NORTHWESTERN UNIVERSITY

Essays in Macroeconomics and Firm Dynamics

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL IN PARTIAL
FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Economics

By

Jane Olmstead-Rumsey

EVANSTON, ILLINOIS

June 2021

Abstract

This dissertation is divided into three chapters. Chapter one studies changes in market concentration and productivity growth in the United States from the 1990s to the 2010s. Chapter two measures the impact of a banking crisis, the British Panic of 1825, on non-financial firms. Chapter three examines how women's employment during pandemic-induced recessions differs from typical recessions and presents a model of the macroeconomic consequences of these differences.

Acknowledgements

I am exceedingly grateful for the advice, insight, encouragement, and support of my committee members Matthias Doepke, Guido Lorenzoni, and David Berger. This work would not be possible without each of them. I also benefited enormously throughout the years from discussions with other Northwestern faculty members and visitors including Luigi Bocola, Nicolas Crouzet, Carola Frydman, Benjamin Jones, Lorenz Kueng, Kiminori Matsuyama, Martí Mestieri, Joel Mokyr, Dimitris Papanikolaou, Giorgio Primiceri, Matthew Rognlie, and Christina Patterson. I also thank Lori Beaman for her work creating the Women in Economics group during my time at Northwestern. Many thanks to all the Northwestern economics faculty members for imparting wisdom during coursework, offering suggestions in seminars, and assisting with preparation for the job market.

I very gratefully acknowledge financial support from the National Science Foundation Graduate Research Fellowship Program Grant No. DGE-1324585.

I made some of the most significant progress on this dissertation during visits to the Centre de Recerca en Economia Internacional in Barcelona in 2019 and the Federal Reserve Bank of Chicago in 2020. My sincere thanks to everyone in both groups for hosting me, especially Jan Eeckhout in Barcelona and Gadi Barlevy, Jason Faberman, Jonas Fisher, François Gourio, and Jing Zhang at the Chicago Fed. I also thank Ufuk Akcigit and Simon Mongey for providing comments and suggestions at an early stage.

It has been a pleasure and a wonderful opportunity to learn alongside all my classmates, but especially Bence Bardóczy, Ricardo Dahis, Ana Danieli, Egor Kozlov, Utsav Manjeer, and Kristina Manysheva. I'm also very grateful to Northwestern alums Titan Alon, Gabriela Cugat, Sasha Indarte, and Abdoulaye Ndiaye for their advice, example, and support.

My undergraduate advisor Paul Carrillo first recommended I undertake a PhD, and I could not be more grateful for his early investments in me and his continued support throughout the years. The same can be said for supervisors and friends at the Federal Reserve Board, especially Shaghil Ahmed and Ozge Akinci. More teachers and mentors than I can count have contributed to my success, curiosity, and knowledge over the years and I'm grateful to each and every one of them.

Last but certainly not least, I'm grateful to my friends and family who loved me, checked in on me, cheered me on, made me laugh, and kept things in perspective. I want to especially thank my parents Mary and Karl and my sister Emma for being my first advisory committee and the most supportive family I could ever ask for, and my husband Alex for seeing this process more intimately than absolutely anyone else, for being my biggest advocate and my own personal comedian, and for always believing in me with absolute certainty.

Contents

List of Figures

List of Tables

Chapter 1

Market Concentration and the Productivity Slowdown

Abstract

Since around 2000, U.S. aggregate productivity growth has slowed and product market concentration has risen. To explain these facts, I construct a measure of innovativeness based on patent data that is comparable across firms and over time and show that small firms make innovations that are more incremental in the 2000s compared to the 1990s. I develop an endogenous growth model where the quality of new ideas is heterogeneous across firms to analyze the implications of this finding. I use a quantitative version of the model to infer changes to the structure of the U.S. economy between the 1990s and the 2000s. This analysis suggests that declining innovativeness of smaller firms can account for about 40 percent of the rise in market concentration over this period and the entire productivity slowdown. Strategic changes in firms' R&D investment policies in response to the decreased likelihood of laggards making drastic improvements significantly amplify the productivity slowdown.

1.1 Introduction

After a boom in the 1990s, U.S. productivity growth began to decline in the early 2000s and industry leaders began capturing an increasingly large share of sales across many sectors of the economy. In this paper I illustrate a new mechanism that explains both of these trends, along with patterns of increasing profitability, increasing productivity differences, and a declining rate of market leadership turnover in U.S. industries. This mechanism is the declining ability of laggard firms to catch and overtake market leaders through innovation.

To support the existence of this mechanism in the data, I show that patent quality, one measure of innovativeness, has fallen sharply among smaller firms since 2000 after a boom in the 1990s. This fact is robust to using market-based measures of patent value or measures of the social value of firms' patents using citation counts. It is also a broad based phenomenon, spanning many sectors of the economy, though it is most pronounced in high tech sectors. I also show that there has been less turnover in market leadership in many industries since 2000.

To understand how leaders (the largest firm in terms of sales in each industry) and laggards (other, smaller firms) respond to diminished opportunities for laggard firms to grow through innovation, and the effect of these responses on aggregate growth, I develop and estimate a general equilibrium, quality ladder model of directed innovation along the lines of Aghion et al. (2001). There are a continuum of industries populated by two incumbent firms producing differentiated goods. Incumbent firms can improve their varieties through innovation. Each industry also contains a competitive fringe of firms with the ability to imitate whichever incumbent has the lower quality variety. Market concentration, measured as the market leader's share of industry sales, is high when the quality difference between the two incumbents' product varieties is large.

Unlike Aghion et al. (2001), the model accommodates the possibility that laggard firms have an "advantage of backwardness," allowing them to improve their variety more drastically than market leaders do when they innovate. A model parameter governs the extent of laggard firms' advantage of backwardness. Other model parameters capture alternative explanations for rising market concentration, slowing productivity growth, or both, that have been suggested in the literature such as slowing knowledge diffusion or entry rates, increasing product differentiation, rising market power, and declining real interest rates.

To infer the relative importance of different changes to the U.S. economy in explaining these trends, I estimate the model parameters for two steady states to match data on concentration, productivity growth, the profit share, patent quality, the rate of turnover in market leadership, aggregate R&D expenditures, and R&D expenditures at the firm level for the U.S. in two separate periods, the 1990s and the 2000s. This exercise suggests a dominant role for the parameter governing laggard firms' advantage of backwardness to explain rising concentration and slowing productivity growth compared to other explanations.

I then use the quantitative model to explore the channels through which laggards' declining patent quality alone can explain trends in concentration and productivity growth. When laggards firms' advantage of backwardness declines, these firms respond to a lower chance of attaining market leadership by investing less in research and development. Facing a diminished probability of being overtaken, market leaders invest slightly more. Together, these decisions lead to larger average quality differences between leaders and laggards in steady state. The model has a direct mapping from quality differences to leaders' market shares and markups such that sales concentration and the profit share of total output also rise. This change explains about 40 percent of the observed rise in concentration between the 1990s and the 2000s, and about 25 percent of the increase in the profit share.

Laggards' declining patent quality also has implications for productivity growth in the model. The source of endogenous growth in the model is quality improvements to the differentiated products that firms produce. Making laggards' quality improvements more incremental generates a productivity slowdown through two channels. One is direct: even if firms devoted the same resources to research and development, the economy would grow more slowly because average quality improvements are smaller than before. Second, there is a strategic effect that amplifies the productivity slowdown: because of the relocation of R&D expenditures in the economy towards market leaders, whose innovations tend to be more incremental, the economy grows even more slowly than before. According to a growth decomposition exercise, this latter force accounts for roughly half of the productivity slowdown in the model. Quantitatively, the estimated decline in laggards firms' advantage of backwardness generates a productivity slowdown in the model of a similar magnitude to the slowdown observed in the U.S.[1]

Finally, I discuss how the model predictions are consistent with patterns of widening productivity gaps within sectors, rising markups, sector level correlations between the productivity slowdown and rising concentration, and the fact that industry leaders conduct a larger share of total industry R&D in the 2000s than in the 1990s.

Related Literature and Contribution This paper contributes a novel mechanism to the large and growing literature linking trends in concentration, productivity growth, and business dynamism using models of endogenous growth. Many of these papers emphasize the increasing importance of intangible assets and information and communications technology (ICT) as a possible explanation (Aghion et al. (2019a), de Ridder (2020), Corhay, Kung, and Schmid (2020)). Non-technological explana-

[1]The fact that changing innovativeness alone can explain the entire productivity slowdown does not rule out other explanations that have been proposed, since there may be forces working to increase productivity growth that the model does not capture such as population growth, entry, improvements in human capital, and globalization.

tions include demographic changes (Hopenhayn, Neira, and Singhania (2018), Peters and Walsh (2019), Karahan, Pugsley, and Sahin (2019), Engbom (2019), Eggertsson, Mehrotra, and Robbins (2019), Bornstein (2018)) or declining real interest rates (Liu, Mian, and Sufi (2019), Chatterjee and Eyigungor (2019)). The most closely related explanation is the one in Akcigit and Ates (2020) and Akcigit and Ates (2019) that diffusion of knowledge from leaders to laggards is slowing down, either because of ICT and the increasing importance of data in firms' production processes or because of anti-competitive use of patents.

Rather than emphasizing particular features of information technology, the theory presented here instead hypothesizes that general purpose technologies (GPTs) may affect firm dynamics and market structure in addition to raising aggregate productivity growth. Past fluctuations in patent quality and productivity growth have been attributed to waves of innovation due to the arrival of new GPTs (Kelly et al. (2018); Kogan et al. (2017)). Bresnahan and Trajtenberg (1995) note that GPTs are applicable in a wide range of sectors and exhibit innovational complementarities, meaning that they increase the productivity of downstream research and development efforts.[2] Given the new evidence presented here on heterogeneity in patent quality across firms and time, I argue that these innovational complementarities appear to be stronger for smaller firms than for market leaders.[3]

[2]See Brynjolfsson, Rock, and Syverson (2018) for further discussion.
[3]See Section 1.2.2 for further discussion.

Most neo-Schumpeterian growth models assume goods within sectors are perfect substitutes so that each sector has just one producer in each period (see Klette and Kortum (2004), Lentz and Mortensen (2008), Acemoglu and Cao (2015), and Akcigit and Kerr (2018), for leading examples). Because of this, these models are not well-suited to address industry-level moments such as sales concentration. Introducing a duopoly (plus a competitive fringe) allows me to make unified predictions both about market concentration at the industry level and firm-level innovation rates, and makes not only markups but also sales concentration within sectors an endogenous outcome of the innovation process.

The duopoly formulation also brings together previously distinct strands of literature in macroeconomics concerned with (i) slowing growth (ii) changes in market structure and potentially market power and (iii) superstar firms. Strands (ii) and (iii) typically rely on opposing assumptions. According to the literature on rising market power, incumbent firms exercise greater pricing power now than in the past and this is reflected in rising markups and profitability (de Loecker, Eeckhout, and Unger (2020), Barkai and Benzell (2018)). On the other hand, the literature on superstar firms contends that greater import competition and greater consumer price sensitivity due to better search technology like online retail have *increased* competitive pressures and reduced the market power of incumbent firms, resulting in reallocation to the most productive (superstar) firms (Autor et al. (2020)). The model resolves this seeming contrast by demonstrating how markups can rise at the same time as there is reallocation to relatively more productive firms without any changes at all to consumer preferences or the aggregate production function. The model is also consistent with the finding of Kehrig and Vincent (2020) that being a superstar firm is a temporary rather than permanent status. In the model, the relative advantage of high value added firms grows in the 2000s and the average duration of these "shooting star" spells increases, but these firms are eventually displaced by competitors.

The model's industry structure with imperfect substitutes makes it possible to quantitatively compare explanations for increased markups and profits in recent years to the superstar firm hypothesis that greater price sensitivity has sparked reallocation to large, productive firms. Within the model, neither story matches the data as well as a decline in laggards' patent quality, though I show that the static superstar firm experiment generates a productivity slowdown alongside rising concentration in the estimated model. To my knowledge, this is the first dynamic version of Autor et al. (2020) with endogenous productivity growth.

The finding that laggards' patent quality has declined since 2000 is consistent with Bloom et al. (2020), who show that despite increasing inputs (expenditures, workers) to R&D, outputs in terms of productivity improvements have declined using a variety of case studies. Anzoategui et al. (2019) also identify a decline in R&D productivity using indirect inference in a DSGE model with endogenous productivity growth. Several empirical papers have also documented that laggard firms are less likely to overtake market leaders in recent years (Bessen et al. (2020), Pugsley, Sedlacek, and Sterk (2020), Andrews, Criscuolo, and Gal (2016)). This paper sheds more light on the channel through which this happens: I estimate a mild *decrease* in the cost per patent to explain rising expenditures on R&D over this period, but also a large decrease in the average contribution of a new patent to the value of the firm for laggard firms.[4]

Finally, many papers have studied rising concentration and the productivity slowdown (Hall (2015), Syverson (2017)) in isolation from one another. Rising concentration is mainly a within-sector phenomenon (Hsieh and Rossi-Hansberg (2020)) that is occurring at the national/product market level rather than at the local level.[5] The finding that market concentration is rising is robust to the inclusion of foreign firms (Covarrubias, Gutierrez, and Philippon (2019)) or more sophisticated methods of identifying firms' direct competitors (Pellegrino (2020), using data from Hoberg and Phillips (2010)).

A variety of explanations for rising sales concentration have been proposed, from the introduction of ICT that creates winner-take-all markets

[4]Contemporaneous work by Cavenaile, Celik, and Tian (2020) estimates an endogenous growth model with incumbents and a competitive fringe with step by step innovations and finds that declining R&D productivity of small firms can explain a large share of the rise in concentration and the productivity slowdown. The advantage of allowing for patent quality heterogeneity and including new data on patent quality as a target for the estimation is that I can separately identify changing costs and changing output of R&D.

[5]In fact Berger, Herkenhoff, and Mongey (2019) and Rossi-Hansberg, Sarte, and Trachter (2020) find evidence that local sales concentration has fallen over this period.

and enables the growth of superstar firms (Bessen (2017), Crouzet and Eberly (2018), van Reenen (2018)), to excessive regulations that erect barriers to entry and create unnatural monopolies (Covarrubias, Gutierrez, and Philippon (2019)), to increased mergers and acquisitions activity, possibly due to weak antitrust enforcement (Grullon, Larkin, and Michaely (2019)). This paper complements these hypotheses by contributing a novel mechanism that, according to the quantitative exercise, explains around 40 percent of the rise in concentration.

1.2 Empirical Motivation

I first review aggregate trends in productivity growth and market concentration to motivate the analysis. I then show that innovativeness has declined relative to the 1990s along various metrics, particularly for laggard firms, and discuss potential causes. Finally, I show that laggard firms are less likely to catch up to the leading firm in their industry to become the sales leader now than in the past.

1.2.1 Market Concentration and Productivity Growth

To establish the main empirical motivation for the paper, Figure 1.1 plots the average market leader's share of total industry sales in Compustat and the total factor productivity growth rate. Among U.S. public companies, market concentration has risen significantly since the late 1990s.[6] The average market leader's sales share within narrowly defined 4-digit Standard Industrial Classification (SIC) industries has risen from around 40% in the 1990s to over 50% in 2017. Total factor productivity growth averaged about 1.7% between 1994 and 2003, but slowed to about 0.5% between 2004 and 2017.

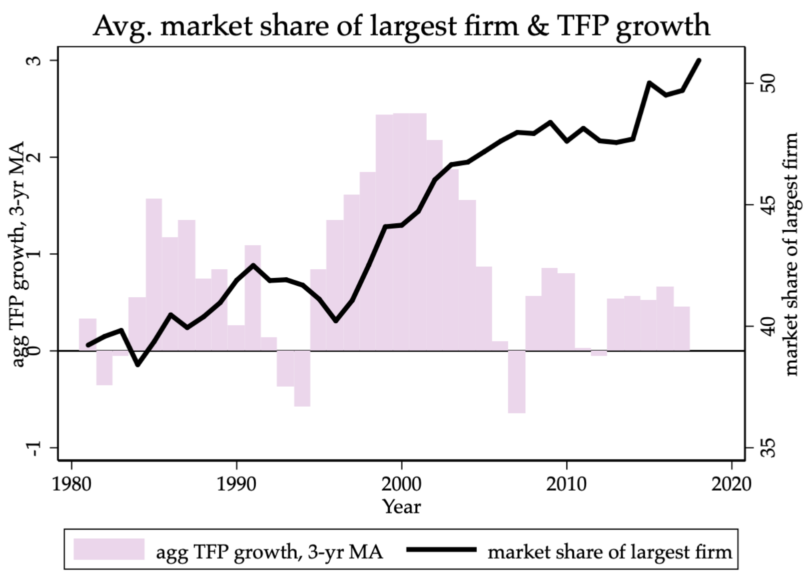

Figure 1.1: Average market share of largest firm (by sales) in 4-digit SIC industries from Compustat (weighted by industry sales); utilization-adjusted total factor productivity (TFP) growth from Fernald (2014), three year moving average.

According to the standard Olley and Pakes (1996) decomposition, aggregate total factor productivity growth could be slowing down for two

[6]See Grullon, Larkin, and Michaely (2019) and Council of Economic Advisers (2016) for overviews of trends in market concentration. More than 75% of U.S. industries have experienced an increase in the Herfindahl-Hirschman index.

reasons. First, average TFP growth across all firms could be slowing down. Second, reallocation to the most productive firms (i.e. the covariance of sales share and productivity) could be slowing down. Baqaee and Farhi (2020) show that within-firm growth has contributed very little to aggregate TFP growth since the late 1990s. Broad-based below-trend productivity growth, not increasing misallocation among U.S. firms, seems to be driving the aggregate slowdown, lending support to explanations focusing on the incentives of existing firms to improve productivity, like the hypothesis I propose here.

1.2.2 Trends in Patent Quality

Economists have long relied on patents as an observable proxy for innovativeness (Griliches (1998)). The most commonly used measure of patent quality, counting the number of forward citations a patent receives from future patents, shows substantial heterogeneity in quality in the cross section of patents, with a few patents receiving many citations and most receiving none or just a few (Akcigit and Kerr (2018)).

Recent evidence using alternative measures of patent quality also points to substantial changes in average quality over time. For example, Kelly et al. (2018) create a text-based measure of patent quality, identifying "breakthrough" patents as those patents where the patent's text differs from the text of past patents but is similar to the text of future patents. This measure has the advantage of covering a longer time series (1860-present) than citation based measures (1940-present). Using this measure, Kelly et al. (2018) find that periods with high average patent quality coincide with the discovery of new general purpose technologies, including the ICT revolution in the 1990s, consistent with Bresnahan and Trajtenberg (1995)'s theory of "innovational complementarities" between general purpose technologies and inventions in other sectors of the economy.[7] The most recent wave of high patent quality driven by ICT began to subside in the late 1990s according to this measure (see Appendix A.1.2).

To explore heterogeneity in the decline in patent quality across firms, I use a measure of patent value from Kogan et al. (2017) that estimates the market value of all patents issued in the U.S. and assigned to public firms from 1926-2010 using firms' excess stock returns in a window around patent approval dates to infer the market value of the patent.[8,9] This measure has the advantage of capturing the private value of the patent to the firm, which will determine firms' investment decisions in the model.

In the model presented in section 1.3, firms make innovations that grow the quality of their product variety by a random amount. I use the

[7] See Helpman (1998) and Aghion, Akcigit, and Howitt (2014) for reviews of the study of GPTs.

[8] I use updated data through 2017 from Noah Stoffman's webpage: https://github.com/KPSS2017/Technological-Innovation-Resource-Allocation-and-Growth-Extended-Data.

[9] Kelly et al. (2018) document the strong correlation between the market- and text-based measures at the patent level as well as the correlation of these measures with forward citation-weighted measures. All three measures show a sharp uptick in average patent quality and in the right tail during the 1990s and a subsequent decline beginning in the late 1990s.

dollar value estimates of Kogan et al. (2017) to construct a measure of each public firm's "patent stock" as the cumulative value of all past patents, intuitively corresponding to the current knowledge or quality embodied in the firm's product(s).[10,11] From 1980 to 2017 this measure covers 1,339,541 patents issued to 4,360 different U.S. public firms. With this measure of the patent stock in hand, I define *patent quality* as the marginal contribution of a new patent to the total value of the firm's existing patent stock. Figure 1.2 plots the average of this measure over time, splitting the sample into market leaders (largest firms by sales in 4-digit SIC industries) and followers (all other firms).

Figure 1.2 illustrates the two key facts for the subsequent analysis:

1. Smaller firms have higher patent quality than market leaders on average.

2. Smaller firms' patent quality rose from 1990 to 2000, but has declined significantly since 2000.

Fact 1 is related to a large debate on the relative innovativeness of large versus small firms (see Akcigit and Kerr (2018)). Typically this debate centers on small startups versus large companies with more than 500 em-

[10] Construction details in appendix A.1.1.

[11] Some depreciation can be applied to the patent stock measure. For example Peters and Taylor (2017) use the Bureau of Economic Analysis' R&D expenditure depreciation rates by sector, ranging from 5-20% per year to construct a measure of firms' intangible capital stock. Applying depreciation rates in this range increases the level of the estimated quality improvements but does not affect the magnitude of the slowdown or the differential decline between leaders and laggards.

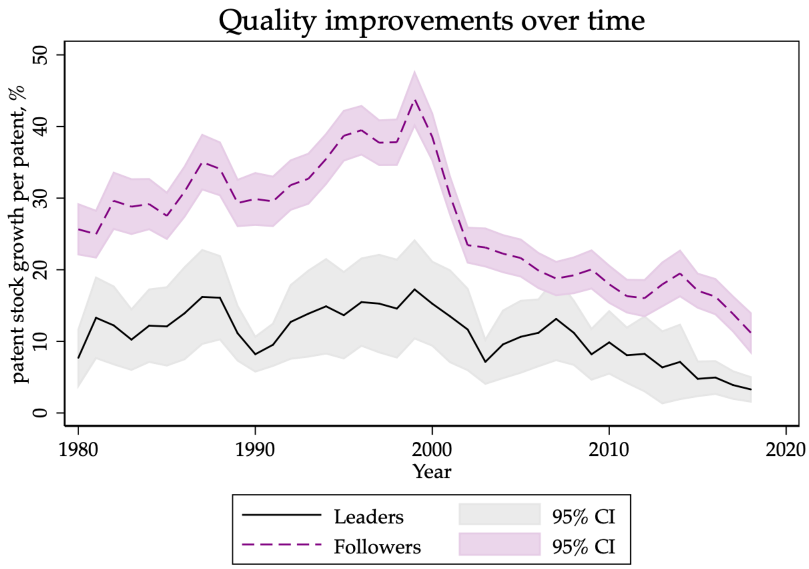

Figure 1.2: Contribution of average new patent to value of filing firm's existing stock of patents, using estimated patent values from Kogan et al. (2017). Leader indicates sales leaders in 4-digit SIC industries and followers are all other firms.

ployees (more than 72% of observations in the sample of patenting firms in Compustat have more than 500 employees). My finding is that even among firms that are large relative to the entire firm size distribution, there are differences in patent quality by size (measured by sales) within industries. Arguments supporting the greater innovativeness of smaller firms should still be relevant even among public firms: managers at smaller firms tend to be more flexible and to be closer both to customers and to researchers within the firm, enhancing their ability to allocate spending to more productive projects (Knott and Vieregger (2016)). Rosen (1991) develops a model to explain why smaller firms within sectors make a disproportionate share of major innovations. In the model, large firms find it optimal to focus on innovations that are complementary to their existing products and processes to avoid Arrow (1962)'s replacement effect. Small firms instead choose to allocate funds toward developing revolutionary technologies, having more to gain in post-innovation rents from doing so.

Fact 2 is new. The patented innovations of relatively smaller firms seem to be more incremental recently than they were in the 1990s. One candidate explanation is that, because this is a market-based measure, investors are internalizing the fact that it is harder for non-leading firms to make a dent in the advantage of their leading competitor than before, perhaps because of anti-competitive practices of market leaders or because of the rise of platform-based technologies. It's not possible to fully rule this out, but constructing a similar measure of patent quality based on forward citation counts instead of dollar values shows a similar decline beginning in 2000 (see Appendix A.1.2). If this was the case and nothing else changed we might expect to see a wedge opening up between the private and social value of laggards' patents, but in fact both declined, perhaps pointing towards technological explanations.

Another possible explanation for fact (2) is that general purpose technologies, or at least ICT, have greater complementarities with some types of firms than others (Jovanovic and Rousseau (2005) find that initial public offerings surge during GPT waves, for example). Smaller firms, with greater flexibility and more incentive to invest in riskier, disruptive ideas, may be better positioned to take advantage of the gains associated with disruptive technologies. After the general purpose technology has diffused through the economy, opportunities for disruption may lessen and laggards' improvements may become more incremental.[12] Consistent with this idea, the pattern of boom and bust in patent quality is more pronounced in high tech sectors than in manufacturing, healthcare, or consumer good sectors, though the trend is present to some extent across all four categories (see Appendix A.1.2).

The very sharp decline in laggards' patent quality between 1999 to 2001 is worth exploring. The only significant change to U.S. patent law in the late 1990s was the American Inventor's Protection Act of 1999 to publish most patent applications 18 months after filing.[13] Previously, only approved patents were published. This change might deter inventors who thought their patent was unlikely to be approved from applying for a patent for fear that their idea would be published but they would not get the exclusive rights to it. In that case one would expect the patent application approval rate to rise. In fact, according to Carley, Hegde, and Marco (2015), the approval rate declined from about 70% in 1996 to 40% in 2005. Moreover, Graham and Hegde (2015) find that firms given the option to opt out of this pre-grant disclosure (U.S. firms that did not file any foreign

[12] Aum, Lee, and Shin (2018) find that the productivity boom from computerization had normalized by 2004.

[13] https://www.uspto.gov/about-us/news-updates/uspto-will-begin-publishing-patent-applications

versions of the same patent) chose to do so less than 10% of the time.

The decline is also not likely driven by the dot-com bubble since the same pattern appears in the citation-based measure of patent quality. It also does not appear to be driven by the increasing age of public firms: this pattern appears even among firms that have been public at least 20 years when the patent is issued (Appendix A.1.2). Nor is it driven by ideas being embodied in multiple patents in recent years: the same pattern is present in the annual patent stock growth rather than the marginal contribution of each individual patent (Appendix A.1.2).

1.2.3 Declining Dynamism and Leadership Turnover

In the model presented in section 1.3, innovations drive growth in market share at the expense of the firm's competitor. Turning to this outcome, Figure 1.3 plots the fraction of U.S. industries with a new sales leader each year to measure the frequency with which smaller firms overtake the largest firm. This fraction has fallen from around 15% per year in the late 1990s to around 9% in recent years (see Bessen et al. (2020) for a detailed empirical analysis of this phenomenon in the U.S. and Andrews, Criscuolo, and Gal (2016) for a cross-country analysis).

Firm-level productivity data also shows that the "advantage of backwardness" has fallen relative to the 1990s, consistent with the idea that it is now harder to catch up through innovation than it was in the 1990s. Andrews, Criscuolo, and Gal (2016) show that in a regression of firm-level

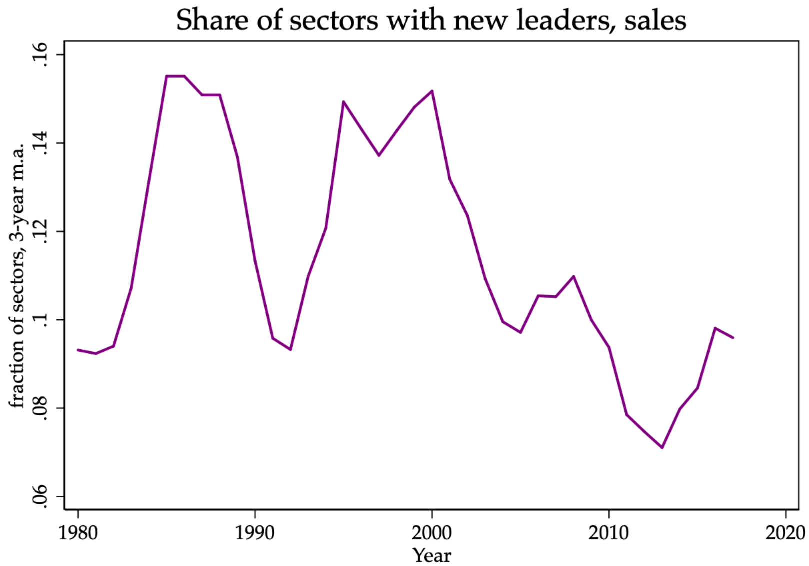

Figure 1.3: Share of 4-digit SIC sectors in Compustat with a new sales leader each year.

productivity growth on a variety of explanatory variables, the coefficient on the lagged productivity gap to the most productive competitor has been declining over the 2000s, suggesting that distance to the productivity frontier is becoming a less important predictor of future productivity growth.[14] Decker et al. (2016) also find that the right skewness of the firm-level productivity growth distribution in the U.S. has declined over this period.

