

Lane, Tom; Zhou, Minghai

Working Paper

Failure of unravelling theory? A natural field experiment on voluntary quality disclosure

CeDEx Discussion Paper Series, No. 2022-17

Provided in Cooperation with:

The University of Nottingham, Centre for Decision Research and Experimental Economics (CeDEx)

Suggested Citation: Lane, Tom; Zhou, Minghai (2022) : Failure of unravelling theory? A natural field experiment on voluntary quality disclosure, CeDEx Discussion Paper Series, No. 2022-17, The University of Nottingham, Centre for Decision Research and Experimental Economics (CeDEx), Nottingham

This Version is available at:

<https://hdl.handle.net/10419/284275>

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.

Discussion Paper No. 2022-17

Tom Lane and Minghai Zhou

December 2022

**Failure of unravelling theory? A
natural field experiment on
voluntary quality disclosure**

CeDEx Discussion Paper Series
ISSN 1749 - 3293



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit <http://www.nottingham.ac.uk/cedex> for more information about the Centre or contact

The full list of CeDEx Discussion Papers is available at

<http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx>

FAILURE OF UNRAVELLING THEORY? A NATURAL FIELD EXPERIMENT ON VOLUNTARY QUALITY DISCLOSURE*

Tom Lane[‡] and Minghai Zhou

December 2022

Abstract

Classic ‘unravelling’ theory holds that buyers should treat with maximal scepticism sellers who withhold verifiable information relating to their quality, as buyers infer from such non-disclosure that the seller possesses the lowest possible quality. This study is the first to use a natural field experiment to test this proposition, and the first to test it in a labour market context. We sent out 12,301 job applications, varying the information on degree classification – a signal of academic quality – that the applicant presented to the employer. Our results do not support unravelling theory. Applications which left degree classification undisclosed were significantly more likely to receive positive responses from employers than those disclosing the lowest possible degree classification. Employers treated non-disclosing applicants similarly to those disclosing mid-scale classifications, suggesting the extent to which adverse inference is drawn from missing information is limited. Evidence is presented against the alternative interpretation that non-disclosure success is driven by recruiters’ usage of software tools.

Keywords: Voluntary Disclosure; Unravelling; Labour Market; Field Experiment

JEL Classifications: C72 – Noncooperative games; D82 – Asymmetric and Private Information; D83 – Communication; J23 – Labor Demand

[‡] **Corresponding author. School of Economics, University of Nottingham Ningbo China, 199 Taikang East Road, Ningbo, China, 315100. Tom.Lane@nottingham.edu.cn**

***Acknowledgements:** We are grateful to Fangqi Ding, Siyang Jiang, Yuanji Jiang, Jiawei Li, Youyan Luo, Qiancheng Ma, Chen Ouyang, Chenyang Pan, Chen Shen, Kaihan Sun, Simiao Sun, Yulong Tang, Yuhui Tian, Chaoxu Xu, Shuning Xu, Xiangyi Yang, Boxiang Zhang, Tianxiang Zheng, Ziyue Zhu for conducting data collection. We received helpful feedback from audience members, including but not limited to Fabrizio Adriani, Abigail Barr, Subir Bose, Robin Cubitt, Sergio Currarini, Matthias Dahm, Sanjit Dhama, Orestis Kopsacheilis, Elke Renner, Devesh Rustagi, Daniel Seidmann, Silvia Sonderegger and Chris Starmer, and valuable written comments from Lawrence Katz. This study received IRB approval from the University of Nottingham Ningbo China. It received no financial support from external sources.

1. Introduction

In a widely-studied class of everyday economic interactions, the seller possesses verifiable information about the quality of the product under negotiation and can choose whether to disclose it to the buyer. Restaurants can choose whether or not to display hygiene ratings to customers. Second-hand car dealers can choose whether or not to post pictures of their vehicles online. And recent graduates applying for jobs can choose whether or not to announce their grades to prospective employers. In each of these circumstances, traditional game theory predicts all but the very lowest-quality sellers will opt, out of self-interest, to voluntarily disclose the information. This outcome, referred to as the ‘unravelling’ of information, rests on the assumptions that buyers follow a sophisticated process of strategic reasoning to infer that no news must be bad news, and that sellers anticipate this (Viscusi, 1978; Grossman and Hart, 1980; Grossman, 1981; Milgrom, 1981).

In the real world, markets do not seem to unravel fully. Many sellers of above the lowest possible quality fail to take opportunities to disclose information demonstrating this. This has been established in numerous contexts, including restaurants and hygiene ratings (Bederson et al., 2018), food products and fat content (Mathios, 2000), hotels and TripAdvisor scores (Butler and Read, 2017), universities and league table rankings (Luca and Smith, 2015; Merika et al., 2018), and academics and their use of professional titles (Harbaugh and To, 2020). While theorists have pointed out a wide range of special circumstances in which we should not expect unravelling to hold even under traditional assumptions of rationality (see Section 2), in many of the identified cases of incomplete disclosure it is unlikely that these exceptions apply.

Why, then, are sellers defying game-theoretic logic by withholding information? In recent years, there has been a proliferation of incentivized laboratory studies exploring this question (e.g.

Benndorf, 2018; Deversi et al., 2021; Jin et al., 2021; Sheth, 2021). This research programme has generally replicated the real-world finding that sellers often avoid disclosing information demonstrating mediocre quality. Buyers, for their part, are not as sceptical about such non-disclosure in these lab games as predicted by traditional economic theory, guessing undisclosed quality levels to be above the minimum. This would suggest that, far from being a foolish strategy, choosing not to reveal quality may in fact be optimal for many sellers. The overall picture is complicated, however, by findings that in some such experiments buyers do become increasingly sceptical with experience and some repeatedly played games evolve towards the unravelling equilibrium (e.g. Forsythe et al., 1989; Jin et al., 2021). This invites one to wonder whether in at least some real-world markets, where buyers are sufficiently experienced, their behaviour might come close to the theoretical prediction.

In this paper, we take this question about buyer behaviour to the field. Several previous studies have attempted to do this by using naturally-occurring data to measure how buyers respond to verifiable information that is not disclosed, and therefore test whether such non-disclosure might benefit some sellers (e.g. Mathios, 2000; Brown et al., 2012, 2013; Carrillo et al., 2013; Zegners, 2016; Zhou and Zhou, 2020). These papers have either compared outcomes between firms which do or do not disclose information, or else compared outcomes for a given firm before versus after a change in its disclosure policy (often driven by a change in the law). These papers have mostly concluded that buyers do not draw the most negative possible inference from missing information, and therefore that it is a profitable strategy for sellers to withhold information revealing quality below a certain threshold.

While these studies provide good suggestive evidence, it is impossible to fully rule out the possibility of endogeneity – that the positive associations found between non-disclosure and seller

outcomes are driven by another, unobservable factor. In the current paper, such potential confounds are eliminated by design. We report a natural field experiment, in which we acted as sellers in a real market and exogenously varied the content of one piece of quality-related information – and whether we disclosed it to prospective buyers – while holding all else constant. This revealed the causal effect of voluntary information disclosure on buyers' responses, and thereby tested the theoretical prediction that they would respond as negatively to silence as to the worst possible revealed information.

The setting we confront is the labour market. In our experiment (a correspondence study) we posed as jobseekers (sellers of labour) and sent out applications in response to real adverts posted on the jobsite [Monster.co.uk](https://www.monster.co.uk) by employers (prospective buyers of our labour). The information we manipulated was the degree classification revealed on the jobseeker's CV (résumé). In the UK, where our experiment was run, the degree an undergraduate earns upon completion is classified as either first class (1st), upper second class (2:1), lower second class (2:2), or third class (3rd). For each vacancy we applied for, we randomly varied which of these degree classes was stated on the CV, or – in a fifth condition – set the CV to state simply that the jobseeker had graduated but without revealing the class. As our fictional applicants purported to be recent graduates, we expected employers to consider academic performance a relevant indicator of their quality and therefore to prefer to hire those of higher classification. Under unravelling theory, employers should assume that a graduate who does not disclose their degree classification in fact achieved a 3rd, and therefore should treat them as unfavourably as those whose CV explicitly stated such.

We monitored the levels of interest employers showed in response to a total of 12,301 applications sent between 2019 and 2021. As expected, applications with higher degree classes

received more positive responses (defined as an invitation to interview or otherwise engage in further discussions with the employer). Our primary interest, however, is in how the applications withholding classification information fared in comparison to those disclosing each possible class. Notably, applications which left degree classification undisclosed received significantly more positive responses than those which stated the applicant had a 3rd – in opposition to the prediction of classic game theory. Indeed, the only disclosed degree class that secured a significantly better response than that achieved by leaving classification unmentioned was the 1st (and only then at a weak significance level).

Overall, our best estimate suggests that applications not disclosing classification do roughly as well as those disclosing classifications in the middle of the available range, suggesting that the extent of negative inferences drawn by employers towards the missing information was limited. As a consequence, we conclude that jobseekers with unimpressive academic credentials would be most successful in securing interviews if they suppress this information on their applications. Our study therefore provides clean and compelling evidence that receivers do not respond to undisclosed information in the way suggested by unravelling theory, and therefore helps to explain why the previous literature has uncovered so many cases where sellers choose not to let information unravel.

