nach den Kreisgebietsreformen 2007–2011, *Raumforschung und Raumordnung*, 70, 49–64.
- Michaelis, J. and Debus, M. (2011) Wage and (un-)employment effects of an ageing workforce, *Journal of Population Economics*, 24, 1493–511.
- Moffat, J. and Roth, D. (2016a) The Cohort Size–Wage Relationship in Europe, *Labour*, forthcoming.
- Moffat, J. and Roth, D. (2016b) Cohort size and youth labour-market outcomes: the role of measurement error, IAB Discussion Paper No. 37/2016, Institute for Employment Research, Nuremberg.
- Morin, L.-P. (2015) Cohort size and youth earnings: Evidence from a quasi-experiment, *Labour Economics*, 32, 99–111.
- Mosca, I. (2009) Population Ageing and the Labour Market, *Labour*, 23, 371–95.
- Shimer, R. (2001) The Impact of Young Workers on the Aggregate Labor Market, *The Quarterly Journal of Economics*, 116, 969–1007.
- Skans, O.N. (2005) Age effects in Swedish local labor markets, *Economics Letters*, 86, 419–26.
- Statistisches Bundesamt (2015) *Bevölkerung Deutschlands bis 2060: 13. koordinierte Bevölkerungsvorausberechnung*, Statistisches Bundesamt, Wiesbaden.

- Stevens, K.** (2007) Adult Labour Market Outcomes: the Role of Economic Conditions at Entry into the Labour Market, available at http://www.iza.org/conference_files/SUMS2007/stevens_k3362.pdf
- Welch, F.** (1979) Effects of Cohort Size on Earnings: The Baby Boom Babies' Financial Bust, *Journal of Political Economy*, **87**, S65–S97.
- Wright, R.E.** (1991) Cohort size and earnings in Great Britain, *Journal of Population Economics*, **4**, 295–305.

Appendix

Table A1: Descriptive statistics

Variable	Observations	Mean	Standard deviation	Minimum	Maximum
Search (3 months)	18,133	63.716	33.848	1	91.250
Search (6 months)	18,133	104.494	72.576	1	182.500
Search (1 year)	18,133	162.522	144.184	1	365
Search (2 years)	18,133	222.391	246.219	1	730
Cohort	18,133	0.009	0.003	0.004	0.020
Unemployment rate	18,133	0.120	0.053	0.024	0.317
German	18,133	0.954	0.209	0	1
<i>Age</i>					
19	18,133	0.144	0.351	0	1
20	18,133	0.292	0.455	0	1
21	18,133	0.278	0.448	0	1
22	18,133	0.175	0.380	0	1
23	18,133	0.111	0.314	0	1
<i>Occupations</i>					
1	18,133	0.041	0.199	0	1
2	18,133	0.495	0.500	0	1
3	18,133	0.223	0.416	0	1
4	18,133	0.023	0.148	0	1
5	18,133	0.034	0.181	0	1
6	18,133	0.084	0.277	0	1
7	18,133	0.068	0.251	0	1
8	18,133	0.023	0.149	0	1
9	18,133	0.010	0.101	0	1
<i>Industries</i>					
1	18,133	0.028	0.166	0	1
2	18,133	0.000	0.018	0	1
3	18,133	0.005	0.067	0	1
4	18,133	0.196	0.397	0	1
5	18,133	0.005	0.069	0	1
6	18,133	0.214	0.410	0	1
7	18,133	0.208	0.406	0	1
8	18,133	0.056	0.229	0	1
9	18,133	0.026	0.158	0	1
10	18,133	0.011	0.106	0	1
11	18,133	0.041	0.197	0	1
12	18,133	0.021	0.144	0	1
13	18,133	0.122	0.327	0	1
14	18,133	0.032	0.176	0	1
15	18,133	0.036	0.185	0	1
16	18,133	0.000	0.013	0	1
17	18,133	0.000	0.011	0	1

Table A2: Classification of occupations and industries

Occupations	
1	Agriculture, forestry, farming and gardening
2	Production of raw materials and goods, and manufacturing
3	Construction, architecture, surveying and technical building services
4	Natural sciences, geography and informatics
5	Traffic, logistics, safety and security
6	Commercial services, trading, sales, the hotel business and tourism
7	Business organisation, accounting, law and administration
8	Health care, the social sector, teaching and education
9	Philology, literature, humanities, social sciences, economics, media, art, culture, and design
Industries	
1	Agriculture and forestry
2	Fishery
3	Mining and quarrying
4	Manufacturing
5	Electricity and water supply
6	Construction
7	Sale, maintenance and repair
8	Tourism
9	Transport
10	Financial and insurance services
11	Real estate
12	Public administration and defence
13	Education
14	Health and social work
15	Other services
16	Households
17	Extraterritorial organisations

The occupation and industry classifications are derived from the *Klassifikation der Berufe 2010* and the *Wirtschaftszweige 1993*, respectively.