[14]This empirical observation is endogenous according to the model, because it may be a result of both structural change to catchup speeds and to the endogenously lower innovation effort by laggard firms since their regression does not control for innovation effort (R&D investment).

1.3 Model

To capture the effect of declining innovativeness of laggards firms, I develop a model along the lines of Aghion et al. (2001) but building on models with heterogeneous patent quality rather than step by step innovations (Akcigit, Ates, and Impullitti (2018), Acemoglu et al. (2018), Akcigit and Kerr (2018)). Relative to Aghion et al. (2001), I also introduce a competitive fringe of firms in each sector that constrains the pricing behavior of the incumbents in order to match observed levels of concentration in the data. The model features endogenous markups and each sector's level of sales concentration evolves over time as the result of innovation. The model will be used to infer changes to the nature of the economy between the 1990s and the 2000s.

The model is of a closed economy in continuous time. There are three types agents: a representative household, a representative final good firm, and firms producing intermediate goods. This section presents the model going through the problem of each type of agent in the economy, then analyzes the equilibrium of the model.

1.3.1 Households

A representative household consumes, saves, and supplies labor inelastically to maximize:

$$U_t = \int_t^{\infty} \exp(-\rho(s-t)) \frac{C_s^{1-\psi}}{1-\psi} ds,$$

subject to:

$$r_t A_t + W_t L = P_t C_t + \dot{A}_t,$$

where ρ is the discount rate, ψ is the inverse intertemporal elasticity of substitution, C_t is consumption at time t, W_t is the nominal wage rate, and P_t is the price of the consumption good C_t. Households' labor supply L will be normalized to 1 and there is no population growth. Households own all the firms, and the total assets in the economy A_t are:

$$A_t = \int_0^1 \sum_{i=1}^{2} \left(V_{ijt} + V_{ijt}^e \right) dj,$$

where V_{ijt} is the value of an incumbent intermediate good firm i in sector j at time t and V_{ijt}^e is the value of an entrant that can displace firm i in sector j at time t. These value functions are explained in greater detail in section 1.3.3. r_t is the rate of return on the portfolio of firms. On a balanced growth path with constant growth rate of output g this yields the standard Euler equation $r = g\psi + \rho$.

1.3.2 Final Good Producers

The competitive final goods sector combines intermediate goods and labor to create the final output good which is used in consumption, research, and intermediate good production. The final good firm operates a constant return to scale technology:

$$Y_t = \frac{1}{1-\beta}\left(\int_0^1 K_{jt}^{1-\beta} dj\right) L^{\beta}, \tag{1.1}$$

where K_{jt} is a composite of two intermediate good firms' products within sector j described below. β determines both the elasticity of substitution across sectors ($\frac{1}{\beta}$) and the labor share. The final good firm's problem of hiring sector composite goods K_{jt} for $j \in [0, 1]$ and labor is:

$$\max_{K_{jt},L} P_t \frac{1}{1-\beta}\left(\int_0^1 K_{jt}^{1-\beta} dj\right) L^{\beta} - \int_0^1 P_{jt} K_{jt} dj - W_t L.$$

The first order condition for sector j's composite good given sector j's composite price index P_{jt} yields the following demand for sector j's good:

$$K_{jt} = \left(\frac{P_{jt}}{P_t}\right)^{-\frac{1}{\beta}} L,$$

and the real wage is equal to the marginal product of labor:

$$\beta \frac{Y_t}{L} = \frac{W_t}{P_t}.$$

To derive the demand curve for each intermediate good producer i within sector j we need to define the sector composite goods K_{jt} explicitly:

$$K_{jt} = \left((q_{1jt}k_{1jt})^{\frac{\epsilon-1}{\epsilon}} + (q_{2jt}k_{2jt})^{\frac{\epsilon-1}{\epsilon}}\right)^{\frac{\epsilon}{\epsilon-1}}, \tag{1.2}$$

where q_{ijt} is the quality of firm i's product at time t (equivalently as firm i's productivity) and k_{ijt} is the output of firm i purchased by the final good producer.[15] The elasticity of substitution between product varieties in the same sector is ϵ.

The first order condition for the final goods firm's problem yields the following demand curve for firm i in sector j's output:

$$k_{ijt} = q_{ijt}^{\epsilon-1} \left(\frac{p_{ijt}}{P_{jt}}\right)^{-\epsilon} \left(\frac{P_{jt}}{P_t}\right)^{-\frac{1}{\beta}} L. \tag{1.3}$$

That is, demand is increasing in the firm's quality, decreasing in its price relative to the sector j price index, and decreasing in the sector's price index relative to the price index in the economy as a whole.

1.3.3 Intermediate Goods Producers

Each intermediate good sector features competition between two large incumbent firms with differentiated products and access to an R&D technology, plus a competitive fringe that constrains the price-setting of the incumbents. Incumbents are periodically hit with exit shocks that cause them to be replaced by a new firm. This section covers the static pricing game played by intermediate good firms and their dynamic R&D invest-

[15] I use quality and productivity interchangeably because final output is homogeneous of degree one in either the qualities or the quantities of the intermediate goods firms' products.

ment decision.

Production and Price Setting

Production Intermediate goods producers purchase final goods to transform them into differentiated intermediate goods. Each unit of intermediate output requires $\eta < 1$ units of the final good to produce. There are no other inputs to intermediate good production.

Competitive fringe Each industry contains a competitive fringe of firms that is able to produce a perfect substitute to the lower quality variety at marginal cost η. I call the incumbent firm with lower quality the *follower*, or laggard, and the incumbent firm with higher quality the *leader*. When $q_{1jt} = q_{2jt}$, the fringe can produce perfect substitutes to both incumbents' varieties. One way to micro-found this assumption is by introducing a cost to filing and maintaining a patent that is sufficiently high that only the leader, who exercises some additional market power by possessing the higher quality and thus earns higher profits in duopoly competition without the fringe, would be willing to pay. The follower then allows its patent to expire and faces imitation by the fringe. Intuitively, this means that sectors in the model feature a high quality variety like a brand name product and competition among other firms to produce a generic version of that sector's product. The competitive fringe firms do not have access to an innovation technology.

This assumption of the presence of a competitive fringe is not necessary to solve the model, but makes it possible to match the average level of sales concentration across sectors in the data and generates plausible predictions for profit shares as a function of market shares (see Appendix A.1.4.) I solve a version of the model without the competitive fringe in Appendix A.2.5 and replicate the main exercise in this setting. The main results in section 1.4 are qualitatively unchanged.

Price setting Firms set prices a la Bertrand at each instant t. The presence of the competitive fringe implies the follower must set its price $p_{ijt} = \eta$.[16] Understanding this, the leader chooses its price as a best response to the price set by the follower.

Dropping the subscript t, the pricing problem of technology leader i in sector j is:

$$\max_{p_{ij}} p_{ij} k_{ij} - \eta k_{ij},$$

subject to the demand:

$$k_{ij} = q_{ij}^{\epsilon-1} \left(\frac{p_{ij}}{P_j} \right)^{-\epsilon} \left(\frac{P_j}{P} \right)^{-\frac{1}{\beta}} L,$$

where

$$P_j = \left(\sum_{i=1}^{2} q_{ij}^{\epsilon-1} p_{ij}^{1-\epsilon} \right)^{\frac{1}{1-\epsilon}}$$

[16] I resolve the indeterminacy of which firm(s) produces the lower quality variety in equilibrium by having the incumbent capture all sales of the lower quality variety so the fringe is not active in equilibrium.

is sector j's price index.

Let $s_{ij} = \frac{p_{ij}k_{ij}}{\sum_{i=1}^{2} p_{ij}k_{ij}}$ denote firm i's market share in sector j. Then the optimal pricing policy for the market leader is:

$$p_{ij} = \frac{\epsilon - (\epsilon - \frac{1}{\beta})s_{ij}}{\epsilon - (\epsilon - \frac{1}{\beta})s_{ij} - 1}\eta. \tag{1.4}$$

The optimal price is the standard one for two-layered constant elasticity of demand structures (nested CES): a variable markup that rises in market share. This is easiest to see for the two extreme cases where market share is 0 or 1. When market share is 0, the firm is atomistic with respect to the sector and charges a markup $\frac{\epsilon}{\epsilon-1}$, the CES solution for an elasticity of substitution equal to ϵ. On the other hand, if the market share is 1, the firm only weighs the elasticity of substitution across sectors and sets a markup $\frac{1}{1-\beta} > \frac{\epsilon}{\epsilon-1}$ since products are less substitutable across sectors than within sectors.

Innovation

Incumbent intermediate goods producers have access to a research and development technology that allows them to choose an amount of research spending R_{ijt} of the final good to maximize the discounted sum of expected future profits. The decision to model R&D as a process of own-product quality improvement by incumbents is consistent with the evidence in Garcia-Macia, Hsieh, and Klenow (2019) that: (i) incumbents are responsible for most employment growth in the U.S., and this share has increased in recent years; (ii) growth mainly occurs through quality improvements rather than new varieties; (iii) creative destruction by entrants and incumbents over other firms' varieties accounted for less than 25% of employment growth from 2003-2013, consistent with earlier evidence from Bartelsman and Doms (2000).

Innovations arrive randomly at Poisson rate x_{ijt} which depends on research spending according to the function:

$$x_{ijt} = \left(\frac{\gamma R_{ijt}}{\alpha}\right)^{\frac{1}{\gamma}} q_{ijt}^{\frac{1-\frac{1}{\beta}}{\gamma}}.$$

That is, since $\beta < 1$, at higher quality levels more research spending is needed to achieve the same arrival rate of innovations x. γ and α are R&D technology parameters.

Innovations improve the quality of the incumbent firm's variety.[17] Conditional on innovating the size of the quality improvement is random. Formally, conditional on innovating,

$$q_{ij(t+\Delta t)} = \lambda^{n_{ijt}} q_{ijt},$$

where $\lambda > 1$ is some minimum quality improvement and $n_{ijt} \in \mathbb{N}$ is a random variable. Note that each competitor improves over their own quality

[17]See Griliches (2001) for a survey of the relationship between R&D and productivity at the firm level and Zachariadis (2003) for a leading empirical test.

when they innovate, rather than over the quality frontier.[18] Initial qualities of all firms at $t = 0$ are normalized to 1. Let $N_{ijt} = \int_0^t n_{ijs} ds$ denote the total number of λ step improvements over a product line i since the beginning of time. The *technology gap* m_{ijt} from firm 1 in sector j's perspective at time t is defined as:

$$\frac{q_{1jt}}{q_{2jt}} = \frac{\lambda^{N_{1jt}}}{\lambda^{N_{2jt}}} \equiv \lambda^{m_{1jt}}.$$

Given λ, m_{ijt} parameterizes the relative qualities of the two firms within sector j from firm $i \in \{1, 2\}$'s perspective, representing the number of λ steps ahead or behind its competitor firm i is. m_{ijt} turns out to be the only payoff relevant state variable for the incumbent firms. For tractability I will impose a maximal technology gap $\bar{m}$, but in calibrating the model I will set the parameters so that this maximal gap rarely occurs in steady state. I assume that the only knowledge spillover between incumbents in the model occurs when a firm at the maximal gap innovates. In that case, both the innovating firm and its competitor's quality increase by the factor λ, keeping the technology gap unchanged but raising the absolute quality of the sector composite good.

The probability distribution of possible quality improvements depends on the firm's current technology gap, consistent with the evidence in section 1.2 that patent quality varies between market leaders and laggards. It is useful to instead imagine firms draw a new position in technology gap

[18] Luttmer (2007) provides an additional rationale for this type of assumption: entrants are typically small and enter far from the productivity frontier, implying that imitation of other firms' technologies is difficult.

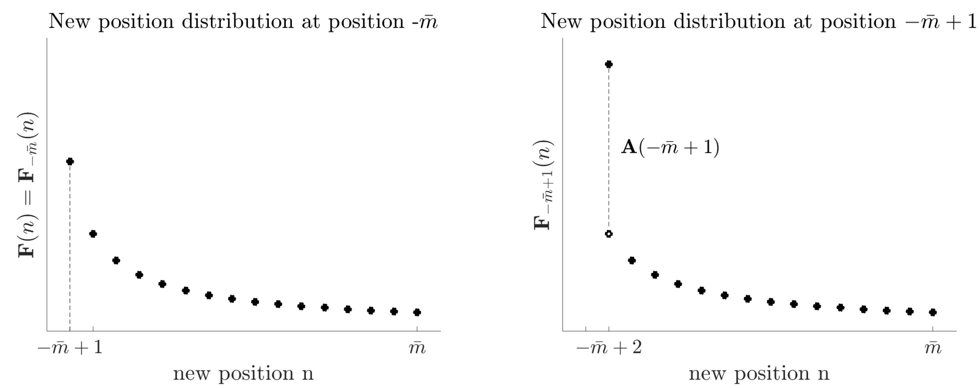

Figure 1.4: Examples of new position distributions for positions $-\bar{m}$ and $-\bar{m}+1$.

space $n \in \{-\bar{m}, \ldots, \bar{m}\}$ when they innovate, rather than an absolute number of λ steps, though given n and m the number of steps can be easily derived as $n - m$. As in Akcigit, Ates, and Impullitti (2018), I assume there exists a fixed distribution $\mathbb{F}(n) \equiv c_0(n + \bar{m})^{-\phi}$ for all $n \in \{-\bar{m}+1, \ldots, \bar{m}\}$ that applies to firms that are the furthest possible distance behind their competitor and describes the probability that they move to each position in technology gap space. An example is shown in the left panel of 1.4. The curvature parameter ϕ is critical in the model and determines the speed of catchup by increasing or decreasing the relative probability of larger innovations. A higher ϕ means a lower probability of these "radical" improvements.[19,20] c_0 is simply a shifter to ensure $\sum_n \mathbb{F}(n) = 1$.

Given this fixed distribution for the most laggard firm, the new posi-

[19] As noted by Akcigit, Ates, and Impullitti (2018), this formulation converges to the less general step-by-step model as $\phi \to \infty$.

[20] The use of "radical innovation" in this paper to describe a relatively large quality improvement differs from some other papers in the literature such as Acemoglu and Cao (2015) who use "radical innovation" to refer to an entrant replacing an incumbent.

tion distribution specific to each technology gap $m > -\bar{m}$ is given by:

$$\mathbb{F}_m(n) = \begin{cases} \mathbb{F}(m+1) + \mathbb{A}(m) & \text{for } n = m+1 \\ \mathbb{F}(s) & \text{for } n \in \{m+2, \ldots, \bar{m}\} \end{cases},$$

where $\mathbb{A}(m) \equiv \sum_{-\bar{m}+1}^{m} \mathbb{F}(n)$. This distribution is shown in the right panel of Figure 1.4 for a firm at gap $-\bar{m}+1$. Simply put, all the mass of the fixed distribution on positions lower than the current position m is put on one-step ahead improvements. This formulation can capture the feature that laggard firms make larger improvements than leaders on average.

Entry and Exit

Incumbent firms face a constant exit risk δ_e. If an incumbent is hit with this shock the incumbent is replaced by an entrant that takes over the product line with the same quality level (and thus technology gap to the other incumbent in the sector) as the incumbent it replaces. This shock captures many reasons why incumbents may exit or be displaced by entrants that are not directly related to the incumbent firms' innovations such as adverse financial shocks, negative taste shocks for the incumbent's brand, expiration of the incumbent's patent or knowledge diffusion as in Akcigit and Ates (2020), or cost shocks to specific inputs used by the incumbent.

Intermediate Goods Firms Value Functions

Turning to the firm value functions, I will show that the technology gap $m \in \{-\bar{m}, \dots, \bar{m}\}$ is sufficient to describe the firms' pricing and innovation strategies, and that firm values scale in some function of their current product quality q_{ijt}.

The proof that pricing decisions and market shares depend only on m_{ijt} (and not on the level of quality q_{ijt}) is in Appendix A.2.1 which shows that we can define $p(m)$ as the price set by a firm at gap m. Next consider the flow profits of an incumbent, denoting the optimal price of the leader at technology gap m as $p(m)$, and dropping subscripts t and j for now:

$$\pi(m, q_i) = \begin{cases} 0 & \text{if } m \leq 0 \\ (p(m) - \eta)k_i & \text{for } m \in \{1, \dots, \bar{m}\} \end{cases}.$$

Plugging in equation 1.3 for k_i and using the definition of the sector price index yields $\pi(m, q_i) = 0$ if $m \leq 0$ and
$\pi(m, q_i) = q_i^{\frac{1}{\beta}-1}(p(m) - \eta)p(m)^{-\epsilon}(p(m)^{1-\epsilon} + (\lambda^{-m})^{\epsilon-1}p(-m)^{1-\epsilon})^{\frac{\epsilon-\frac{1}{\beta}}{1-\epsilon}}$ for $m \in \{1, \dots, \bar{m}\}$.

For the dynamic problem, I will use a guess and verify method to verify that firms' strategies depend only on m and that firm values scale in some function of q_{ijt}. Dropping the subscript ij and given an interest rate r_t, the value function of a firm with technology gap m_t to its competitor and quality level q_t can be written:

$$\begin{aligned}
r_t V_{mt}(q_t) - \dot{V}_{mt}(q_t) = \max_{x_{mt}} \{ & \pi(m, q_t) - \alpha \frac{(x_{mt})^\gamma}{\gamma} q_t^{\frac{1}{\beta}-1} \\
& + x_{mt} \sum_{n_t=m+1}^{\bar{m}} \mathbb{F}_m(n_t)[V_{nt}(\lambda^{n_t-m} q_t) - V_{mt}(q_t)] \\
& + x_{(-m)t} \sum_{n_t=-m+1}^{\bar{m}} \mathbb{F}_{-m}(n_t)[V_{(-n)t}(q_t) - V_{mt}(q_t)] \\
& + \delta_e(0 - V_{mt}(q_t)\}. \qquad (1.5)
\end{aligned}$$

The firm chooses the arrival rate of innovations x_{mt}. The first line denotes the flow profits and the research cost R_{ijt} given the choice of x_{mt}. The second line denotes the probability the firm innovates and sums over the possible states the firm could move to using the new position distribution and the firm's new value function with higher quality and a larger quality advantage over its rival. The third line denotes the chance the firm's rival innovates and the change in the firm's value because its relative quality falls when the rival innovates. The final line denotes the chance the entrant displaces the incumbent. The slightly altered equations for firms at the minimum and maximum gaps because of knowledge spillovers are given in Appendix A.2.2.

A guess and verify approach verifies that $V_{mt}(q_t) = v_{mt} q_t^{\frac{1}{\beta}-1}$. Thus one can focus on a Markov Perfect equilibrium where firms' strategies depend only on the payoff-relevant state variable m, which characterizes the technology gap between incumbents.

The firm's optimal innovation rate x_{mt} is the solution to the first order condition of equation (1.5), which gives:

$$x_{mt} = \begin{cases} \left(\frac{\sum_{n=m+1}^{\bar{m}} \mathbb{F}_m(n_t)[(\lambda^{n_t - m})^{\frac{1}{\beta}-1} v_{nt} - v_{mt}]}{\alpha} \right)^{\frac{1}{\gamma-1}} & \text{for } m < \bar{m} \\ \left[\frac{1}{\alpha}(\lambda^{\frac{1}{\beta}-1} - 1) v_{\bar{m}t} \right]^{\frac{1}{\gamma-1}} & \text{for } m = \bar{m} \end{cases}.$$

Intuitively, firms choose a higher arrival rate of innovations when the cost of R&D α is low, and when the expected gain from innovating is high, captured by the probability of moving to different positions in technology gap space upon innovating $\mathbb{F}_m(n)$, the value v_n of being at gap n, and the minimum size of quality improvements λ. All else equal, greater expected innovativeness of laggards (more weight on states where they catch up to or overtake the leader), should encourage more innovation by laggard firms. However, the v_n terms also capture the probability of being displaced in the future, so these values are endogenously determined along with the chance of displacement by rivals due to innovation or the chance of being hit with an exit shock δ_e. At t, the value of a potential entrant in product line i in sector j is simply $V^e_{ijt} = \delta_e V_{ijt}$.

1.3.4 Equilibrium Output

Plugging in the intermediate goods firms' pricing decisions yields the following expression for final output Y_t, derived in Appendix A.2.3:

$$Y_t = \frac{1}{2}\frac{L}{1-\beta}P^{\frac{1-\beta}{\beta}}\sum_{m=-\bar{m}}^{\bar{m}} Q_{mt}, \tag{1.6}$$

where Q_{mt} is defined as:

$$\begin{aligned} Q_{m,t} &= \int_0^1 \left(q_{it}^{\epsilon-1}p(m)^{1-\epsilon} + q_{-it}^{\epsilon-1}p(-m)^{1-\epsilon}\right)^{-\frac{(1-\beta)}{\beta(1-\epsilon)}} \mathbb{1}_{\{i\in\mu_{mt}\}}di \\ &= \left(p(m)^{1-\epsilon} + (\lambda^{-m})^{\epsilon-1}p(-m)^{1-\epsilon}\right)^{\frac{1-\beta}{\beta(\epsilon-1)}} \int_0^1 q_{i,t}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i\in\mu_{mt}\}}di. \end{aligned} \tag{1.7}$$

Here, μ_{mt} is the measure of firms at each technology gap m at time t (normalizing measure of firms to one) and Q_{mt} is a particular index of the qualities of all firms at gap m. The change in output between t and $t+dt$ will therefore depend on the changes $\dot{Q}_{mt}$ for each technology gap m which in turn depend on the innovation arrival rates x_{mt} chosen by firms and the exogenous distribution of quality improvement sizes $\mathbb{F}(n)$. The term $(p(m)^{1-\epsilon}+(\lambda^{-m})^{\epsilon-1}p(-m)^{1-\epsilon})^{\frac{1-\beta}{\beta(\epsilon-1)}}$ weights the change in qualities of firms at gap m depending on the prices set by firms at gap m and $-m$, capturing static distortions from firms' markups. Note that entry and exit are not a source of growth in the model because they have no impact on the qualities of the intermediate goods in the economy or on markups. The final component determining output will be the measure of firms at each technology gap μ_{mt} that is itself an endogenous object. The next section describes how to solve for the measures μ_{mt}.

1.3.5 Distribution Over Technology Gaps

Firms move to technology gap n through innovation from a lower technology gap, or because their competitor innovates to gap $-n$. The distributions $\mathbb{F}_m(n)$ and $\mathbb{F}_{-m}(-n)$ respectively determine these probabilities, combined with the innovation efforts of firms at m and $-m$, for all $m < n$ and $-m < -n$. The outflows from gap n are due to the firm at n or its competitor at $-n$ innovating. Putting this together into the Kolmogorov forward equations for the evolution of the mass of firms at each gap:

$$\dot{\mu}_{nt} = \sum_{m=-\bar{m}}^{n-1} x_m \mathbb{F}_m(n)\mu_{mt} + \sum_{m=n+1}^{\bar{m}} x_{-m}\mathbb{F}_{-m}(-n)\mu_{mt} - (x_n + x_{-n})\mu_{nt}. \quad (1.8)$$

The highest and lowest technology gaps are special cases because of spillovers: if the firm at the highest gap innovates both firms remain at the same gap in the next instant:

$$\dot{\mu}_{-\bar{m}t} = \sum_{m=-\bar{m}+1}^{\bar{m}} x_{-m}\mathbb{F}_{-m}(\bar{m})\mu_{mt} - x_{-\bar{m}}\mu_{-\bar{m}t} \quad (1.9)$$

$$\dot{\mu}_{\bar{m}t} = \sum_{m=-\bar{m}}^{\bar{m}-1} x_m \mathbb{F}_m(\bar{m})\mu_{mt} - x_{-\bar{m}}\mu_{\bar{m}t}. \quad (1.10)$$

On a balanced growth path, $\mu_{mt} = \mu_m$ for all m, t. Replacing the left hand side of the above equations with zero change in equilibrium and the measures on the right hand side with the constants μ_n, μ_m defines a system of $2\bar{m}+1$ equations in $2\bar{m}+1$ unknowns that determine the steady state distribution of firms over possible technology gaps. There are several additional restrictions on the solution to this system. First, for each firm at m there is a firm at $-m$ (that is, the stationary distribution is symmetric). Second, I impose the restriction that the measure of all incumbent firms sums to one.

1.3.6 Output Growth

Differentiating equation 1.6 with respect to time yields the following expression for the growth rate:

$$\frac{\dot{Y}_t}{Y_t} = g_{Yt} = \frac{1}{2}\frac{1}{1-\beta}\sum_{m=-\bar{m}}^{\bar{m}}\frac{\dot{Q}_{mt}}{Y_t}.$$

It's useful to define:

$$\tilde{Q}_{mt} = \int_0^1 q_{m,t,i}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i\in\mu_{mt}\}} di \tag{1.11}$$

So that:

$$g_{Yt} = \frac{1}{2}\frac{1}{1-\beta}\sum_{m=-\bar{m}}^{\bar{m}}(p(m)^{1-\epsilon} + (\lambda^{-m})^{\epsilon-1}p(-m)^{1-\epsilon})^{\frac{1-\beta}{\beta(\epsilon-1)}}\frac{\dot{\tilde{Q}}_{mt}}{Y_t}.$$

The subsequent analysis focuses on a balanced growth path where $\frac{\dot{\tilde{Q}}_{mt}}{Y_t}$

is constant for all m. On this balanced growth path consumption and output grow at a constant growth rate g and the mass of firms at each technology gap μ_m is constant. In general it is not possible to solve for this growth rate in closed form, but for a given set of model parameters it is possible to check the existence and uniqueness of such a balanced growth path and find the value of g as the solution to a system of equations. A more detailed derivation of these results is provided in Appendix A.2.4.

1.3.7 Equilibrium Definition

Let $R_t = \int_0^1 \sum_{i=1}^2 R_{ijt} dj$ denote total research and development spending by incumbents, C_t total consumption, and $K_t = \int_0^1 \sum_{i=1}^2 \eta k_{ijt} dj$ total purchases of final goods for production of intermediate goods.

A Markov-Perfect equilibrium is an allocation
$\{k_{ijt}, K_t, x_{ijt}, R_t, Y_t, C_t, L, \mu_{mt}, Q_{mt}, A_t\}_{i\in\{1,2\},j\in[0,1],m\in[-\bar{m},\bar{m}]}^{t\in(0,\infty)}$
and prices $\{r_t, W_t, p_{ijt}\}_{i\in\{1,2\},j\in[0,1]}^{t\in(0,\infty)}$ such that for all t:

1. Households choose C_t and A_t to solve the problem described in section 1.3.1.

2. Final goods firms solve their problem to hire labor L and purchase intermediate goods k_{ijt} optimally according to the problem in section 1.3.2.

3. Intermediate good firms choose p_{ijt} and x_{ijt} to solve their innovation and price-setting problems described in section 1.3.3.

4. The final goods market clears: $Y_t = C_t + R_t + K_t$.

5. The asset market clears, pinning down r_t via the household's Euler equation.

6. Labor market clears, pinning down the wage rate from the final good producer's problem.

7. μ_{mt} and Q_{mt} are consistent with firms' choices of x_{ijt}.

This completes the description of the model. The next section develops further intuition about the model under reasonable model parameter values for the U.S. economy.

1.4 Model Estimation

The quantitative analysis precedes as follows. I estimate an initial steady state for the model by matching various moments for the U.S. economy in the period of high patent quality between 1994 and 2003 ("1990s") using data on U.S. public firms from Compustat as well as aggregate moments. Using this initial calibration I describe firms' pricing and innovation strategies to develop intuition about the model. I then re-estimate the model parameters for 2004-2017 ("2000s") in order to infer changes to the economy between these two periods.

1.4.1 Baseline Calibration for the 1990s

Four parameters are calibrated outside the model. The inverse intertemporal elasticity of substitution ψ is set to 1. The labor share, β, is set to 0.6. This implies an elasticity of substitution across sectors of $\frac{1}{\beta} = \frac{5}{3}$, within the range of upper-level elasticities of substitution estimated in Hobjin and Nechio (2019). The curvature of the R&D cost function, γ, is calibrated outside to match the empirical evidence on the elasticity of patenting to R&D expenditures, discussed in Acemoglu et al. (2018). The maximum technology gap, $\bar{m}$, is set to 16.

