

Dinlersoz, Ermin; Dunne, Timothy; Haltiwanger, John C.; Penciakova, Veronika

**Working Paper**

## Business formation: A tale of two recessions

Working Paper, No. 2021-5

**Provided in Cooperation with:**

Federal Reserve Bank of Atlanta

*Suggested Citation:* Dinlersoz, Ermin; Dunne, Timothy; Haltiwanger, John C.; Penciakova, Veronika (2021) : Business formation: A tale of two recessions, Working Paper, No. 2021-5, Federal Reserve Bank of Atlanta, Atlanta, GA,  
<https://doi.org/10.29338/wp2021-05>

This Version is available at:

<https://hdl.handle.net/10419/244308>

**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.*

## Business Formation: A Tale of Two Recessions

Ermin Dinlersoz, Timothy Dunne,  
John Haltiwanger, and Veronika Penciakova

Working Paper 2021-5  
January 2021

**Abstract:** The trajectory of new business applications and transitions to employer businesses differ markedly during the Great Recession and the COVID-19 recession. Both applications and transitions to employer startups decreased slowly but persistently in the post-Lehman crisis period of the Great Recession. In contrast, during the COVID-19 recession new applications initially declined but have since sharply rebounded, resulting in a surge in applications during 2020. Projected transitions to employer businesses also rise, but this projection is dampened by a change in the composition of applications in 2020 toward applications that are more likely to be nonemployers.

JEL classification: D2, E65, G33

Key words: COVID-19, business failures, liquidity, small business

<https://doi.org/10.29338/wp2021-05>

---

The authors thank Shawn Klimek and Scott Ohlmacher for helpful comments. Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau, the Federal Reserve System, the Board of Governors, or their staff. The Census Bureau's Disclosure Review Board and disclosure avoidance officers have reviewed this data product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release (approval number CBDRB-FY21-062). Any remaining errors are the authors' responsibility.

Please address questions regarding content to Ermin Dinlersoz, U.S. Census Bureau, [ermin.m.dinlersoz@census.gov](mailto:ermin.m.dinlersoz@census.gov); Timothy Dunne, University of Notre Dame, [tdunne1@nd.edu](mailto:tdunne1@nd.edu); John Haltiwanger, University of Maryland, [haltiwang@econ.umd.edu](mailto:haltiwang@econ.umd.edu); or Veronika Penciakova, Federal Reserve Bank of Atlanta, [veronika.penciakova@atl.frb.org](mailto:veronika.penciakova@atl.frb.org).

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed's website at [www.frbatlanta.org](http://www.frbatlanta.org). Click "Publications" and then "Working Papers." To receive e-mail notifications about new papers, use [frbatlanta.org/forms/subscribe](http://frbatlanta.org/forms/subscribe).

How have recent recessions affected business formation? The startup rate of new employer businesses declined sharply in the Great Recession and rose slowly in the subsequent recovery. Indeed by 2018, the startup rate remained 24% lower than in 2006 (see online appendix). There is evidence that the COVID-19 Recession is leading to widespread business closure ([Leland Crane, Ryan Decker, Aaron Flaaen, Adrian Hamins-Puertolas and Christopher Kurz, 2020](#)), but less is known about its impact on business formation.

In this paper, we use a new Census Bureau data series on applications for Employer Identification Numbers (EIN), the Business Formation Statistics (BFS), to describe the path of business formation in the Great Recession and the COVID-19 Recession. There are three advantages of the BFS. First, it can shed light on the path of current and future business formation because new employer businesses must obtain an EIN in order to file payroll taxes. Second, it is currently published on a weekly basis, enabling a high frequency and timely analysis. Finally, [Kimberly Bayard, Emin Dinlersoz, Timothy Dunne, John Haltiwanger, Javier Miranda and John Stevens \(2018\)](#) shows that new business applications can be used to track the fluctuations in actual new employer business formation.