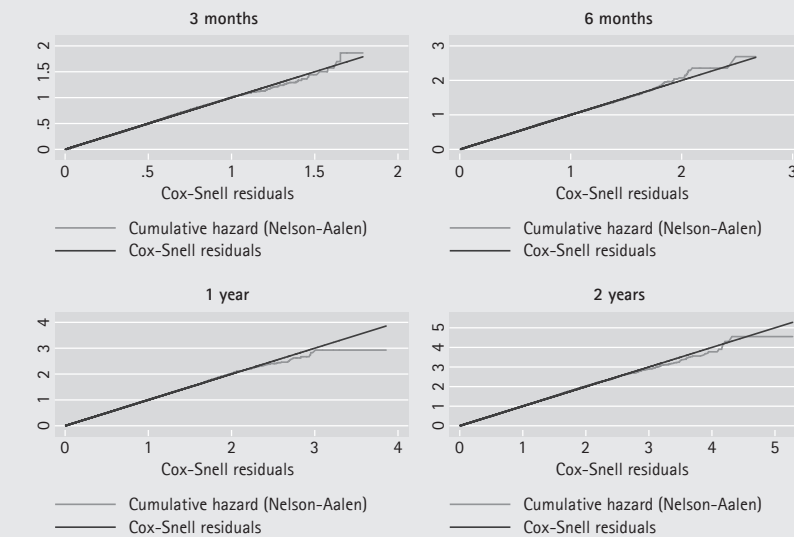
Supplementary material

This document's purpose is to assess the validity of the paper's empirical model as well as the robustness of the results. This is done by first performing a set of residual-based tests concerning the specification of the Cox model. Second, the sensitivity of the results is analysed by estimating different variations of the initial Cox model as well as by presenting the results from different fully parametric specifications and comparing them with those of the paper's semi-parametric Cox model.

S1 Model specification

The Cox model gives rise to different kinds of residuals which form the basis for testing the adequacy of the specified model (for further detail on the tests and the different residuals see Box-Steffensmeier and Jones, 2004 or Cleves et al., 2010). The first specification test assesses how well the model proposed in Equation 1 of the paper fits the data (this is done separately for each of the common observation periods of 3 months, 6 months, 1 year and 2 years, respectively). This test is based on the Cox-Snell residuals, which can be derived from the estimated coefficients and the estimated cumulative baseline hazard rate and which can be interpreted as the number of transitions that an individual is expected to have experienced (assuming he can repeatedly experience transitions) within the time that it actually takes the individual to find employment. If the model is correctly specified, these residuals will follow a unit-exponential distribution with a hazard rate that is equal to 1. To assess whether this is the case, the cumulative hazard function of the Cox-Snell residuals is estimated (using the Nelson-Aalen estimator) and is then plotted against the residuals. For the correct model specification this estimate will be close to the 45-degree line.

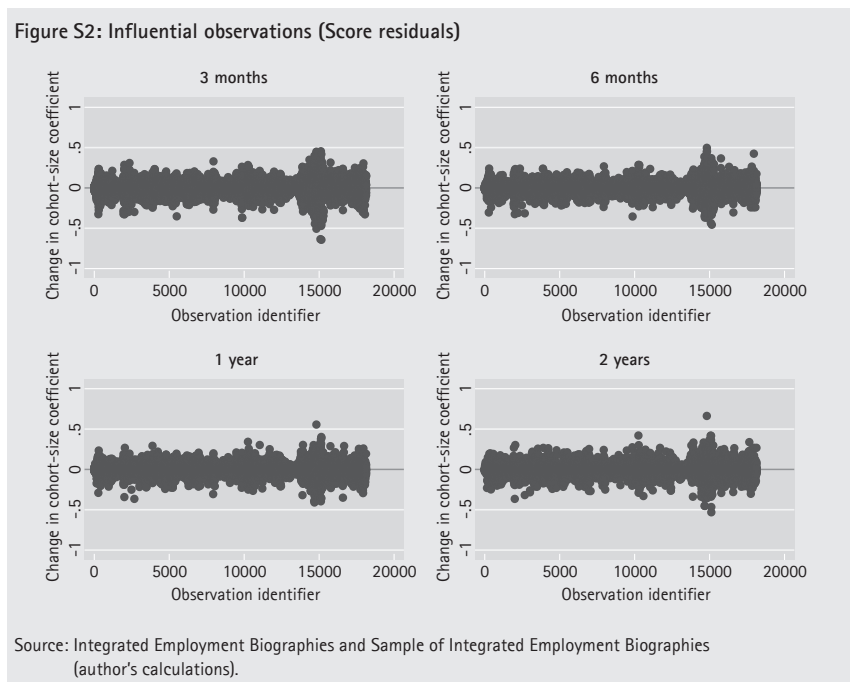
Figure S1: Overall model fit (Cox-Snell residuals)



Source: Integrated Employment Biographies and Sample of Integrated Employment Biographies (author's calculations).