The rest of the parameters for the baseline model, shown in Table 1.1, are estimated using a simulated method of moments approach described in appendix A.3.2 to match targets for the 1990s equilibrium (1994-2003). These targets are given in Table 1.2. The data sources and computation methods for the data moments are given in appendix A.1.1. Appendix A.3.1 describes the solution method for finding the model steady state.

The moments include the main phenomena of interest: aggregate productivity growth, average market leader's share of industry sales, the profit share of total output, average patent quality, and the rate of leadership turnover from either entry or being overtaken by an incumbent rival. In addition to average patent quality across all firms, I include the average patent quality of market leaders to help identify λ and ϕ separately. The other two moments, R&D as a share of output and R&D as a share of sales at the firm level, are included to help discipline the R&D cost parameter

Parameter	Value	Meaning/source
ψ	1	Inverse intertemporal elasticity of substitution
ρ	**0.026**	Rate of time preference (annual)
β	0.6	Labor share/Nechio & Hobijn (2017)
ϵ	**4.21**	Elasticity of substitution within sectors
η	**0.64**	Marginal cost of intermediate producers
δ_e	**0.089**	Exogenous entry/exit rate (annual)
λ	**1.059**	Min. qual. improvement
γ	2	Curvature of R&D function
α	**4.18**	R&D cost parameter
$\bar{m}$	16	Maximum number of steps ahead
ϕ	**0.88**	Curvature of patent quality distribution

Table 1.1: Model parameters (estimated parameters in bold), 1990s.

and the discount rate.

The model performs relatively well in fitting the data, particularly for productivity growth and concentration. Intuitively, the minimum step size λ and ϕ govern the average patent quality, with λ acting as a level shift in patent quality for all types of firms and ϕ shifting the probability that laggards make drastic or incremental improvements, holding patent quality of leaders fixed. The R&D cost parameter α influences the amount all firms spend on R&D and helps match aggregate expenditures as a share of output and R&D as a share of firms' sales. The entry/exit shock δ_e helps match leadership turnover. One problem with the model fit is for the R&D as a share of sales at the firm level. This can be attributed to the fact that productivity growth is purely due to R&D in the model, whereas in the reality productivity may improve for other reasons, such as management practices or improved human capital.

The estimated parameters are reasonable: a discount rate ρ of 2.6% annually implies a real interest rate in the model of 4.4%. An elasticity of substitution ϵ of 4.21 results in an average markup of 1.24, in line with the evidence summarized in Mongey (2017), particularly de Loecker, Eeckhout, and Unger (2020), that suggests markups for U.S. public firms in the 1990s ranged from 1.2 to 1.3. The entry/exit rate of about 9% per year is in line with entry and exit rates for the U.S. reported by Decker et al. (2016). The model also matches non-targeted heterogeneity in R&D intensity (R&D as a share of sales) well, as shown in Table 1.3. Section 1.4.4 describes the model's fit for additional non-targeted moments.

Targeted moments, 1994-2003	Data	Model
Avg. TFP growth, %	1.74	1.75
Avg. leader market share, %	43.34	44.62
R&D share of GDP, %	1.8	1.91
Profit share of GDP, %	5.24	6.02
Avg. R&D/sales, %	2.56	5.18
Avg. patent stock growth per patent, %	23.52	22.26
Avg. patent stock growth per patent, leaders, %	9.08	10.52
Avg. leadership turnover, %	13.74	13.26

Table 1.2: Model fit for targeted moments from estimation of 7 parameters for 1990s.

Non-targeted moments, 1994-2003	Data	Model
Avg. R&D/sales, followers, %	4.80	6.98
Avg. R&D/sales, leaders, %	1.66	2.51

Table 1.3: Model fit for R&D heterogeneity, 1990s equilibrium.

1.4.2 Properties of the Baseline Model

Before turning to how the estimated parameters change when targeting the same moments for the 2000s, I use the 1990s steady state to develop intuition about the quantitative model, particularly the incumbent firms' strategies.

Using the parameters from Table 1.1, the market shares and prices for market leaders as a function of the leader's technology gap are plotted in Figure 1.5. The leader's optimal price $p(m)$ rises as the technology gap widens (that is, as the leader's relatively quality improves). The leader's market share rises from around 30% of sales when the leader is one step ahead (that is, when the leader's quality is 5.9% higher than the laggard's) to 80% of the market at the maximum 16 steps ahead. The follower, which must set price equal to marginal cost because of the presence of the competitive fringe, has a large market share due to its relatively low price, and its market share is increasing in its relative quality.

The competitive fringe assumption also plays an important role in determining the shape of the innovation policy as a function of technology differences depicted in Figure 1.6a, specifically the hump shape. This shape has been suggested theoretically in the work of Harris and Vickers (1987), Aghion et al. (2001), and Akcigit, Ates, and Impullitti (2018), and found in a variety of studies including Aghion et al. (2005), Aghion et al. (2014), Aghion et al. (2019b), and Zhang (2018). The hump shape appears in this model because the competitive fringe assumption means that

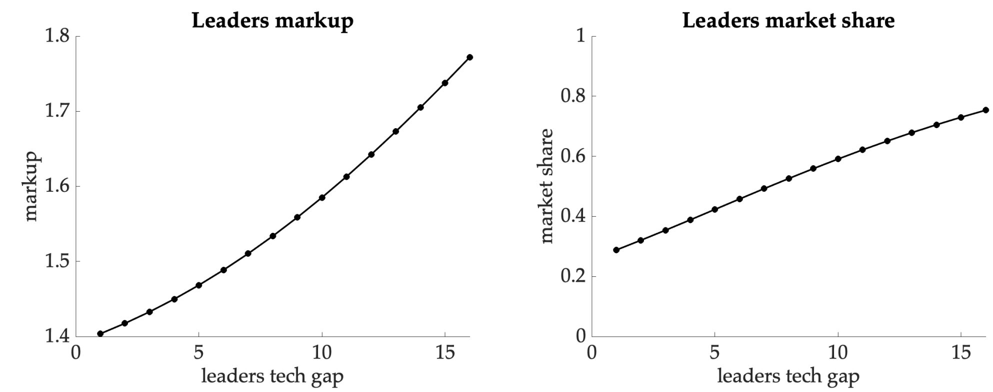

Figure 1.5: Markups and resulting market shares as a function of the leader's technology gap m.

the greatest incremental gain in flow profits comes from obtaining quality leadership (and thus escaping competition with the fringe), so the arrival rate of innovations will be highest when firms have equal quality.[21]

Figure 1.6: Baseline Innovation Policies and Stationary Distribution

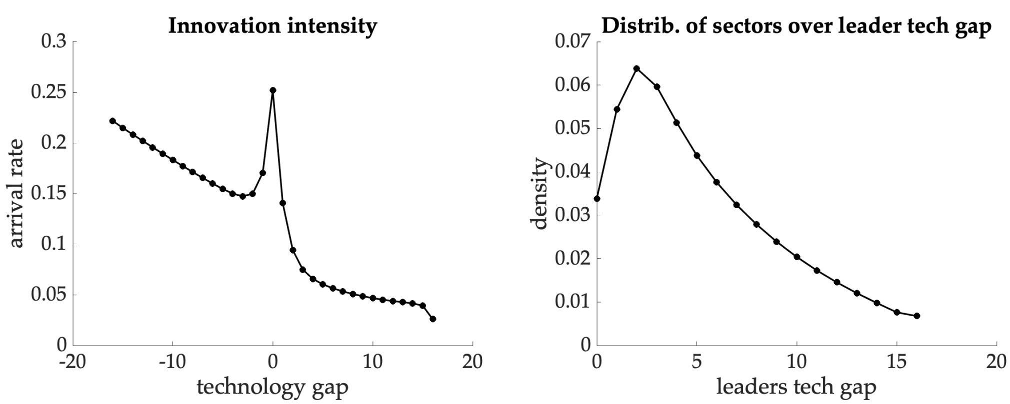

(a) Innovation policies (b) Stationary distribution

Finally, Figure 1.6b shows the stationary distribution of sectors over

[21] See Appendix A.2.5 for the version without the competitive fringe. The mechanism and main results are qualitatively unchanged.

the market leader's technology gap to its rival. Because followers innovate more frequently than leaders and have a high chance of catching their competitor when they do, there is a high rate of turnover in market leadership and technology gaps do not grow very large on average. Most sectors feature a leader that is just a few steps ahead of its rival, but there is a right tail of sectors with a large and dominant "superstar" that has a much higher quality product than its rival and captures a large share of industry sales.

1.4.3 Re-estimation for the 2000s

Re-estimating the model for the 2000s uses the model to infer the role of different channels suggested in the literature to explain changes in concentration and productivity growth and compare the strength of these other channels to the strength of declining laggard patent quality to explain these trends.

Changing the discount rate ρ captures the interest rate channel proposed by Liu, Mian, and Sufi (2019).[22] A decrease in entry and exit shocks δ_e can capture declining knowledge diffusion from incumbent firms to new entrants or rising entry costs (Akcigit and Ates (2020), Corhay, Kung, and Schmid (2020)). An increase in the research cost parameter α implies that more R&D spending is needed to achieve the same arrival rate of in-

[22] However, the model economy is not close to the very low interest rate environment discussed in Liu, Mian, and Sufi (2019) where strategic effects can dominate. See Goldberg, Lopez-Salido, and Chikis (2020) for further discussion.

novations, capturing the cost side of the hypothesis of Bloom et al. (2020) that ideas are getting harder to find. A decrease in the elasticity of substitution within sectors ϵ captures increased market power over the leader's variety, in line with Jones and Philippon (2016). On the other hand, an increase in ϵ captures the superstar firm hypothesis of Autor et al. (2020) that competitive pressures within industries have risen, causing the most productive firms to capture a larger share of total industry sales.[23] Finally, changes in ϕ govern the expected patent quality for different types of firms by changing the distributions $\mathbb{F}_m(n)$. Changing ϕ represents the research *output* side of Bloom et al. (2020)'s hypothesis, capturing the possibility that the quality of new ideas, particularly for laggard firms, is falling.

Table 1.4 shows the targeted moments and model fit for the 2000s estimation. Productivity growth slowed substantially compared to 1994-2003, while the average market leader's sales share grew by about 4 percentage points. Both aggregate and firm level research and development expenditures grew, as noted by Bloom et al. (2020). As discussed in detail in section 1.2, patent quality and leadership turnover declined. The estimation has some trouble matching the decline in the growth rate alongside an increase in R&D expenditure, but otherwise performs well.

Table 1.5 compares the estimated parameters to fit the two steady states. The households' discount rate declines slightly in the 2000s. Consistent

[23] Autor et al. (2017) speculate that such pressures may have risen because of increasing competition from foreign firms or greater price sensitivity due to better search technology such as online retail.

Targeted moments, 2004-2017	Data	Model
Avg. TFP growth, %	0.49	0.74
Avg. leader market share, %	48.12	48.89
R&D share of GDP, %	1.89	1.32
Profit share of GDP, %	6.61	6.71
Avg. R&D/sales, %	3.8	3.54
Avg. patent stock growth per patent, %	11.71	11.88
Avg. patent stock growth per patent, leaders, %	5.38	7.3
Avg. leadership turnover, %	9.27	9.37

Table 1.4: Targeted moments from estimation of 7 parameters for 2000s.

with Decker et al. (2016), the entry rate of new firms declines (alternately, incumbents are less likely to be displaced, consistent with the hypothesis of Akcigit and Ates (2019) that the rate of knowledge diffusion is slowing down). To match the fact that R&D expenditures as a share of GDP rose between the 1990s and the 2000s, the cost α of performing R&D declines. However, the expected *output* of R&D (patent stock growth per patent) conditional on innovating declines substantially due to the decrease in the probability of radical innovations, driven by the substantial increase in ϕ. The elasticity of substitution within sectors ϵ rises slightly. The marginal cost of the intermediate goods firms rises modestly. I explore these results in more detail in Section 1.5.

1.4.4 Model Validation

The model performs well in matching not just the average level of concentration but the entire distribution of market leaders' market shares across

Parameter	1990s	2000s	Meaning/source
ρ	**0.026**	**0.025**	Rate of time preference (annual)
ϵ	**4.21**	**4.32**	Elasticity of substitution within sectors
η	**0.64**	**0.70**	Marginal cost of intermediate producers
δ_e	**0.089**	**0.081**	Exog. entry/exit rate
λ	**1.059**	**1.063**	Min. qual. improvement
α	**4.18**	**3.33**	R&D cost parameter
ϕ	**0.88**	**1.52**	Curvature of patent quality distribution

Table 1.5: Comparison of estimated parameters, 1990s vs. 2000s model equilibria.

sectors of the economy in both periods. Figure 1.7 compares the empirical distribution of leader market shares in the two study periods (1994-2003 and 2004-2017) in the data and in the model. This shift is mainly due to increased average quality differences between leaders and followers in steady state, consistent with the findings of Andrews, Criscuolo, and Gal (2015) and Andrews, Criscuolo, and Gal (2016) that productivity differences within industries have grown over this period. They also find that this divergence is particularly pronounced in ICT intensive sectors, and that sectors with wider productivity gaps have experienced deeper productivity slowdowns. In Figure A.7 of Appendix A.1.5 I also show that rising market concentration and the productivity slowdown are correlated at the sector level.

The model also predicts that the average leader's share of total industry R&D expenditures rises from 20% to 53%, I find that among Compustat firms, market leaders now perform a larger share of total R&D expendi-

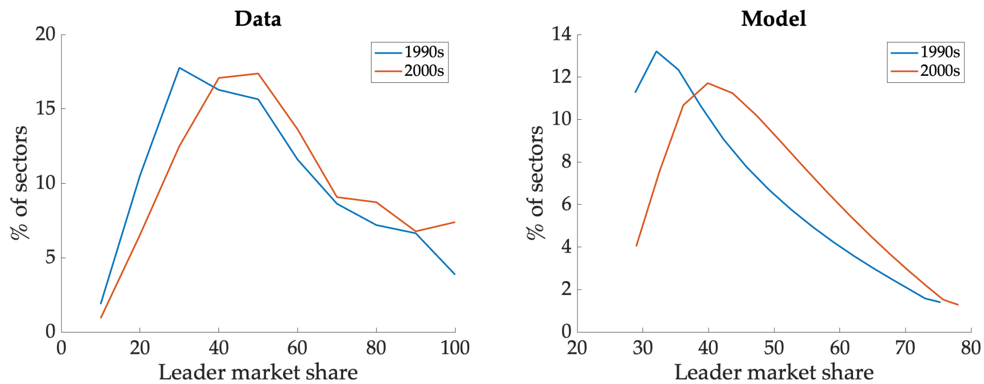

Figure 1.7: Shift in market concentration across sectors, data vs. model. Leader market share in the data is for 4-digit SIC industries in Compustat.

tures in their industries than during the 1990s (Figure A.8 in Appendix A.1.5), though the increase in the data is less dramatic than in the model, from around 38% in 1999 to a peak of 50% in 2010. Anderson and Kindlon (2019) also find a decline in R&D intensity among companies with fewer than 250 employees and an increase among larger firms in the National Science Foundation's Business R&D and Innovation Survey covering both public and private firms over this period. Akcigit and Ates (2019) also document increasing concentration of patents among the top 1% of patenting firms and increasing flows of R&D employees from small to large firms.

1.5 Results

This section obtains the main results of the paper by decomposing the role of different parameters changes estimated in Section 1.4.3 in explaining rising concentration and the productivity slowdown. The decomposition suggests the greatest scope for declining patent quality of laggards, compared with other explanations such as declining real interest rates, slowing knowledge diffusion from incumbents to entrants, and declining entry rates, to explain the observed trends in productivity growth and concentration. Holding the other parameters fixed at their initial levels, the model-implied change to the patent quality distribution explains just over 100 percent of the productivity slowdown and about 40 percent of the rise in concentration observed in the data, and is consistent with the decline in patent quality documented in section 1.2. Two different decomposition exercises suggest that between 25 and 60 percent of the productivity slowdown generated by an exogenous change in the patent quality distribution is due to firms' responses to this change, in particular a relocation of innovation effort from laggard firms towards market leaders.

Only one other estimated parameter change moves both market concentration and productivity growth in the directions suggested by the data. This is an increase in the elasticity of substitution between product varieties within sectors. I discuss how such a change represents a dynamic version of the exercise in Autor et al. (2017) modelling the rise of superstar firms. This experiment in the quantitative model suggests that on impact this change raises measured TFP, but has a negative effect on growth dynamically through a standard Schumpeterian channel that lower markups reduce incentives for innovation. A change in patent quality, on the other hand, generates a modest rise in markups and the profit share that is consistent with the data.

1.5.1 Decomposition

To understand the contribution of each estimated parameter change to matching the trends in the data, Table A.4 in Appendix A.3.3 reports the effect of changing each parameter from its 1990s value to its 2000s value, holding the other model parameters fixed at their 1990s estimated values. Note that these are not the marginal effects of each parameter on each moment since the moments are endogenously determined in steady state.

The decline in the discount rate ρ and the exit rate δ_e play a similar role in increasing incumbents' R&D expenditures in order to match the rise in the R&D as a share of GDP, since all incumbent firms discount expected future profits less which increases incentives for innovation. A decrease in the cost of R&D α also helps match the rise in R&D expenditures in the data. However, because they result in more R&D, these changes all have the additional effect of *raising* the TFP growth rate absent the other parameter changes. They do not substantially change the average level of concentration. Only the estimated changes in ϕ, governing relative patent quality of leaders and laggards, and ϵ, governing product substitutability, push both concentration and productivity growth in the same direction as in the data. I next explore these two parameter changes in more detail.

1.5.2 Role of Changing Patent Quality

Table 1.6 summarizes the role of the model-implied change in ϕ compared to changes in all the other parameters at once to match the moments of interest. To decompose the effect of a change in the patent quality distribution, I use the values in Table A.4 to compute the share of the changes in the data that are explained by a change in ϕ as follows:

$$\frac{M_j(\boldsymbol{\theta}_{1990s}, \phi_{2000s}) - M_j(\boldsymbol{\theta}_{1990s}, \phi_{1990s})}{D_{j,2000s} - D_{j,1990s}} \times 100$$

Where M_j is moment j in the model steady state with the other parameters $\boldsymbol{\theta}$ held fixed at their estimated 1990s values and $D_{j,t}$ denotes the moment's value in the data at time $t \in \{1990s, 2000s\}$.

A change in the patent quality distribution alone, consistent with lower probability that the followers catch up to leaders through innovation, can explain 102% of the productivity slowdown and about 46% of the rise in concentration in the data.[24] It explains about a quarter of the rise in the profit share and more than three quarters of the decline in turnover in market leadership. As can be seen in Table 1.6, the model-implied change in ϕ from the re-estimation of the model also closely matches the observed decline in average patent quality documented in section 1.2.2. Figure 1.8 shows the expected quality improvement from innovation for a firm at

[24] Transition dynamic analysis in Appendix A.3.4 suggests the productivity slowdown occurs within a few years, while concentration takes a long time to reach its new steady state level, consistent with Figure 1.1.

Moment	Data change (pp)	Decomposition	
		Innov. (%)	All others (%)
TFP	-1.28	102.4	-60.0
Concent.	4.78	45.6	31.0
R&D/GDP	0.09	-1322.2	1056.6
Profits/GDP	1.37	23.4	16.1
R&D/Sales	1.24	-262.1	221.0
Pat. qual.	-11.81	95.5	-15.7
New leader	-4.47	76.3	-7.6

Table 1.6: Share of changes in moments between 1990s and 2000s explained by estimated parameter changes in Table 1.5. Column labelled "Innov." holds other parameters fixed at 1990s values while ϕ changes to its estimated 2000s value. "All others" column holds ϕ fixed and allows the six other estimated parameters to change to 2000s values. Positive sign in the second and third columns indicates same direction of change as in the data.

each technology gap for the two different model-implied values of ϕ, holding the minimum quality improvement λ fixed at its 1990s value.

Turning to the mechanisms behind these changes, firms with lower quality than their rival ($m < 0$) respond strongly to a decline in the expected return from innovating by choosing a lower arrival rate of innovations (Figure 1.9a). Facing a lower probability of being overtaken, market leaders discount future gains to innovation less and choose a slightly higher rate of innovations.

Taken together, this relocation of innovative activity from followers to leaders causes the stationary distribution of sectors over the leader's technology gap to shift right: more sectors now feature a leader that is further

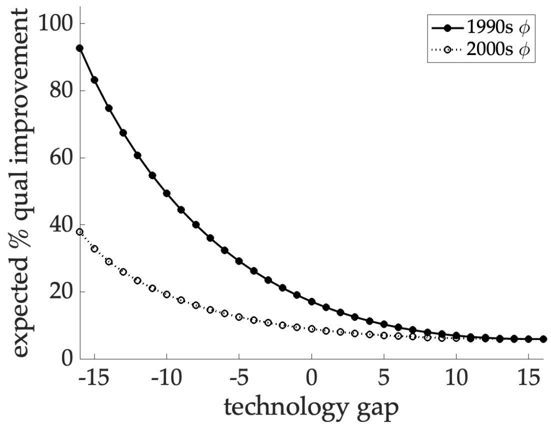

Figure 1.8: Expected quality improvement from innovation as a function of firm's current technology gap, two different values of patent quality parameter ϕ.

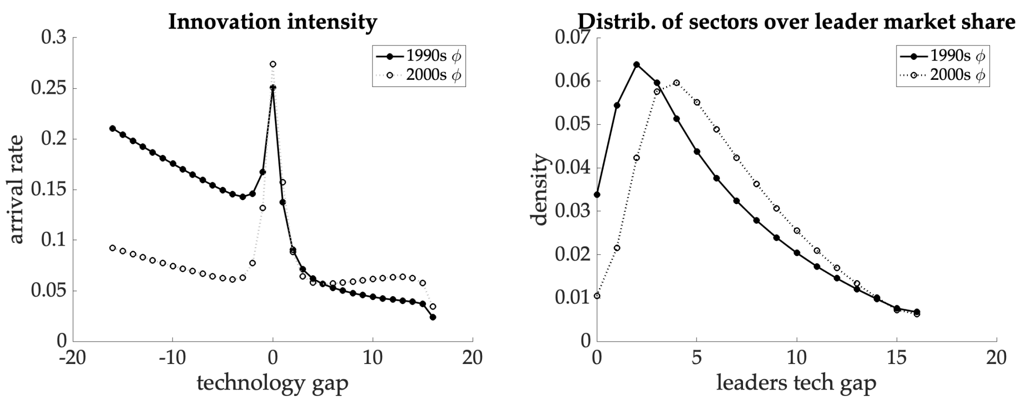

(a) Innovation policies

(b) Stationary distribution

Figure 1.9: Innovation Policies and Stationary Distribution, Role of Patent Quality

ahead (Figure 1.9b). The rise in concentration, average markups (from 23.8% to 25%), and the profit share in the model equilibrium with lower patent quality is driven purely by this composition effect. Note that fixing innovation effort at its 1990s level in the model but increasing ϕ (de-

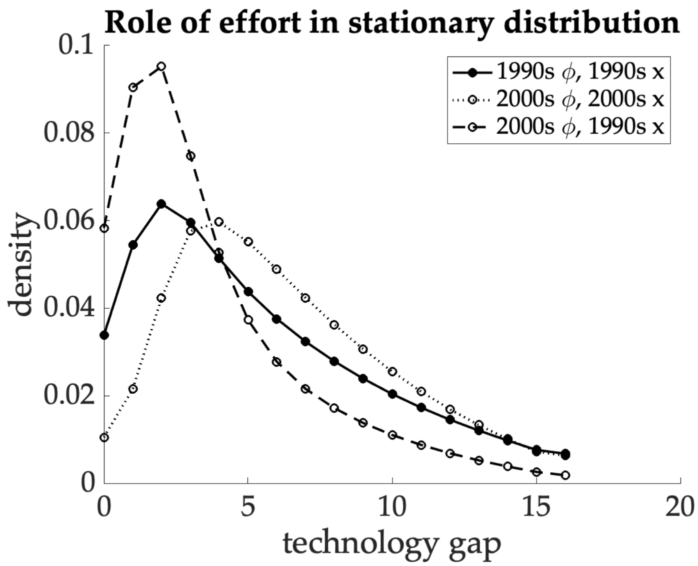

Figure 1.10: Stationary distribution of firms over technology gaps in 1990s and 2000s equilibria, plus distribution assuming firms innovation policies x_m are fixed at their 1990s values while ϕ changes to its estimated value in the 2000s equilibrium (counterfactual).

creasing the average innovation size) would result in a *tighter* distribution around the neck-and-neck state, because competitors pull away from each other less frequently under a higher ϕ regime. This counterfactual is plotted in Figure 1.10).

The growth rate declines when laggards' patent quality is lower for two reasons. First, there is an endogenous effect that comes from changes in firms' innovation policies x (Figure 1.9a). R&D expenditures as a share of output decline from 1.9% to 0.7% (Table 1.6). The decline in R&D expenditures is concentrated among industry laggards, whose average R&D intensity declines from 7% in the 1990s to 1.5% in the 2000s. Leaders' average R&D intensity declines just slightly from 2.5% to 2%. As a result, the average leader's share of total industry R&D rises from 20% in the equilibrium corresponding to the 1990s to 51% in the equilibrium with lower patent quality. Leaders' quality improvements are more incremental on average, so both the level effect of reduced R&D expenditures and the reallocation effect contribution to slower productivity growth.

Second, even if firms' innovation policies were unchanged, lowering the average patent quality exogenously lowers the growth rate. To decompose the importance of these two channels, I conduct two different counterfactual exercises (Table 1.7). First, I solve the model holding ϕ fixed at its 1990s estimated value but allowing firms' innovation policies x to change to their 2000s values. The growth rate under this counterfactual is 0.9% per year, which accounts for 60% of the decline in productivity growth due to changing patent quality in the model. The other decomposition fixes firm innovation policies at their 1990s values and reduces patent quality exogenously. The growth rate in this counterfactual is 0.8%, accounting for about 74% of the productivity slowdown in the model.

Decomposition	% of slowdown explained
Role of effort (ϕ fixed, x changes)	60.2
First order effect (x fixed, ϕ changes)	74.4

Table 1.7: Growth decompositions. Details in the text.

1.5.3 Role of Elasticity of Substitution

The elasticity of substitution within sectors ϵ plays an important role in the determination of the level of concentration and the growth rate.[25] Though the re-estimation exercise suggests only a small change in ϵ, I next explore larger changes in both directions in ϵ because each can capture (in a very reduced form) different structural changes in the U.S. economy that have been suggested in the literature recently to explain rising concentration or rising markups. These exercises illustrate how the model can be used to unify the neo-Schumpeterian endogenous growth literature with the literature on superstar firms and rising market power. Neither change matches the direction of all the moments of interest that declining patent quality of laggard firms does, though the superstar firm experiment gets closer to the data than increased market power. This is because the calibrated model has the standard Schumpeterian feature that increased market power gives a greater incentive for innovation.

Increasing Market Power?

Recent research has focused on the potential costs of rising market power and markups (see de Loecker, Eeckhout, and Unger (2020), Eggertsson, Robbins, and Wold (2018) and Edmond, Midrigan, and Xu (2018) for example) for growth and welfare. Can an increase in market power generate

[25] Following the same decomposition as in the previous section, the estimated change in ϵ explains 1.6% of the productivity slowdown and 8% of the rise in concentration.

the same predictions for the macroeconomic changes experienced in the U.S. in recent years as a change in the probability of radical innovations in the model? I model an increase in market power as a decrease in the substitutability of products in the same sector, ϵ, making the incumbents' varieties more differentiated and increasing the markup the leader charges for the same level of quality differences.

The calibration remains the same as in Table 1.1. I decrease ϵ from 4.2 in the baseline to 3 as an illustration. The model-generated moments for this exercise are compared to the 1990s baseline results in Table 1.8. Average markups rise by about 10 percentage points, about a third of the total rise estimated by de Loecker, Eeckhout, and Unger (2020). With the exception of markups, R&D, and the profit share, the results from this exercise are the opposite of what has happened in the data. Because of greater market power, the leader's markups and profits are higher for the same level of the technology gap (see Figure 1.12b) and this induces more innovation effort by laggard firms as they try harder to overtake the market leader (R&D/GDP rises from 1.9% to 2.8%, see Table 1.8 and Figure 1.11a). This results in a higher growth rate. There is also greater turnover in market leadership and average quality differences between leaders and followers go down (Figure 1.11b), contrary to the data.