By combining the BFS data on applications with records of employer businesses from the Longitudinal Business Database (LBD), we can measure both the evolution of business applications (BA) and the observed or modeled transition of BA to employer businesses. Our findings show marked differences in the path of BA between the Great Recession and the COVID-19 Recession. The fall in BA is slow and persistent during the Great Recession, whereas BA experience a sharp V-shaped recovery in the COVID-19 Recession. However, the composition of BA shifts substantially during the COVID-19 Recession toward those with a lower propensity to transition to employer businesses. This composition shift implies that the surge in new BA in 2020 is likely to yield a more substantial increase in new non-employer businesses relative to new employer businesses.

## 1 Measuring Business Formation

The analysis uses the BFS, which tracks the weekly number of BA by state. The BFS is derived from administrative data from the IRS on EIN applications. All employer businesses in the United States are required to have an EIN to file payroll taxes. New non-employer businesses also file for an EIN if forming a partnership or an incorporated business. Even new sole proprietor non-employers often file for an EIN to facilitate their business activity (e.g., working with other businesses or opening a business bank account). The application form includes the name and address of the applicant and business, business start date, type of business entity, principal industry, and planned date of initial wage payments. The filing date and business location information are used to aggregate individual applications to the state level at a weekly frequency. The main data series analyzed, BA, represents the number of EIN applications

received in a state in a specific week of the year.

We complement the confidential microdata underlying the weekly BFS with the LBD. The LBD contains firm and establishment-level information on age, location, industry, number of employees, quarterly payroll, and EIN for the near-universe of employer businesses in the United States. Using the EIN, the BFS is matched to the LBD to identify the incidence and timing of transitions to employer businesses. In particular, we focus on transitions that occur within the first four quarters of the application date, which are formally referred to in the BFS as BF4Q. [Bayard et al. \(2018\)](#) find that the majority of the transitions from application to employer business occur within the first year of the EIN filing.

The matched data are employed in two ways to obtain the number of applications that transition, or are likely to transition, to employer businesses. For the Great Recession, the count of realized transitions is aggregated to the state-week level, as this is directly observable in the matched data. For the COVID-19 Recession, transitions to employer business are not yet measured in the LBD, which is currently available through 2018. Instead, we project the number of transitions during 2019 and 2020 using a linear probability model (LPM) following the approach discussed in [Bayard et al. \(2018\)](#). The model generates a predicted probability of transitioning to an employer business for each application, including the out-of-sample period from 2019 through 2020. The probabilities are aggregated to yield an estimate of the number of applications that will transition to employer status at the state-week level. This approach allows for the number of estimated transitions to vary by both the number of applications received and the composition of those applications (see online appendix).

We use a regression framework to evaluate whether the Great Recession and COVID-19 Recession significantly affected the trajectory of BA and transitions. In doing so, we first construct measures of cumulative applications and transitions for crisis and normal time benchmark (reference period), forwards and backwards relative to a week 0 defined below. This approach smooths out seasonality observed in the weekly flow of these data. Second, we construct our outcome variables as the growth rate in cumulative applications (or transitions) during the crisis period relative to the reference period. Finally, we estimate application-weighted (or transition-weighted) regressions of this growth rate on distance from week 0 dummies and state fixed effects to capture how deviations from normal times evolve within states during each crisis. Note that each distance dummy contains two weeks, with the week 0 dummy containing weeks zero and one.

## 2 Evolution of Business Applications and Transitions

We begin by comparing the paths of BA and transitions for the two recessions. Figure 1a depicts the growth rates of cumulative applications and transitions during the crisis period relative to the reference period of the Great Recession. All statistics should be interpreted as

deviations relative to the two-week period containing week 0, which is normalized to zero.<sup>1</sup> Focusing on deviations relative to week 0 allows us to highlight high frequency cyclical variation, abstracting from lower frequency trends. The bands represent 95 percent confidence intervals. Application growth rates (the black line) decline after week 0 and remain persistently lower 52 weeks after the Lehman collapse. The second line, depicting transition growth rates, exhibits a sharp initial decline and a subsequent flattening out.