Figure S1 shows the estimated cumulative hazard functions of the Cox-Snell residuals for each observation period against the 45-degree line. In each case the estimate lies close to this line, which indicates that the model provides a reasonable fit. While deviations from the 45-degree line can be found for higher values of the Cox-Snell residuals, this is likely to reflect the fact that the number of individuals for whom many transitions are expected will be relatively small.

Next, it is tested to what extent the coefficient of the cohort-size variable is driven by any single observation. Instead of estimating the model separately after successively excluding one observation, the effect on the coefficient is approximated by multiplying the matrix of score residuals with the variance-covariance matrix of the estimated coefficients. As shown in Figure S2, the largest (absolute) change in the cohort-size coefficient is about 0.5 units when a single observation is dropped, which is a small change in light of an estimated coefficient that ranges between 26 (3-month period) and 6 (2-year period).



The final test concerns the proportional-hazards property. The Cox model belongs to a class of models for which the hazard rate can be decomposed into one component that depends only on time (i.e. the duration of search) and which is given by the baseline hazard $h_0(t)$ and another component which is a function of the model's covariates

and their coefficients given by the term $\exp(\delta'x_{irp})$. Changes in the covariates are therefore expected to shift the baseline hazard up (or down) in a parallel way. This assumption is testable on the basis of yet another type of residuals, the Schoenfeld residuals. The former are observation-specific and covariate-specific and can be interpreted as the difference between the observed and the expected value of a covariate. The test can be performed globally (i.e. for the whole set of covariates) as well as locally for individual regressors. Table S1 shows the test-statistics and p-values associated with the cohort-size variable as well as with the whole set of regressors. Looking at the 3-month period it can be seen that the null hypothesis of the proportional-hazards assumption being satisfied is not rejected locally for the cohort-size variable. However, the test statistic for the global test is sufficiently large that the null hypothesis is rejected at the 5%-level. For all other observation periods, the results of the tests suggest that the proportional-hazard assumption is violated for the cohort-size variable as well as for the whole set of covariates. A possible response to the null hypothesis of the proportional-hazards assumption being violated is to allow for the effects of the regressors to vary with time by including interactions with search duration (Box-Steffensmeier and Jones, 2004). In order to allow as flexible an approach as possible and to avoid that certain variables pick up the effect of other regressors for which no interactions have been included, a model should be specified that contains interactions with search duration for every regressor. However, given the relatively large number of control variables, the estimation procedure for such a model does not converge. If a model is estimated that only interacts the size of the graduation cohort with the duration of search, the former variable's effect on the hazard rate is found to decrease in magnitude with time spent searching. Due to the above-mentioned concerns about models in which interactions are only included for a subset of regressors (and in this case only for a single regressor) this approach is not pursued any further. In order to justify the use of the paper's results I argue that the main conclusions about the effects of entry-cohort size are based on the 3-month period for which the proportional-hazard assumption appears to be locally satisfied.

Table S1: Proportional-hazard test (Schoenfeld residuals)

	3 months	6 months	1 year	2 years
Local test				
Cohort size	0.50 (0.48)	15.71 (0.00)	15.39 (0.00)	23.93 (0.00)
Global test	170.59 (0.04)	454.85 (0.00)	226.01 (0.00)	252.17 (0.00)
Tests statistics are based on Harrell's rho for the local test and on the Grambsch and Therneau method for the global test. Test statistics are χ^2 -distributed. P-values are in parentheses.				

S2 Sensitivity analysis

This subsection consists of two parts. In the first part, the robustness of the Cox model's results is assessed against various changes in the sample as well as in the empirical specification. The aim of the second part is to compare the results from the Cox model, which does not make an assumption regarding the specific distribution of the search durations, with a number of models that assume that the search durations follow a particular distribution.