Moment	Model	
	$\epsilon = 4.2$	$\epsilon = 3$
TFP growth, %	1.75	2.12
Leader market share, avg, %	44.62	40.6
R&D share of GDP, %	1.91	2.77
Profit share of GDP, %	6.02	6.54
Pat stock growth/patent, avg, %	22.26	22.08
R&D intensity, avg, %	5.18	6.76
Leadership turnover, %	13.26	14.18
Markup, avg, %	23.84	33.25

Table 1.8: Model comparison, market power experiment.

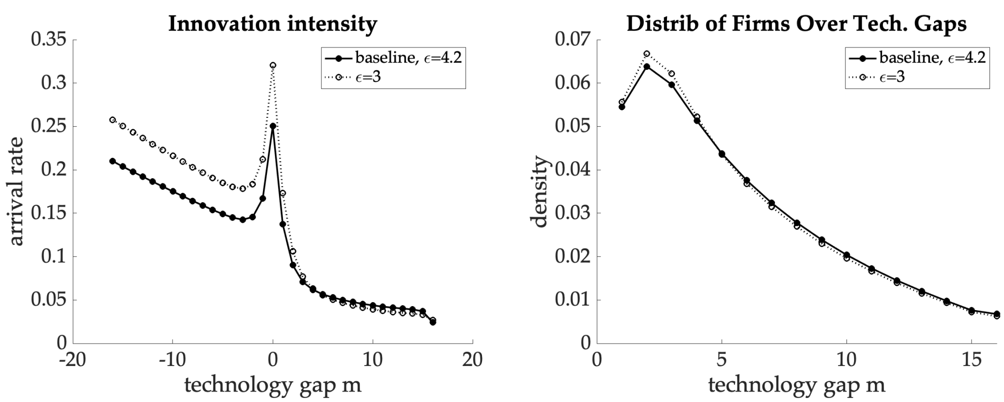

(a) Innovation policies

(b) Stationary distribution

Figure 1.11: Innovation Policies and Stationary Distribution, Market Power

Superstar Firms?

Seminal work on the macroeconomic effects of superstar firms is Autor et al. (2020), who show that the rise of superstar firms can explain the declining labor share of GDP. In their static industry model (Autor et al. (2017)), firms draw labor productivities from an exogenous distribution and then make an entry decision.[26] Firms that decide to enter produce differentiated varieties of the sector good and set prices a la Bertrand.

The force for reallocation to more productive firms in Autor et al. (2017)'s model is an increase in product substitutability between varieties. The authors argue that this increase could represent more fierce import competition from abroad, particularly China, in recent years or increased price sensitivity due to better search technology such as online retail. Keeping the exogenous productivity distribution fixed, an ancillary result of their analysis is that a sector's measured TFP will rise unambiguously when substitutability increases because of two forces: first, the minimum productivity threshold for entrants rises, and second, more productive firms increase their sales share.

I first show that this static reallocation result is present in my two-firm industry model with a competitive fringe for the estimated parameter values. Figure 1.12a plots market shares as a function of the technology gap for different values of ϵ for the baseline parameterization of the model given in Table 1.1. For most values of the technology gap, increasing the substitutability of the incumbents' varieties statically increases the leader's market share which raises sale-weighted sector TFP.[27]

On impact, therefore, increasing the substitutability of the firms' prod-

[26] The model in Autor et al. (2020) generalizes this formulation. The exercise presented here studies the specific shock to market toughness suggested in Autor et al. (2017).

[27] Without the competitive fringe this is always true. But it is not necessarily true under the assumption that the follower sets price equal to marginal cost: if relative quality differences are small, increasing ϵ can cause a drop in the leader's market share.

Figure 1.12: Markups and Market Shares, Role of Elasticity of Substitution

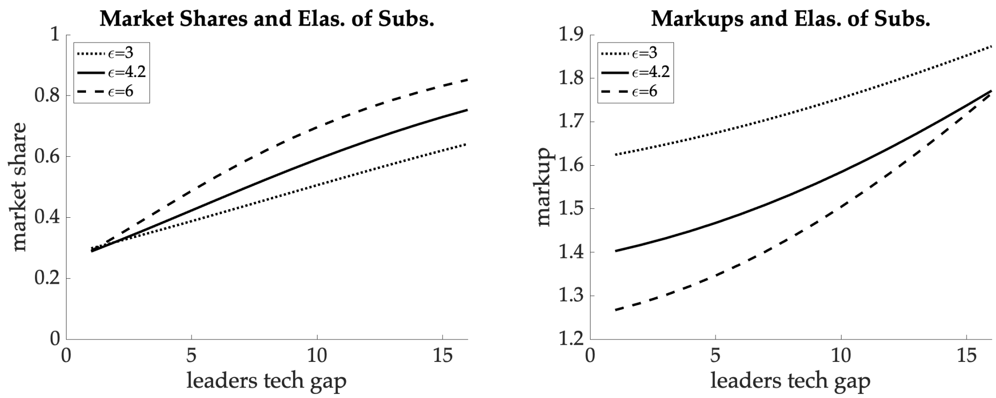

(a) Market shares (b) Markups

ucts tends to raise measured TFP as in Autor et al. (2017). To analyze the dynamic effect of this change I compare the steady state of the model with higher ϵ to the 1990s equilibrium of the model in Table 1.9. Under this parameterization, raising ϵ lowers the growth rate while dramatically increasing concentration. The rise in concentration comes from two forces. First, the static reallocation force operates: even if technology gaps were unchanged from one steady state to another, these same gaps would generate a higher average leader market share according to Figure 1.12a. Second, changes in patenting frequency across different types of firms (Figure 1.13a) cause the average technology gap to grow (Figure 1.13b).

This latter effect is due to the fact that, for the given parameter values, the markup the leader charges as a function of its technology gap is lower at all possible values of the technology gap when ϵ changes from 4.2 to 6 (Figure 1.12b). This reduces the post-innovation gains to attaining market

Moment	Model	
	$\epsilon = 4.2$	$\epsilon = 6$
TFP growth, %	1.75	1.56
Leader market share, avg, %	44.62	51.12
R&D share of GDP, %	1.91	1.52
Profit share of GDP, %	6.02	6.05
Pat stock growth/patent, avg, %	22.26	22.54
R&D intensity, avg, %	5.18	4.75
Leadership turnover, %	13.26	12.72
Markup, avg, %	23.84	18.79

Table 1.9: Model comparison, superstar firm experiment.

leadership, reducing the innovation effort of laggard and neck and neck firms (Figure 1.11a). On the other hand, markups and profits become more elastic in the technology gap when ϵ is higher, and the likelihood of being overtaken falls, so leaders choose a higher arrival rate of innovations than before.

Similar to the experiment reducing laggards' patent quality, there is both a reduction in the level and a shift in the location of R&D expenditures that generates a productivity slowdown, though this slowdown (from 1.75% annual productivity growth to 1.56%) is much smaller than the slowdown driven by changing patent quality. Moreover, the average markup falls, contradicting the findings of de Loecker, Eeckhout, and Unger (2020). Compared to laggards' declining patent quality, the main exercise suggests only a very modest scope for the superstar firm channel.

Figure 1.13: Innovation Policies and Stationary Distribution, Superstar Firms

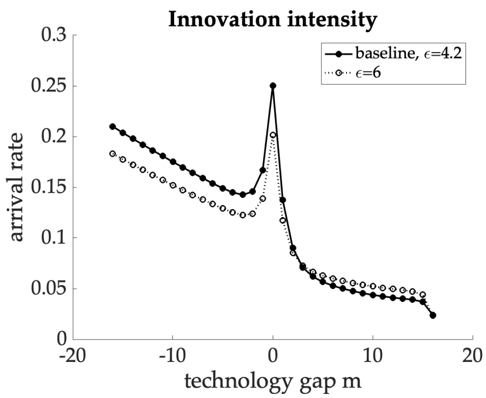

(a) Innovation policies

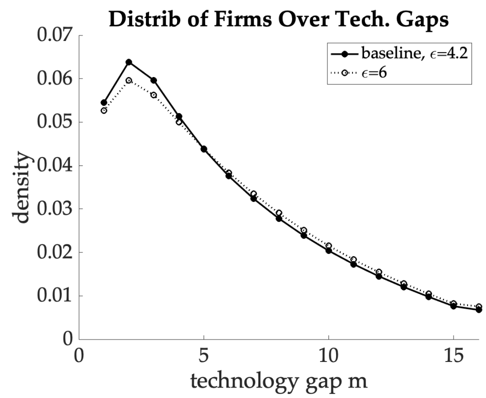

(b) Stationary distribution

1.6 Conclusion

This paper documents a striking decline in the contribution of new patents to firms' existing portfolios of patents since 2000, raising the possibility that the quality of new ideas has declined relative to the 1990s. Smaller firms within industries drove the boom in innovation quality that has been attributed to the arrival of information and communication technology as a general purpose technology, as well as the bust that began in the early 2000s. This finding contributes to the debate on whether ideas are getting harder to find, emphasizing heterogeneity in this phenomenon across firms. Further empirical work should investigate heterogeneity in the complementarities of general purpose technologies with different types of firms' R&D investments, and determine whether similar patterns of firm dynamics and productivity growth were present in the wake of previous general purpose technology waves.

To understand the consequences of this empirical fact I develop a general equilibrium model of innovations and growth where multiple firms are active in each sector in each period and goods within sectors are imperfect substitutes. A quantitative version of the model estimated for the U.S. in two different periods, the 1990s and the 2000s, points to declining patent quality as the main driver of rising concentration and the productivity slowdown over this period. Rising concentration in the 2000s is driven by a decline in the research effort of laggard firms and an increase in research effort by large firms, which causes average quality differences between competitors to grow, consistent with the rise of superstar firms. Because leaders make more incremental improvements on average, the economy grows more slowly as a result.

Through the lens of the model I unify the Schumpeterian endogenous growth literature with the growing literature on the rise of superstar firms. The estimated model points to only a modest rise in the elasticity of substitution within sectors, though I show that larger changes also have the potential to explain rising concentration and the productivity slowdown, providing a dynamic complement to the experiment in Autor et al. (2017), one that rationalizes the emergence of superstar firms alongside a productivity slowdown for a standard Schumpeterian reason: in this environment, laggard firms' incentives to innovate fall because varieties are less differentiated and the value of market leadership is lower. However, this is inconsistent with patterns of markups and the profit share, which have risen rather than fallen over this period.

Analyzing welfare and optimal policy is left to future research, though the preceding discussion offers some insight into the relevant trade-offs between reducing static markup distortions and providing dynamic incentives to innovate. Such analysis should ensure that knowledge spillovers between firms are properly accounted for. Another area for future research is how policy can incentivize the development of new general purpose technologies. Past research suggests this is difficult because there are significant positive externalities for other sectors that the inventor of the general purpose technology does not internalize. This paper suggests there may also be winners and losers within other industries, further complicating this problem.

Chapter 2

Country Banks and the Panic of 1825

Abstract

The Panic of 1825 was one of the world's first international financial crises. I document how this crisis spread from London banks to England's real economy. England's correspondent banking network propagated trouble in sovereign debt markets to small banks outside of London and ultimately to non-financial firms. Using exogenous variation in town-level exposure to the crisis, I show that bank failures led to a substantial number of bankruptcies among non-financial firms, particularly in non-tradable sectors. These findings highlight the costs of disruptions to the payment system: country bank notes were the primary means of payment during the first industrial revolution.

2.1 Introduction

The Panic of 1825 was a British financial panic that followed a credit and speculative boom from 1821-1825. The crisis involved the first cases of sovereign default on government bonds in international capital markets and has been called the first emerging-market induced financial crisis (Dawson (1990); Bordo (1998); Morgan and Narron (2015)). Following a series of bad news shocks in the fall of 1825, the London money market seized suddenly in mid-December and runs occurred on many London banks, causing several major banks to temporarily stop payment and others to fold entirely. Through the correspondent banking network of relationships between London banks and small, so-called "country" banks in English towns outside of London, the crisis exerted significant financial stress on the country banks.

More than 10% of England's country banks went bankrupt during the Panic and real activity declined dramatically. Among Britain's banking crises over the past 200 years, Turner (2014) puts only the Panic of 1825 on par with the Great Recession of 2007-8 in terms of financial distress and output costs. Construction activity, measured by brick production, fell by 30% from 1825 to 1826 (Shannon (1934)). The value of exported cotton manufactures fell 20% and the quantity of raw wool imports fell by more than half over the same period (Gayer, Rostow, and Schwartz (1975)). Bankruptcies more than doubled from 1,141 in 1825 to 2,590 in 1826 (Marriner (1980)). In this paper I provide the first causal evidence that bank failures during the crisis contributed to the decline in real activity. The main channel appears to be through a disruption to the payment system rather than a credit supply shock, which stands in contrast to much of the literature on the real effects of financial crises.

Like many modern financial crises, including the global financial crisis of 2008, the Panic of 1825 occurred after a large number of new securities appeared in financial markets. Surprising news about the low quality of these assets caused runs on the unregulated financial institutions that were exposed to them. Credit and financial intermediation both within and outside of the financial center (London) contracted. However, several features distinguish the Panic of 1825 from more recent crises. First, towns outside of London used bank notes issued by local country banks as currency so that bank failures directly affected household cash balances through a sharp devaluation of bank notes held by households. Second, the banking network was not diversified: most country banks relied on a single correspondent London bank.

In this paper I explore the impact of country bank failures during the crisis on local economic activity. Economic historians have debated the role of country banks in supporting industrialization and economic growth during the first industrial revolution but lack of data and identification have both been major barriers to answering this question. To address these challenges, I use new, hand-collected data on the universe of English country banks from 1820-1830 and data on local bankruptcies to compare the number of non-financial firm bankruptcies in towns with different numbers of bank failures. I instrument for bank failures using town-level exposure to the sovereign debt crisis through the correspondent banking network. I show that transmission of financial stress occurred geographically and through the banking network.

I find that towns with country banks that failed due to the crisis experienced a higher number of non-financial firm bankruptcies than comparable towns that were not exposed. The direct, partial equilibrium effect of bank failures can explain about 27% of the increase in bankruptcies during the crisis. I discuss two possible ways financial stress on the country banks was transmitted to local economic conditions: (i) a drop in aggregate demand due to lost household wealth; (ii) a negative credit supply shock, particularly to working capital lending.

The paper contributes to two principal strands of literature. First, the paper is related to the literature measuring the cost of financial crises for non-financial firms. Chodorow-Reich (2014) and Fernando, May, and Megginson (2012) use a similar identification strategy to study the effects of particular institutions' collapse on the availability of credit during the Great Recession. Amiti and Weinstein (2018) and Mian and Sufi (2008) use matched firm-lender data in Japan and Pakistan, respectively, to identify bank shocks in a more reduced form way, using firm and bank fixed effects to identify loan supply shocks from individual banks. A concern in many of these studies is the extent to which firms can switch banks to avoid loan supply shocks. The highly localized nature of English lending and rigid bank-firm relationships during this period provide an ideal setting to isolate the effects of financial stress when switching is not possible.

The focus of much of this literature has been on identifying the effects of a credit supply shock. Yet financial crises that disrupt the payment system or otherwise negatively household wealth may also feature reductions in aggregate demand. One study that focuses on the local demand effects of banking crises is Huber (2018), though the demand effects he documents for Germany are essentially second-round effects of the credit supply shock due to employment losses. I instead highlight the role of country banks in the payment system and study the first-order effects of payment disruptions for local aggregate demand. This channel bears a resemblance to the recent demonetization episode in India, where Chodorow-Reich et al. (2018) show that districts experiencing more severe cash shortages suffered greater reductions in economic activity in the short run, and is also consistent with the Friedman and Schwartz (1963) hypothesis that the money supply has first-order effects on output.

Several papers have examined the impact of bank distress on firm outcomes in a historical context. In Britain, Kenny, Lennard, and Turner (2017) use data on various banking crises in Britain in a VAR framework to identify industrial production contractions of about 8% from these crises. For the U.S, Frydman, Hilt, and Zhou (2015) identify a large causal effect of runs on shadow banks on aggregate investment in the Panic of 1907 using exposure to a financial scandal. Calomiris and Mason (2003) show that counties and states with more financially distressed lenders saw greater slowdowns in construction activity during the Great Depression. Hansen and Ziebarth (2017) use geographic variation in financial distress within Mississippi to show that financial distress caused firm exit but not bankruptcy during the Great Depression. I provide further evidence that financial crises feature large effects on non-financial firms, including during one of the earliest modern financial crises, the Panic of 1825, and focus on how historical banking crises disrupted the functioning of the payment system in the absence of a single national currency.

The source of exogenous variation in financial stress in this setting merits special attention. Several recent papers have documented the effects of foreign financial crises on the credit supply decisions of domestic lenders (Bottero, Lenzu, and Mezzanotti (2017), Ongena, Peydró, and Van Horen (2015), and Huber (2018)). To the best of the author's knowledge, only in one other study (Acharya et al. (2018)) is the case where *foreign* sovereign debt devaluation was the driving force behind deterioration of the domestic banking system's balance sheet considered, focusing on the recent European debt crisis. The authors show that banks in EU countries contracted lending due to the balance sheet effects of sovereign debt devaluation (see Bocola (2016) for a model of this channel) but the effects were concentrated in banks from the five countries (Greece, Ireland, Italy, Portugal, Spain) most affected by the sovereign debt crisis. The present analysis provides further evidence of this particular cost of sovereign default: contraction of financial intermediation in the investing country, even by domestic banks. The existence of this channel in early 19th century capital markets is a novel finding.

Second, the role of finance, and of country banks in particular, during England's first industrial revolution is hotly contested. Contemporaries pointed to the high failure rate of country banks as evidence that these banks harmed the towns they served and excessive note issuance by country banks was blamed for credit boom and bust cycles in the early 19th century. Some historians have argued that restrictions on the maximum number of partners and the Usury Law restricting the maximum interest rate banks could charge caused the industrial revolution to be "financed out of the pockets of tinkerers and manufacturers, not through bank lending" (Calomiris and Haber (2014)). Others argue that public finance of the Napoleonic wars largely crowded out private finance (Murphy (2014); Temin and Voth (2013)). Crouzet (1972) summarizes the conventional view of the relationship between banks and industry during the industrial revolution: "they lived in two separate worlds and that the contribution of the banking system to the industrial revolution was therefore quite insignificant."

Those arguing for the importance of private finance include Crouzet (1972), Mathias (1973), and Pollard (1964), who point out that banks' provision of short term credit freed up internal profits for reinvestment in longer term capital. Indeed, these so-called plough-backs were the primary source of fixed capital formation during this period, and would not have been possible had firms been obliged to meet their short-term needs with profits.[1] Heblich and Trew (2019) find that employment in the financial sector in 1817 was associated with structural transformation and industrialization by 1881. This paper complements those findings by focusing on the short term consequences of a loss of financial services and by highlighting the role of country banks in particular.

The rest of the paper is organized as follows: Section 2.2 describes the history of country banks, their role in the towns they served, and the structure of the English financial system during the period of study, as well as providing background on the Panic of 1825. Section 2.3 describes the two new datasets I collected. Section 2.4 presents the main estimation results of the paper and discusses the validity of the instrumental variable approach, then discusses possible transmission mechanisms from bank failures to firm bankruptcies. Section 2.5 considers robustness checks.

[1] Brunt (2006) suggests country banks also played a role in industrialization by funding fixed capital investment but evidence on this point is sparse. Policy changes after the Panic (described in detail at the end of section 2.2.2) make it difficult to study the longer-term causal effects of bank failures on town-level outcomes.

2.2 Historical Context

2.2.1 Country Banks in England

The banking system in England during this period was a three-tiered structure comprised of the Bank of England, London banks, and country banks. Country banks began to appear in England in the mid-18th century, around the time the first industrial revolution began. By 1815 it was estimated that there was a country bank within 15 miles of anywhere in England, according to Pressnell (1956), a comprehensive text on country banks during this period.

Legal restrictions capped the number of partners in a bank at six, so country banks were small and served a very limited geographic area (usually they were unit banks but were not legally prohibited from branch banking). These wealthy partners provided the initial capital for the bank. Country banks also took deposits, but these tended to only be from wealthy individuals and large firms. The main employments for bankers' resources, other than cash reserves, were local loans and purchases of "London assets": British government securities and interest-bearing balances with London bankers. Country bankers tended to match their assets to their liabilities. Private country bank note issues were backed by liquid cash reserves and government securities in London; deposits were used to discount local or London bills of exchange, and longer term loans came from bankers' capital (Pressnell (1956)).

Each country bank maintained an account with a bank in London, referred to as its London agent. While most of the country banks had at most one or two branches so that they appeared quite isolated, the London agent system served to connect banks across the country with one another. The London agent performed important functions for country banks, transferring excess capital from one part of the country to another, particularly from agricultural areas to industrial areas (Pressnell (1956)). London agents also settled transactions among different country banks. For these services, country banks compensated their agents by promising to leave a large permanent deposit at their London bank (Pressnell (1956)). This scheme meant that if a country bank's London agent failed, as many did during the Panic of 1825, the country bank could face the substantial capital loss of their London balance. In line with recent findings for the U.S.'s correspondent banking network of the 1920s and 1930s (see Calomiris, Jaremski, and Wheelock (2019)), I will show that this network was an important source of propagation of financial shocks.

2.2.2 Panic of 1825

The key to my identification strategy will be that the Panic of 1825 crisis originated outside the small town economies I study. In this section I first discuss the causes of the crisis, then provide a timeline of events during the crisis, and finally discuss the conduct of the Bank of England and other policy responses.

The Panic of 1825-26 has been called the first Latin American debt crisis (Dawson (1990)). The success of Barings' French bond offerings in 1820, combined with the prospect to invest in metal-rich newly independent Latin American governments, and the low return on British government consols, created a huge demand for Latin American securities in the early 1820s (Neal (1998)). As John Horsley Palmer, the Governor of the Bank of England in 1832, put it: "the excitement of that period was further promoted by the acknowledgement of the South American republics by this country, and the inducements held out for engaging in mining operations, and loans to those governments" (Great Britain (1832)). According to Gayer, Rostow, and Schwartz (1975), these Latin American issuances constituted the "largest single category of new investment" in the lead up to the crisis.

Information on the quality of these assets emerged in London only slowly over the years following the first issuances in 1822. Because little was known about each individual government, all Latin American bonds were priced at a heavy discount, and prices for all countries tended to move together before 1825 (see Figure 2.1). So little information was available about Latin America at the time that Scottish explorer Gregor McGregor was able to issue bonds for the fictional Central American government of Poyais on similar terms as bonds issued for Chile (Morgan and Narron (2015)). Latin yields surged during the fall of 1825 as the Poyais and other schemes were revealed.[2] A November 20, 1825 article in The Examiner (1825b) questioned Buenos Aires' ability to meet its next coupon payment. A December 8 report (The Morning Chronicle (1825)) noted that "every description of Foreign Security continues under a cloud–but more especially South America." By January 1826 it was clear that most, if not all, Latin American borrowers were insolvent (Dawson (1990)), though it was not until April 1826 that the first country, Peru, formally defaulted (that is, completely stopped making interest payments).

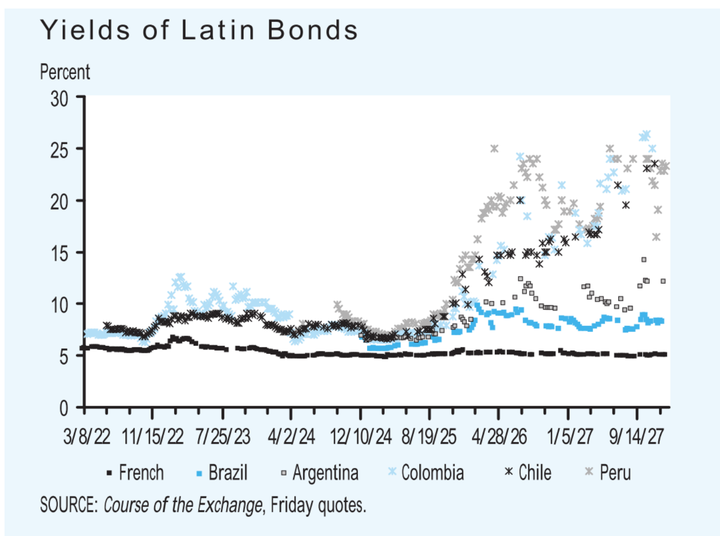

Figure 2.1: Source: Neal (1998).

London banks could assume two possible roles in sovereign debt issuances during this period. First, banks could underwrite foreign government debt, assuming liability in case of default. Second, banks could

[2]The Hereford Journal (1825) reported MacGregor's arrest on December 8th, 1825, just before the first London bank failed on December 12th.

assume a "payee" role, responsible for collecting the payments of loan subscribers and then acting as a "window" on behalf of the issuing country and paying out coupon payments. Flandreau and Flores (2009) make the distinction between underwriting and being the window/payee on sovereign debt in the 1820s. They argue that "the risks and revenues of the last two operations were much smaller than those from the first, but leads and lags could cause trouble." For example, Dawson (1990) describes a case as early as 1823 when the London bank Everett & Co., the payee for Peru, temporarily froze payments it had collected from loan subscribers to Peru because of uncertainty about regime stability in Peru. The Peruvian government in turn needed these payments to make coupon payments back to bondholders as it had no other gold or Bank of England notes on hand, and ended up suing Everett & Co. to release the money. Early incidents like these undermined confidence in both Peru's and Everett's ability to meet their obligations and also suggest that payee banks were sometimes expected to supply the coupon payments when the country could not come up with the money itself.

Rumors of impending defaults put immense pressure on the London banks involved in debt issuances of these countries. The failures of sovereign borrowers associated with a bank could have devastating reputational consequences for the bank, according to Flandreau and Flores (2009) and Indarte (2016). Often the only information published about sovereign debt issuances in the 1820s were the amount of debt, the interest rate, the underwriter, and which bank would make the coupon payments.

Three of the six largest London agents (by number of country bank correspondents) that failed during the crisis had been the party responsible for paying interest on Latin American debt issuances, according to Dawson (1990). Perring & Co. paid interest on Poyais, Everett & Co. on Peru, and Fry and Chapman on Mexico and on a portion of the Peruvian debt. These three London agents accounted for nearly half of the country banks exposed to the failure of their London agent during the crisis.

Other London banks that had no country bank correspondents but were also involved in underwriting and issuing South American securities also failed during the panic: Goldschmidt failed in February 1826 and Barclay, Herring, Richardson & Co. in April 1826. Flandreau and Flores (2009) cite evidence that underwriting banks, particularly the market leaders Barings and Rothchilds, would intervene to support securities' prices by buying them up during selloff periods. No surviving evidence proves that other banks involved in issuing these securities also did this, but it is certainly possible. Holding these assets on one's balance sheet would be extremely costly: Peruvian bonds earned a negative return of 15 percent during this period, for example (Flandreau and Flores (2009)).

Other causes of the Panic of 1825 have been suggested. Like the lemons in the sovereign debt markets, Neal (1998) identifies 624 companies floated in England from 1824 to 1825, only 127 of which survived to 1827. Many of these were international companies, particularly mining companies in

South America. The Bank of England had also created highly accommodative monetary conditions since 1819 that may have created a credit boom ready to burst by 1825. The Bank lowered its discount rate from 5 percent to 3 percent during this period. Pressnell (1956) argues that a fall in the rate of interest on their London balances encouraged country banks to search for higher yield by increasing their own note issues from December 1823 to December 1825, but also explains that these increases were matched by increases in the demand for credit due to good harvests and increased foreign trade. All three of these causes, the sovereign debt crisis, the decline in foreign private stock prices, and accommodative monetary policy, were more or less external to real economic activity in the interior of England.

A concrete timeline of the crisis helps to clarify my identification strategy. Against the backdrop of increasing tightness in the London money market caused by the sovereign debt crisis, the panic began on December 12, 1825 when the London bank Pole, Thornton & Co. stopped payment.[3] News of the crisis spread rapidly, even to the countryside, and runs began the next day on country banks known to be Pole's correspondents. Runs also occurred on other London banks suspected to be in trouble. By the end of that week four major London banks with a total of 65 country correspondents had failed. By the end of December, 30 country banks had

[3] On the causes of Pole, Thorton & Co.'s failure, The Examiner (1825a) writes "The decline of this house is generally attributed to the anxiety felt by the partners at the time when the rate of interest was low, to make a profitable use of their capital, and hence they were led to employ it on securities capable of being realized only at a distant period, or of an inferior degree of credit."

been declared bankrupt and 41 more would follow suit from January to May of 1826 (see Figure 2.2).