We also address an alternative explanation of our findings, which supposes the success of non-disclosure is not driven by the reactions of recruiters reading CVs but instead by their usage of software tools to pre-screen applications. If correct, this would rather muddy the waters, as human behaviour would not be (directly) responsible for the game-theoretic prediction failing. However, we present evidence to exclude this interpretation. Responses to a follow-up questionnaire distributed among employers from our field experiment rule out that a significant

proportion use filtering tools in such a way as would result in a CV not disclosing degree classification being more likely to be read than one disclosing a 3rd. Furthermore, audit tests show that reliance on AI-based CV-rating software, which some employers do report using, would also fail to lead to this eventuality.

This paper advances our understanding of interactions with asymmetric verifiable information; as we discuss in the conclusion, our results suggest unravelling theory provides neither descriptive accuracy nor prescriptively useful guidance for market participants. A supplementary contribution is that our study is the first to test unravelling theory within a real-world labour market. Arguably, one might expect the theory to have a greater chance of success in such a setting, compared to a product market, as employers are, unlike consumers, acting in a professional capacity and are likely to be highly experienced in handling information sent by jobseekers. Indeed, the first paper to introduce the term ‘unravelling’ (Viscusi, 1978) used the example of a labour market interaction. It is, therefore, particularly notable that the theory does not succeed in this setting.

More generally, our findings add to the broad literature on the signalling roles of education credentials and grades in the labour market, wherein employers use such information to screen jobseekers with innate ability that cannot be directly observed (Spence, 1973; Arrow, 1973; Stiglitz, 1975; Daley and Green, 2014). Following a longstanding general interest in interpreting degree effects through either market signalling or human capital models (e.g., Weiss, 1995; Chevalier et al., 2004; Flores-Lagunes and Light, 2010), burgeoning empirical studies have recently begun to examine further the effects of university degree classifications on labour market outcomes (McGuinness, 2003; Walker and Zhu, 2011; Di Pietro, 2016; Feng and Graetz, 2017, for the UK; Freier et al., 2015, for Germany). While only some of these studies have employed causal

identification designs, all have found large returns to achieving higher degree classes, in terms of such graduates finding better job placements and receiving higher wages. Our result is consistent with theirs in that our response rates are significantly higher for CVs explicitly stating a 1st or 2:1, rather than a 3rd. While some empirical strategies in the literature using student records data have struggled to distinguish between signalling and human capital theories of education (as discussed in Feng and Graetz, 2017), our natural field experiment setting provides a straightforward test for a pure signalling effect, and finds evidence for this.¹

In the following section, we provide background on the theoretical concept of unravelling and relevant empirical literature. We then describe our study's methods in Section 3, and present its results in Section 4, before discussing whether they are influenced by employer software usage in Section 5. Finally, Section 6 offers our conclusions.

2. Conceptual framework and related literature

Unravelling theory was introduced almost simultaneously by a combination of related papers around 40 years ago (Viscusi, 1978; Grossman and Hart, 1980; Grossman, 1981; Milgrom, 1981). Under consideration are exchanges with asymmetric information between a seller and buyer, where the seller possesses information pertaining to the quality of the good under negotiation that the buyer is not party to. The seller can truthfully reveal the information or remain silent, but

¹ Note, however, that our results do not imply that human capital effects are unimportant in general. They are simply ruled out of our study by design.

cannot lie (i.e. the information is verifiable). In the context we consider in this paper, a graduate can state on their CV the classification of their degree or omit the information.²

The standard analysis concludes that apart from those whose information reveals the quality level to be as low as it is possibly could be, all sellers have an incentive to voluntarily disclose the information. The argument relies on sellers drawing adverse inference from the choice not to disclose. The logic proceeds as follows. Sellers whose information demonstrates the highest possible quality will clearly reveal this information, otherwise buyers would underestimate their quality. Buyers should anticipate this and, therefore, if they observe no signal from the seller they must conclude the seller's quality is at best of the second-highest level. Having understood this, sellers whose quality is of the second-highest level should realise their best strategy is also to reveal this, to avoid being misidentified as having even lower quality. The logical process repeats, with buyers assuming no signal means a seller is at best third rate, third-rate sellers therefore revealing their quality, and so on. In equilibrium, all sellers of the second-lowest quality level or above will reveal their information, while sellers of the lowest possible quality level will be indifferent between revealing and not revealing; buyers will conclude that any seller who does not disclose has the lowest possible quality level. This analysis can easily be transferred to our labour market setting. All applicants with a 1st will want to convey this fact; employers will deduce that if degree class is unstated it cannot be higher than a 2:1; all applicants with a 2:1 will want to reveal that; employers will realise that no information means a 2:2 at best; all those with a 2:2 therefore

² In reality, they could also lie about their degree classification. However, for the types of positions we apply for in our experiment, employers are likely to ask for proof of the applicant's qualification before they start the job, which would render lying a very poor strategy.

disclose; employers therefore understand that non-disclosure guarantees the applicant has a 3rd. As such, the information can be said to unravel without any external interference. An important conclusion of this theory is that, in markets where governments might consider legally mandating sellers to disclose product quality information to customers, the result of natural unravelling implies this would be unnecessary.

Theorists have established many exceptions which should prohibit complete unravelling even when the ‘standard’ economic assumptions of perfect rationality and self-interest are retained. Most notably, if sellers must incur a cost to disclose the information, only those with quality above a certain threshold will do so (Jovanovic, 1982). Information may also fail to unravel where buyers hold different prior beliefs about seller quality which differ across seller type (Harbaugh and To, 2020), where there is uncertainty over how quality will be interpreted upon receiving the information (Suijs, 2007), where there are strategic interactions between sellers (Board, 2008), where buyers are unsure whether sellers possess the relevant information (e.g. Dye, 1985; Shin, 1994) and where disclosing information today may force a seller to continue disclosing it in future (Grubb, 2011). While these, and other, identified exceptions to the conditions theoretically leading to unravelling may apply in a significant array of real-world markets, it is also not difficult to find settings where they should not, including the one studied in our experiment.³ The predictions of unravelling theory are best tested under circumstances where these ‘standard’ exceptions do not apply, as an empirical refutation under such conditions would appear to be driven by failure of one

³ Harbaugh and To (2020) list previous theoretical papers that have provided exceptions to the prediction of unravelling. A less recent but very readable account can also be found in Dranove and Jin (2010).

of the fundamental *homo economicus* principles of traditional theory (e.g. full rationality or self-interest).

Many empirical studies have indeed found insufficient levels of information disclosure by sellers, relative to the theoretical prediction. Mathios (2000) found many sellers in the salad dressing market did not voluntarily disclose the fat content of their products. Bederson et al. (2018) established the limited extent to which restaurants revealed their hygiene ratings to diners. Jin (2005) uncovered incomplete disclosure of service quality accreditations by health maintenance organisations. Luca and Smith (2015) showed many American business schools did not announce their ranking positions on their websites, a result replicated by Merika et al. (2018) for UK universities. Conlin et al. (2013) found that applicants often chose not to reveal SAT scores to colleges, while Harbaugh and To (2020) estimated that a large proportion of academic staff do not announce their professional titles on voicemail greetings and course syllabi. Several studies have identified incomplete disclosure in housing markets, in terms of photos (Bian et al., 2021), local school quality (Carrillo et al., 2013) and energy information (Fronzel et al., 2020). In hospitality, Butler and Read (2017) demonstrated that many properties did not take the opportunity to reveal TripAdvisor scores. It seems unlikely that all of these results can be explained by the aforementioned ‘standard’ exceptions to the theoretical expectation of unravelling.

A better understanding is needed, therefore, as to why more sellers are not using voluntary disclosure. Here, the possibility needs to be considered that some players are not following the strategic logic of the game. Our experiences suggest that attempts to explain unravelling theory can initially draw blank expressions on the faces of listeners, even those trained as professional economists. People can easily understand that those with the highest quality will choose to reveal this, but iterating all the remaining logical steps to conclude that all but the very worst will also

do so may not come naturally.⁴ Sellers of average-to-low quality may choose not to disclose their information because the strategic logic of the unravelling argument has not occurred to them. However, such sellers may also prefer not to disclose even if this logic has occurred to them, but they expect it not to have occurred to buyers – or, indeed, that they expect buyers not to believe that it will have occurred to sellers⁵ – which in either case would lead sellers to expect buyers not to draw the most negative possible inference from non-disclosure, and would lead sellers with quality below a certain level to believe they would be better off not revealing this. The potential that such breakdowns in principles of rationality may explain non-disclosures has attracted more interest in recent years, with several scholars attributing such outcomes to models of cursed equilibria (Eyster and Rabin, 2005) or limited attention (Hirshleifer et al., 2004).⁶

In our experiment, we address the question of whether sellers would be correct to believe that buyers do not equate a lack of information with the worst information that could possibly be

⁴ In other words, people may be categorised according to the theory of level-k reasoning (Stahl and Wilson, 1994, 1995; Nagel, 1995). Some may not consider strategic motivations in the game at all (level-0), while others may follow only the first logical step (level-1), and others may iterate the process further (level-2 and above). Brown et al. (2012) explain this in more detail in the context of information (not) unravelling.