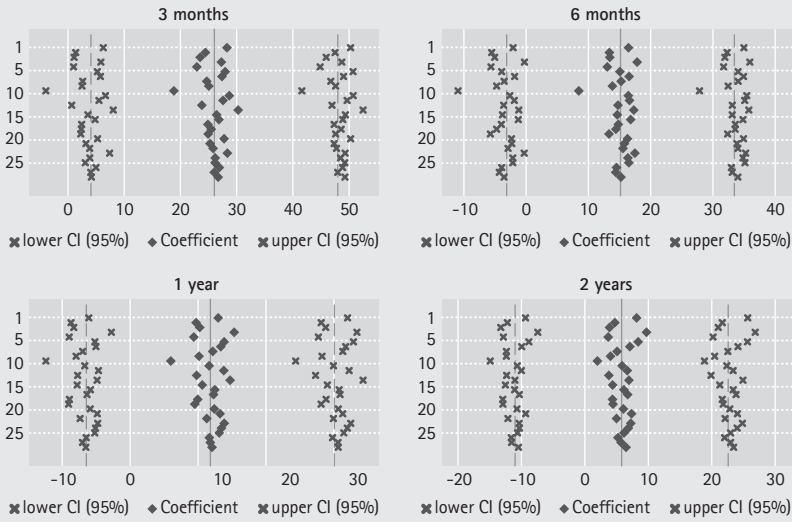
S2.1 Robustness checks on the Cox model

The first set of robustness checks continues with assessing to what extent the results are driven by influential observations. However, in contrast to the above approach in which individual observations were excluded from the sample (see Figure S2) this analysis successively drops all observations from a given graduation period. A similar analysis is then performed in which observations from individual graduation regions are excluded from the sample.

Figure S3 shows the cohort-size coefficient and the corresponding 95% confidence interval for each case in which one of the 28 graduation periods is excluded. In order to allow for a comparison with the results from the full sample, the former's coefficient and 95% confidence interval is included in form of vertical lines. In each case the estimated cohort-size coefficient lies within the full model's confidence interval (represented by the dashed lines) and is typically close to the coefficient of the full model. An exception is period 10, which corresponds to the graduation period May–October 2003, where the change is more pronounced and the coefficient becomes insignificant.

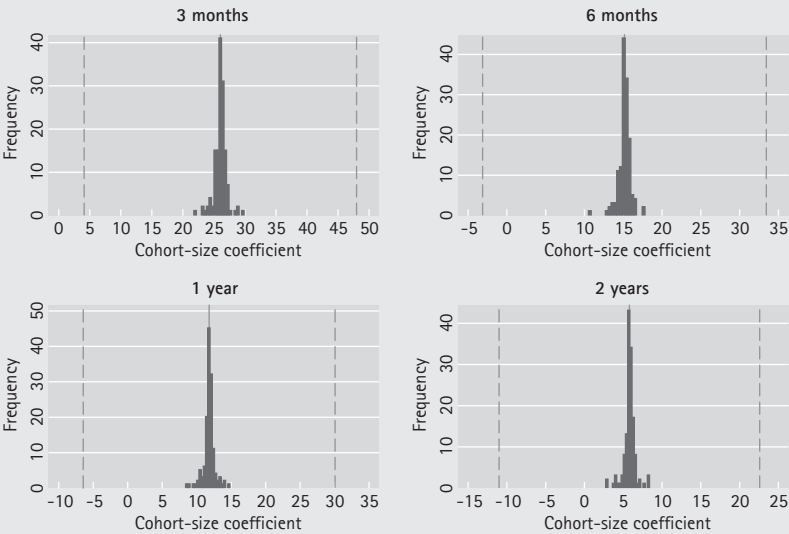
Due to the large number of regions (141) the results from omitting individual graduation regions are shown in form of a histogram which illustrates the distribution of the resulting cohort-size coefficients. As can be seen from Figure S4, there is a certain degree of variation around the coefficient from the full model for each observation period, but these differences are small when compared to the confidence interval of the full model's coefficient. Moreover, for the 3-month observation period all cohort-size coefficients are significant at the 5% level, while in the 6-month period some of the coefficients become significant at the 10% level.

Figure S3: Exclusion of single graduation periods from the sample



Source: Integrated Employment Biographies and Sample of Integrated Employment Biographies (author's calculations). Cohort-size coefficients are estimated as in Equation 1. The solid line represents the cohort-size coefficients from the full model, the dashed lines the corresponding 95% confidence interval.

Figure S4: Exclusion of individual graduation regions from the sample



Source: Integrated Employment Biographies and Sample of Integrated Employment Biographies (author's calculations). Cohort-size coefficients are estimated as in Equation 1. The solid line represents the cohort-size coefficients from the full model, the dashed lines the corresponding 95% confidence interval.

In the paper the sample is homogenised by only including individuals that are aged between 19 and 23 at the time of graduation. As can be seen from Table S2, comparable results are obtained when this restriction is not imposed: the cohort-size coefficient remains positive and significant for the first observation period; moreover, it falls in size and becomes insignificant when the observation period is increased. For the first three periods of observation the coefficients are between 15% and 20% smaller than those of the full model, though the absolute changes are always considerably smaller than the size of the estimated standard errors.