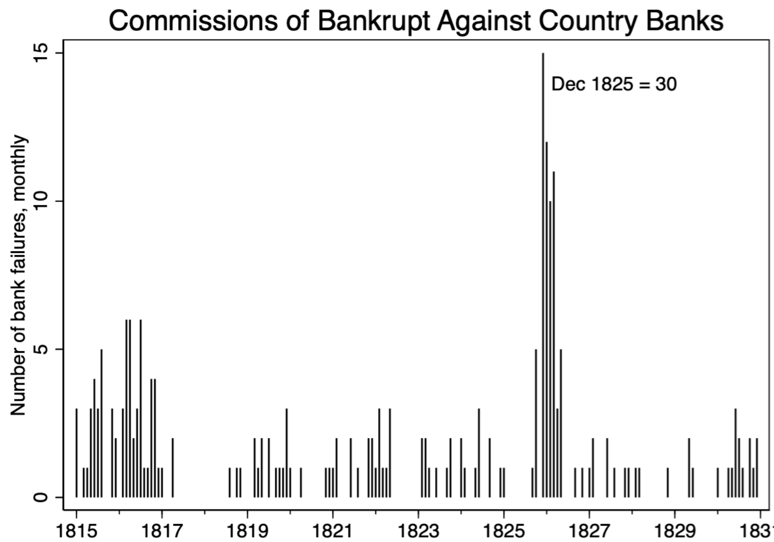

Figure 2.2: Source: Report from the Committee of Secrecy on the Bank of England Charter, 1832, Appendix No. 101.

How did the Bank of England respond to the panic in the money market? James (2012) called the Bank of England's response to panics from 1790 until 1825 "limited, episodic and inconsistent", but noted the more active role the Bank played in addressing the 1825 crisis through liberal discounting (p. 299). However, at the same time, the Bank increased the discount rate from 4% to 5% on December 14, 1825 at the height of the crisis, likely due concerns that its bullion reserves were getting too low (the Bank had resumed convertibility of its notes into gold in 1821).

The Bank of England also began to issue one pound notes for a brief period during the crisis as an emergency measure (all small notes had been removed from circulation in 1821). Parliament set up special boards to make advances to economically distressed country towns. In a previous financial crisis during 1810, the Bank of England had issued Exchequer bills specifically to support "Manufacturers, in different parts of Great Britain, who, having in great degree suspended their works, were enabled to resume, and to afford employment to a number of workmen who must otherwise have been thrown on the Public for support" (Pressnell (1956) p. 468). Similar concerns were at work in 1825.

The Panic of 1825 eventually caused significant reforms to the English banking system. In 1826 Parliament lifted the prohibition on joint-stock banking outside London that had been in place since 1708 (Black (1995)). These larger banks slowly absorbed or out-competed the country banks. At the same time, the Bank of England sought to expand its sphere beyond London by opening branches in seven major English cities. Because small note issues were blamed for the crisis, in March 1826 Parliament declared that all private notes below five pounds had to be withdrawn by 1829 (Black (1995)). These additional changes to the banking system after 1826 make it difficult to study the longer term effects of the crisis, so I focus my analysis on quite a narrow period, as described in further detail in the next section.

2.3 Data

To undertake this study I created two novel datasets.[4] One database records the universe of country banks in the United Kingdom, matching each country bank to a London agent. The other dataset includes all bankruptcy records in England for the period December 1, 1824 to June 30, 1826. This section describes the datasets. Additional details are provided in Appendix B.1.

2.3.1 Banking Network

I construct the banking network characterizing the English financial system from 1820-1830 using five years of data from the Post-Office London Directory: 1820, 1823, 1825, 1827[5], and 1830. For the main results of the paper I identify country bank branch failures on the basis of their disappearance from the 1827 London Directory relative to 1825. This introduces some measurement error since I am not able to learn when a branch's existence was recorded for publication in the Post-Office Directory, and some may be mistakenly omitted in 1827, resulting in a misidentified failure.

Table 2.1 describes some of the features of the data. Consistent with narrative evidence about the banking system at the time, the number of country bank branches in England peaked in 1825, and began to decline af-

[4]I supplement these two databases with parish/town-level population data from Census of Great Britain (1821).

[5]I thank a researcher at Reed College for providing access to the 1827 volume.

terwards. The greatest number of failures per year came between 1825 and 1827, my window of study. However, a large number of disappearances occurred between 1827 and 1830, likely driven by other changes like the prohibition of small note issues, the introduction of joint stock banking, and competition from the branches of the Bank of England that opened during this period.

Table 2.1: Post-Office London Directories, English Banks Only

	1820	1823	1825	1827	1830
Number of towns	319	343	357	302	268
Number of country banks	465	465	470	418	370
Number of country bank branches	532	565	591	487	447
Number of London agents	61	57	56	50	48
Average bank (branch) failures per year	21	26	64	58	

Source: Post-Office London Directories, 1820, 1823, 1825, 1827, 1830. Average bank (branch) failures denotes the average number of failures per year in the years between the given year and the next year in which data is available.

2.3.2 Firm Bankruptcies

The second set of data I collect is individual bankruptcy statistics from the Edinburgh Gazette.[6] The records include the date of the announcement of

[6] Bankruptcy notices for all of Britain had to be printed in the London, Dublin, and Edinburgh editions of the Gazette. For English bankruptcies the Edinburgh Gazette prints the information most readably. However, using Edinburgh Gazette entries means I don't capture bankruptcies in Scotland and Ireland well, so despite having information on the London agents of many Scottish and Irish banks, I exclude these countries from the subsequent analysis. Appendix B.1.2 includes a further discussion of the decision to focus on England.

the bankruptcy proceedings to the general public, the bankrupt individual's name, location of residence, and occupation. To my knowledge, this is the first time these records have been collected at the town level rather than at the national level.

Bankruptcy commissions could seize an individual's assets, determine which creditors would be paid, and how much each creditor would receive. To be eligible for bankruptcy, an individual's total debt had to exceed 100 pounds, a large sum at the time, and the individual had be classified as a trader rather than a farmer or a professional (Duffy (1973)). These criteria remained fixed over the period of study. This means that the data I collect omits gentlemen, farmers, professionals like attorneys and doctors, and merchants owing amounts under one hundred pounds. Private businesses were not entitled to limited liability during this period because of the Bubble Act of 1720, so I treat individual and firm bankruptcies as equivalent and refer to bankruptcies as firm failures throughout the paper.

The sample I collect covers December 1, 1824 to June 30, 1826 and includes 1,440 bankruptcies in 488 towns, 208 of which had country bank branches. Bankruptcies increased substantially across all occupation classes (see Table 2.2). Note that I discover just 42 bankruptcy notices for bankers during the crisis, compared to the 99 branch failures I identify using the Post-Office directories. In some cases bankers were also merchants or industrialists so I exclude from the bankruptcy database of non-financial firms any individual with the same last name as a named partner in a failed bank. Another reason why bankruptcy notices understate bank failures is that each bankruptcy notice lists a single location, where the individual actually lived. In many cases multi-branch banks failed corresponding to just one individual named in bankruptcy proceedings.

Table 2.2: Bankruptcies in England from Edinburgh Gazette

	Pre-Crisis	Crisis	Total
Bankers	1	42	43
Other financial	11	37	48
Trade	74	202	276
Manufacturing	32	255	287
Retail	28	78	106
Food	71	237	308
Clothing	67	187	254
Construction	34	88	122
Total	318	1122	1440

Source: Edinburgh Gazette. Pre-crisis: Dec. 1, 1824-June 30, 1825. Crisis: Dec 1, 1825-June 30, 1826.

A final concern with the bankruptcy statistics is that many troubled debtors may not appear in the statistics at all due to the inefficacy of bankruptcy commissions during this period.[7] Hearings on bankruptcy laws conducted in 1818 suggested that in cases of debts less than £1000 the costs of bankruptcy commissions usually exceeded the amount recovered

[7] Duffy (1973) argues that "faulty laws and administration encouraged dishonesty and prevented speedy collection of estates" and that bankruptcy laws were unpopular as a result (p. 153).

from bankrupts' estates (Duffy (1973)) and were rarely initiated as a result, so I am likely measuring truly large firms. Because of concerns like these, Silberling (1919) and others have used bankruptcies as a barometer of economic activity rather than a measure of activity in itself. Gayer, Rostow, and Schwartz (1975) show that bankruptcies strongly comove with many other cyclical indicators like trade volumes, indices of goods production, inflation, and the money supply at the national level.

2.3.3 Summary Statistics

Summary statistics for the variables used in the town-level regressions in the rest of the paper are shown in Table 2.3. Combining all towns with country banks with all towns with at least one bankruptcy over the study period, there are 616 total towns in the sample. Just over half the towns in the sample had at least one country bank before the crisis,[8] and the maximum number of country banks in any town before the crisis was 10 (in Bristol). The number of banks per town whose London agent failed ("Exposed banks") and the number of bank failures per town both range from 0 to 3. Because there is little variation in the number of bank failures I will use a Poisson count model to model bank failures in the IV setup (described in more detail in section 2.4.3). There is more variation in the outcome variable, number of bankruptcies per town during the crisis,

[8] In section 2.5 I show that the IV results are roughly the same in the subsample of 328 towns with at least one bank before the crisis.

where the number ranges from 0 to 73 (in Manchester).

Table 2.3: Summary Statistics for Main Variables

	Source	Mean	Std. dev.	Min.	Max.
Bank failures, 1825-1827	POD	0.19	0.48	0	3
Exposed banks, 1825	POD	0.17	0.43	0	3
Total banks, 1825	POD	0.91	1.17	0	10
Has bank, 1825	POD	0.53	0.50	0	1
Population, 1821, thousands	C	6.60	12.10	0.07	119
Firm bankruptcies, pre-period	EG	0.47	1.76	0	27
Firm bankruptcies, crisis	EG	1.64	5.19	0	73
Number of towns: 616					

Pre-crisis: Dec. 1, 1824-June 30, 1825. Crisis: Dec 1, 1825-June 30, 1826. POD denotes Post-Office Directories, EG denotes Edinburgh Gazette, C denotes 1821 Census.

2.4 Results

In this section I first present ordinary least squares results for regressions of individual (firm) bankruptcies on bank failures for 616 English towns. Next, to control for the possibility that bank failures are endogenous to local economic conditions, or that bank failures are mismeasured, I use town exposure to failed London agents as an instrument for town-level bank failures. Before presenting the IV results I discuss the validity of the instrument. Then I use my preferred instrumental variable specification and the bankruptcy occupation data to isolate one possible channel: a

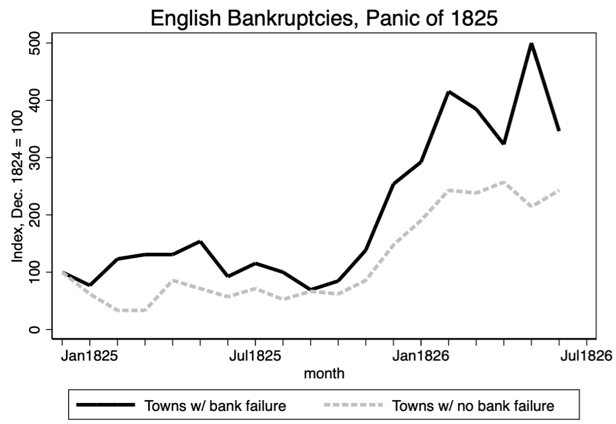

Figure 2.3: Source: Author's calculations from London Post-Office Directory and Edinburgh Gazette. 99 towns with bank failures and 229 towns without bank failures, indexed to their December 1824 total.

negative aggregate demand shock from lost household wealth. Using the historical record, I discuss an additional channel: a credit supply shock to the availability of working capital loans.

2.4.1 Ordinary Least Squares Results

First, to visually compare the difference in firm bankruptcies between towns with and without bank failures, Figure 2.3 compares towns with at least one bank failure to towns with banks but no bank failures and shows that these two groups had similar patterns of firm bankruptcies until the Panic of 1825, when towns with bank failures began to experience much higher bankruptcy rates.

Table 2.4 shows ordinary least squares (OLS) regressions of firm bankruptcies on bank failures. Letting f_i denote the number of firm failures in town i, bf_i the number of bank failures, and X_i a set of town-level controls (town's population, number of pre-period bankruptcies, and an indicator for whether the town had at least one bank before the crisis), the OLS model is:

$$f_i = \beta_0 + \beta_1 bf_i + X_i'\theta + \varepsilon_i$$

The relationship, measured by β_1, is positive as expected, but not statistically significant after controlling for pre-period bankruptcies. Ex ante, my main concern with the OLS results was the possible omitted variable bias inherent in regressing firm bankruptcies on bank failures, as many unobserved local economic shocks could push these variables in the same direction, biasing the estimates of the effects upwards. However, since the OLS results suggest that bank failures had little to no effect on local bankruptcies, the IV strategy can also be used to correct for measurement error in the bank failure variable that may have attenuated the OLS coefficient on bank failures toward zero. This correction (partially) addresses measurement error as long as the country bank's London agent is uncorrelated with the chance it is mis-reported in the London Directories.

Table 2.4: Ordinary Least Squares Estimates for Firm Bankruptcies Dec. 1, 1825-Jun. 30, 1826

	(1)	(2)	(3)	(4)
Bank failures	1.806^{+}	0.820	0.265	0.255
	[0.948]	[0.584]	[0.351]	[0.346]
Has bank	0.079	-0.669*	-0.349	-0.489*
	[0.412]	[0.295]	[0.212]	[0.224]
Population, 1821, thousands		0.320**	0.116**	0.129**
		[0.065]	[0.032]	[0.029]
Firm bankruptcies, pre-period			1.977**	1.942**
			[0.293]	[0.324]
County FE				Yes
Observations	616	616	616	616
R^2	0.029	0.568	0.786	0.813

Source: Post-Office London Directories, 1820-1830; Edinburgh Gazette; (Census of Great Britain 1821). Pre-period for firm bankruptcies is Dec. 1 1824-Jun. 30 1825. Firm bankruptcies exclude banks. Robust standard errors in parentheses. + p<0.1, * p<0.05, ** p<0.01.

2.4.2 Instrumental Variable Validity

Using country towns' exposure to the financial crisis of 1825 through their banks' connection with London agents that failed during the crisis as an instrument relies on several assumptions. This section discusses each condition that must be satisfied for the identification strategy to be valid and provides evidence in favor of the identification strategy. I find that interbank relationships were conduits of financial stress from large London banking houses to smaller country banks.[9]

Agent Failures Affected Country Banks

The first condition is that the failure of a country bank's London agent had a material effect on the chance that the country bank itself would fail, that is, the relevance of the instrument. Evidence from the banking network data shows that bank-agent relationships in the English banking system during the 1820s were sticky: even 10 years later, 76% of country banks that survived until 1830 had the same London agent in 1830 as they did in 1820. This suggests that switching London agents likely involved some cost that country banks were unwilling to pay, and that problems at the London bank would therefore be transmitted to the country bank.

Even if such relationships were sticky, it is still not clear a priori that an agent failure would put financial stress on their country bank clients; it could be that London balances and transactions were an unimportant part of a country bank's balance sheet, in which case the instrument would be weak. Pressnell (1956) uses surviving bank balance sheet data to argue that the London account was the best-managed and most important part of a typical country bank's balance sheet and the first resource in times of liquidity crisis.

[9] Empirical studies of propagation are rare, as data on interbank exposures have been difficult to obtain, according to Iyer and Peydro (2015). While I also lack data on balance sheet exposures, the environment of 1825 is well-suited to studying propagation in a network with few connections since 98% of country banks had just one London agent.

Moreover, narrative evidence about the Panic suggests a second channel was at play during this crisis that did not exist in other financial crises studied in this literature. Following the failure of the first London bank on December 12, 1825, "That Monday saw runs upon banks known to be correspondents of Pole's, upon other banks in the same towns, and upon banks in nearby towns." (Pressnell (1956) p. 486). Thus, the failure of a London agent could affect its correspondents through a news channel even if the financial impact of its failure on its correspondents was small. As Duffy (1973) writes, "the failure of a London bank could, by arousing panic in provincial areas, cause the stoppage one after the other of banks which were completely solvent" (p. 249). Even if country banks could costlessly switch to another London agent, financial contagion spread on rumor and relationships as much as actual solvency concerns.

To demonstrate that agent failures during the Panic of 1825 had material effects on their country bank clients Table 2.5 shows bank branch-level probit models for bank failures for 561 English bank branches (includes only banks in towns where population data is available).[10][11] Across all specifications, agent bankruptcy has the expected positive association with bank bankruptcy. It seems that banks with multiple branches were more likely to fail than unit banks, which is somewhat surprising given

[10] I treat 13 bank branches with 2 London agents as separate branches in the regressions.

[11] Linear probability models that avoid the incidental parameter problem for the regressions including county fixed effects show qualitatively similar results and are available upon request.

that these banks were better positioned to insure themselves across space. Banks that had been founded more recently were much more likely to fail than older banks.

Table 2.5: Probit Models for Bank Failure

	(1)	(2)	(3)	(4)
Bankruptcy				
Agent bankruptcy	0.181	0.234	0.383*	0.389*
	[0.150]	[0.151]	[0.158]	[0.178]
Number of other branches		0.075*	0.059+	0.039
		[0.032]	[0.032]	[0.041]
Founded 1821-1825		0.333*	0.290*	0.481**
		[0.135]	[0.137]	[0.153]
CB bankruptcies in same city			0.161+	-0.208
			[0.091]	[0.128]
Agent's number of clients			-0.003	0.002
			[0.009]	[0.009]
CB bankruptcies of same agent			0.032+	0.020
			[0.019]	[0.020]
County FE				Yes
Observations	561	561	561	506

Source: Post-Office London Directories, 1820-1830. Robust standard errors in parentheses. + p<0.1, * p<0.05, ** p<0.01.

Column 4 provides evidence that the financial crisis was propagated not just through agent failures but through two additional channels. First, there is the within-town contagion effect: banks in towns where other banks failed were more likely to fail themselves, though this effect disappears once I control for county fixed effects, perhaps suggesting that contagion occurred at the county level rather than at the town level.[12] There is also evidence of the within-agent contagion effect: failures of other country banks' connected to a bank's London agent increased that bank's probability of failure, conditional on whether the agent itself failed or not. See Calomiris, Jaremski, and Wheelock (2019) for a similar study of propagation in correspondent banking networks in the U.S.

The magnitude of the effect of an agent failure is large. Using the model in column 3 of Table 2.5 at the mean of the other covariates, branches whose London agent did not fail had a 19% chance of failing during the crisis while banks whose agent failed had a 31% chance. Given this, plus the finding that there was a within-town contagion effect, the first stage regression of town bank failures on town exposure to London agent failures, reported in Section 2.4.3, is expected to be a strong instrument.

No Selection on London Agents

A second assumption necessary for instrumental variable validity is that banks with London agents that failed were not systematically different from other banks. Irresponsible, insolvent country bankers who were more likely to fail ex ante may have chosen irresponsible London agents who were also more likely to fail, creating an upward endogeneity bias in the previous results. Not much is known about how agents were cho-

[12] Certain counties experienced no bank failures, so including county fixed effects omits banks in those counties, thus decreasing the number of observations.

sen. Pressnell argues that the choice of a particular London banker was affected largely by the nature of the business of the country banker and of his clients, but family ties also played a role. The fact that relationships were so sticky, as already demonstrated, makes it unlikely that more savvy banks were able to foresee and avert risks related to which London agent they used.[13]

Table 2.6 compares banks with London agents who failed during the crisis to those whose agent survived and shows few differences. Banks in these two groups were equally likely to have more than one London agent, be founded in the last five years, had roughly the same number of competitors in their town, and their towns had roughly the same number of firm bankruptcies in the pre-crisis period. The only statistically significant differences between the two groups are in the number of bank branches, with exposed banks having fewer bank branches on average, and exposed banks tending to be in towns with lower populations. In the regression analysis I control for town size to account for this difference.

Historical evidence also supports the fact that the exposed banks were no more risky than other banks ex ante. Many country bank failures during the 1825 panic were caused by illiquidity rather than insolvency. By 1828 23 out of 63 banks that declared bankruptcy during the crisis had resumed payment, and records from the same year show that an additional

[13] I also show in a placebo test in Section 2.5 that having a London agent that failed in 1825 does not predict bank failure between 1823-1825.

31 of these 63 were still attempting to resume operation (Pressnell (1956) p. 491).[14] Still, payment stoppages that were successfully resolved several years later could have large consequences in the short run, as discussed later in this section.

Table 2.6: Comparison of Exposed vs. Not Exposed Banks, 1825

	Exposed	Not Exposed	Difference	
	Mean	Mean	Diff.	t-stat
Number of bank branches	1.49	2.00	-0.51**	-3.85
Has more than one agent	0.04	0.05	-0.01	-0.40
Founded 1821-1825	0.21	0.25	-0.04	-0.94
Number of banks in town	2.19	2.45	-0.27	-1.42
Population, 1821, thousands	9.25	14.10	-4.85**	-2.65
Firm bankruptcies, pre-per.	1.18	1.49	-0.32	-0.77
Observations	102	459	561	

Source: Post-Office London Directories, 1820-1830; Edinburgh Gazette; (Census of Great Britain 1821). Pre-period for firm bankruptcies is Dec. 1 1824-Jun. 30 1825. Firm bankruptcies exclude banks. Standard errors in parentheses. + $p<0.1$, * $p<0.05$, ** $p<0.01$.

Exclusion Restriction

The final requirement for the instrument to be valid is the exclusion restriction that failures of the London agents serving a town's country banks did not affect local economic conditions, especially firm bankruptcies, in any way other than through financial stress on the town's country banks. Lon-

[14] Surviving bankruptcy records from three of these banks show that two were solvent and the third was short only £6,000 on a debt of £71,000 (Pressnell (1956)).

don agents did occasionally lend to and take deposits from firms outside London, and may have been more likely to lend to firms in towns where they had a country bank client. Few balance sheets have survived to shed light on this concern. On the liability side, Duffy (1973) reproduces the claims of major claimants against Brickwood & Co., a London agent that failed in an earlier banking crisis in 1810. For this particular bank, with liabilities of £621,117, only 6% of those were owed to traders outside of London, and just three individuals made up these claims (p. 381). The country bank with the largest balance at Brickwood, Bowles Bank, accounted for about 20% of all outstanding claims on Brickwood and folded a few months later.

It turns out that the reduced form estimates of the town exposure instrument on the firm bankruptcies at the town level show a *negative* correlation between town exposure to agent failures and firm bankruptcies during the crisis after controlling for firm bankruptcies in the pre-period which is the strongest predictor of firm failures during the crisis (see Table 2.7). This suggests that the exclusion restriction is valid and the true channel is through country bank failures.

Table 2.7: Reduced Form Estimates for Firm Bankruptcies Dec. 1 1825-Jun. 30, 1826

	(1)	(2)	(3)	(4)
Exposed banks	0.978	0.620	-0.366	-0.327
	[1.123]	[0.656]	[0.296]	[0.289]
Has bank	0.424	-0.576^{+}	-0.144	-0.138
	[0.466]	[0.337]	[0.222]	[0.217]
Population, 1821, thousands		0.323**	0.115**	0.114**
		[0.064]	[0.032]	[0.031]
Firm bankruptcies, pre-period			2.005**	2.000**
			[0.294]	[0.289]
County type FE				Yes
Observations	616	616	616	616
R^2	0.010	0.566	0.787	0.792

Source: Post-Office London Directories, 1820-1830; Edinburgh Gazette; (Census of Great Britain 1821). Pre-period for firm bankruptcies is Dec. 1 1824-Jun. 30 1825. Firm bankruptcies exclude banks. Robust standard errors in parentheses. + p<0.1, * p<0.05, ** p<0.01.

2.4.3 Instrumental Variable Results

Because the endogenous bank failure variable is a count variable ranging from zero to three, the appropriate first stage is a count model. To account for this, I follow the three stage IV method for non-linear first stages described by Wooldridge (2002) (p. 623) which I summarize below. Doing so avoids using the "forbidden regression" with an endogenous count variable and a linear second stage. In section 2.5 I show the results carry through in a garden-variety two stage least squares (2SLS) setup.

The drawback of that approach is significantly worse model fit and less explanatory power of the exposed banks instrument when trying to fit a linear model to count data for bank failures.

Under Wooldridge's approach, the usual 2SLS standard errors and test statistics are asymptotically valid if the standard IV assumptions hold. If the first stage is correctly specified (that bank failures follow a Poisson distribution in my case[15]) and the errors are homoskedastic (a strong assumption), the estimator is asymptotically efficient in the class of IV estimators.

The estimation procedure is as follows:

Step one, count model for bank failures at the town level:

$$bf_i = \exp(\delta_0 + \delta_1 eb_i + X_i' \boldsymbol{\delta} + \ln(tb_i + 1) + \eta_i) \tag{2.1}$$

where eb_i is exposed banks (number of country banks in town whose London agents failed during the crisis), tb_i is total banks, and X_i is a vector of town controls. η_i is a random error term. Note that I control for a town's exposure to the possibility of bank failures using the total number of banks in the town. Since many towns had no banks but did have firm failures, I use total banks plus one as the control.

I report first stage results with different sets of controls in Appendix

[15]I also estimated a negative binomial model but the estimated over-dispersion parameter α was very close to 0, suggesting Poisson is an appropriate fit. Estimating a zero-inflated Poisson model to account for the fact that many towns in this sample will have no bank failures because they had no banks to begin with also doesn't change the results much. For more on these models see Long and Freese (2014).

B.2. Here I discuss the results with no controls (column 1 in Table B.1). The relationship between the number of exposed banks and the number of bank failures is positive and significant at the 1% level. Interestingly, the incidence rate ratio $e^{\delta_1} = e^{0.393} \approx 1.481$, meaning that having an additional exposed bank increases the rate of bank failures by more than a factor of one. This may capture the contagion effect of runs not just on clients of failed London agents but on other banks in town. Another way to interpret the results is using predicted counts. "Exposed banks" ranges from 0 to 3, and the mean predicted bank failures (conditional on total banks) for each value of exposed banks is 0.15, 0.37, 0.66, and 2.1, respectively.

Step two, first-stage IV regression using predicted bank failures $\widehat{bf}_i$ from the previous step as an instrument for actual failures:

$$bf_i^{IV} = \gamma_0 + \gamma_1 \widehat{bf}_i + X_i'\boldsymbol{\gamma} + \nu_i \tag{2.2}$$

Obtain predicted bank failures from this regression, denoted $\widehat{\widehat{bf}}_i$.

Step three, second-stage IV regression for firm failures at the town level:

$$f_i = \beta_0 + \beta_1 \widehat{\widehat{bf}}_i + X_i'\boldsymbol{\beta} + \varepsilon_i \tag{2.3}$$

as in the OLS model I expect β_1 to be positive, measuring the cost of a contraction in financial intermediation for firm survival.

Table 2.8: Instrumental Variable Estimates for Firm Bankruptcies Dec. 1 1825-Jun. 30, 1826

	(1)	(2)	(3)
Bank failures	6.939**	2.291*	0.981+
	[1.645]	[1.091]	[0.510]
Has bank	-2.781**	-1.053**	-0.749**
	[0.495]	[0.337]	[0.199]
Population, 1821, thousands	0.294**	0.115**	0.128**
	[0.067]	[0.031]	[0.028]
Firm bankruptcies, pre-period		1.904**	1.935**
		[0.321]	[0.316]
1st Stage F-Stat	59.75	46.12	111.54
County FE			Yes
Observations	616	616	616

Source: Post-Office London Directories, 1820-1830; Edinburgh Gazette; (Census of Great Britain 1821). Pre-period for firm bankruptcies is Dec. 1 1824-Jun. 30 1825. Firm bankruptcies exclude banks. Robust standard errors in parentheses. + p<0.1, * p<0.05, ** p<0.01.

The main results are shown in Table 2.8.[16] I find that in my preferred specification controlling for local economic conditions using firm bankruptcies in the pre-period and town size using population data, each bank failure causes about 2.3 additional firm failures (column 2 of Table 2.8). This effect is large: about half a sample standard deviation (5.2) in the town bankruptcy variable, and about one and a half times the sample mean (1.6).