Figure 1b focuses on the COVID-19 Recession.<sup>2</sup> In March and April 2020, the growth rate of applications drops sharply, but subsequently experiences a strong rebound. Indeed within 18 weeks, applications have recovered and in subsequent weeks rise sharply above the reference week level of activity. The differences in projected transition growth rates fall even more sharply than applications and recover more slowly. As will become clear below, to ascertain the quantitative implications for projected transitions in 2020, we need to consider patterns without normalizing week 0 differences to zero.

Figures 2a and 2b highlight the differences in the evolution of transitions by depicting the implied differences in transition rates from Figure 1. They represent the difference between the crisis and reference period in the transition rate of cumulative applications. Figure 2a focuses on the Great Recession, and shows a decline post-Lehman, followed by a recovery in transition rates within 30 weeks.<sup>3</sup> During COVID-19, depicted in Figure 2b, the projected transition rate shows a sharp decline at the onset of the crisis and no recovery 30 weeks later.

The contrast between the Great Recession and the COVID-19 Recession is stark. Applications dropped and remained low in the Great Recession, but after an initial decline, transition rates recovered. While actual transitions fell, their decline was driven by the fall in applications, rather than a persistent decline in transition rates. In contrast, during the COVID-19 Recession applications initially fell sharply and rebounded several weeks later, while the transition rate also fell but remain depressed through week 30. Differences between the two recessions in the evolution of transition rates reflect differences in the compositional changes in applications.

---

<sup>1</sup>Week 0 is the two-week period ending September 20, 2008 in the crisis period, which coincides with the Lehman Brothers collapse, and the equivalent week in 2006 for the reference period.

<sup>2</sup>Week 0 is the two-week period ending March 14, 2020, and the equivalent week in 2019 for the reference period. The growth rate of applications during the crisis relative to the reference period are shown for weeks 0 through 40 using the public domain BFS. The growth rate of transitions during the crisis relative to reference period is reported for weeks 0 through 30.

<sup>3</sup>The differences are normalized to zero in the two-week period containing week 0. Some caution is required in interpreting the path of the transition rates post-Lehman. As shown in the appendix there is a decline in transition rates from 2006 to 2008 pre-Lehman – much of this occurs in 2007, before the Great Recession.




Figure 1: Business Applications and Transitions




Figure 2: Transition Rates

### 3 Changing Composition of Business Applications

The decline in observed and projected transition rates during both crises is indicative of compositional changes in applications playing a role in business formation during recessions. To explore the relevance of these compositional changes, we first disaggregate applications into two groups – business applications with planned wages (WBA) and non-wage business applications (non-WBA). WBA are submissions that indicate the intent to pay wages. [Bayard et al. \(2018\)](#) report the transition rate for WBA to be substantially higher than non-WBA.

Figures 3a and 3b contrast the differences in WBA and non-WBA growth rates in cumulative differences for the two recessions. In the Great Recession, WBA initially fell more rapidly than non-WBA. However, after about 20 weeks, the gap between the decline in WBA and non-WBA stabilized.

During the COVID-19 Recession, both types of applications initially declined sharply. After around 4 weeks, non-WBA began to rebound, while WBA continued to decline. Though WBA started to rebound after about 8 weeks, by week 40, the gap between the cumulative differences in the two types of applications remains substantial. This pattern is consistent with the fall in the projected transition rate.

Other compositional changes are also important for understanding the decline in projected transition rates during the COVID-19 Recession. A special projects release from the weekly BFS reports that BA among non-store retailers (NAICS 454) account for the vast majority of the observed surge in BA in 2020. This is not surprising, given the pandemic has provided strong incentives for shifting retail activity to online transactions. Yet, in 2018 over 90% of non-store retailers are non-employer businesses (see combined NES-CBP data). Therefore, this concentration of applications among non-store retailers also helps explain the declining transition rate in Figure 2b.

To further highlight the importance of the changing composition of applications during the COVID-19 Recession, we evaluate the number of new businesses that would be formed during this period if applications filed during the COVID-19 Recession had the same composition as those filed during 2019. The projected number of transitions would be 17.2 percent higher during COVID-19 if the composition of applications remained unchanged from the reference period.