⁵ To extend this further, the unravelling argument rests upon common knowledge of rationality. If, instead, there is common knowledge of the possibility of errors, we might consider a model of quantal response equilibrium (McKelvey and Palfrey, 1995), which would not make such stark predictions.

⁶ Fishman and Hagerty (2003) also consider the possibility that buyers may simply not understand the information presented to them.

revealed. In other words, we will study how buyers actually respond to non-disclosure. The recently burgeoning literature on laboratory signalling games with verifiable information offers insights. These experiments – in which sellers with above-minimum but below-average quality very frequently pursue a strategy of non-disclosure – generally find that buyers do *not* apply full scepticism to non-disclosing sellers and often act towards such sellers in a manner similar to how they act towards those who reveal quality levels below average but well above the minimum (Benndorf, 2018; Montero and Sheth, 2019; Li and Schipper, 2020; Deversi et al., 2021; Jin et al., 2021; Sheth, 2021). As such, it is indeed a profitable strategy for many sellers to leave their information undisclosed.⁷ However, some of these experiments have shown that patterns of behaviour change when the game is played repeatedly. In the experimental asset markets of Forsythe et al. (1989) and King and Wallin (1991), both buyer and seller behaviour started far away from the unravelling equilibrium, but moved towards it as the market was played repeatedly. Similar results were yielded by the sender-receiver games of Jin et al. (2021) – though only when salient feedback was available between experimental periods – and the worker-employer-framed interactions of Benndorf (2018). It strikes us as an open question, therefore, how buyers will behave away from the unfamiliar confines of a lab game and when operating in markets they are

⁷ Whether *all* non-disclosing sellers in these experiments are following a profitable strategy depends on how well-calibrated buyers are in their assessments of the average quality of non-disclosing sellers. If these assessments are close to the true average, the distribution of non-disclosing sellers must contain plenty of sellers who could profitably deviate to disclosing their quality. However, several experiments have found buyers are somewhat overoptimistic in their assessments of non-disclosing sellers (Montero and Sheth, 2019; Li and Schipper, 2020; Deversi et al., 2021; Jin et al., 2021); if such overoptimism is strong enough, it may be that even the best non-disclosing sellers are following an optimal strategy.

highly acquainted with. In our experiment, buyers are employers who have professional experience assessing the information submitted to them by jobseekers.

The response of buyers to non-disclosed information has also been addressed by a number of studies observing data in real-world markets.⁸ This literature has tended to suggest buyers exhibit the weak scepticism characteristic of the early rounds of the lab games, rather than the strong scepticism observed in the later rounds of some experiments. Most of these studies have examined how a seller's decision to withhold information correlates with a particular measure of seller success (e.g. price, units sold), while controlling for observable measures of quality and other characteristics. If buyers are fully sceptical, this correlation should be null for sellers of the lowest quality and negative for all others; however, the correlation has generally been found to be positive. Brown et al. (2012, 2013) found that movies which cold opened (i.e. did not disclose reviews to customers) achieved greater box office success than movies of equivalent quality which disclosed their reviews. Conlin and Dickert-Conlin (2017) concluded that college applicants who withheld poor SATI scores improved their admission chances relative to those who revealed such results. Zegners (2016) found that e-books which did not reveal their quality through making free samples available sold for more than those which did. Lewis (2010) determined that used cars whose sellers made no statement about rust sold better than those whose sellers admitted them to be rusty. Other studies have conducted the analysis on a within-firm basis, identifying that some sellers perform better when they do not disclose information than when they do (for instance,

⁸ Note that there also exists a literature on the distinct but related phenomenon of price shrouding. This addresses how consumers respond when certain price components are hidden. Price shrouding generally seems to increase sales (e.g. Chetty et al., 2009), but not always (Brown et al., 2010).

because they are forced to by a change of the law), controlling for other observable features of the situation (Mathios, 2000; Frondel et al., 2020; Zhou and Zhou, 2020).

Overall, these papers suggest that buyers do not draw negative inferences from missing information to the extent predicted by unravelling theory. However, as noted in Section 1, none of these observational studies can guarantee they are unimpeded by endogeneity. In a cross-sectional study, the correlation between a seller's disclosure decision and their success could yet be driven by a third, unobservable factor. In a before/after study, the change to a seller's success following the change in their disclosure strategy might be the independent consequence of the event that prompted the seller to change their strategy (e.g. the change in law). We argue, therefore, that it is worth testing buyer behaviour in a context where such endogeneity is excluded by design. We argue further that the labour market provides a worthwhile setting for such a test, as buyers in this market may be more experienced than those in consumer markets, where the literature has tended to focus, and this may provide us with conditions more favourable to the prospect of finding support for unravelling theory. In the next section, we turn to our experimental design and depict how, through the controlled manipulation of information conveyed to buyers in our market of interest, it allows us to cleanly measure the causal effect of non-disclosure on buyers' reactions to sellers.

3. Experimental Design

To test how employers respond to the non-disclosure of quality-related information by applicants, we ran a natural field experiment in which we posed as jobseekers, and sent out fictional CVs in response to job adverts. This type of experiment, also known as a correspondence study, has a rich

history in its application to labour market research. Traditionally, such studies were mostly used to measure employer discrimination by systematically varying characteristics such as the race or gender of the name on the CV (see Bertrand and Duflo, 2017, for a review). More recently, the method has been applied to a wide range of questions by varying other aspects of the CV, including duration of non-employment (e.g. Oberholzer-Gee, 2008; Kroft et al., 2013; Farber et al., 2016), college major (Nunley et al., 2016) and alma mater (Deming et al., 2016), and volunteering experience (Baert and Vujić, 2018).

In our experiment, the information we vary relates to the jobseeker's degree classification. Our study is run in the UK, where it is widely understood that a degree certificate will identify an undergraduate's completed academic performance as having been either first class, second class (upper), second class (lower), or third class. The categorisations are known universally as a 1st, 2:1, 2:2 and 3rd.⁹ In common parlance, UK graduates typically refer to the classification they achieved,

⁹ In very unusual circumstances, universities award degrees which are of even lower status than third class, known as 'pass degrees' and 'ordinary degrees'. In our experiment, the lowest degree class we disclose on a CV is a 3rd. Under a literal interpretation of the unravelling theory, we should predict that CVs with undisclosed classifications will perform even worse than degrees which disclose a 3rd, as employers should assume them to be concealing an ordinary or pass degree, the lowest existing categories. However, as ordinary and pass degrees are vanishingly rare and we suspect many employers may not even be aware of their existence, we believe a more reasonable interpretation of the prediction of unravelling theory is that employers should assume non-disclosed classifications to be third class, the lowest commonly used classification.

while a precise score (i.e. GPA) is very rarely discussed.¹⁰ Informally, degrees are generally considered to be of good quality if they are a 2:1 or above.¹¹

We created a total of 22 fictional jobseekers' CVs and sent them out to job adverts. Each was given a different name and background information, with different jobseekers purporting to have graduated in different disciplines from different universities; we also used divergent formatting styles across CVs, so that if one recruiter happened to receive CVs from more than one of our jobseekers, they would not suspect they in fact came from the same source (See Figure I for an example of one of our CVs).¹² Of each CV, we created five versions, which constitute our five experimental treatments. The versions were identical in all respects except for whether they stated that the applicant had graduated with a 1st, a 2:1, a 2:2, a 3rd, or – in the case of non-disclosure – simply stated that the applicant had graduated but did not mention degree class. If included, degree classification was reported in parentheses following the CV's statement on the university and course from which the applicant had graduated. When an advertised position was identified for a

¹⁰ Several audience members have enquired whether the choice of jobseekers to only rarely disclose exact GPA on CVs is itself a failure of unravelling theory. We believe not, since this information is not generally printed on degree certificates and is therefore not easily verifiable.

¹¹ This, perhaps, is reflected in the proliferation of commonly understood pejorative epithets bestowed upon degrees below the level of a 2:1, such as a 'drinker's degree' to refer to a 2:2, and a 'gentleman's degree' for a 3rd.

¹² In early 2021, Monster changed its system so that, instead of allowing jobseekers to upload pre-made CVs, it automatically constructed the CVs on the basis of information jobseekers uploaded. Naturally, from this point onwards the CVs we submitted were stylistically identical.

given jobseeker to apply for, we would randomise which of the five versions of the jobseeker's CV to submit. Since our experiment used multiple jobseekers with distinct CVs, in our analysis of treatment effects we will include fixed effects for each fictional jobseeker to control for between-individual differences in quality.