Table S2: Dropping the age restriction

	3 months	6 months	12 months	24 months
Cohort	22.09** (9.31)	11.95 (8.10)	9.64 (8.20)	3.66 (7.76)
<i>Dummies</i>				
Occupation	Yes	Yes	Yes	Yes
Industry	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes
Log pseudo-likelihood	-105,389.80	-134,403.74	-163,907.26	-187,558.95
Observations	22,828	22,828	22,828	22,828
Clusters	141	141	141	141
ME(std)	1.06**	1.03	1.03	1.01
All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The Breslow method is used to handle tied observations. ME(std) shows the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation.				

Search durations are measured in days and therefore transitions into employment are only observed at discrete points in time (even though the underlying process that generates the search series may be continuous). Under these conditions observations might be tied, i.e. there may be two or more observations with the same observed search duration. In order to account for tied observations the partial likelihood function has to be adjusted. Specifically, an adjustment has to be made to the definition of the risk set, i.e. the set of individuals that are at risk of experiencing a transition into employment at any given value of search duration. Assuming that the true search duration is indeed continuous, two individuals with the same observed search duration will in fact not have experienced transition into employment at the same point in time. If this is the case, both individuals will be in the risk set at the time of the first transition, but at the time of the second transition the individual that found employment earlier will no longer be part of the risk set. The Breslow method for handling tied observations, which has been

used in the empirical analysis up to this point, does not make this distinction and instead assumes that the risk set is the same for all individuals sharing the same observed search duration. In contrast, the Efron method takes into account that sequential transitions give rise to different risk sets. Table S3 shows that employing the Efron method yields estimated coefficients and standard errors for the cohort-size coefficient that are very similar to those derived from the Breslow method.

Table S3: Efron method

	3 months	6 months	12 months	24 months
Cohort	26.12** (11.25)	15.20 (9.37)	11.85 (9.35)	5.82 (8.60)
<i>Dummies</i>				
Occupation	Yes	Yes	Yes	Yes
Industry	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes
Log pseudo-likelihood	-81,407.13	-103,453.55	-126,157.62	-145,272.21
Observations	18,133	18,133	18,133	18,133
Clusters	141	141	141	141
ME(std)	1.07**	1.04	1.03	1.02
All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region. The Efron method is used to handle tied observations. ME(std) shows the proportional change in the hazard rate for an increase in the size of the graduation cohort by one standard deviation.				

S2.2 Parametric model specifications

The empirical analysis of the paper uses the Cox model, which represents an example of a semi-parametric estimation approach as the functional form of the relationship between the hazard rate and the covariates is parameterised, whereas the actual distribution of failure times is left unspecified. Despite not specifying a distribution the Cox model is able to consistently estimate the effect that changes in the covariates have on the hazard rate. Alternatively, a fully parametric approach can be employed, which makes an assumption about the type of distribution from which search durations are drawn. A drawback of this approach is that the validity of the results depends on having chosen the correct distribution function.

In the following, the robustness of the Cox model's results is assessed against specifying a particular distribution. The first three distributions considered – the exponential, the Weibull and the Gompertz distribution – are compatible with the metric in which the Cox model is formulated, i.e. they can also be specified

in form of a model of the hazard rate and, moreover, the first two also share the proportional-hazards property of the Cox model. The hazard functions that can be derived from these distributions take the following form:

$$h_i(t) = h_0(t) e^{(\gamma \text{cohort}_{rp} + \delta' x_{irp})} = e^{(\alpha)} e^{(\gamma \text{cohort}_{rp} + \delta' x_{irp})} \quad [S1]$$

$$h_i(t) = h_0(t) e^{(\gamma \text{cohort}_{rp} + \delta' x_{irp})} = p t^{p-1} e^{(\alpha)} e^{(\gamma \text{cohort}_{rp} + \delta' x_{irp})} \quad [S2]$$

$$h_i(t) = h_0(t) e^{(\gamma \text{cohort}_{rp} + \delta' x_{irp})} = e^{(\gamma t)} e^{(\alpha)} e^{(\gamma \text{cohort}_{rp} + \delta' x_{irp})} \quad [S3]$$

These specifications differ from the Cox model in that the baseline hazards are fully parameterised and depend on the constant of the model α as well as, in the case of the Weibull and the Gompertz specifications, on a distribution-specific shape parameter that also has to be estimated. These parameters determine the shape of the baseline hazard, i.e. the predicted hazard rate when all covariates take on a value of zero. Under the Weibull and the Gompertz specification the baseline hazard may be a flat, monotonically increasing or monotonically decreasing function of search duration. In contrast, the exponential model is less flexible in this respect as the baseline hazard is invariably flat, which may not be a realistic prior assumption if the probability of finding employment decreases with the duration of search.

For each of the above models and observation period, Table S4 shows the coefficient of the cohort-size variable as well as the estimated auxiliary parameters. The size of the coefficients can be compared directly with the results from Table 2 in the paper. The results show that the estimated effect of cohort size on search duration is robust to the use of a fully parametric specification, with each of the three distributions yielding coefficients that are similar in size and significance to those of the Cox model. Moreover, the negative sign of the auxiliary parameters in the Weibull and the Gompertz model suggests that the baseline hazard is decreasing with the duration of search, implying that, *ceteris paribus*, the hazard of finding employment is lower the longer the duration of search.