[16]The first-stage Poisson regressions with each set of controls can be found in Appendix B.2. The F-statistic reported in Table 2.8 is the F-statistic for the regression of actual bank failures on the predicted failures and controls. That is, from equation (2.2).

Controlling for pre-crisis bankruptcies seems important given potential omitted variable bias: poor local economic conditions before the crisis could be correlated both with pre-crisis bankruptcies and with bank failures. Or, for local institutional reasons, firm bankruptcies may have been more likely at all times in some towns than in others regardless of size and pre-crisis bankruptcies can capture this difference across towns. The inclusion of pre-period bankruptcies reduces the estimated number of firm bankruptcies due to a bank failure from about 7 to 2.3 and I prefer this more conservative estimate due to the concerns above. Between the second two models, I prefer the model without county fixed effects. The average number of towns per county in my sample is just 15 so there probably is not enough power to identify the effects using variation within counties.

The number of firm failures caused by each bank failure might appear modest, but the discussion in Section 2.3 about using bankruptcies as a barometer of economic activity overall suggests that towns with bank failures were likely to be suffering greater output losses than towns without failures. Moreover, I find that having a local bank played a mitigating role in the crisis overall: holding the number of bank failures fixed, towns with at least one bank had fewer firm failures than towns without a bank (coefficient on "Has bank" is negative). This could suggest that banks made better quality loans to firms in their own towns, since monitoring was easier, or that surviving banks provided liquidity in the form of country bank notes that helped their towns weather the crisis better than towns without easy access to a local means of payment.

To understand the aggregate implications of bank failures during the crisis, I perform a simple counterfactual analysis using the results in Table 2.8. Taking the estimated coefficients from equation 2.3, I instead assume a world with no bank failures, setting the value of the bank failures variable to zero for all towns. Comparing this predicted value to the model-predicted value with actual bank failures, firm failures rise by 27% less when bank failures are zero.

2.4.4 Discussion

The above results show a causal relationship between bank failures and firm failures in the short run, suggesting a role for banks in promoting economic activity in the towns they served. In this section I describe two possible ways for bank failures to cause firm failures.

Aggregate Demand

Towns outside London did not usually use Bank of England notes until after 1826, and instead relied on country bank notes.[17] Day laborers' wages were usually paid with country bank notes and these notes subsequently circulated through purchases of local goods and services. Failures of the

[17] Testimony given to Parliament by the Governor of the Bank of England in 1832 suggested that the public had preferred private notes issued by a banker they knew and trusted to notes issued by the Bank of England that were subject to forgery to a greater degree than country notes (Great Britain (1832)).

banks backing these notes could wipe out the household wealth of noteholders as well as wealthier depositors. For example, the failure of the bank Turner, Turner, & Morris, was said to have caused "much alarm and difficulty among the middling and lower orders, as the circulation of their notes was very great" (The Examiner (1825a)).

Whether the value of bank notes issued by bankrupt bankers was wiped out entirely or whether they continued to circulate at a fraction of their face value is unclear and seems to have varied. There were some cases where a particular local merchant would accept bank notes of a defunct bank in exchange for goods at a fraction of their face value, hoping to recover some of the value in bankruptcy proceedings according to parliamentary testimony given by (Pease (1848)).[18] At other times, small noteholders themselves were forced to participate in bankruptcy proceedings to recover the value. Banks that failed in 1825 eventually paid an average of 85% on their obligations (a figure of 17 shillings on the pound, worth 20 shillings, was cited by Pease (1848).) However, this dividend on the bankrupt's estate was paid out several years after payment was stopped (Duffy (1973)). Whether the noteholder received a large fraction of the value in some years' time, a smaller fraction immediately, or nothing at all, all three constitute a drop in household money balances in the short

[18] A notice posted in the Northampton Mercury (1826) asserted that anyone who purchased or accepted bank notes of already bankrupt banks had no legal claim to recover the value in bankruptcy proceedings, though the legality of this behavior one way or the other is unclear.

run.

Recent work by Chodorow-Reich et al. (2018) studies the effect of India's demonetization on economic activity at the local level and finds that local cash scarcity generated substantial output losses. In their model this happens because of a cash-in-advance constraint on a portion of household consumption. Because of this constraint, a negative shock to money balances directly translates to a negative local demand shock as households are forced to cut consumption and the market-clearing level of output drops.[19] Such a constraint likely held for English households in 1825 as well considering the many credit market imperfections (lack of effective enforcement and legal protection for creditors, costly monitoring, information asymmetries) that existed at the time. In a similar setting to the one examined here, recent work by Kenny and Turner (2019) uses narrative evidence to argue that money supply shocks from bank failures in Ireland's banking panic of 1820 resulted in similar contractions in economic activity.

Chodorow-Reich et al. (2018)'s model of demonetization and other models of local aggregate demand shocks (for example Mian and Sufi (2014)) predict that firms producing non-tradables will be differentially more affected by local aggregate demand shocks than firms producing tradable

[19]The size of differences in output losses between locations in their model depends both on the size of local money balance shocks and on the size of the non-tradable sector. With a larger non-tradable sector there is less opportunity to smooth idiosyncratic money balance shocks across space. Moreover, in this model, cumulating the estimated local effects like I do in the previous section to obtain the 27% figure is a lower bound on the true aggregate effect because of trade between locations.

goods. This is because tradable goods may be exported to areas that are less affected and thus the prices of these goods will fall by less than the prices of non-tradable goods and services.

I use the preferred IV model to investigate whether non-tradable firms did indeed experience higher bankruptcy rates than tradables firms during the Panic of 1825.[20] As expected, I find no statistically significant effect of bank failures on tradable firm failures in Table 2.9, while non-tradable industries were significantly affected. Each bank failure results in 3.5 non-tradable firm bankruptcies. The results are suggestive that a drop in local aggregate demand was an important channel and the magnitude of the difference in the effects on the two sectors is very large, but the results are not conclusive since the coefficients are statistically indistinguishable.

Table 2.9: IV Estimates for Bankruptcies, Tradables vs. Non-tradables Dec. 1, 1825-Jun. 30, 1826

	T	NT	Other
Bank failures	-1.266	3.546*	0.011
	[1.635]	[1.562]	[0.305]
Has bank	0.223	-1.246*	-0.030
	[0.528]	[0.507]	[0.106]
Population, 1821, thousands	0.091**	-0.004	0.028^{+}
	[0.026]	[0.014]	[0.015]
Firm bankruptcies, pre-period	0.787**	1.041**	0.076
	[0.289]	[0.118]	[0.088]
1st Stage F-Stat	46.12	46.12	46.12
Observations	616	616	616

Source: Post-Office London Directories, 1820-1830; Edinburgh Gazette; (Census

[20]The first stage results are the same as column 2 of Table B.1 in Appendix B.2.

of Great Britain 1821). Pre-period for firm bankruptcies is Dec. 1 1824-Jun. 30 1825. Firm bankruptcies exclude banks. Robust standard errors in parentheses. + $p<0.1$, * $p<0.05$, ** $p<0.01$.

Credit Supply Shock

The devaluation of existing bank notes held by households was just part of a larger drop in the availability of country bank notes that were the primary means of payment in British towns. Banks supplied their bank notes to firms as a form of working capital in exchange for longer term promissory notes. Pressnell writes that "the bankrupted banks represented a reduction of the means of payment and an immobilization of much capital. The survivors contracted their lending... contraction enforced by caution was reinforced by reduced confidence in the ordinary banks" (p. 491). Thus, much like the recent financial crisis, the Panic of 1825 was characterized by a contraction in bank lending to firms. James, McAndrews, and Weiman (2013), studying similar disruptions to the payment system in correspondent banking networks in the U.S., argue that these stoppages act as severe adverse supply shocks, mainly by preventing firms from being able to make payroll. They find that payment stoppages by New York banks (analogous to London banks) from 1866-1914 were associated with declines in real activity of 10-20%.

Figure 2.4 approximates the contraction in short term lending by showing the volume of *new* bank notes issued, differentiating small and large

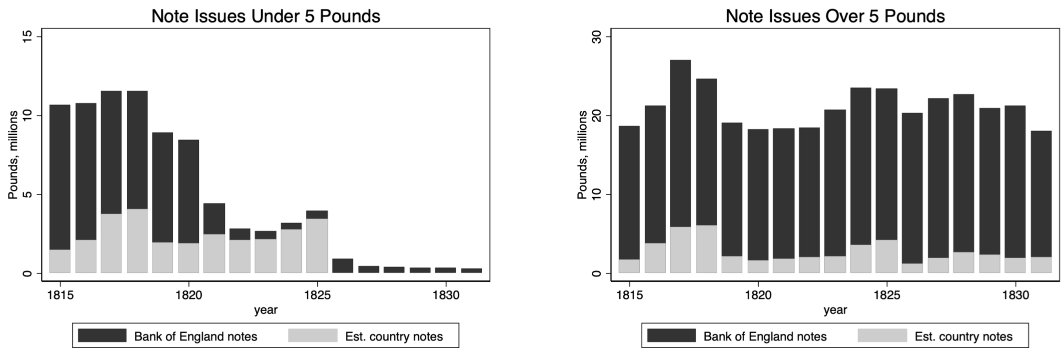

Figure 2.4: Source: Report from the Committee of Secrecy on the Bank of England Charter, Appendix No. 99 (stamp duties) and Appendix No. 82 (Bank of England notes). I estimate the volume of country bank notes using tax rates and the amount of stamp duties collected.

denominated bills. Small denominated bills were commonly used to pay workers' wages.[21] The Bank of England responded to the credit crunch by issuing its own notes[22], but not enough to completely offset the contraction in private notes. The fact that small private bank note issues never recovered was due to the 1826 banking reform requiring that notes under £5 had to be withdrawn by 1829. Unfortunately, no data on note issues at the local level is available so it's not possible to identify the causal effect of credit contractions on firm failures during this period.

The above discussion provides support for the arguments of Crouzet

[21]The cash in advance model of Sargent and Velde (2002) provides an example of how the composition of the money supply can play a role in determining output. Smaller denominations provide greater liquidity services in the model.

[22]The Examiner (1825a) quotes a local Birmingham paper: "the failure of the house of Smith and Gibbins created a good deal of local inconvenience from the quantity of their paper which was in circulation...It appears that the issue of £1 Bank of England notes in Birmingham, has been very considerable, and by no means unwelcome."

(1972), Pollard (1964), and Mathias (1973) that country banks did indeed contribute, at least modestly, to industrialization, but largely not through long-term lending for fixed capital formation, as Brunt (2006) argued more recently. The above reexamination of the historical evidence suggests that country banks provided liquidity to provincial economies and greased the wheels of nascent factory systems by supplying a means of payment through working capital loans.[23] This is why I can detect effects of country bank failures in the very short run, six month period during and after the crisis. If long term lending were the main driver, so many failures of large firms would likely not have occurred so rapidly.

2.5 Robustness

This section explores various robustness checks for the results already presented. I check that the firm bankruptcy results hold in the subsample of towns with at least one bank. I test an alternative instrument that uses only exposure to London agents directly connected to the Latin American debt crisis. I use a placebo test to show that the instrument only predicts country bank failures during the crisis period. Finally, I show that the results hold up to assuming a linear model rather than a Poisson model in the first stage.

One concern with the main IV results is that I pool towns with and

[23] According to Crouzet (1972) "short-term credit to finance increases in inventories was quantitatively by far the largest need of industry" during the industrial revolution.

without banks together. While these towns experienced bankruptcies (hence their inclusion in the dataset), recording zero bank failures for these towns is misleading in the sense that they had no bank to begin with. Other unobserved differences between towns with and without banks, such as financial integration with London or level and type of of economic activity, may correlate with bankruptcies and bias the results in one way or another. So, in this first robustness check, I consider the subsample of towns with at least one bank at the beginning of the crisis, leaving 328 towns total. The results in Table 2.10 suggest that the main IV results broadly carry through to the subsample of towns with banks, though the effect of bank failures is estimated to be smaller than in the main results. Precision is a concern here since the sample size has been reduced substantially. Since the results are mainly driven by non-tradables, I check the effect of bank failures on non-tradables only in the bank town subsample in columns 4-6. Here I find similar results to the full sample.

Table 2.10: IV Estimates for Firm Bankruptcies in Bank Towns Only

	All bankruptcies			Non-tradables only		
	(1)	(2)	(3)	(4)	(5)	(6)
Bank failures	4.254*	0.699	0.371	5.282*	3.028*	0.402
	[1.661]	[0.908]	[0.509]	[2.399]	[1.466]	[0.369]
Pop, 1821, 1000s	0.399**	0.179**	0.178**	0.122**	-0.021	0.002
	[0.074]	[0.039]	[0.039]	[0.033]	[0.023]	[0.019]
Firm bankr., pre		1.901**	1.965**		1.231**	1.097**
		[0.335]	[0.367]		[0.175]	[0.147]
1st Stage F-Stat	52.21	42.73	51.14	52.21	42.73	51.14
County FE			Yes			Yes
Observations	328	328	328	328	328	328

Source: Post-Office London Directories, 1820-1830; Edinburgh Gazette; (Census of Great Britain 1821). Pre-period for firm bankruptcies is Dec. 1 1824-Jun. 30 1825. Firm bankruptcies exclude banks. Robust standard errors in parentheses. + $p<0.1$, * $p<0.05$, ** $p<0.01$.

As already mentioned, only three of the London banks that failed during the crisis were directly connected with the Latin American debt crisis through their payee role, though it is likely that the others that failed were exposed through holding these securities on their balance sheet. Using the failure of the other banks ignores reverse causality concerns that country bank failures may have caused their London agent to fail, overstating the strength of the instrument and potentially violating the exclusion restriction. One can see in Table 2.11 that the F-statistics are still quite large and only slightly smaller than the baseline F-statistics. The main IV results are not at all sensitive to using this alternative instrument.

Table 2.11: IV Estimates With LA Agents Only

	(1)	(2)	(3)
Bank failures	7.357**	2.694*	1.020*
	[1.866]	[1.166]	[0.517]
Has bank	-2.925**	-1.193**	-0.763**
	[0.559]	[0.347]	[0.198]
Population, 1821, thousands	0.292**	0.115**	0.127**
	[0.066]	[0.031]	[0.028]
Firm bankruptcies, pre-period		1.889**	1.934**
		[0.322]	[0.316]
1st Stage F-Stat	45.50	39.44	115.75
County FE			Yes
Observations	616	616	616

Source: Post-Office London Directories, 1820-1830; Edinburgh Gazette; (Census of Great Britain 1821). Pre-period for firm bankruptcies is Dec. 1 1824-Jun. 30 1825. Firm bankruptcies exclude banks. Robust standard errors in parentheses. + $p<0.1$, * $p<0.05$, ** $p<0.01$.

Next, to further check the validity of the instrument and ensure that there were no systematic differences between London agents that did and did not survive the 1825 Panic, I use a placebo test. The probit models reported in Table 2.12 show that agent failure during the 1825 crisis did not predict earlier county bank failures between 1823 and 1825, after controlling for agent characteristics like the total number of clients the agent had and the agent's total number of correspondent failures over the same period. Interestingly, in normal times (1823-1825), the failure of one of the bank's local competitors reduced the probability it would fail, presumably because this expanded its business prospects. However, during the crisis the within-town contagion effect dominated and reversed the sign on "CB bankruptcies in same city." This is a novel finding relative to the analysis of Calomiris, Jaremski, and Wheelock (2019) who find a consistently negative effect of the failure of other banks in town during the Great Depression in the U.S., perhaps because their analysis covers a protracted period with elevated bank run risk, whereas my sample covers both the boom and bust parts of the credit cycle.

Table 2.12: Placebo Probit Models for Bank Failure, 1823-1825

	(1)	(2)	(3)	(4)
Bankruptcy, 1823-1825				
Agent bankruptcy, 1825	0.261	0.331^{+}	0.294	0.220
	[0.179]	[0.185]	[0.188]	[0.208]
Number of other branches		0.007	0.018	0.023
		[0.043]	[0.042]	[0.052]
Founded 1821-1823		0.687**	0.698**	0.750**
		[0.167]	[0.169]	[0.186]
CB bankr. in same city, 1823-1825			-0.124	-0.462*
			[0.201]	[0.191]
Agent's number of clients, 1823			-0.002	-0.000
			[0.007]	[0.007]
CB bankr. of same agent, 1823-1825			-0.028	-0.032
			[0.045]	[0.050]
County FE				Yes
Observations	578	578	578	409

Source: Post-Office London Directories, 1820-1830. Robust standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Finally, some may be skeptical of the non-standard three-step procedure I employ for the main IV results and prefer a standard two-stage linear IV approach. The main problem with this approach is that a count model fits the bank failure data much better, as discussed in Section 2.4.3. Using the town's number of exposed banks as an instrument in a linear first stage is weak, whereas it is a strong and statistically significant instrument when using the count model. Nevertheless, the main results are largely unchanged by using the standard two-stage linear setup (Table 2.13).

Table 2.13: IV Estimates: 2SLS

	(1)	(2)	(3)
Bank failures	7.371**	3.219**	3.174**
	[1.945]	[1.159]	[1.105]
Has bank	-2.930**	-1.376**	-1.537**
	[0.577]	[0.329]	[0.349]
Population, 1821, thousands	0.292**	0.115**	0.122**
	[0.065]	[0.030]	[0.027]
Firm bankruptcies, pre-period		1.870**	1.912**
		[0.318]	[0.330]
1st Stage F-Stat	21.69	17.21	15.79
County FE			Yes
Observations	616	616	616

Source: Post-Office London Directories, 1820-1830; Edinburgh Gazette; (Census of Great Britain 1821). Pre-period for firm bankruptcies is Dec. 1 1824-Jun. 30 1825. Firm bankruptcies exclude banks. Robust standard errors in parentheses. + p<0.1, * p<0.05, ** p<0.01.

2.6 Conclusion

This paper contributes to two strands of literature. The first is on the output costs of banking crises. Recent studies have generally focused on firms directly connected with financially stressed banks and have not attempted to identify spillovers to local demand. I instead focus on the role of banks in the payment system and show that payment suspensions and bank failures affected local aggregate demand directly during the Panic of 1825.

In general, the effects I find demonstrate that the first modern financial crisis in Britain looked somewhat like financial crises in the twenty-first century, but with important institutional caveats. In particular I argue that the destruction of household wealth when private bank notes lost value was likely the most important channel of transmission from financial shocks to the real economy, but classical features like a contraction in bank lending and a loss of market liquidity for previously safe assets may have also been important.

These findings update arguments in the second strand of literature on the importance of finance in the industrial revolution in England. The rapid spread of bank failures to bankruptcies of non-financial firms suggests that banking services, particularly the means of payment they provided to households and firms, were important for the normal functioning of local economies at a short-term frequency. This point may be useful for understanding the potential consequences of rapid demonetizations like

India's in 2016 as well as the potential costs of disruptions to cryptocurrency payment systems as these currencies become more widely used.

Finally, I have found that integration in the form of the correspondent banking network played a critical role in transmitting financial stress induced by the Latin American debt crisis in 1825 to provincial economies. This finding is consistent with models of failures in networks of interdependent financial organizations (see Elliott, Golub, and Jackson (2014), for example). These models predict non-monotonic effects of financial integration: integration initially allows contagion to travel farther, but eventually reduces individual organizations' exposure to their own idiosyncratic shocks. For example Scotland, with its more mature and well-integrated banking system, experienced much milder effects of the Panic of 1825 compared to England.[24] Geographic integration via correspondent banking may well have had positive effects in normal times; for example I find that towns with banks but no bank failures weathered the Panic with fewer non-financial firm failures than towns without banks. Policy reforms in England in response to the Panic allowed banks to grow larger and expanded branch banking significantly, eventually enabling banks to more effectively smooth idiosyncratic local shocks.

[24]Calomiris and Haber (2014) provide a comparison of the two countries' banking systems during this period. Scotland allowed joint-stock banking and Scottish banks were much larger than English country banks.

Chapter 3

This Time It's Different: The Role of Women's Employment in a Pandemic Recession

with Titan Alon, Matthias Doepke & Michèle Tertilt

Abstract

In recent US recessions, employment losses have been much larger for men than for women. Yet, in the economic downturn caused by the Covid-19 pandemic the opposite is true: women's employment declined much more than men's. Why does a pandemic recession have a disproportionate impact on women's employment, and what are the wider repercussions of this phenomenon? We argue that more women lost jobs because their employment is concentrated in contact-intensive sectors such as restaurants and because increased childcare needs during school and daycare closures prevented many from working. We analyze the macroeconomic implications of women's employment losses using a model that features heterogeneity in gender, marital status, childcare needs, and human capital. A pandemic recession is qualitatively different from a regular recession because women's labor supply behaves differently than men's. Specifically, our quantitative analysis shows that a pandemic recession features a stronger transmission from employment to aggregate demand and results in a persistent widening of the gender wage gap. Many of the negative repercussions of a pandemic recession can be averted by prioritizing opening schools and daycare centers during the recovery.

3.1 Introduction

Economic fluctuations display a number of regularities, such as comovement of output across sectors and higher volatility in aggregate investment than in aggregate consumption. These observations motivated Robert Lucas to famously claim that "business cycles are all alike" Lucas (1977), and business cycle theory has been devoted to accounting for these regularities ever since.

As a consequence of the Covid-19 pandemic, in 2020 the United States and other countries entered the sharpest contraction in economic activity since the Great Depression. While this contraction displays some of the regularities of other economic downturns, in other ways it is unlike any other in recent history. Understanding the differences between regular and pandemic recessions is important both to further our understanding of what the recovery from the current downturn will look like, and to inform policy responses to possible pandemic recessions in the future.

In this paper, we show that a crucial difference between regular recessions and the current downturn lies in the role of women's employment. In recent recessions preceding the current crisis, men were more severely affected by employment losses. This disproportionate impact was particularly pronounced in the Great Recession that followed the financial crisis of 2007–2008, which gave wide currency to the term "mancession" for this and earlier downturns.

The Covid-19 recession reverses the usual pattern. This time, it is women, rather than men, who have experienced larger employment losses and higher unemployment. Figure 3.1 displays the difference between the rise in women's and men's unemployment in every recession in the United States since 1948.[1] Unlike in all previous recessions, women's unemployment has risen much more than that of men during the current recession—a 2.9 percentage point gap between February and April of 2020. Moreover, the gap in the rise in unemployment is larger in absolute value than during any other recession; there never has been a recession that has affected workers of one gender so much more than the other.

Our analysis aims to answer two questions about the role of women's employment in a pandemic recession. First, why is it that women's employment has declined much more than men's in the current downturn? Second, what are the wider repercussions of the fact that this pandemic recession is a "shecession" rather than a "mancession"?

Our answer to the first question builds on the observation that the current recession has in large part been triggered by "lockdowns," i.e., the social distancing measures, business shutdowns, and stay-at-home orders implemented during the pandemic. The need for social distancing impacted different sectors of the economy than those usually hardest hit in

[1]For pre-Covid-19 recessions, we use the difference in the seasonally adjusted unemployment rate between the first and last months of each recession based on recession dates from the NBER Business Cycle Dating Committee. For the current recession, we use the difference between unemployment in February 2020 (the trough in unemployment before lockdown measures were taken) and April 2020 (the peak in unemployment).

Figure 3.1: Difference between Rise in Women's and Men's Unemployment, US Recessions from 1948 to 2020

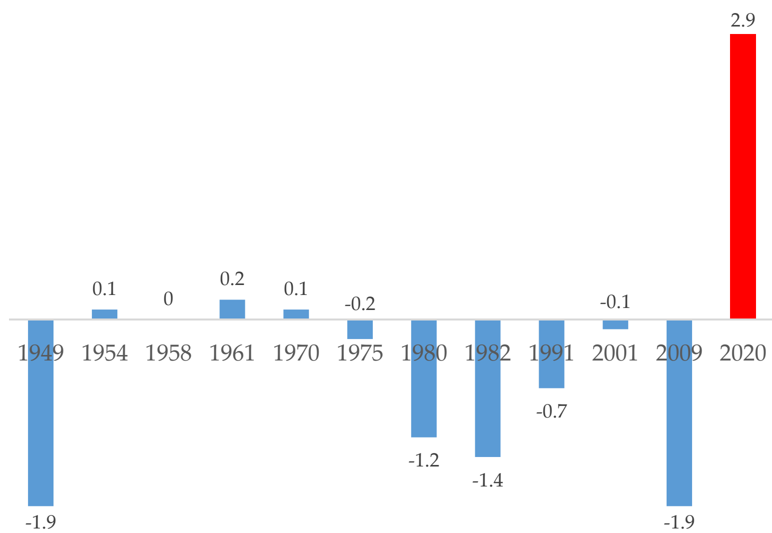

Notes: Data from Bureau of Labor Statistics. Each bar is the rise in the women's unemployment rate minus the rise in the men's unemployment rate from the first to the last month of each recession according to NBER business cycle dates. For the Covid-19 recession, change in unemployment from February to April 2020 is displayed. The underlying series are seasonally adjusted monthly unemployment rates by gender.

recessions. Regular recessions lead to large employment losses especially in construction and manufacturing, both sectors with a high male employment share. In contrast, in the current recession employment losses have been greatest in high-contact service sectors such as restaurants, hospitality, and travel.[2] These are sectors where women represent a large share of the workforce.

[2]The largest increases in unemployment have been in the "personal care and service" as well as "food preparation and serving" occupations, with June unemployment rates close to 30 percent in both cases. See *BLS Labor Force Statistics from the Current Population Survey*, Table A-30.

Another set of social distancing measures consisted of closures of schools and daycare centers, usually for a much longer period than business shutdowns. With children at home, parents had to spend more time on childcare, which reduced time available for market work and led to additional employment losses beyond layoffs. As Dingel, Patterson, and Vavra (2020) show, a third of the US workforce has children aged 14 or younger in the household, implying that the employment effects from school and daycare closures are potentially large. We argue that women's employment is more strongly affected than men's by increased childcare needs. There are many more single mothers than single fathers. Among couples raising children together, before the crisis women spent a lot more time on childcare than men, even if both mother and father were working full time. Survey evidence shows that during the crisis this lopsided division of childcare was sustained, implying that more women than men have been unable to work because of childcare obligations.

Our second question regarding the wider repercussions of the impact of the pandemic recession on women's employment hinges on gender asymmetries. Is a shecession the same as mancession, just with signs reversed, or does it make a qualitative difference which gender is more affected by a downturn? We argue that a shecession is indeed qualitatively different from a mancession, because of the different roles women's and men's labor supply play in many families. In married couples, women are more likely than men to be the secondary earner, and their labor supply is more elastic than that of men. One consequence of married women's more elastic labor supply is that their earnings can serve as a shock absorber when men lose employment in a typical recession. When a husband becomes unemployed, married women become more attached to the labor force and may increase their labor supply on the intensive margin (see Bardóczy 2020). In a shecession, married men cannot provide the same kind of insurance for women's job losses, because most have little room to increase their labor supply. For this reason, the reaction of aggregate labor supply is qualitatively different depending on whether a negative shock to employment is initially concentrated on women or men.

Differences in women's and men's flexibility of labor supply also imply qualitatively different long-run repercussions of a recession. Women's labor supply is responsive to within-family relative wages. If women's future earnings prospects decline because of job loss in a recession, they become more likely to further specialize in non-market work, reducing future earnings even more. In contrast, when men become unemployed in a recession they usually continue to seek future full-time employment.

Building on these insights, we assess the repercussions of regular versus pandemic recessions with the help of a quantitative model of the household sector in the economy. The model features women and men, single and married households, households with and without kids, and workers who can telecommute and those who cannot. Households decide on consumption, labor supply, and savings, and households with kids have to decide on how to meet childcare needs. The labor market is subject to search frictions: workers may lose jobs and unemployed workers must wait for job offers. Workers who receive job offers decide whether to accept or reject the offer and, if they accept, whether to choose full-time or part-time work. The skills of employed workers increase over time due to returns to experience, whereas the skills of workers who are out of employment depreciate. The ability of workers to combine work with childcare responsibilities depends on their occupation: telecommuters have an easier time meeting childcare needs. The division of labor within the household is in part governed by a social norm: there is a fraction of "traditional" households that prefer that childcare be provided by the mother rather than the father.

We calibrate this model to the pre-pandemic US economy. Among other statistics, the calibrated model matches observed labor market flows, married women's labor supply, the division of childcare in dual-earner couples, estimates of returns to experience and skill loss in unemployment, and the gender wage gap. By reproducing how joint decision-making in families generates the different structure of women's and men's labor supply, the model captures the gender asymmetries that are the root cause of qualitative differences between shecessions and mancessions.