(FIGURE I ABOUT HERE)

Our empirical strategy relies on degree classification being taken by employers as a meaningful signal of applicant quality. In order to achieve this, we designed our fictional applicants to be recent graduates with limited job experience. Moreover, we applied for jobs that were specifically targeted at recent graduates. All applications were made on the jobsite Monster.co.uk; job adverts were identified by searching for 'graduate jobs' on the site.¹³ We did not apply for any jobs on which the advert specifically stated a requirement of a minimum degree classification. Demanding, for instance, a 2:1 at minimum is a fairly common specification on graduate job adverts; we avoided applying in such cases as we expected employers might assume that any applicant failing to specify degree classification was at least at the minimum requested level for the job.

Beyond this, we placed a few limitations on the types of jobs we would apply for. We did not apply for any jobs which specifically requested applicants had a degree in a subject different from the one on our fictional CV, as we presumed the success rate from any such applications

¹³ Eventually, this search would run out of available results, in which case we modified the search by randomly adding one of the job categories listed on Monster to the search term, so that it became, for instance, 'graduate administration jobs'.

would be close to zero, regardless of degree classification. In order to save time, we only responded to adverts that required applications to be made via Monster directly, as opposed to those requiring an external application. Applications sent within Monster do not require an accompanying cover letter, and in our applications we never attached one.

A total of 12,301 applications were sent out between May 2019 and September 2021. Applications were administered by a team of research assistants, with a single assistant in charge of any given fictional jobseeker (and therefore all versions of this jobseeker's CV). We limited each jobseeker to a maximum of five applications to different jobs posted by the same organisation. Where a given jobseeker made subsequent applications to an organisation they had already applied to, we ensured we sent the same version of the CV that had been randomly selected to be sent on the first occasion¹⁴. This was done to prevent employers becoming suspicious that the CVs were not genuine, which would have been likely if they found themselves in possession of two otherwise identical CVs reporting different qualifications. As such, our unit of randomisation is the organisation applied to; our analysis will adjust for this by clustering standard errors accordingly.

¹⁴ There could still have been some circumstances under which we ultimately sent different versions of the CV to the same organization. This could occur if the listed recruiter on Monster was actually an agent, who passed on CVs to an unspecified end user; if a single end user employed multiple agents, they could end up receiving our CVs via multiple sources. Alternatively, if the listed recruiter on Monster was the employer but they passed on received CVs to an unspecified agent who conducted the actual recruitment process, the agent could receive multiple versions of a CV via different listed recruiters. We believe either circumstance is unlikely. If it did occur, the likely consequence would be a slight dampening of our treatment effects.

The outcome variable of interest is whether or not an application receives a positive response. We define a positive response as an invitation for an interview or for the applicant to otherwise communicate further with the employer about the vacancy (e.g. the recruiter requesting a phone conversation or directly asking the applicant questions by email). The only way in which employers could communicate with us was via email (the CVs did not include phone numbers)¹⁵; we monitored the inboxes of the email accounts associated with each applicant's Monster profile, and tallied each applicant's positive responses. When a positive response was received, we sent a reply withdrawing the applicant's interest in the position, in order to minimise the time wasted by recruiters. We prepared a maximum of five versions of the withdrawal email for each applicant to send to a given organization, to prevent employers becoming suspicious from repeatedly receiving identically-worded withdrawals of interest. Our experiment received approval from our university's ethics committee.

Before moving to our results, we highlight one important point. Implicit in unravelling theory is the assumption that the seller will benefit from being perceived by buyers as high quality. In principle, this assumption could fail in some labour market contexts where applicants are penalised for being overqualified. An employer might prefer to interview a candidate with a mid-rank degree than one with a 1st if they believe retaining the higher quality candidates to be unrealistic. While this is likely to be the case for *some* employers in our experiment, we doubted it would hold true on average. For applicants to benefit on average from being perceived as higher quality, we require

¹⁵ Our CVs also included fictional postal addresses. Following Koellinger et al. (2015), we used real postcodes and street names but a non-existent house number. As Koellinger et al. (2015) argue, it is highly unlikely any recruiter would use postal communication as their first means of contact with an applicant.

positive response rates to be monotonically increasing in revealed degree classification. Our results will verify that this is indeed the case.¹⁶

4. Results

The average rate of positive responses across all applications was 12.5%. The raw treatment differences are visible in Figure II. As anticipated, among those disclosing degree classifications, CVs stating a 1st are most successful, followed successful by those with a 2:1, 2:2 and then 3rd. The positive response rate is 15% for those disclosing a 1st and 10.3% for a 3rd. Although only a few percentage points in difference, this is a fairly substantial effect as it implies 1st class candidates will be offered almost three interviews for every two achieved by the 3rd classers. It confirms that degree classification is considered a relevant and quite important quality attribute within this section of the labour market.

¹⁶ A more complex version of this objection could argue that rational applicants should be able to identify for which of the vacancies in our experiment they would be penalized for being overqualified, and therefore should not be expected by rational employers to want to reveal themselves as high quality for these particular vacancies. However, the theoretical consequence of this would just be to alter the ranking of classifications that applicants prefer to be perceived as holding – for instance, resulting in a 2:2 being the most highly regarded classification and a 1st the least so. This does not change the logic of unravelling theory that a non-disclosing applicant will be perceived as belonging to the least preferred category – whichever that may be. It therefore would not be able to explain how – as our experiment will find – a strategy of non-disclosure outperforms, on average, a strategy of disclosing one particular classification, since in theory non-disclosure should be at best an equally successful strategy (on the occasions where that degree classification is the least preferred one), and should be a less successful strategy on other occasions.

The success rate of CVs with undisclosed degree classification falls in between these two extreme categories, at 12.6%. This is somewhat below the rate for CVs with 2:1s (13.1%) but above that for CVs with 2:2s (11.3%). That this positive response rate outperforms that for CVs disclosing a 3rd suggests employers do not treat non-disclosers equivalently to those revealed to be of the lowest quality.

(FIGURE II ABOUT HERE)

While these positive response rates are averaged across our 22 different CVs, the fact that we randomised classification information for each application means these treatment differences should be interpretable as the causal effects of varying classification information. However, it is possible that by chance particular classification information could have been sent out disproportionately often by CVs that were more impressive on other dimensions. To address this, we subject our data to regression analysis in which we include fixed effects for each jobseeker. The dependent variable is binary: whether or not an application receives a positive response. We use a logistic model with dummy variables for each type of classification information, excluding the 3rd as the baseline category. The regression estimates the effects of disclosing particular degree classes – or not disclosing one at all – on the positive response rate, for a given jobseeker. It also includes control variables for how many days ago a job advert was posted; we use dummy variables to capture this because, if a job was posted more than 6 days ago, Monster only specifies this fact rather than the precise number of days.¹⁷ The regression allows us to estimate the statistical

¹⁷ We attempted to collect other information listed on the website about each vacancy, including industry type and career level, but the categorisations employed on the adverts for these variables were inconsistent over time and were largely left empty after Monster's interface underwent a revamp in early 2021.

significance of our treatment differences, having made the necessary adjustment to the standard errors by treating each hiring organisation as provided one cluster of observations.^{18,19}

Odds ratios are presented in Table I. An odds ratio above 1 indicates a variable exerts a positive effect on the likelihood of receiving a positive response, while an odds ratio below 1 implies a negative effect. Of most importance, the positive result on *Undisclosed* demonstrates that a given CV which leaves degree classification undisclosed has 1.265 times higher odds of achieving a positive response than if it discloses a 3rd (the odds represent the ratio of the probability of receiving a positive response to the probability of not receiving one). The result is significant, with a p-value of 0.022. This, of course, is inconsistent with unravelling theory, which suggests the odds should be the same for both.

Regarding other pairwise treatment comparisons, we find unsurprisingly that both the 1st and 2:1 outperform the 3rd with very high significance levels, while the effects of a 2:2 and 3rd are not statistically distinguishable. We use Wald tests to compare the effects of non-disclosure versus disclosing a 1st, 2:1 or 2:2. Non-disclosure is only significantly inferior to disclosing a 1st, and only

¹⁸ Note that 23 observations are omitted from the regression analysis due to data entry errors which meant the hiring organisation was unidentifiable or the posted date information was missing in our dataset. These exclusions make no qualitative difference to the output.

¹⁹ In addition to adding control variables into our regression, we also performed a series of balance checks, where we regressed the dummies for jobseekers and elapsed time against the treatment variables. These only very rarely identified significant treatment differences, as would be expected by chance, and indicates that our randomization was successful. The output of these regressions can be replicated from our do-file code (freely available upon publication).

weakly so ($p=0.079$). We also find that the 1st is significantly more successful than the 2:2 but not than the 2:1, while the 2:1 outranks the 2:2 at weak significance.²⁰

(TABLE I ABOUT HERE)

5. Alternative explanation: recruitment software usage

Our interpretation is that our key experimental result stems from employers viewing CVs with non-disclosed degree classification and not always inferring from the absence of information that such candidates are of the lowest possible quality. A potential objection to this interpretation derives from the fact that employers nowadays often use software tools to assist their handling of applications. Such tools, known generally as Applicant Tracking Systems (ATSs), can result in a large proportion of CVs succumbing to a preliminary filtering process and not ultimately being read by recruiters (see Mukherjee et al., 2014; Laumer et al., 2015).