While the compositional change implies a reduction in the transition rate, our findings still imply an increase in employer startups in 2020, relative to 2019, from the increase in applications. To understand this implication, it is instructive to consider the patterns of transitions in the COVID-19 Recession without normalizing the week 0 difference to zero. In online appendix Figure A1, we find that by week 30, cumulative transitions are about 8% higher in 2020 relative to 2019. Consistent with this finding, Figure A4 in the online appendix shows that applications with a high propensity to transition to employer business (HBA) are 17% higher by week 40 in 2020 relative to 2019. At the national level, this implies about 200 thousand more HBA by week 40 in 2020 relative to 2019. These findings suggest that the surge in applications in 2020 will translate into an increase in actual startups in 2020 relative to 2019 – it just will not be as large as implied by the increase in overall applications.

The compositional changes imply an even larger surge in non-employer businesses in 2020. We find that the cumulative increase in applications for likely non-employers (Figure A.4, non-HBA) is 30% in week 40 of 2020 relative to 2019. At the national level, this implies about 600 thousand more non-HBA by week 40 in 2020 relative to 2019. Non-employer businesses are an important source of business activity. Counting each non-employer establishment as one job and combining with County Business Patterns data, non-employers accounted for about 17 percent of jobs in 2018. Many non=employers generate a small amount of revenue, but EIN non=employers are three times as large in terms of revenue as SSN non=employers (see online appendix for more discussion).




Figure 3: WBA versus non-WBA

## 4 Conclusion

Our findings show that the trajectory of business applications and transitions to employer businesses differ markedly during the Great Recession and COVID-19 Recession. The persistent reduction in actual employer business formation during the post-Lehman crisis period of the Great Recession is primarily accounted for by a decrease in applications, as transition rates to employer businesses recover after an initial decline. In contrast, during the COVID-19 Recession new applications initially declined but have since sharply rebounded, resulting in a surge in applications during 2020. Although our analysis predicts a rise in employer startups in 2020 relative to 2019, this rise is dampened by a reduction in the projected transition rate of applications to employer businesses. Compositional changes in applications underlie this reduction, including a decrease in the share of applications with intent to pay wages and an increase in the share of applications in industries dominated by non-employers (e.g., non-store retailers). These compositional changes suggests some caution regarding the observed surge in new applications in 2020 in terms of their implications for new employer businesses. On the other hand, the surge in applications of likely non-employers in 2020 has been quite dramatic.



# A Appendix

**Applications and Transitions Data:** The data for all business application series (BA, HBA, WBA) are obtained from the public domain Business Formation Statistics (BFS). The BFS publishes quarterly data on business applications, as well as realized (BF4Q) and projected (PBF4Q) transitions. Starting soon after the onset of the COVID-19 Recession, the BFS began releasing weekly data on business applications at the national, regional, and state level. In our analysis, we use the public domain weekly BFS business application data. Since actual and projected transitions are only available at a quarterly frequency in the public domain BFS, the actual and projected transitions and transition rates reported in this paper are based on the integration of the BFS and Longitudinal Business Database (LBD) micro data, as described in the main text. In addition, we used the newly redesigned LBD in this process which enabled us to project transitions for 2019 and 2020 based on actual transitions through 2018.

**Projecting Transitions:** The empirical model used to predict transitions is described in detail in [Bayard et al. \(2018\)](#). Briefly, the projections are based on a linear probability model (LPM), which relates observed transitions of business applications to employer businesses in the LBD to a rich set of predictors from the EIN application form. The variables include indicators of the business start date, type of entity, industry, limited liability status, reason for applying, and wage date. The model also controls for location (state), week of application submission within the year, presence of prior EIN, and whether the application indicates a trade name, an executor’s name, or a distinct business address. The covariates of the empirical model also include a rich set of interactions between industry, wage date, type of entity, and reason for applying. Figure 6 of [Bayard et al. \(2018\)](#) shows that projected transitions track actual transitions closely, both within and out of sample.