A prominent feature of the distributions in Table S4 is that they could be expressed in the same metric as the Cox model, i.e. in terms of the hazard rate. However, obtaining such an expression is not possible for all parametric models and the following set of examples – the lognormal, the loglogistic and the generalised Gamma distribution – are instead expressed in the accelerated failure time (AFT) metric: instead of estimating the effect of a change in a specific covariate on the hazard rate, its effect on the survivor function is estimated, which shows the share of observations that have not experienced transition into employment for each value of search duration. Estimates from the hazard rate can be derived

from the survivor function and in contrast to the above group of distributions the former may take on a non-monotonic shape. The generalised Gamma distribution in particular allows for more flexibility as its shape is determined by two auxiliary parameters. Since these models are parameterised in terms of the survivor function, the coefficients cannot be directly compared with the results of Table 2. However, if the relationship between cohort size and search durations in these models has the same sign as estimated by the Cox model, the cohort-size coefficients should have the opposite sign as in Table 2: if an increase in the size of the graduation cohort increases the hazard of finding employment, the corresponding effect on the survivor function should be negative since a larger instantaneous probability of finding employment should lead to a smaller share of individuals not having experienced a transition at any given time.

Table S4: Parametric models (Exponential, Weibull, Gompertz)

	3 months	6 months	12 months	24 months
<i>Exponential</i>				
Cohort	27.02** (11.85)	16.71 (10.50)	12.90 (11.12)	5.61 (11.28)
Log pseudo-likelihood	-24,167.63	-29,295.60	-33,706.08	-36,401.95
<i>Weibull</i>				
Cohort	26.25** (11.38)	15.39 (9.61)	11.82 (9.61)	5.63 (8.90)
Auxiliary parameter: ln(p)	-0.22*** (0.01)	-0.31*** (0.01)	-0.36*** (0.01)	-0.37*** (0.01)
Log pseudo-likelihood	-23,906.67	-28,519.32	-32,258.50	-34,407.30
<i>Gompertz</i>				
Cohort	26.14** (11.28)	15.31 (9.33)	12.49 (9.43)	7.08 (8.81)
Auxiliary parameter: γ	-0.01*** (0.00)	-0.01*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)
Log pseudo-likelihood	-23,852.43	-28,331.70	-32,351.38	-34,842.75
<i>Dummies</i>				
Occupation	Yes	Yes	Yes	Yes
Industry	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes
Observations	18,133	18,133	18,133	18,133
Clusters	141	141	141	141
All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed as proportional hazard estimates. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region.				

Table S5 contains the coefficients of the cohort-size variable as well as the auxiliary parameters estimated from a model based on the lognormal, the loglogistic and the generalised Gamma distribution, respectively. In order to compare the size of the estimated effects to those of the Cox model and to those for the first set of distributions, Table S5 also reports the results from a Weibull model, which may also be specified in terms of the survivor function and which, in the hazard-rate metric, yielded coefficients that were very close to those of the Cox model. As hypothesised, the coefficients of the cohort-size variable turn out negative for each period of observation, which implies that, as was the case with the Cox model, increases in the size of an individual's graduation cohort are associated with shorter search durations. Moreover, the size of the coefficient decreases as the observation period becomes longer, which also corresponds to the results from the Cox model. The results from the lognormal, the loglogistic and the generalised Gamma specifications differ, however, in that the coefficients tend to be larger than in the Weibull model, which for the hazard-rate metric produced coefficients that were very close to those of the Cox model: for the 3-month observation period the former are between 15% and 30% larger than those from the Weibull model, with the difference being larger for longer periods of observation. Moreover, the estimated effects of cohort size are also significant in the 6-month and the 1-year period of observation.

Table S5: Parametric models (Weibull, lognormal, loglogistic, generalised Gamma)

	3 months	6 months	1 year	2 years
<i>Weibull</i>				
Cohort	-32.62** (14.16)	-21.05 (13.14)	-17.00 (13.83)	-8.19 (12.94)
Auxiliary parameter $\ln(p)$	-0.22*** (0.01)	-0.31*** (0.01)	-0.36*** (0.01)	-0.37*** (0.01)
Log pseudo-likelihood	-23,906.67	-28,519.32	-32,258.50	-34,407.30
<i>Lognormal</i>				
Cohort	-42.01*** (16.08)	-33.24** (14.86)	-30.13** (14.71)	-24.45* (13.64)
Auxiliary parameter: $\ln(\sigma)$	0.64*** (0.01)	0.64*** (0.01)	0.62*** (0.01)	0.58*** (0.01)
Log pseudo-likelihood	-23,817.50	-28,301.08	-32,032.23	-34,322.21
<i>Loglogistic</i>				
Cohort	-37.02*** (15.60)	-28.23** (14.64)	-26.27* (14.43)	-22.63* (13.63)
Auxiliary parameter: $\ln(\gamma)$	0.06*** (0.01)	0.09*** (0.01)	0.07*** (0.01)	0.03*** (0.01)
Log pseudo-likelihood	-23,847.60	-28,349.18	-32,074.06	-34,370.44