We use the quantitative model to compare the repercussions of a regular recession and those of a pandemic recession. We model regular recessions as a temporary shift in job destruction rates and job finding probabilities, calibrated to capture the large impact of regular recessions on men's employment. In contrast, a pandemic recession has an equally large direct impact on women's and men's employment, and also brings about an increase in parents' childcare needs, which generates additional employment losses as some parents reduce labor supply to look after their children.

A first finding from the model is that due to endogenous decisions on the allocation of childcare within households, a pandemic recession lowers women's employment much more than men's employment. The pattern of household specialization persists in the recession, implying that mothers shoulder the majority of the increased childcare load and consequently bear more of the employment consequences.

A second finding is that the transmission of income shocks to consumption is qualitatively different in a pandemic versus a regular recession. A pandemic recession has a large impact on households with children, whose marginal propensities to consume (MPCs) differ from those of average households. Single parents are especially limited in their ability to offset income shocks, meaning that such shocks have a large impact on their consumption. Distinct macro implications of regular and pandemic recessions also arise from the role of within-family insurance among married couples. In a regular mancession, many wives partially compensate for their husband's lost earnings by joining the labor force or working more. In aggregate terms, within-family insurance serves as a shock absorber that lowers the transmission of aggregate income shocks to aggregate consumption. In contrast, we find that within-family insurance is more limited during a pandemic recession. Men have a more limited ability to compensate for a job loss of their spouse because their attachment to the labor force is already high. Moreover, in families with children, increased childcare needs during the recession limit the ability of secondary earners to increase labor supply. The loss of within-family insurance together with the large impact on single-parent households with high MPCs imply a stronger transmission from income to consumption in a pandemic recession. To the extent that aggregate demand partly determines output, this finding results in a greater amplification of the initial shock and thus a deeper recession and a delayed recovery.[3]

The third finding from our quantitative analysis is that a pandemic recession has long-run repercussions for gender inequality in the labor market. Workers who lose employment lose skills. Given that regular recessions are mancessions, they primarily reduce men's skills and therefore moderately reduce the gender wage gap, consistent with the evidence in Solon, Barsky, and Parker (1994). In contrast, a pandemic recession depreciates the skills of women who reduce their hours or drop out of the labor force, leading to a substantial widening of the wage gap. A qualitative

[3] In our analysis, we focus on the household sector and do not spell out such an aggregate demand channel explicitly. It would be straightforward to include such a channel by adding a production sector subject to nominal frictions, as in, e.g., Hagedorn, Manovski, and Mitman (2019).

difference between regular and pandemic recessions is that the impact of a pandemic recession on the gender wage gap is much more persistent. Women who lose employment during a pandemic recession become more likely to leave the labor force permanently or to only seek part-time work. In contrast, most men who lose employment in a regular recession ultimately return to full-time work. In our baseline model, it takes almost twenty years until women's relative wages return to their previous level after a pandemic recession.

We also account for forces that may ultimately reduce gender gaps in the labor market. Inspired by evidence that working from home is here to stay, we impose that the pandemic permanently increases the fraction of flexible jobs that allow telecommuting.[4] We also allow the pandemic to have a persistent effect on social norms, based on evidence from "daddy months" showing that short-term changes in fathers' involvement in childcare lead to a more equal division of childcare in the long-term.[5] Our model indicates that fathers, even though they do less than mothers, still substantially increase the time they spend on childcare during a pandemic recession.[6] We conjecture that this sudden change will gradually increase the share of "modern" couples with gender-equal social norms. In our baseline model, these changes generate a long-run rise in

[4]See Barrero, Bloom, and Davis (2020) and Appendix C.3.1.

[5]See Appendix C.3.2 for a description of the evidence.

[6]Quantitatively, we find as schools close, the fraction of fathers who do any childcare increases from 53 to 76 percent and that the number of married couples in which the husband is the primary childcare provider rises by 2 percentage points.

women's labor force participation and a decline in the gender wage gap, with increased job flexibility and changing social norms each accounting for about half of the change. However, it does take a long time for these effects to dominate the direct impact on women's skills: after the gender wage gap reaches its previous level after 20 years, it takes an additional 20 years for the the gender wage gap to shrink by two percentage points.

Our results on how a pandemic recession is different from a regular recession are important for guiding policy. Higher MPCs during a pandemic recession imply that fiscal policy will be more effective compared to a regular recession, especially if directed at families with children.[7] We also show that school and daycare reopenings, if they can be safely done, can have a sizeable impact on the recovery. Interestingly, even though childcare needs are larger for small children, reopening schools has a larger effect on the economy than reopening daycare centers. The reason is that a larger percentage of the workforce has school-age children, and that these parents are more likely to work full time than those of smaller children.

Related Literature

Our work contributes to the literature on the role of women's employment in economic fluctuations. In December 2019, women accounted for the majority of the US labor force for the first time, capping a decades-

[7]During the lockdown period consumption demand may be restricted for other reasons, such as the impossibility to travel or go to the mall. Higher MPCs will then emerge after the lockdown ends.

long convergence between male and female employment. Yet, for a long time most business cycle models have been "unisex" models that do not allow for gender differences, while many macroeconomic studies of labor supply have been calibrated to data on men's employment only. More recently, studies such as Albanesi (2020) and Fukui, Nakamura, and Steinsson (2019) argue that the role of women in aggregate fluctuations has changed substantially over time due to rising female labor force participation. Albanesi (2020) provides evidence that women's employment plays a crucial role in phenomena such as jobless recoveries, the productivity slowdown, and the great moderation. Bardóczy (2020) argues that joint household decision-making is an important determinant of the transmission of macroeconomic shocks. Other contributions to the literature on women's employment and household decision-making within macroeconomics include Greenwood, Seshadri, and Yorukoglu (2005), Ortigueira and Siassi (2013), Doepke and Tertilt (2016), Mankart and Oikonomou (2017), Borella, De Nardi, and Yang (2018), Mennuni (2019), Olsson (2019), and Wang (2019).[8] In addition, Albanesi and Şahin (2018) and Coskun and Dalgic (2020) note the impact that the gender breakdown of employment in various industries has on the contrasting cyclicality of male and female employment, which is a key element of how we model the impact of regular recessions.

[8]Macroeconomic studies of the policy implications of joint household decisions include Guner, Kaygusuz, and Ventura (2012), Guner, Kaygusuz, and Ventura (2020), Bick (2016), and Krueger and Wu (2019).

One of the central mechanisms in our theory is within-family insurance of job loss and income shocks. In the labor literature, Lundberg (1985) introduced the notion of the "added worker effect," i.e., a worker joining the labor force in response to their spouse's job loss. More recent studies supporting the important role of within-family insurance include Attanasio, Low, and Sánchez-Marcos (2005), Blundell, Pistaferri, and Saporta-Eksten (2016, 2018), Birinci (2019), García-Pérez and Rendon (2020), Pruitt and Turner (2020), and Guner, Kulikova, and Valladares-Esteban (2020). Meanwhile Guler, Guvenen, and Violante (2012) and Pilossoph and Wee (2020) analyze the impact of within-family insurance on job searches. Ellieroth (2019) uses a joint-search model similar to our setting to characterize the quantitative importance of within-household insurance over the business cycle. Unlike existing search models with within-family insurance, our model allows for the accumulation and depreciation of human capital, incorporates single and married households, accounts for childcare needs, and allows for different occupations and social norms. All of these features play a central role in our analysis.

Our analysis also contributes to a rapidly growing body of work on the macroeconomic consequences of the Covid-19 recession. Much of this literature combines epidemiological and economic modeling to examine how policy interventions and endogenous behavioral adjustments shape the evolution of the pandemic and its macroeconomic consequences (see Eichenbaum, Rebelo, and Trabandt 2020, Berger, Herkenhoff, and Mongey 2020, Glover et al. 2020, and Brotherhood et al. 2020, among others). Our paper departs from such studies as it does not model the pandemic explicitly, but rather focuses on the economic consequences of the employment losses and increased childcare needs brought about by the pandemic.[9] In this regard, our approach is more similar to Guerrieri et al. (2020), Gregory, Menzio, and Wiczer (2020), and Danieli and Olmstead-Rumsey (2020), who also focus on the macroeconomic transmission of the lockdown shock in models that abstract from epidemiology. These papers focus on different mechanisms than our study, namely the role of incomplete markets and liquidity constraints, employment stability, and the sectoral distribution of the downturn. Hence, our focus on the differential impacts on women and men provides a novel contribution to this literature.

3.2 Why the Role of Gender is Different in Pandemic Recessions

The social distancing measures and stay-at-home orders imposed in many US states and other countries during the Covid-19 crisis have resulted in a drop in employment, a rise in unemployment, and an economic contraction. In this section, we discuss why this pandemic recession differs from

[9] The pandemic itself also has a gender dimension, as men appear to be at higher risk of death than women. However, to date vastly more people are affected by the economic repercussions of the pandemic than by Covid-19 itself.

earlier recessions in its implications for women's versus men's employment.

3.2.1 Gender Differences in Regular Recessions

In recent economic downturns preceding the current crisis, including the Great Recession of 2007–2009, men's employment was affected more strongly than women's. Doepke and Tertilt (2016) summarize the evidence on how employment varies over the business cycle for women and men. Table 3.1 shows that women's aggregate labor supply is less volatile overall than men's, as measured by the percentage standard deviation of the Hodrick-Prescott residual of average labor supply per person. For cyclical volatility, i.e., the component of overall volatility that is correlated with aggregate economic fluctuations, the gap between women and men is even larger. Over the period 1989–2014, men account for more than three quarters of overall cyclical fluctuations in employment, and women for less than one quarter.

One reason why women's employment varies less over the cycle is insurance within the family, i.e., some married women increase their labor supply in a recession to compensate for their husband's unemployment or higher unemployment risk.[10] An indication of the importance of this channel is that the cyclical volatility of labor supply illustrated in Table 3.1

[10] See Ellieroth (2019) for a study documenting the quantitative importance of this mechanism.

Table 3.1: Volatility of Hours Worked by Gender and Marital Status

	All			Married		Single	
	Total	Women	Men	Women	Men	Women	Men
Total Vol.	1.15	0.87	1.47	0.79	1.16	1.30	2.25
Cyclical Vol.	0.91	0.51	1.23	0.38	0.95	0.70	1.82
Hours Share		42.64	57.36	25.89	39.83	16.75	17.53
Vol. Share		23.68	76.32	10.80	41.51	12.88	34.81

Notes: All data from Current Population Survey, March and Annual Social and Economic Supplements, 1989 to 2014. Total volatility is the percentage standard deviation of the Hodrick-Prescott residual of average labor supply per person in each group. Cyclical volatility is the percentage deviation of the predicted value of a regression of the HP-residual on the HP-residual of GDP per capita. Hours share is share of each component in total hours. Volatility share is share of each group in the cyclical volatility of total hours. See Doepke and Tertilt (2016) for further details.

is much lower for married women (to whom the family insurance channel applies) than for single women.

Additional channels also contribute to differences in the volatility of women's and men's labor supply. This is apparent from the large volatility gap between single women and single men, to whom the within-family insurance channel does not apply. The second crucial channel is the different sectoral composition of female and male employment. In typical recessions, sectors such as manufacturing and residential construction are more severely affected compared to, say, education and health care. Men's employment is more concentrated in sectors with a high cyclical exposure, whereas women are more represented in sectors with relatively stable employment over the cycle. These facts are documented in a recent paper by Coskun and Dalgic (2020). The authors find that employment in the "Government" and "Education and Health Services" sectors is actually countercyclical. These two sectors account for 40 percent of women's employment, but only 20 percent of men's employment. Conversely, the highly cyclical sectors of "Manufacturing," "Construction," and "Trade, Transportation, Utilities" account for 46 percent of male employment but only 24 percent of female employment.

These two channels are neither exhaustive nor independent—for example, some women may choose to work in a countercyclical sector to compensate for their husbands' cyclical employment risk. But the bottom line is clear: past downturns have affected men's employment more severely than women's.

3.2.2 Why a Pandemic Recession is Different

In Alon et al. (2020a), we predicted that unlike a regular recession, the current pandemic recession would reduce women's employment more than men's employment. This prediction, which has since been confirmed by the evidence, was based on two channels. The first consists of the impact of social distancing measures in a pandemic across sectors and occupations. To quantify this channel, in Alon et al. (2020a) we combine data from the American Community Survey (ACS), the American Time Use Survey (ATUS), and the Current Population Survey (CPS) to rank occupations by the ability to work from home (meaning that work during the lockdown is possible) and by whether an occupation is critical during the lockdown (such as healthcare workers). We document that women are underrepresented in the occupations with the highest ability to telecommute and in critical occupations, implying that women's employment has a stronger exposure to the pandemic recession shock.

The second channel is increased childcare needs due to closures of schools and daycare centers. This channel is further amplified by the reduced availability of other means of childcare provision, such as from relatives, neighbors, nannies, or babysitters, during a lockdown with minimal social contact. To quantify the childcare channel, in Alon et al. (2020a) we combine CPS and ATUS data to document that women provide a much larger share of overall childcare than men. There are many more single mothers than single fathers, and many more married mothers than fathers who work part-time or are a stay-at-home parent with their spouse working full-time. Even among married parents who both work full time, mothers provide about 40 percent more childcare than fathers.[11] Taken together, these observations suggest that women will end up shouldering most of the increased childcare needs during a pandemic recession, and thus face reduced opportunities for employment.[12]

Since the onset of the current recession, a number of studies have provided additional evidence on the importance of these channels. Mongey, Pilossoph, and Weinberg (2020) use O*NET data on occupational charac-

[11]The gap between women's and men's provision of childcare is even larger during regular working hours (8 a.m. to 6 p.m. on weekdays; see Schoonbroodt 2018).

[12]Women provide the majority of childcare in all industrialized countries, though there is considerable variation between countries in the gap between women's and men's contributions (Doepke and Kindermann 2019).

teristics to examine the burden of social distancing policies based on the ability to work from home and a measure of physical proximity at work in different occupations. In contrast to the time-use data used by Alon et al. 2020a, they find that women are more likely to be able to work from home, but that they are also over-represented in occupations requiring physical proximity. Combining these factors, the authors expect the overall impact on women's and men's employment to be similar, and hence qualitatively different from regular recessions in which the most adversely affected occupations have a higher share of male employment. Albanesi et al. (2020) also examine the gender breakdown in employment between occupations that are high and low in personal contact, and find that women account for 74 percent of employment in high-contact occupations.

Dingel, Patterson, and Vavra (2020) quantify the extent to which childcare obligations will hold back the recovery. Based on ACS data, they document that 32 percent of the US workforce has a child under the age of 14 in their household, and that two-thirds of these households do not include an adult who is out of the labor force (e.g., a stay-at-home parent). In 30 percent of households with children, all offspring are under the age of 6, meaning that these households will be relieved of additional childcare needs when daycare centers reopen. These numbers underscore that childcare obligations have been a major driver of reduced employment during the recession, and that a strong recovery will not be possible until these needs are met.

To assess the implications of this key distinction between regular and pandemic recessions for macroeconomic dynamics, gender inequality, and welfare, we now turn to our macroeconomic model.

3.3 A Model to Assess the Wider Repercussions of a Pandemic Recession

Our quantitative model focuses on the household side of an economy with search frictions. Macroeconomic shocks affect households primarily through changes in job-loss and job-finding probabilities. In our analysis, we take the impact of aggregate shocks on these labor-market variables as given, and focus on the question of how the household sector will respond in terms of labor supply, consumption demand, and the accumulation of skills.[13]

3.3.1 Demographics and State Variables

The economy is populated by a continuum of three types of households: single women, single men, and couples. Every period, a new cohort of singles and couples enters the economy. The household type is permanent. Singles and couples face a constant probability ω of death. Couples stay

[13]It would be conceptually straightforward to expand towards a full general equilibrium analysis by modeling job creation and destruction by firms in the usual way and, if desired, adding additional features such as nominal rigidities.

together and die together, and hence there are no widows, widowers, or divorcees in the economy.

The state variables of a household include assets/savings a and the labor market productivity h of each (adult) household member. Additional discrete state variables are kids $k \in \{0, s, b\}$ (no kids, small kid, big kid), employment of each member $e \in \{E, U\}$ (employed or unemployed), and the occupation of each household member $o \in \{TC, NT\}$ (can telecommute or cannot). The unemployed state $e = U$ in the model corresponds to both unemployment and being out of the labor force in the data. For couples, a final state variable is a social norm $m \in \{0, 1\}$ where $m = 0$ denotes a "traditional" social norm that values a within-household division of labor in which the mother provides the majority of childcare, whereas a couple with $m = 1$ has the "modern" view that no childcare arrangement is inherently superior.[14] The aggregate state variable for the economy is denoted by X, which captures whether the economy is or is not currently in a recession.

New singles and couples start out with zero assets. The initial human capital levels for singles are drawn from gender-specific distributions $F^g(h)$ and for couples from the joint distribution $F(h^f, h^m)$. The initial probability of each occupation and each social norm is given by the stationary distribution over these states implied by the current aggregate

[14]One indication for the relevance of social norms is that men raising children in same-sex couples provide more childcare than men in different-sex couples do (Prickett, Martin-Storey, and Crosnoe 2015).

state. Singles or couples may already have a small or large child when they enter the economy. The probabilities of having a job offer in the initial period are identical to the offer probabilities for an unemployed individual of the same gender.

After the initial period, the level of assets is determined by a household's consumption-savings decision. Labor market productivity evolves as a function of shocks and labor supply. Employment status and occupation type evolve as a function of shocks—individuals can get laid off, and finding a job in a particular occupation is random. People can also decide to reject a job offer or quit a job. Labor supply (conditional on having a job) is either part-time or full-time, chosen by the worker.

For singles, the transition probabilities for kids are given by $\pi^g(k'|k)$, and for couples these probabilities are given by $\Pi(k'|k)$. The transition probabilities for employment are given by $\pi^g(e'|e, X)$. Naturally, employment transition probabilities depend on the aggregate state X, which captures that in a recession jobs are easier to lose and harder to find. The transition probabilities also depend on the current employment state e and gender g. The employment state e' at the beginning of the next period denotes whether the worker receives a job offer. If a job offer is received, the worker can still decide whether to accept the offer and, if so, whether to work full-time or part-time. The transition probabilities for human capital $\pi(h'|h, n)$ are independent of gender and only depend on current human capital h and labor supply n. People also face constant probabilities of switching occupations and social norms, given by $\pi(o'|o, X)$ and $\pi(m'|m, X)$.

3.3.2 The Decision Problem for Singles

We use v to denote the value functions of singles, while V denotes the value functions of couples. Similarly, $\tilde{v}$ and $\tilde{V}$ denote the value functions at the beginning of the period before job offers are accepted or rejected. The value function for an employed single is given by:

$$v_E^g(a, h, k, o, X) = \max_{a', c, l, n, t} \left\{ u^g(c, l) + \omega \beta E \left[\tilde{v}_{e'}^g(a', h', k', o', X') \right] \right\}.$$

Here β is the time discount factor, c denotes consumption, l denotes leisure, $n \in \{0, 0.5, 1\}$ is labor supply (part time or full time), and t is time spent on childcare. The period utility function is given by:

$$u^g(c, l) = \log(c) + \alpha^g \log(l).$$

We allow leisure preference to depend on gender to facilitate matching labor supply to the data. The social norm does not apply to singles because it only affects the time allocation of couples. The constraints for employed singles are as follows:

$$c + a' = w^g h n^\theta + (1 + r)a,$$

$$t + \phi(k)\, n\, I(o = TC) \geq \gamma(k, X),$$

$$l + n + t = T.$$

The first constraint is the budget constraint. The parameter $\theta > 0$ allows for increasing or decreasing returns in labor supply. For example, part-time workers (who supply half as much labor as full-time workers) may be less than half as productive because of commuting time, or more than half as productive because workers get tired. The second constraint is the childcare constraint, which says that total childcare time has to be at least as large as the childcare need $\gamma(k, X)$, where $\gamma(s, X) > \gamma(b, X) > \gamma(0, X) = 0$. The term $\phi(k)\, nI(o = TC)$ reflects the fact that in a telecommuting job ($o = TC$), fraction $\phi(k)$ of work time can be used to simultaneously provide childcare. Intuitively, workers with TC jobs can supervise a child at home while still getting some work done, and they do not have to take an entire day off of work if a child is sick at home. This matters a lot when childcare requirements rise during a pandemic recession. The ability of a worker in a TC occupation to simultaneously work and provide childcare depends on the age of the child. Specifically, a younger child requires more full-time attention than does an older child. The remaining childcare time is denoted as t. The final constraint is the time constraint, where T is the time endowment.

The value function and constraints for unemployed singles are:

$$v_U^g(a,h,k,o,X) = \max_{a',c,l,t} \left\{ u^g(c,l) + \omega\beta E\left[\tilde{v}_{e'}^g(a',h',k',o',X')\right]\right\}.$$

$$c + a' = zw^g h + (1+r)a,$$

$$t = \gamma(k,X),$$

$$l + t = T.$$

Here z denotes the unemployment benefit replacement rate relative to potential productivity $w^g h$. Notice that even when unemployed, occupation o is defined, because the current occupation defines the probability distribution of receiving job offers in each possible occupation.

The value function at the beginning of the period for a single with a job offer is:

$$\tilde{v}_E^g(a,h,k,o,X) = \max\left\{ v_E^g(a,h,k,o,X), v_U^g(a,h,k,o,X)\right\}.$$

Without a job offer, there is no choice to be made, so we have:

$$\tilde{v}_U^g(a,h,k,o,X) = v_U^g(a,h,k,o,X).$$

3.3.3 The Decision Problem for Couples

We now turn to married households. The overall structure of the decision problem is the same as for singles. The spouses act cooperatively with bargaining weights λ for the wife and $1-\lambda$ for the husband. Here, the household decision problem also reflects the role of the social norm. If $m=0$ (the traditional social norm applies), the household suffers a utility loss of ψ per unit of time if the father provides more childcare than the mother, and a utility benefit if she does more. The value function for two working spouses is given by:

$$V_{EE}(a, h^f, h^m, k, o^f, o^m, m, X) = \max\left\{\lambda u^f(c^f, l^f) + (1-\lambda)u^m(c^m, l^m)\right.$$
$$\left.-(1-m)\psi(t^m - t^f) + \omega\beta E\left[\tilde{V}_{(e^f)',(e^m)'}(a', (h^f)', (h^m)', k, (o^f)', (o^m)', m', X')\right]\right\}.$$

The budget and time constraints are:

$$c^f + c^m + a' = w^f h^f (n^f)^\theta + w^m h^m (n^m)^\theta + (1+r)a,$$
$$t^f + t^m + \phi(k)\left(n^f I(o^f = TC) + n^m I(o^m = TC)\right) = \gamma(k, X),$$
$$l^f + n^f + t^f = T, \tag{3.1}$$
$$l^m + n^m + t^m = T. \tag{3.2}$$

If only the woman has a job, the decision problem is:

$$V_{EU}(a, h^f, h^m, k, o^f, o^m, m, X) = \max \big\{ \lambda u^f(c^f, l^f) + (1-\lambda) u^m(c^m, l^m)$$
$$-(1-m)\psi(t^m - t^f) + \omega\beta E \left[\tilde{V}_{(e^f)', (e^m)'}(a', (h^f)', (h^m)', k, (o^f)', (o^m)', m', X') \right] \big\}$$

subject to (3.1) and:

$$c^f + c^m + a' = w^f h^f (n^f)^\theta + z w^m h^m + (1+r)a,$$
$$t^f + t^m + \phi(k)\, n^f I(o^f = TC) \geq \gamma(k, X),$$
$$l^m + t^m = T.$$

The reverse case is analogous. If both are unemployed, the decision problem is:

$$V_{UU}(a, h^f, h^m, k, o^f, o^m, m, X) = \max \big\{ \lambda u^f(c^f, l^f) + (1-\lambda) u^m(c^m, l^m)$$
$$-(1-m)\psi(t^m - t^f) + \omega\beta E \left[\tilde{V}_{(e^f)', (e^m)'}(a', (h^f)', (h^m)', k, (o^f)', (o^m)', m', X') \right] \big\}$$

subject to (3.1), (3.2), and:

$$c^f + c^m + a' = z(w^f h^f + w^m h^m) + (1+r)a,$$
$$t^f + t^m = \gamma(k, X).$$

At the beginning of the period, if both spouses have a job offer, we get:

$$\tilde{V}_{EE}(a, h^f, h^m, k, o^f, o^m, m, X) = \max\big\{V_{EE}(a, h^f, h^m, k, o^f, o^m, m, X),$$
$$V_{EU}(a, h^f, h^m, k, o^f, o^m, m, X), V_{UE}(a, h^f, h^m, k, o^f, o^m, m, X),$$
$$V_{UU}(a, h^f, h^m, k, o^f, o^m, m, X)\big\}.$$

The initial value functions for the other permutations are analogous.

3.3.4 The Stochastic Process for Labor Productivity

Human capital h evolves as a function of shocks and captures both random shocks to productivity and the returns to experience. There is a finite grid $h \in H = \{h_1, h_2, \ldots, h_I\}$ of possible human capital levels, where the ratio of subsequent points is constant, i.e., $\log(h_{i+1}) - \log(h_i)$ is constant across i. There are returns to experience to working full time, meaning that full-time workers upgrade to the next human capital level with a fixed probability η:

$$\pi(h_{i+1}|h_i, 1) = \eta, \quad \pi(h_i|h_i, 1) = 1 - \eta.$$

Individuals who do not work face possible skill depreciation with probability δ:

$$\pi(h_{i-1}|h_i, 0) = \delta, \quad \pi(h_i|h_i, 0) = 1 - \delta.$$

The human capital of part-time workers is constant: $\pi(h_i|h_i, 0.5) = 1$.

3.3.5 The Aggregate State

The aggregate state X takes four possible values: $X \in \{N, NN, R, P\}$. Here N denotes normal times, before a recession hits. R denotes a regular recession, modeled as a large decline in job-finding probabilities and large rise in job-loss probabilities for men and smaller changes in the same direction for women, with unchanged childcare requirements. P denotes a pandemic recession, where there are considerable changes in labor market flows for both men and women, as well as a large increase in childcare requirements. Finally, NN denotes the "new normal," or the state of the economy after a pandemic recession is over. This state allows us to model the consequences of permanent transformations brought about by a pandemic, such as a rise in the share of TC jobs and a shift in social norms.

The transition matrix between these four states is parameterized as follows:

$$\pi(S'|S) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1-\rho_R & 0 & \rho_R & 0 \\ 0 & 1-\rho_P & 0 & \rho_P \end{pmatrix}.$$

Note that the N and NN states are absorbing; in either the normal or the new-normal state, people expect to stay in that state forever. Recessions therefore arise as unexpected "MIT shocks" (this could be easily generalized). Once in a regular R recession, the economy returns to normal N with probability $1 - \rho_R$. If in a pandemic P recession, then the economy switches to the new normal with probability $1 - \rho_P$.

3.3.6 The Stochastic Processes for Occupation and Social Norm

The transition probabilities of occupation and the social norm depend only on the state variable itself and on the aggregate state. Hence, the transition probabilities for occupation are given by numbers $\pi(o'|o, X)$, while the transition probabilities for the social norm are denoted as $\pi(m'|m, X)$. The dependence on the aggregate state captures the possibility that a pandemic recession can promote the spread of TC jobs and the modern social norm. The transition matrix for $o \in \{TC, NT\}$ is given by:

$$\pi(o'|o, X) = \begin{pmatrix} \rho_{TC}(X) & 1 - \rho_{TC}(X) \\ 1 - \rho_{NT}(X) & \rho_{NT}(X) \end{pmatrix}$$

and similarly for the social norm $m \in \{0, 1\}$ we have:

$$\pi(m'|m, X) = \begin{pmatrix} \rho_0(X) & 1 - \rho_0(X) \\ 1 - \rho_1(X) & \rho_1(X) \end{pmatrix}.$$

For both transition matrices, we assume that one set of coefficients applies to the aggregate states N and R (normal and regular recession), and a second set of coefficients applies to the states P and NN (pandemic recession and new normal).

3.4 Calibrating the Model to Normal Times and Recessions

We aim to quantify the impact of regular versus pandemic recessions on different households and on the aggregate behavior of the household sector. To this end, we first calibrate the normal state $X = N$ of the economy to match a number of characteristics of the US economy before the onset of the current pandemic recession. We then pin down the properties of regular versus pandemic recessions using data on earlier recessions and on the current one. Finally, we calibrate the new normal after a pandemic recession based on changes in telecommuting during the crisis and evidence on the dynamics of social norms.