The concern is that the greater success of non-disclosing applicants, relative to those disclosing a 3rd, may arise simply because ATS usage results in non-disclosing CVs being more likely to be read. This could happen in two particular ways. First, recruiters could use filtering tools to exclude from the pool of readable applications any CV mentioning a 3rd class degree. Secondly, they might rely on AI tools which scan CVs and numerically score them; if such tools provide a higher score for a non-disclosing CV than an otherwise identical CV revealing a 3rd, and if recruiters are more likely to read CVs with higher AI scores, non-disclosing CVs would be read more often than those disclosing a 3rd. While it may be argued this would merely relocate responsibility for the failure of the game-theoretic prediction onto the filtering strategies of

²⁰ Estimations of marginal effects on the CV dummy variables lead to qualitatively similar p-values and the same levels of significance as those reported on the odds ratios in Table I.

recruiters and the design choices of algorithm developers, we would still like to exclude this possibility, as it would bring into question the generalizability of our findings to contexts absent of automation.²¹

To address this alternative explanation, we ran a follow-up study in which we explored the ways ATSS are used by the employers in our experiment. We were able to distribute a questionnaire to recruiters who had provided positive responses in the field experiment, since we had their contact email addresses from this correspondence. In the questionnaire, we asked respondents specifically about their usage of ATSS for graduate job applications received via Monster. Through our own explorations of the Monster recruiter interface, we had also already verified there were no AI or filtering tools available within the platform itself that would result in the patterns we sought to rule out; these outcomes could occur only if recruiters downloaded the CVs and entered them into external ATSS, and therefore our questions focused on such behaviour.

We sent out questionnaire invitations to 584 individuals and received 41 complete responses between April and July 2022.²² The low response rate was likely due to the time lag of

²¹ For instance, ATS usage might fail to follow the sceptical strategy towards non-disclosure advocated by unravelling theory simply because of technological limitations which make doing so impractical.

²² We were unable to invite all those who provided positive responses because some of the email accounts we created were inaccessible by 2022, meaning we could no longer access their correspondence to identify recruiters' contact details. We also sent invitations only to email addresses which appeared to belong to a single person, since we wanted respondents to be those who we actually interacted with in the field experiment, rather than a colleague sharing their company email account.

up to three years between the experiment and questionnaire; many of the invitations bounced back, indicating the recruiter was no longer employed in the same position. 10 respondents reported that they did not use Monster for recruitment (i.e. their organization's choice of recruitment platform had evidently changed since the time of the experiment) and therefore did not provide answers to the questions of interest, leaving us with 31 usable responses. Of these, 26 reported downloading CVs received via Monster onto external ATSs. When asked directly, 24 of these respondents stated that they had never used tools to filter out CVs revealing low degree classifications. The remaining two reported that they had sometimes done this, but when prompted to provide more information on exactly how they did it, one explained it was done using questions via their online application form (which could not have been used for any applications in our experiment), demonstrating they had not understood the question referred specifically to applications received directly via Monster. The other person also appeared to have misunderstood the question, explaining why filtering tools were useful in general; we wrote to this respondent for a clarification, in which they confirmed their organization in fact filtered CVs by degree subject rather than classification. Overall, then, it appears that none of the 31 usable respondents employ filtering tools in such a way that would result in 3rd class CVs being less likely to be read than non-disclosing ones.²³

While the sample size is small, the proportion of 0/31 gives us confidence that the type of filtering strategy of concern was not sufficiently common within our experimental sample to meaningfully impact the results. One might worry that our respondents were not a representative sample of employers from the experiment, as we only surveyed positive respondents. We argue, however, that this sample is good enough for ruling out the pattern of behaviour of concern. If this filtering strategy had been responsible for the difference in outcomes between the 3rd class and

²³ The full wording of the survey questions and all responses are available in the online supplementary materials.

non-disclosure treatments, this would necessarily mean that some of those recruiters employing the strategy (especially those randomly assigned to receive 3rd class CVs) would have to end up as negative responses, while others (especially those randomly assigned to receive undisclosed class CVs) would end up as positive responses. Therefore, there would need to be some evidence for use of the strategy just among the positive response sub-sample – but we do not find this.

Six respondents did, however, report using AI tools to score CVs. As a result, this type of behaviour could have a meaningful impact on our results, if the algorithms tended to score non-disclosure CVs more favourably than those disclosing a 3rd. To explore whether this was indeed the case, we ran a series of audit tests. We identified nine online AI-based CV rating tools, and fed into them the non-disclosure and 3rd class versions of otherwise identical CVs which had been used in the field experiment. The tools on which we ran these checks were varied, including both free-to-use and premium software, and including some which ranked a CV's suitability against a specific job description and others which just provided a general score.²⁴ In every case, we found that the algorithm gave an identical score to the two versions of the CV, indicating that it was not sensitive to degree classification information. This speaks against the possibility that the usage of such tools resulted in different filtering rates for CVs belonging to the different treatments.

6. Conclusion

The results of our experiment are quite straightforward. Employers do not treat applicants who fail to disclose the quality of their academic credentials equivalently to those who disclose the lowest possible level of quality. On the contrary, our best estimate is that those who do not disclose degree

²⁴ Specifically, the tools we ran our tests on were those offered by cvscan.uk, enhancv.com, jobscan.co, livecareer.co.uk, resume.shine.com, resumeworded.com, rezscore.com, skillsyncer.com and vmock.com.

classification are treated similarly to those who disclose themselves as being around the middle of the scale.

Our research sheds new light on job market interactions. Jobseekers with low-quality qualifications would be better served in securing interviews by remaining silent about this than by fronting up to employers. As with any correspondence study we only observe the first stage in the hiring process, and we cannot confirm what the ultimate effect on the job-offer success rate of a third class graduate suppressing their classification information would be. However, an invitation to interview is a foot in the door for the applicant; while we cannot empirically test whether a third class graduate who suppresses the classification on their CV increases their chance of a job offer, given that they increase their chance of an interview the most plausible conclusion is that they do so.

Our primary focus, moreover, is on the employer rather than the jobseeker. Our study has shown plainly that employers do not react to non-disclosure with the maximal scepticism unravelling theory predicts they should adopt. In our design, we were careful to ensure we focus on a setting where, as discussed in Section 2, none of the exceptions to the prediction of unravelling resulting from ‘standard’ theoretical analysis apply. As such, the incorrectness of the game-theoretic prediction appears to be driven by failure of one of its fundamental *homo economicus* building blocks – relating, most likely, to complete rationality and/or the common knowledge of its existence.

One might offer the conjecture that employers are in fact sceptical but playing a game of wait and see, holding onto applicants who do not disclose in order to resolve the uncertainty about their qualification standard at the interview. However, even this would not be consistent with unravelling theory – the theory holds that employers should be certain that non-disclosure implies

minimal quality, and an employer should need to be sufficiently optimistic about the probability of the applicant turning out to be of a high quality in order for it to be worth incurring the cost of interviewing them. Furthermore, we suspect that this conjecture would be incorrect in any case. If employers are giving positive responses to non-disclosing candidates on the off-chance that they will turn out to belong to a high class, then given the apparent importance of degree classification (according to our data, employers prefer applicants with a 1st over those with a 3rd by a margin of 50%) one would expect many of them to try to resolve the uncertainty immediately, i.e. by asking, in their initial response email, for the applicant to disclose their classification. However, from the 2,421 applications we made which left degree classification undisclosed, we encountered just one solitary case of an employer requesting this information in their immediate response email. In contrast, employers' emails often asked for clarification on other pieces of information related to the CV – for instance, the applicant's current location.

This raises the possibility that one reason for the lack of scepticism towards non-disclosure in the experiment was that some employers simply failed to notice the missing information. Since traditional economic theory, from which the unravelling concept stems, places no limits on human ability to process information, this does not let the theory off the hook. When we have presented this paper in seminars, some listeners have questioned whether we are subjecting the theory to a fair test, given that a CV contains many different pieces of information and it is difficult for readers to pay attention to all of them. Technically, the original versions of unravelling theory model an environment in which there is a single, measurable dimension of quality. The logic of the theory will, however, still apply if there are multiple dimensions of quality, as long as attention costs are not built in. Since it is almost impossible to think of a real-world environment in which product quality can be compressed into a single dimension, it seems reasonable to believe the proponents

of unravelling theory intended for it to apply in multi-dimensional cases. Indeed, they illustrated their arguments with such examples (e.g. Grossman, 1981). If we take the view that the theory abstracts from attention costs for the sake of simplicity, it would then be reasonable to expect it to best predict outcomes when we focus on relatively important dimensions of quality, which buyers should have the strongest incentives to pay attention to. As our experiment shows, degree classification appears to be one such important characteristic, leading to substantial variation in an applicant's attractiveness to employers.²⁵

Readers might then consider what role there is left for unravelling theory in light of our results, which support the conclusions of existing studies from the lab and from real-world naturally-occurring data that buyers do not draw the most negative adverse inference from sellers withholding verifiable information. As has by now been well established in a range of contexts,