**Crisis and Reference Periods:** For the reference period (or normal time benchmark) for the Great Recession, we use 2006 to avoid any overlap between the crisis and reference periods between weeks 0 and 52. For the COVID-19 Recession, at the time the projection model was estimated, BFS data were available through the week ending October 3, 2020 (which is 30 weeks after the start of the crisis period). Since the BFS currently releases new data weekly, in the main text and appendix figures we use applications through the week ending December 19, 2020 in our analysis. The estimates in the figures for 40 weeks after the crisis period reflect the two week period that ends December 19, 2020.

**Growth Rate Calculation:** For the analysis, we compute growth rates for cumulative applications and transitions using the DHS ([Steven Davis, John Haltiwanger and Scott Schuh \(1996\)](#)) growth rate, defined as  $(X_{cw} - X_{rw}) / (0.5 * (X_{cw} + X_{rw}))$  for variable  $X$ , where  $r$  is the reference period,  $c$  is the crisis period and  $w$  is the week through which  $X$  has been calculated. The DHS growth rate is a second order approximation of the log first difference and accommodates zeroes.

**Weighting:** In the main text, we report application-weighted (or transition-weighted) regression results, where the weights are the relevant state-level DHS denominator. In unreported results, we estimated both population-weighted and unweighted regressions. These regressions yield results very similar to those reported in the main text.

**Startup and Non-employer Statistics:** The discussion in the main text of the decline in actual new employer startups draws from the [Business Dynamic Statistics](#). Further analysis of the decline in employer startups during the Great Recession can be found in [Ryan Decker, John Haltiwanger, Ron Jarmin and Javier Miranda \(2014\)](#) and [Steven Davis and John Haltiwanger \(2019\)](#). The discussion in the main text on the importance of non-employers draws from [Steven Davis, John Haltiwanger, Ron Jarmin, C.J. Krizan, Javier Miranda, Alfred Nucci and Kristin Sandusky \(2009\)](#) and the [2018 Non-employer Statistics](#). [Davis et al. \(2009\)](#) find that more than ten percent of EIN non-employers ultimately transition to employer businesses. Most of these transitions occur at least one year after the non-employer business started. Consequently, such transitions are not included in the application transition rates discussed in the main text.

**Additional Figures:** Figures A1-A3 are analogous to Figures 1-3, but without normalizing the week 0 difference to zero.<sup>4</sup> Figure A4 depicts the growth rate of cumulative high-propensity business applications (HBA) and non high-propensity business applications (non-HBA) in the crisis relative to the reference period, without normalizing week 0 to zero. As discussed in [Bayard et al. \(2018\)](#), non-HBA have a very low probability of transiting to employer businesses in the next 4 quarters (less than 3%), while HBA have a 27% transition rate. HBA are those applications that have either (i) planned wages, (ii) are an application for a corporation, or (iii) are in small number of industries that have notably higher transition rates. These characteristics are among those used in the much richer LPM model to project transitions of applications to employer businesses. HBA are intended as a transparent way to classify applications on core characteristics.

Figure A1 shows that there is a much larger gap at week 0 between applications and transitions in the Great Recession (Figure A1a) compared to the COVID-19 Recession (Figure A1b). Indeed, in the Great Recession, the gap is large and negative, while in the COVID-19 Recession the gap is smaller and positive. For the Great Recession, this means that the growth rate in applications from reference week 0 to crisis week 0 is close to zero, while the analogous growth rate for transitions is very negative. This large negative gap, observed during the Great Recession, is consistent with the large negative decline in the transition rate in Figure A2a in week 0. Importantly, these patterns indicate that the transition rate declined before the Lehman collapse (between 2006 and 2008). Consequently, our inference of the decline in transitions during the Great Recession being mostly due to the decline in applications refers specifically to the

---

<sup>4</sup>In practice, we run the same regressions, but rather than reporting the coefficients and 95% confidence interval of the distance from week 0 dummies, we report the predictive margins and their 95% confidence interval evaluated at the distance from week 0 dummies.