	3 months	6 months	1 year	2 years
<i>Generalised Gamma</i>				
Cohort	-39.61** (15.70)	-31.28** (14.64)	-26.48* (14.45)	-16.75 (13.18)
Auxiliary parameter: $\ln(\sigma)$	0.55*** (0.02)	0.61*** (0.01)	0.57*** (0.01)	0.51*** (0.01)
Auxiliary parameter: κ	0.25*** (0.05)	0.15*** (0.04)	0.26*** (0.03)	0.43*** (0.03)
Log pseudo-likelihood	-23,804.83	-28,293.44	-31,997.53	-34,184.58
<i>Dummies</i>				
Occupation	Yes	Yes	Yes	Yes
Industry	Yes	Yes	Yes	Yes
Period	Yes	Yes	Yes	Yes
Region	Yes	Yes	Yes	Yes
Observations	18,133	18,133	18,133	18,133
Clusters	141	141	141	141

All variables refer to the time of graduation from apprenticeship training. Coefficients are expressed in the accelerated failure time metric. ***/**/* signifies significance at the 0.01/0.05/0.1 level of significance. Standard errors are clustered at the level of the labour-market region.

References

- Box–Steffensmeier, J.M. and Jones, B.S. (2004) *Event History Modeling: A Guide for Social Scientists*, Cambridge University Press, Cambridge.
- Cleves, M., Gutierrez, R.G., Gould, W. and Marchenko, Y.V. (2010) *An Introduction to Survival Analysis Using STATA (3rd edition)*, STATA Press, College Station.

Kurzfassung

Diese Arbeit befasst sich in vier Essays mit dem Zusammenhang zwischen regionalen Bevölkerungsstrukturen und verschiedenen Arbeitsergebnissen. Grundlage der empirischen Analysen bilden dabei entweder europäische Mikrodaten der *European Union Statistics on Income and Living Conditions* (EU-SILC) oder deutsche Mikrodaten der *Stichprobe der Integrierten Arbeitsmarktbiografien* (SIAB) und der *Integrierten Erwerbsbiografien* (IEB).

In jedem der vier Papiere stellt die Kohortengröße die zentrale erklärende Variable dar. Dabei handelt es sich um ein Maß für die Größe einer Gruppe, deren Mitglieder ein ähnliches Alter oder ein vergleichbares Maß an Arbeitserfahrung aufweisen, womit der Annahme Rechnung getragen wird, dass Personen unterschiedlichen Alters nur unvollständig miteinander substituierbar sind. Eine Veränderung in der Größe einer solchen Kohorte sollte zunächst deren Grenzproduktivität beeinträchtigen. Im Fall abnehmender Grenzproduktivität sollte ein Anstieg in der Kohortengröße auf Wettbewerbsmärkten dazu führen, dass die Löhne innerhalb der Kohorte sinken. Erfolgt jedoch keine vollständige Lohnanpassung, sind darüber hinaus Auswirkungen auf Beschäftigung und Arbeitslosigkeit in der Kohorte denkbar.

Das erste Papier befasst sich mit dem Zusammenhang zwischen der Größe einer Kohorte, deren Mitglieder neben einer ähnlichen Berufserfahrung auch ein vergleichbares Ausbildungsniveau aufweisen, und den gruppenspezifischen Löhnen. Der Beitrag des Papiers besteht darin, eine Identifikationsstrategie zu verwenden, die nicht nur Endogenität aufgrund von Ausbildungsentscheidungen, sondern auch Selektion durch regionale Migration adressiert. Die Ergebnisse zeigen, dass diese Identifikationsstrategie im Fall der größten Ausbildungsgruppe zu höheren negativen Lohnwirkungen führt.

Auch im zweiten Papier werden die Auswirkungen auf Löhne behandelt, jedoch werden andere Schwerpunkte gesetzt. Zunächst wird argumentiert, dass die Verwendung administrativer räumlicher Einheiten zu Messfehlern in der Kohortenvariable führen kann, da diese typischerweise keine Arbeitsmärkte abbilden. Die Ergebnisse legen nahe, dass die geschätzten Koeffizienten bei Verwendung administrativer statt funktionaler Einheiten tatsächlich nach unten verzerrt sind. Darüber hinaus weisen die Ergebnisse darauf hin, dass ein Teil des negativen Effekts auf Selektion in Berufe und Wirtschaftszweige zurückzuführen ist, in denen niedrigere Löhne gezahlt werden.