²⁵ Given that multiple dimensions of applicant quality exist, another possibility that has been put to us is that employers believe the decision to non-disclose is correlated with an unobservable dimension of an applicant's quality. Under this interpretation, it is possible that employers might actually deduce that non-disclosing applicants have 3rd class degrees but treat them relatively favourably because they take non-disclosure as a positive signal of something else. While we are unable to empirically rule out this possibility, we do not find it convincing. It is difficult to think of unobservables that are likely to be correlated with the decision to non-disclose, with one important exception. In fact, we believe the decision to disclose a 3rd class degree is a signal of rather poor strategic thinking, and that non-disclosure is therefore likely to be associated with higher intelligence. However, the only reason we think disclosing a 3rd is a signal of poor strategic reasoning is because it is likely to lower employers' beliefs about the applicant's academic performance (i.e. because employers do not respond to non-disclosure by inferring the lowest classification level, as our results suggest). If, instead, employers actually do infer from non-disclosure that the applicant has a 3rd, this argument becomes self-defeating and there is then no good reason to believe non-disclosure is correlated with better strategic reasoning.

sellers do not generally act in accordance with the theory by disclosing information in instances it suggests they should. The possibility still remained, however, that this was merely down to the naivety of sellers – that they *should* be fully disclosing, but were failing to anticipate that buyers would draw maximally negative conclusions when they did not. However, we have shown that such conclusions are not drawn, even by buyers in labour markets who are likely to have more contextual experience than those in consumer markets. It seems, therefore, that unravelling theory is not only a descriptively inaccurate model in predicting the outcomes that will occur in seller-buyer interactions, but is also an unhelpful prescriptive model in guiding market participants on how they should act. Indeed, the finding that buyers do not take no news to unequivocally imply bad news suggests that much of the non-disclosure observed throughout the literature reflects sellers profitably ignoring the guidance of standard game theory.

An alternative position is that, despite its shortcomings, the empirical evidence does not render unravelling theory fully irrelevant. In lab experiments, buyers appear to believe that sellers who do not disclose quality information are below average in quality (Montero and Sheth, 2019; Li and Schipper, 2020; Jin et al., 2021). This suggests there is some degree of adverse inference drawn from silence, even if not to the extent the theory predicts. Our results may be consistent with this, although we cannot state this with certainty due to the size of the confidence intervals on our estimates of the positive response rates for each treatment. If, as our results suggest, employers treat non-disclosing applicants equivalent to those revealing themselves to be around the middle of the classification scale, this would actually imply they infer non-disclosing applicants to be below the mid-point in the quality distribution. This is because, following years of grade inflation, a large majority of students at UK universities graduate with a 2:1 or better; the proportion reached 76% in 2018-19 (HESA, 2021). In regarding an applicant who leaves

classification undisclosed as worse than one who discloses a 2:1, employers may consider non-disclosure as a bad signal.

Considering the whole literature, it does appear that some degree of strategic reasoning is relevant in seller-buyer interactions with verifiable information. Models which retain this feature but limit the extent to which we assume (common knowledge of) complete rationality appear better suited than traditional game theory in explaining behaviour in such interactions. Previous researchers in this literature have noted – correctly, we believe – that models of level-k thinking (Stahl and Wilson, 1994, 1995; Nagel, 1995) and cursed equilibrium (Eyster and Rabin, 2005) can offer insights. Any successful theory will likely need to account for heterogeneity in players' strategic sophistication and beliefs about others. We encourage future researchers to devise and implement tests of such theories.

Author affiliations

Tom Lane: School of Economics, University of Nottingham Ningbo China

Minghai Zhou: School of Economics, Zhejiang University

References

Arrow, K. J. (1973). Higher education as a filter. *Journal of Public Economics*, 2(3), 193-216.

Baert, S., & Vujčić, S. (2018). Does it pay to care? Volunteering and employment opportunities. *Journal of Population Economics*, 31(3), 819-836.

- Bederson, B. B., Jin, G. Z., Leslie, P., Quinn, A. J., & Zou, B. (2018). Incomplete disclosure: Evidence of signaling and countersignaling. *American Economic Journal: Microeconomics*, 10(1), 41-66.
- Benndorf, V. (2018). Voluntary disclosure of private information and unraveling in the market for lemons: An experiment. *Games*, 9(2), 23.
- Bertrand, M., & Duflo, E. (2017). Field experiments on discrimination. *Handbook of Economic Field Experiments*, 1, 309-393.
- Bian, X., Contat, J. C., Waller, B. D., & Wentland, S. A. (2021). Why Disclose Less Information? Toward Resolving a Disclosure Puzzle in the Housing Market. *The Journal of Real Estate Finance and Economics*, 1-44.
- Board, O. (2009). Competition and disclosure. *The Journal of Industrial Economics*, 57(1), 197-213.
- Brown, A. L., Camerer, C. F., & Lovallo, D. (2012). To review or not to review? Limited strategic thinking at the movie box office. *American Economic Journal: Microeconomics*, 4(2), 1-26.
- Brown, A. L., Camerer, C. F., & Lovallo, D. (2013). Estimating structural models of equilibrium and cognitive hierarchy thinking in the field: The case of withheld movie critic reviews. *Management Science*, 59(3), 733-747.
- Brown, J., Hossain, T., & Morgan, J. (2010). Shrouded attributes and information suppression: Evidence from the field. *The Quarterly Journal of Economics*, 125(2), 859-876.
- Butler, D., & Read, D. (2017). Unravelling & strategic disclosure: Evidence from the hospitality industry. *Available at SSRN 3068601*.

- Carrillo, P., Cellini, S. R., & Green, R. K. (2013). School quality and information disclosure: Evidence from the housing market. *Economic Inquiry*, 51(3), 1809-1828.
- Chetty, R., Looney, A., & Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, 99(4), 1145-77.
- Chevalier, A., Harmon, C., Walker, I., & Zhu, Y. (2004). Does education raise productivity, or just reflect it?. *The Economic Journal*, 114(499), F499-F517.
- Conlin, M., & Dickert-Conlin, S. (2017). Inference by college admission departments. *Journal of Economic Behavior & Organization*, 141, 14-28.
- Conlin, M., Dickert-Conlin, S., & Chapman, G. (2013). Voluntary disclosure and the strategic behavior of colleges. *Journal of Economic Behavior & Organization*, 96, 48-64.
- Daley, B., & Green, B. (2014). Market signaling with grades. *Journal of Economic Theory*, 151, 114-145.
- Deming, D. J., Yuchtman, N., Abulafi, A., Goldin, C., & Katz, L. F. (2016). The value of postsecondary credentials in the labor market: An experimental study. *American Economic Review*, 106(3), 778-806.
- Deversi, M., Ispano, A., & Schwardmann, P. (2021). Spin doctors: An experiment on vague disclosure. *European Economic Review*, 103872.
- Di Pietro, G. (2017). Degree classification and recent graduates' ability: Is there any signalling effect?. *Journal of Education and Work*, 30(5), 501-514.
- Dranove, D., & Jin, G. Z. (2010). Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4), 935-63.

- Dye, R. A. (1985). Disclosure of nonproprietary information. *Journal of Accounting Research*, 123-145.
- Eyster, E., & Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73(5), 1623-1672.
- Farber, H. S., Silverman, D., & Von Wachter, T. (2016). Determinants of callbacks to job applications: An audit study. *American Economic Review*, 106(5), 314-18.
- Feng, A., & Graetz, G. (2017). A question of degree: The effects of degree class on labor market outcomes. *Economics of Education Review*, 61, 140-161.
- Fishman MJ, Hagerty KM. 2003. Mandatory Versus Voluntary Disclosure in Markets with Informed and Uninformed Customers. *Journal of Law, Economics, and Organization*, 19 (1), 45-63.
- Flores-Lagunes, A., & Light, A. (2010). Interpreting degree effects in the returns to education. *Journal of Human Resources*, 45(2), 439-467.
- Forsythe, R., Isaac, R. M., & Palfrey, T. R. (1989). Theories and tests of "blind bidding" in sealed-bid auctions. *The Rand Journal of Economics*, 214-238.
- Freier, R., Schumann, M., & Siedler, T. (2015). The earnings returns to graduating with honors—Evidence from law graduates. *Labour Economics*, 34, 39-50.
- Fronzel, M., Gerster, A., & Vance, C. (2020). The power of mandatory quality disclosure: Evidence from the German housing market. *Journal of the Association of Environmental and Resource Economists*, 7(1), 181-208.
- Grossman, S. J. (1981) "The informational role of warranties and private disclosure about product quality", *Journal of Law and Economics*, 24(3): 461-483.