post-Lehman period. To help understand the pre-Lehman decline in the transition rate, a few additional findings are relevant. First, from the BFS, the composition of applications changed significantly in 2007 even before the beginning of the recession in December 2007. In January 2007, the share of applications that are HBA is 62% and WBA is 38%. In November 2007, these shares are 51% and 25% respectively. These compositional changes imply a significant decline in the transition rate pre Great Recession which is consistent with the patterns in Figure 6 of Bayard et al. (2018). The latter shows both actual and projected transitions declined substantially in 2007 pre Great Recession. In the main text we intentionally abstract from these pre Great Recession changes in applications and transitions as this reflects lower frequency factors underlying the well-known negative trend in employer startups in the U.S. that accelerated in the post-2000 period. As discussed in the main text, we abstract from these trend effects by focusing on deviations from week 0.

In contrast, in the COVID-19 Recession, the transition rate at week 0 in 2020 is slightly higher than in week 0 in 2019. There are not sizable pre-recession trend effects in 2019 compared to early 2020. Figure A3 shows similar patterns to Figure A1 for the WBA and non-WBA during both recessions. Importantly, during the COVID-19 Recession, both WBA and non-WBA show marked increases by week 40 in 2020, relative to 2019.

For the COVID-19 Recession, Figure A4 shows that cumulative HBA are 17% higher by week 40 in 2020, relative to 2019. It is instructive to compare this increase of 17% in HBA to the 25% increase in overall applications in Figure A1. This difference is consistent with our finding that during the COVID-19 Recession, the transition rate has declined as the share of HBA in overall applications has declined. However, HBA are increasing substantially in 2020, relative to 2019, which implies that employer startups are likely to be higher in 2020 than in 2019. At the same time, as discussed in the main text, the increase in employer startups is likely to be smaller than suggested by the overall increase in business applications.




Figure A.1: Business Applications and Transitions (non-normalized)




Figure A.2a: Great Recession — Transition Rate (actual, non-normalized)




Figure A.2b: COVID — Transition Rate (projected, non-normalized)

Figure A.2: Transition Rates (non-normalized)




Figure A.3a: Great Recession — WBA vs. non-WBA (non-normalized)




Figure A.3b: COVID — WBA vs. non-WBA (non-normalized)

Figure A.3: WBA versus non-WBA (non-normalized)




Figure A.4a: Great Recession — HBA vs. non-HBA (non-normalized)




Figure A.4b: COVID — HBA vs. non-HBA (non-normalized)

Figure A.4: HBA versus non-HBA (non-normalized)

## References

- Bayard, Kimberly, Emin Dinlersoz, Timothy Dunne, John Haltiwanger, Javier Miranda, and John Stevens.** 2018. "Early Stage Business Formation: An Analysis of Applications for Employer Identification Numbers." NBER Working Paper No. 24364.
- Crane, Leland, Ryan Decker, Aaron Flaaen, Adrian Hamins-Puertolas, and Christopher Kurz.** 2020. "Business Exit during the COVID-19 Pandemic: Non-Traditional Measures in Historical Context." FEDS Working Paper No. 2020-089.
- Davis, Steven, and John Haltiwanger.** 2019. "Dynamism Diminished: The Role of Housing Markets and Credit Conditions." NBER Working Paper No. 25466.
- Davis, Steven, John Haltiwanger, and Scott Schuh.** 1996. *Job Creation and Destruction*. MIT Press.
- Davis, Steven, John Haltiwanger, Ron Jarmin, C.J. Krizan, Javier Miranda, Alfred Nucci, and Kristin Sandusky.** 2009. "Measuring the Dynamics of Young and Small Businesses: Integrating the Employer and Nonemployer Universes." In *Producer Dynamics: New Evidence from Micro Data*. 329–366. NBER and University of Chicago Press.
- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda.** 2014. "The Role of Entrepreneurship in U.S. Job Creation and Economic Dynamism." *Journal of Economic Perspectives*, 24(3): 3–24.