Thema des dritten Papiers ist der Zusammenhang von Kohortengröße und Arbeitslosigkeit sowie Beschäftigung. Der Beitrag besteht darin zu zeigen, dass sich die Effekte – in Bezug auf Vorzeichen und Größe – stark zwischen den be-

trachteten Altersgruppen unterscheiden. Um diesen Befund zu erklären, wird das Argument entwickelt, dass die für den Arbeitsmarkt relevante Größe einer Kohorte für junge Altersgruppen aufgrund umfangreicher Teilnahme an Ausbildungsmaßnahmen falsch gemessen wird.

Das vierte Papier befasst sich damit, wie sich die Größe der Kohorte beim Eintritt in den Arbeitsmarkt auf die Dauer bis zum Beginn der ersten Beschäftigung auswirkt – ein Zusammenhang, der in der bestehenden Literatur noch nicht untersucht worden ist und für den verschiedene Hypothesen denkbar sind. Die Ergebnisse zeigen, dass innerhalb relativ kurzer Beobachtungszeiträume die Suchdauer in größeren Eintrittskohorten kürzer ist. Darüber hinaus wird gezeigt, dass dieses Ergebnis nicht mit Selektion in schlechter bezahlte Beschäftigung einhergeht und auch nicht durch regionale Selektion oder durch eine veränderte Zusammensetzung größerer Kohorten hinsichtlich der Produktivität ihrer Mitglieder erklärt werden kann.

Abstract

This thesis consists of four essays which address the relationship between regional population structures and various labour-market outcomes. Each of the essays contains an empirical analysis which uses either European microdata from the *European Union Statistics on Income and Living Conditions* (EU-SILC) or German microdata from the *Sample of Integrated Employment Biographies* (SIAB) and the *Integrated Employment Biographies* (IEB).

The main explanatory variable in each of the four papers is a cohort-size variable, which measures the size of a group whose members are of a similar age or have acquired a comparable amount of work experience. This type of variable is therefore based on the assumption that members of different age groups are only imperfectly substitutable. A change in the size of a cohort should have an effect on the marginal product of that group. In the case of diminishing marginal productivity an increase in cohort size should lead to wage decreases if labour markets are perfectly competitive. However, in the absence of a complete wage adjustment, it is also possible that there are effects on cohort-specific employment and unemployment.

The first paper deals with the wage effects of the size of a cohort whose members not only have a comparable amount of work experience but also a common level of education. The paper's contribution consists of employing an identification strategy that not only addresses the possibility that cohort size may be endogenous due to educational choices but also due to regional migration. In the case of the largest educational group this identification strategy indeed produces estimated coefficients that are larger in magnitude.

The second paper also addresses the relationship between cohort size and wages but focuses on other aspects. First, the argument is proposed that a cohort-size variable that is based on administrative regional entities is potentially subject to measurement error because these entities typically do not represent labour markets. This argument appears to be supported by the results as the estimated coefficients turn out to be smaller when the cohort-size variable is derived from administrative rather than functional units. Second, the paper aims to shed light on the mechanisms through which the negative wage effect manifests itself. The results suggest that a considerable part of this negative effect can be ascribed to members of larger cohorts being more likely to be employed in lower-paying occupations and industries.

The topic of the third paper is the relationship between cohort size and employment as well as unemployment. It is shown that the estimated coefficients vary considerably across different age groups. To explain this finding the argument

is developed that cohort-size measures of young age groups potentially suffer from measurement error because a substantial part of these groups are likely to be enrolled in education and therefore unavailable to the labour market.

The fourth paper analyses the effect that cohort size at the time of labour-market entry has on the subsequent duration of search for employment – a relationship that has so far not been addressed by the extant literature and for which different hypotheses are conceivable. Within a relatively short period following labour-market entry, the results show that members of larger cohorts have shorter search durations. It is further shown that this result is neither due to selection into lower-quality jobs nor to regional selection nor to changes in the composition of the entry cohort in terms of the productivity of its members.

How does the size of young age cohorts affect the labour-market outcomes of these groups? Employing different microeconomic methods in an empirical analysis at the regional level, Duncan Roth addresses this question in the four essays contained in this book. The analysis deals with the effects on wages, employment and unemployment as well as the duration of search for employment following labour-market entry. Each chapter contributes to the extant literature by focussing on issues that, in the opinion of the author, have so far been insufficiently dealt with.

W. Bertelsmann Verlag



Institut für Arbeitsmarkt- und Berufsforschung der Bundesagentur für Arbeit
Regensburger Str. 100 ■ 90478 Nürnberg ■ www.iab.de