- Grossman, S. J. & Hart, O. D. (1980) "Disclosure laws and takeover bids", *Journal of Finance*, 35: 323-334.
- Grubb, M. D. (2011). Developing a reputation for reticence. *Journal of Economics & Management Strategy*, 20(1), 225-268.
- Harbaugh, R., & To, T. (2020). False modesty: When disclosing good news looks bad. *Journal of Mathematical Economics*, 87, 43-55.
- HESA (2021). What are HE students' progression rates and qualifications? <https://www.hesa.ac.uk/data-and-analysis/students/outcomes> (page accessed September 29, 2021)
- Hirshleifer, D. A., Lim, S. S., & Teoh, S. H. (2004). Disclosure to an audience with limited attention. *Available at SSRN 604142*.
- Jin, G. Z. (2005). Competition and disclosure incentives: An empirical study of HMOs. *Rand Journal of Economics*, 93-112.
- Jin, G. Z., Luca, M., & Martin, D. (2021). Is no news (perceived as) bad news? An experimental investigation of information disclosure. *American Economic Journal: Microeconomics*, 13(2), 141-73.
- Jovanovic, B. (1982). Truthful disclosure of information. *The Bell Journal of Economics*, 36-44.
- King, R. R., & Wallin, D. E. (1991). Market-induced information disclosures: An experimental markets investigation. *Contemporary Accounting Research*, 8(1), 170-197.
- Koellinger, P. D., Mell, J. N., Pohl, I., Roessler, C., & Treffers, T. (2015). Self-employed but looking: A labour market experiment. *Economica*, 82(325), 137-161.

- Kroft, K., Lange, F., & Notowidigdo, M. J. (2013). Duration dependence and labor market conditions: Evidence from a field experiment. *The Quarterly Journal of Economics*, 128(3), 1123-1167.
- Laumer, S., Maier, C., & Eckhardt, A. (2015). The impact of business process management and applicant tracking systems on recruiting process performance: An empirical study. *Journal of Business Economics*, 85(4), 421-453.
- Li, Y. X., & Schipper, B. C. (2020). Strategic reasoning in persuasion games: An experiment. *Games and Economic Behavior*, 121, 329-367.
- Luca, M., & Smith, J. (2013). Salience in quality disclosure: Evidence from the US News college rankings. *Journal of Economics & Management Strategy*, 22(1), 58-77.
- Mathios, A. D. (2000). "The Impact of mandatory disclosure laws on product choices: An analysis of the salad dressing market." *Journal of Law and Economics*, 43(2): 651–77.
- McKelvey, R. D., & Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1), 6-38.
- McGuinness, S. (2003). University quality and labour market outcomes. *Applied Economics*, 35(18), 1943-1955.
- Merika, A., Kapareliotis, I., & Triantafyllou, A. (2018) TEF disclosure in UK higher education: Countersignaling and ranking consistency. *Mimeo*.
- Milgrom, P. R. (1981) "Good news and bad news: Representation theorems and applications", *Bell Journal of Economics*, 12(2): 380–391.

- Montero, M., & Sheth, J. D. (2019). *Naivety about hidden information: An experimental investigation* (No. 2019-11). CeDEx Discussion Paper Series.
- Mukherjee, A. N., Bhattacharyya, S., & Bera, R. (2014). Role of information technology in human resource management of SME: A study on the use of applicant tracking system. *IBMRD's Journal of Management & Research*, 3(1), 1-22.
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5), 1313-1326.
- Nunley, J. M., Pugh, A., Romero, N., & Seals Jr, R. A. (2016). College major, internship experience, and employment opportunities: Estimates from a résumé audit. *Labour Economics*, 38, 37-46.
- Oberholzer-Gee, F. (2008). Nonemployment stigma as rational herding: A field experiment. *Journal of Economic Behavior & Organization*, 65(1), 30-40.
- Sheth, J. D. (2021). Disclosure of information under competition: An experimental study. *Games and Economic Behavior*.
- Shin, H. S. (1994). News management and the value of firms. *The RAND Journal of Economics*, 58-71.
- Spence, M. (1973). Job Market Signaling. *The Quarterly Journal of Economics*, 87(3), 355-374.
- Stahl, D. O., & Wilson, P. W. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization*, 25(3), 309-327.
- Stahl, D. O., & Wilson, P. W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1), 218-254.

Stiglitz, J. E. (1975). The theory of "screening," education, and the distribution of income. *The American Economic Review*, 65(3), 283-300.

Suijs, J. (2007). Voluntary disclosure of information when firms are uncertain of investor response. *Journal of Accounting and Economics*, 43(2–3), 391–410.

Viscusi, W. K. (1978). A note on "lemons" markets with quality certification. *Bell Journal of Economics*, 9(1): 277–79.

Walker, I., & Zhu, Y. (2011). Differences by degree: Evidence of the net financial rates of return to undergraduate study for England and Wales. *Economics of Education Review*, 30(6), 1177-1186.

Weiss, A. (1995). Human capital vs. signalling explanations of wages. *Journal of Economic Perspectives*, 9(4), 133-154.

Zegners, D. (2016). Voluntary disclosure of product information: The case of e-book samples. *Available at SSRN 2380216*.

Zhou, F. S., & Zhou, Y. (2020). The dog that did not bark: Limited price efficiency and strategic nondisclosure. *Journal of Accounting Research*, 58(1), 155-197.

Tables and Figures

Table I: Logistic Regression

Dependent Variable: Positive Response (1 if yes, 0 otherwise)	Odds Ratio	Robust Standard Error
1 st	1.514***	(0.163)
2:1	1.394***	(0.141)
2:2	1.146	(0.130)
Undisclosed	1.265**	(0.130)
Posted 1 Day Ago	1.021	(0.0885)
Posted 2 Days Ago	0.790**	(0.0767)
Posted 3 Days Ago	0.723***	(0.0873)
Posted 4 Days Ago	0.763**	(0.105)
Posted 5 Days Ago	0.937	(0.130)
Posted 6 Days Ago	0.676**	(0.112)
Posted 7+ Days Ago	0.717	(0.184)
Baseline Odds	0.144***	(0.0252)
Jobseeker Fixed Effects	Yes	
Observations	12,278	
Wald Tests (p-values)		
Undisclosed = 1 st	0.079	
Undisclosed = 2:1	0.309	
Undisclosed = 2:2	0.375	
1 st = 2:1	0.414	
1 st = 2:2	0.005	
2:1 = 2:2	0.072	

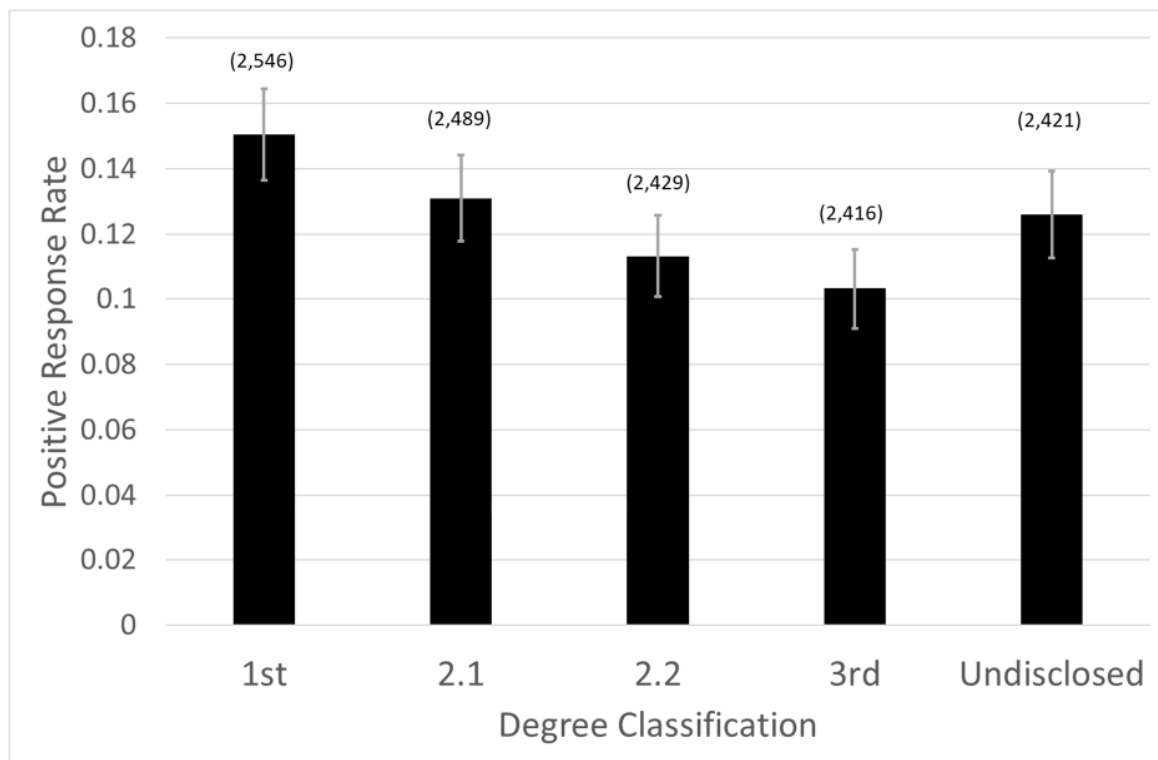
Notes: Table I reports odds ratios from a logistic regression. A ratio greater than 1 implies a positive effect; smaller than 1 implies a negative effect. The omitted degree class is 3rd. Robust standard errors, in parentheses, are clustered by hiring organisation (2,084 clusters). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure I: Example CV (1st and Undisclosed treatments)

<p>PROFILE</p> <p>I am a fresh graduate with a well-rounded personality and broad range of capabilities and interests. Although my degree was in maths, I also speak Chinese, have a strong interest in business and have been involved in journalism. I am seeking the right organization to begin my post-university career.</p> <p>CONTACT</p> <p>EMAIL: simonlittle19@outlook.com</p> <p>27 Tring Road, Northchurch, Berkhamsted HP4 3SY</p> <p>HOBBIES</p> <p>Writing Go karting Hiking Travelling</p>	<h1>SIMON LITTLE</h1> <p>EDUCATION</p> <p>Durham University 2016 - 2019</p> <p>Graduated in BSc Mathematics (Degree Class: 1st)</p> <ul style="list-style-type: none"> Final year modules included optimization and heuristics, multivariate analysis, stochastic processes, advanced numerical methods Dissertation on network theory Studied Chinese in first two years <p>WORK EXPERIENCE</p> <p>Sainsbury's Store Assistant 2016-2019</p> <ul style="list-style-type: none"> I worked part time as a shelf-stacker in order to help fund my studies <p>KPMG Summer Intern Summer 2018</p> <ul style="list-style-type: none"> Participated in intern program in the advisory department. Shadowed senior staff and performed administrative duties Developed key business skills and received on-the-job training workshops <p>OTHER ACTIVITIES</p> <ul style="list-style-type: none"> Reporter and editor for university newspaper Committee member of Durham University hiking society Member of student business society and winner of society's product design competition <p>SKILLS</p> <ul style="list-style-type: none"> Very strong oral and written communication skills Logical approach to problem solving Excellent computer skills, particularly Microsoft Office, MATLAB, SPSS
<p>PROFILE</p> <p>I am a fresh graduate with a well-rounded personality and broad range of capabilities and interests. Although my degree was in maths, I also speak Chinese, have a strong interest in business and have been involved in journalism. I am seeking the right organization to begin my post-university career.</p>	<p>EDUCATION</p> <p>Durham University 2016 - 2019</p> <p>Graduated in BSc Mathematics</p> <ul style="list-style-type: none"> Final year modules included optimization and heuristics, multivariate analysis, stochastic processes, advanced numerical methods Dissertation on network theory Studied Chinese in first two years

Note: The top picture displays the entire CV for the treatment disclosing a 1st. The bottom picture displays the education section of the CV in the treatment with undisclosed degree classification. The rest of the CV was identical between treatments.

Figure II: Positive Response Rates by Treatment



Note: Number of applications sent out in each treatment reported in parentheses. Error bars represent 95% confidence intervals.

Online Supplementary Materials

A. Survey Wording

The questions in this survey specifically refer to recruitment that your organization conducts by placing adverts on the website Monster.

Furthermore, these questions refer only to cases where applicants for jobs you are advertising must apply **directly** via Monster. Please note that these questions do not refer to cases where you place a job advert on Monster which requires applicants to apply on your own platform.

Q1: In cases where applicants have applied for a position directly via Monster, do you always rely on the CV management and evaluation tools available on Monster's platform, or do you ever download applicants' CVs and put them into another CV management/evaluation tool/platform (e.g. an "applicant tracking system" or similar software)?

- a) I always rely on the CV management and evaluation tools available on Monster's platform.
- b) I sometimes or always download applicants' CVs and put them into another CV management/evaluation tool/platform (e.g. an "applicant tracking system" or similar software).
- c) My organization does not place job adverts which require applicants to apply directly via Monster.

Q2 [*Does not appear if 1a or 1c are answered*]: In cases where applicants have applied for a position directly via Monster, do you always rely on the CV management and evaluation tools available on Monster's platform, or do you ever download applicants' CVs and put them into another CV management/evaluation tool/platform (e.g. an "applicant tracking system" or similar software)?

- a) Very rarely (less than 20% of CVs received via Monster)
- b) Rarely (20-40% of CVs received via Monster)
- c) Somewhat often (40-60% of CVs received via Monster)
- d) Often (60-80% of CVs received via Monster)
- e) Very often (more than 80% of CVs received via Monster)

Q3 [*Does not appear if 1a or 1c are answered*]: In your use of CV management/evaluation tools or platforms, have you ever used **filtering** tools to **automatically** screen in (i.e. include for consideration) or screen out (i.e. exclude from consideration) applicants on the basis of their **degree status** (i.e. whether the applicant has a degree or not)?

- a) Yes
- b) No

Q4 [*Does not appear if 1a or 1c are answered*]: In your use of CV management/evaluation tools or platforms, **when evaluating the CVs of graduates**, have you ever used **filtering** tools to **automatically** screen in (i.e. include for consideration) or screen out (i.e. exclude from consideration) applicants on the basis of their degree **classification** (i.e. which class a graduate achieved in their undergraduate degree: 1st Class, 2:1, 2:2 or 3rd class)?

- a) Yes
- b) No

Q5 [Does not appear if 1a, 1c or 4b are answered]: Please briefly describe the way you use tools to automatically filter applicants on the basis of degree classification.

Q6 [Does not appear if 1a or 1c are answered]: When evaluating CVs that applicants have sent directly via Monster, do you ever put CVs into software which reads the CV and uses artificial intelligence to provide a score or rank for the CVs?

- a) Yes
- b) No

Q7 [Does not appear if 1a, 1c or 6b are answered]: When evaluating CVs that applicants have sent directly via Monster, how often do you put CVs into software which reads the CV and uses artificial intelligence to provide a score or rank for the CVs?

- a) Very rarely (less than 20% of CVs received via Monster)
- b) Rarely (20-40% of CVs received via Monster)
- c) Neither rarely nor often (40-60% of CVs received via Monster)
- d) Often (60-80% of CVs received via Monster)
- e) Very often (more than 80% of CVs received via Monster)

Q8 [Does not appear if 1a, 1c or 6b are answered]: When you do this, which software do you use?

Q9: Do you have any additional remarks to add in relation to any of the questions in this survey?

B. Survey Responses

Respondent	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9
1	b	e	No	No		No			
2	a								no
3	b	e	No	No		No			no
4	a								
5	b	e	No	No		No			
6	b	b	No	No		No			no
7	c								
8	b	c	No	No		No			
9	c								
10	b	e	No	No		Yes	A	N/A	Automation is definitely becoming more present in the world of recruitment, however from personal experience CV review is very much a manual/human activity. I am aware for larger recruitment campaigns, automation is a necessity due to high volumes of applicants.

11	b	e	No	No		No			CVs are not written by AI so I wouldn't trust AI to screen which CVs I should consider are relevant. You may end up wrongly discounting candidates due to typos or an irregular formatting that the software doesn't anticipate (for instance 2:1, 2:i, 2.1, 2-1 etc).
12	b	a	Yes	No		No			Monster is by far my least favorite job board. I recruit in the Life Sciences/Pharma industry and I never get the right applicants through the Job Board. In honesty, I find it is the Job Board most used by Job Centre candidates and so get a lot of irrelevant applications. Also it is not as user friendly for recruiters compared to Reed and CV Library
13	b	d	Yes	No		No			
14	b	b	No	No		No			No
15	b	e	No	No		No			Our clients very very rarely ask us about whether a candidate has a degree. If they do its usually a degree catered to that specific industry they're hiring. For instance a Marketing Manager for a biotech company - they may require a degree in a biotech or related field.
16	b	d	No	No		No			n/a
17	c								We use multiple platforms to advertise so it's easier for us to track internally via excel.
18	c								We don't use Monster anymore - This used to be our most valued candidate-job board. The investment and response has dropped off so much we have cancelled our subscription. We now invest heavily in Indeed as this is where our candidate response comes from. Thanks
19	b	c	No	No		No			No, thank you.
20	b	e	Yes	Yes	Questionnaire as part of the application	No			
21	b	e	No	No		Yes	e	Access	None
22	b	e	No	No		No			
23	c								
24	b								

25	b	e	Yes	No		No			
26	b	d	No	No		No			No
27	c								Never even heard of CV evaluation software
28	c								No additional remarks
29	a								
30	c								No
31	b	e	No	No		Yes	e	RDB	no
32	b	b	Yes	No		Yes	c	Our own Gravity System	Why is this so focused on Monster - there are other better boards out there.
33	a								No
34	b	d	Yes	No		Yes	d	JobAdder	I think the rise of AI within the process of recruitment and hiring is a compelling part of the puzzle, however I don't believe there is a silver bullet and whilst the most obvious benefits (automation of admin heavy processes) will add value, I'm not convinced that human intervention can ever or should ever be eradicated.
35	b	a	Yes	Yes	for graduate roles this is important	No			no
36	a								No
37	b	e	Yes	Yes	We use boolean searches to filter the candidates we are looking for i.e only including 1st or other similar terms if the role we are recruiting for requires this as a must have	Yes	e	Vincere, our ATS/CRM	N/A
38	b	b	No	No		No			
39	c								
40	b	b	No	No		No			
41	c								We always ask for a covering letter to accompany a CV