3.4.1 Externally Calibrated Parameters

The model economy operates at a quarterly frequency. Newly born people in the model correspond to singles and couples at age 25 in the data. A number of model parameters directly correspond to specific empirical observations and can be pinned down individually. These parameters are listed in Table 3.2. The survival probability ω determines life expectancy in the model. Given that we do not model retirement, we interpret the lifespan in the model as corresponding to active working life. As an increasing number of people retire starting around age 55 in the data, we set ω to match a life expectancy of 55 years.[15] We set the discount factor to $\beta = 0.98$ at a quarterly frequency. The discount factor primarily drives asset accumulation. In addition, because assets determine how financially constrained households are, the discount factor drives the distributions of marginal propensities to consume (MPCs) and save in the economy. Macroeconomic models have typically been calibrated to match overall asset accumulation in the economy, but a recent literature documents that such models imply counterfactually low MPCs (e.g., Kaplan and Violante 2014). We therefore choose a lower value than in typical macroeconomic calibrations, which in turn raises the average MPC of single and married households in the economy to a more realistic value.[16] The interest rate is set to $r = 0.02$, a relatively high value allowing for the fact that households are not compensated for accidental bequests left at their death. We normalize the time endowment to $T = 1.5$. Since we interpret a labor supply of $n = 1$ as a full-time job of 40 hours, this corresponds to a time endowment of 60 hours per week.[17] The childcare parameters $\gamma(s, N)$ and $\gamma(b, N)$ are calibrated based on information on time spent on childcare in families with younger and older children from the American Time Use

[15]Explicitly modeling retirement would primarily affect asset-accumulation decisions in the model. However, given that death is modeled as a shock, people still accumulate a substantial amount of assets and leave accidental bequests.

[16]Kaplan and Violante (2014) and Auclert, Bardóczy, and Rognlie (2020) report a quarterly MPC of about 0.25 for the US economy.

[17]We interpret our model as allocating fungible time during a typical weekday. Thus, we subtract sleep and personal care time and weekends to arrive at a time endowment of 60 hours per week.

Survey. The returns to experience parameter η is set to match a return to labor market experience of 1.1 percent per quarter, which is computed using the NLSY97 data set. The skill-depreciation parameter δ matches a quarterly depreciation of skills of 2.5 percent, which matches observations by Davis and von Wachter (2011) on the earnings implications of job loss during recessions. Further details on the calibration and the underlying data sources are provided in Appendix C.1.

Table 3.2: Externally Calibrated Parameters

Parameter	Value	Interpretation
ω	0.99	Expected retirement at age 60
β	0.98	Discount factor
r	0.02	Interest rate
T	1.5	Time endowment
$\gamma(s,N)$	0.34	Younger kids require 13.7 hours of childcare per week
$\gamma(b,N)$	0.11	Older kids require 4.2 hours of childcare per week
η	0.03	Return to labor market experience
δ	0.06	Skill depreciation in unemployment
ρ_{NT}	0.999	8.2% of pre-pandemic jobs are telecommuting

Notes: Hours are converted into fractions based on our assumptions that one unit of time corresponds to 40 hours per week.

In addition to the parameters listed in Table 3.2, we calibrate the initial distributions of human capital $F^g(h)$ and $F(h^f, h^m)$ to match evidence on the distribution of earnings of singles and couples at age 25 (see Appendix C.1.4). We match the transition probabilities for children $\pi^g(k'|k)$ and $\Pi(k'|k)$ with evidence on the distribution of different types of households (having younger children, older children, or neither; see Appendix C.1.3). The calibration yields a stationary distribution in which

59 percent of households are married, 51 percent are parents, 7 percent of households are single mothers, and 3 percent are single fathers. Among households with children, 45 percent have young kids under the age of six. Similarly, we initialize telecommuting status to match occupational patterns by gender and marital status observed in the data. Couples are jointly initialized so as to reflect the extent of occupational correlation between spouses (which, according to Malkov 2020, is quantitatively important for couples' exposure to risk in the current pandemic). Transitions between telecommuting and non-telecommuting jobs are then chosen such that the stationary equilibrium matches the prevailing level of telecommuters just before the pandemic, as documented in Bick and Blandin (2020).[18] The resulting fraction of telecommuters in the labor force is 12.9 percent. The share of telecommuters is substantially higher among married than single workers. Finally, we set the pre-pandemic share of married couples with traditional social norms to 30 percent, to match evidence from the General Social Survey. Appendix C.1 provides additional details on these parameter values and the data sources.

3.4.2 Jointly Calibrated Parameters

The remaining parameters are jointly calibrated to match a set of target moments that characterize the US economy before the onset of the cur-

[18]Specifically, we normalize the persistence of telecommuting jobs to 0.99, and choose the persistence of non-telecommuting jobs to match the target.

rent recession. Table 3.3 displays the calibrated parameter values, and Table 3.4 shows the model fit. Though the parameters are jointly chosen, in most cases there is a fairly direct mapping from a particular parameter to a particular moment.

Table 3.3: Jointly Calibrated Parameters

Description	Parameter	Value
Exogenous gender wage gap	w^f	0.91
Wife's bargaining power in married couples	λ	0.40
Diminishing returns to market work	θ	0.55
Women's leisure preference	α^f	0.64
Men's leisure preference	α^m	0.43
Telecommuters' childcare bonus for young kids	$\phi(s)$	0.07
Telecommuters' childcare bonus for older kids	$\phi(b)$	0.14
Job offer probability for employed women	$\pi^f(E\|E,N)$	0.93
Job offer probability for non-employed women	$\pi^f(E\|U,N)$	0.40
Job offer probability for employed men	$\pi^m(E\|E,N)$	0.93
Job offer probability for non-employed men	$\pi^m(E\|U,N)$	0.40
Utility cost of violating social norms	ψ	0.23

We normalize men's wage per efficiency unit of labor to one, $w^m = 1$. We then choose the exogenous part of the gender wage gap (women's wage per efficiency unit of labor w^f) to match an overall gender wage gap of 0.81 (see Appendix C.1.2 for details on how we compute this target). The resulting parameter is $w^f = 0.91$, implying that about half of the gender wage gap is due to this exogenous gap, with the remainder accounted for by differences in labor supply and in the accumulation of experience over the life cycle between women and men.

The parameters for leisure preference and for women's bargaining power primarily determine the distribution of labor supply across women and men and within couples. The social-norm parameter also helps match labor supply, as this parameter specifically affects the labor supply of married women with children. With regard to the childcare bonus for telecommuters, we impose that the bonus is twice as large for older compared to younger kids, based on the notion that older children require less supervision and therefore interfere less with working from home. The level of the childcare bonus for telecommuters is pinned down based on the observation that, in the ATUS data, men who telecommute do 50 percent more childcare than those who do not work from home (conditional on being married to women who do not telecommute, see Alon et al. 2020a). The returns to scale parameter θ for market work helps to match the breakdown between part-time and full-time work.

For labor-market flows, we impose that job-offer probabilities are identical for women and men in normal times. This assumption makes our results easier to interpret, in that it implies that gender differences in job flows in the model are entirely due to endogenous behavior (i.e., job-acceptance decisions) rather than hard-wired differences. Furthermore, as Table 3.4 shows, the observed job flows are still matched fairly well. The higher persistence in the model of non-employment for women compared to men arises because women reject more offers, primarily due to childcare obligations.

As Table 3.4 shows, the calibrated model matches the target moments

Table 3.4: Model Fit for Target Moments

	Data	Model
Gender wage gap	0.81	0.81
Childcare division, full-time couples, men-to-women	0.65	0.66
Men who telecommute do 50% more childcare	1.50	1.48
Relative labor supply, men-to-women	1.19	1.17
Labor supply of married women without kids	0.72	0.73
Labor supply of married women with younger kids	0.56	0.59
Labor supply of married women with older kids	0.64	0.70
Share of married mothers not employed	0.30	0.26
Share of married mothers working part-time	0.18	0.19
Share of married mothers working full-time	0.52	0.55
Women's Labor Market Flows: E-to-E	0.91	0.92
Women's Labor Market Flows: U-to-U	0.77	0.73
Men's Labor Market Flows: E-to-E	0.93	0.92
Men's Labor Market Flows: U-to-U	0.66	0.66

Notes: See Appendix C.1 for further details and data sources. Labor market state U here refers, as in the model, to all individuals who are either unemployed or out of the labor force. For telecommuters, childcare time in the model is computed as $t^g + 0.5\phi(k)\, n^g$, that is, time that is spend on childcare and work simultaneously is counted as 50 percent childcare. Counting all of the combined time as childcare leads to similar results.

well. Even though we use relatively few parameters to match these moments (nine degrees of freedom to match 14 moments), the model provides a good fit for the distribution of married women across employment states and for the impact of having children on women's labor supply. Generally, as in the data, women's labor supply in the model is more responsive to having children than is that of men. While the social norm does matter for traditional couples, the main driver behind specialization in childcare is wage differences between wives and husbands (as in Alon, Coskun, and Doepke 2020). The exogenous part of the gender wage gap implies that among a majority of couples, the wife is the secondary earner when the first child arrives, making it more likely that she will reduce her employment to meet childcare needs. As reducing employment means forgoing returns to labor market experience and potentially suffering skill loss, the within-couple wage gap will tend to grow, leading to even more childcare specialization as time passes.

3.4.3 Fit for Non-Targeted Moments

Table 3.5: Model Fit for Non-Targeted Moments

	Data	Model
Composition of single fathers by employment state:		
– not employed	0.16	0.15
– part-time	0.07	0.08
– full-time	0.77	0.77
Composition of married fathers by employment state:		
– not employed	0.07	0.19
– part-time	0.04	0.05
– full-time	0.89	0.75
Composition of single mothers by employment state:		
– not employed	0.24	0.15
– part-time	0.17	0.37
– full-time	0.59	0.48
Share of full-time dual earner couples by kids' age:		
– no kids	0.61	0.53
– younger kids	0.43	0.21
– older kids	0.49	0.47

Notes: See Appendix C.1 for further details and data sources for the data moments.

Table 3.5 shows how well the model performs in terms of matching a larger set of moments that were not explicitly targeted in the calibration.

While we focused on matching the overall women-to-men labor supply ratio and specific patterns of married women's labor supply in the calibration procedure, Table 3.5 shows that the model nevertheless matches the employment breakdown for men and single women fairly well (and remarkably well for single fathers). The model accounts for the observation that most married fathers work full time, and that single fathers are more likely to work than single mothers. Even though the model underpredicts the share of dual full-time earner couples couples with small children, it does capture the overall variation in this share with fertility, and matches well the fraction of dual full-time earners among couples with either older kids or without kids.

3.4.4 Modeling Regular versus Pandemic Recessions

The calibration described thus far pins down the economy in the normal state $X = N$, before a recession takes place. We now turn to the parameters that characterize the aggregate changes when the economy enters a regular recession R or a pandemic recession P. We impose that regular and pandemic recessions have the same expected duration of six quarters, i.e., $\rho_R = \rho_P = \frac{5}{6}$. We model the aggregate changes during recessions in a stylized way so as to allow for a transparent comparison of the different types of recessions. Specifically, to capture the larger impact of regular recessions on men's employment, we impose that in a regular downturn the job-offer probabilities for men are reduced twice as much as those for women. This scaling allows for a simple decomposition of which employment changes are due to shocks (i.e., job loss) versus changes in behavior (i.e., probability of accepting job offers). In a pandemic recession, we instead impose that both women and men experience the same change in job offer probabilities as men in a regular recession. The different impacts on women versus men are thus primarily accounted for by changing childcare obligations (which only occur in a pandemic recession) rather than hard-wired differences in job flows.

Table 3.6 summarizes all the parameter values that differ across aggregate states. The pandemic recession leads to a substantial increase in childcare obligations, from 13.7 to 42 hours per week for younger kids, and from 4.2 to 26 hours per week for older kids. The underlying assumption is that small children need near-constant supervision, meaning that the time cost of childcare is just as a large as working full time. While older kids require less time, there is still a large increase, in part due to the need to homeschool them. These values can be compared to the findings of Adams-Prassl et al. (2020b), who show that in a typical work week during the pandemic, US parents working from home spent roughly 22.5 (men) and 30 (women) hours doing childcare and homeschooling, for a total of 52.5 hours. Given that there are also single parents and married couples where only one parent works from home, the childcare burden in the model for younger kids roughly corresponds to the half-way point between the total childcare burden of 52.5 hours provided by a couple and the 30 hours a mother provides on her own during the pandemic.

Table 3.6: Parameters Varying across Aggregate States

Parameter	Interpretation	Normal N	Recession R	Pandemic P	New Norm. NN
$\gamma(s, X)$	Childcare time, younger kids	0.34	0.34	1.05	0.34
$\gamma(b, X)$	Childcare time, older kids	0.11	0.11	0.65	0.11
$\rho_1(X)$	Persistence modern norms	0.99	0.99	0.99	0.99
$\rho_0(X)$	Persistence traditional norms	0.98	0.98	0.94	0.94
$\rho_{TC}(X)$	Persistence TC occupations	0.99	0.99	0.99	0.99
$\rho_{NT}(X)$	Persistence NT occupations	0.999	0.999	0.996	0.996
$\pi^m(E\|E, X)$	Job offer, employed men	0.93	0.91	0.91	0.93
$\pi^m(E\|U, X)$	Job offer, unemployed men	0.40	0.38	0.38	0.40
$\pi^f(E\|E, X)$	Job offer, employed women	0.93	0.92	0.91	0.93
$\pi^f(E\|U, X)$	Job offer, unemployed women	0.40	0.39	0.38	0.40

The job offer probabilities during regular recessions were chosen to match employment flows during previous US recessions, as described in Appendix C.1.2 (see Table C.1). While this facilitates comparisons of regular and pandemic recessions in the model, it also means that our model somewhat understates the direct employment impact of the current pandemic recession (e.g., Kahn, Lange, and Wiczer 2020 report that there were 30 percent fewer vacancy postings in April 2020 than at the beginning of the year).

We allow for a one-time jump in the share of telecommutable jobs at the beginning of a pandemic recession, which captures the immediate rise in telecommuting at the beginning of the lockdown. Bick, Blandin, and Mertens (2020) report that in May 2020 more than 30 percent of the labor force worked from home, up from less than 10 percent in February. To match this increase, at the start of a pandemic recession, workers in NT occupations (who cannot telecommute) experience a one-time probability that their job switches to TC (telecommutable), where this probability is chosen to move the share of TC workers to 30 percent. After this one-time shock, the transition probabilities displayed in Table 3.6 apply, and the share of telecommuters remains at 30 percent throughout the pandemic.

Our model assumes that after a pandemic recession, rather than returning to its previous state, the economy approaches a new normal NN due to permanent changes brought about by the pandemic. We allow for such permanent effects along two dimensions: work organization and social norms. There is ample evidence by now that the "working-from-home experiment" caused by the pandemic has led to permanent changes in the organization of work. We therefore expect telecommuting to stay elevated in a post-Covid world. We summarize the existing evidence in Appendix C.3.1. We thus impose that the occupational transition probabilities during the pandemic recession continue to apply during the new normal. This implies that the fraction of telecommutable jobs will stay elevated, at about 30 percent.

With regard to social norms, we conjecture that the share of traditional couples will ultimately decline by half, from 30 to 15 percent. This is motivated by empirical evidence that short temporary changes in the division of labor in the household have lasting effects – not only on the families themselves but also on peers.[19] The transition probabilities that apply both

[19]See in particular Dahl, Loken, and Mogstad (2014) exploiting a paternity leave reform in Norway in 1993. Further evidence is summarized in Appendix C.3.2.

during the pandemic recession P and the new normal NN were chosen such that the modern state is highly persistent (0.99 probability of staying modern), and such that the persistence of the traditional state results in the desired long-run share of traditional couples of 15 percent. In addition, new cohorts also display these new long-run shares of 85 percent modern and 15 percent traditional couples.

Clearly, the future evolution of social norms is difficult to predict. Our calibration here should be regarded less as an empirical estimate and more as an "if-then" scenario. In other words, our simulations answer the question of how the economy will evolve if the current pandemic ends up having a substantial impact on the evolution of gender norms. Below, we also provide a decomposition analysis that examines different outcomes where social norms fail to respond. Still, in the past, gender norms have often evolved rapidly in response to economic changes (e.g., Fernández 2013 and Fogli and Veldkamp 2011). In our simulation, the change in social norms is slower than that implied by the learning model of Fernández (2013) during the rise of female labor force participation in the United States from the 1960s to the 1980s. The data already plainly show that the Covid-19 recession has led to a historically unprecedented increase in men's participation in childcare. Based on past experiences, we believe that such transformations are bound to have a substantial impact on social norms. Hence, while our assumptions on shifting social norms are necessarily more speculative than other aspects of our analysis, we believe a shift towards more gender-equal norms is the most likely scenario.

3.5 Regular versus Pandemic Recessions in the Quantitative Model

We now use our quantitative model to compare the consequences of regular versus pandemic recessions for macroeconomic aggregates and changes in gender inequality. We display outcomes for recessions that last for six quarters (the expected duration of a recession given $\rho_R = \rho_P = \frac{5}{6}$), and then revert to the normal state N in the case of a regular recession or the new normal NN in the case of a pandemic recession. A duration of six quarters places the end of the Covid-19 recession in the third quarter of 2021, which lines up with the expected wide availability of vaccines by the summer of 2021.[20] We start with an analysis of the division of childcare before studying the impact of the recession on labor supply and earnings.

3.5.1 Division of Childcare and Leisure during the Pandemic Recession

We find that mothers are more affected than fathers by the large increase in childcare needs during a pandemic recession. This can be seen in Fig-

[20]In some countries schools reopened in the fall of 2020, but in the United States many schools are likely to remain closed for the school year given persistently high infection rates.

ure 3.2, which compares the increase in childcare time during the pandemic for mothers and fathers. Naturally, the increase in childcare is largest for single parents, whose weekly time spent on childcare increases by about 23 hours. Single mothers and fathers react in a similar way. Among married couples, the increase in childcare hours is much larger for women than for men. This difference is due to endogenous decisions on the allocation of childcare within households, and is a crucial driver of the large impact of a pandemic recession on women's employment in the model. The model implications align well with empirical findings that both women and men are spending more time on childcare during the crisis, but that this increase is much larger for women (see, e.g., Adams-Prassl et al. 2020b for evidence on the US, UK, and Germany).[21]

The increase in childcare comes partly at the expense of reduced leisure. Among married couples, the reduction in leisure is larger for women than for men (see Figure C.3 in Appendix C.2.2). The reduction in leisure implies that women experience a larger welfare loss during the pandemic than men do (see Figure C.4 in Appendix C.2.3). This finding may help explain the observed increase in the gender gap in mental health during the pandemic.[22] Yet, reductions in leisure are only part of parents' reaction

[21]Researchers have documented that women are taking over the majority of increased childcare needs in a wide range of countries; see, e.g., Costoya et al. (2020) for evidence on Argentina.

[22]See Adams-Prassl et al. (2020a) for the United States and Oreffice and Quintana-Domeque (2020) for the UK. In addition, Biroli et al. (2020) document an increase in reported tensions in families in Italy, the UK, and the US. See also Wozniak (2020), who reports that households with school-age children indicated a greater decline in well-being

Figure 3.2: Childcare Provided by Single and Married Parents

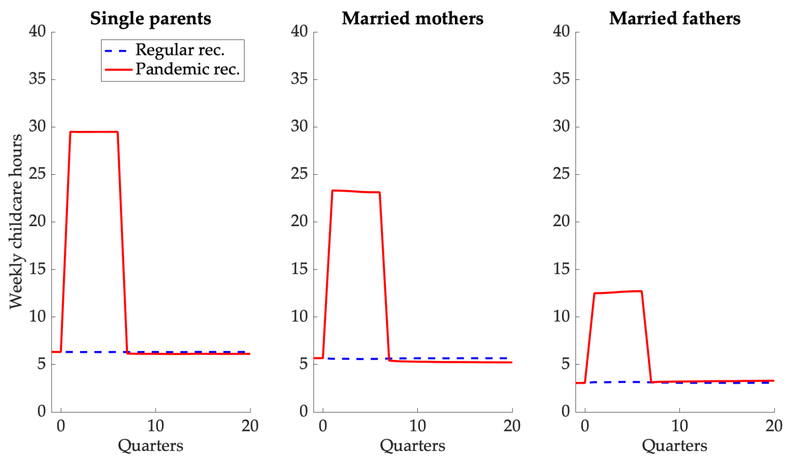

Notes: For telecommuters, childcare time in the model is computed as $t^g + 0.5\phi(k)\, n^g$*, that is, time that is spend on childcare and work simultaneously is counted as 50 percent childcare. Counting all of the combined time as childcare leads to similar results.*

to the sudden increase in childcare needs. We also find large reductions in parents' labor supply: many women switch from full-time to part-time work or drop out of the labor force entirely to meet the extra childcare needs. We will discuss these findings in detail in the next section.

While Figure 3.2 clearly shows that women are taking over the majority of the increase in childcare hours, the impact is large for men as well. In fact, in relative terms (compared to childcare during normal times) the increase in childcare is slightly larger for married fathers than for married mothers (see Figure 3.2). In some families, this leads to a substantial change in the division of childcare time. In our quantitative model, the fraction of couples in which both parents do at least ten percent of

during the shutdown than other households.

childcare increases from 31 percent in normal times to 43 percent in the first period of the pandemic. Similarly, the fraction of fathers who do any childcare at all rises from 53 percent to 76 percent.

Some families even experience a complete reversal. We find that a pandemic recession increases the share of couples in which the husband is the main provider of childcare. In normal times, specialization in the household is primarily driven by the within-couple gender wage gap and, for traditional couples, by gender-unequal social norms. Both factors push toward a division of labor that makes mothers the main provider of childcare. Although these factors remain present during a pandemic recession, the parents' occupations begin to play a major role—specifically, whether or not they can be carried out remotely. When a husband can telecommute while his wife cannot, the husband often becomes the primary childcare provider, since he can more easily combine childcare with work.[23] In the model, as the fraction of telecommutable jobs increases during the pandemic recession, the fraction of men who are main childcare providers immediately rises from 24 to 26 percent.[24]

The model predictions of a rise in shared childcare and a rise in men who are primary childcare providers are consistent with the evidence.

[23]One example of such a couple would be a wife who is a doctor or nurse working in a hospital married to an office worker who can work from home during the crisis. Alon et al. (2020a) document that there are millions of such couples in the United States (about 12 percent of married couples with children).

[24]In a regular recession, there is also a rise in the number of men who are the main childcare providers as more men lose their jobs and take on childcare responsibilities, but this increase is smaller and disappears in the recovery.

Carlson, Petts, and Pepin (2020) find that in the United States 28 percent of women reported sharing childcare equally prior to the pandemic, which increased to 34 percent during the pandemic. This increase was even larger for families with older children: from 29 to 42 percent. Biroli et al. (2020) find that the proportion of families that divide childcare responsibilities equally increased by 8 percentage points in the UK and 17 percentage points in Italy.[25] For Germany, Möhring et al. (2020) report that in April 2020 fathers were the main childcare provider in over 20 percent of families. von Gaudecker et al. (2020) find that in 30 percent of Dutch couples where the mother works in a critical occupation fathers were the sole childcare provider in April. The central role of telecommuting in driving these changes is supported by the findings of Adams-Prassl et al. (2020b), who observe that fathers working from home in the United States in April 2020 spent 4.8 hours per day on childcare and homeschooling, while fathers who could not work from home but still had a job spent less than half as much (2.3 hours).

We expect that this increase in fathers' involvement during the pandemic will ultimately lead to more gender-equal norms in terms of the division of childcare, in spite of the overall gender gap we observe. Arguably, having to do a lot of childcare is a bigger shock for most men than for most women. Many men learn for the first time how much work child-

[25]Del Boca et al. (2020) also documents that many Italian men report an increase in childcare responsibilities during April, especially if they were working from home or not working at all due to the crisis.

care entails and the full range of tasks that it involves. Men's increased awareness of the challenges of combining childcare and work may erode gender norms that work against men contributing equally to childcare. To be sure, this may not apply to every individual case. Indeed, some men may be even more hesitant to provide childcare after their pandemic experience. However, existing evidence from policy-induced increases in father's contributions to childcare (e.g., through paternity leave) does suggest that the rise in men's engagement during the crisis will result in a higher involvement of fathers in childcare in the future, and a corresponding greater ability of mothers to pursue their careers (see., e.g., Farré and González 2019 for evidence from Spain, Tamm 2019 for evidence from Germany, and Appendix C.3.2 for further evidence). Furthermore, fathers who are the main providers of childcare can be role models and thus affect social norms in other families as well. Such peer effects among fathers have been documented in the context of paternity leave taking (see Dahl, Loken, and Mogstad 2014). We explore the implications of such potential shifts in social norms in Section 3.7.

3.5.2 Labor Supply During Pandemic Recessions

Figures 3.3a and 3.3b compare the impact of regular and pandemic recessions on total labor supply and on total labor earnings in the economy. Hours worked decline by more than 20 percent in the pandemic recession, versus less than 3.3 percent in the regular recession.

Figure 3.3: Hours Worked and Aggregate Labor Earnings, Pandemic vs. Regular Recessions

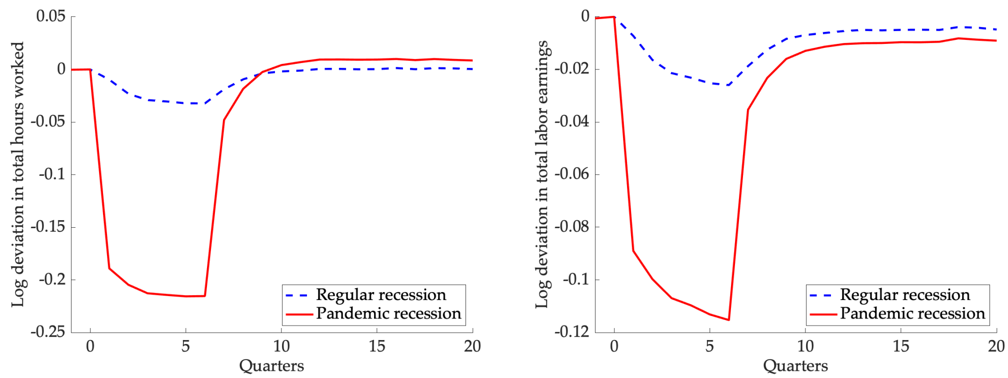

(a) Total Hours Worked (b) Aggregate Labor Earnings

The decline in total labor earnings (which measures the decline in labor supply in efficiency units) is only about half as large as the decline in hours in the pandemic recession. This reflects the fact that the workers who drop out of employment or reduce their hours during the recession tend to have low earnings to begin with. Given that job separation and job finding probabilities do not depend on earnings, this selection effect is entirely due to endogenous decisions on accepting or rejecting job offers. Selection on earnings is less pronounced in a regular recession.

In our model, labor supply quickly rebounds during the recovery following the end of the recession. However, total labor earnings remain lower than before the recession, and particularly so after a pandemic recession. The persistent decline in total labor earnings reflects the depreciation of human capital suffered by many workers who lose employment during the recession.

Figure 3.4: Women's vs. Men's Labor Supply, Pandemic vs. Regular Recessions

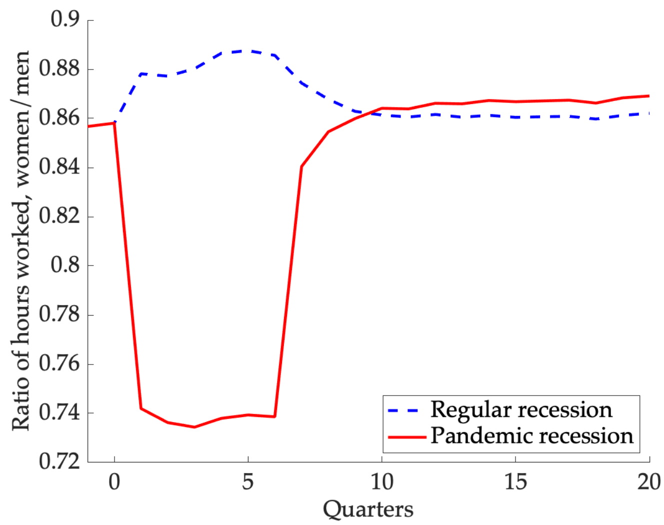

When considering the different implications for women versus men, we observe (Figure 3.4) that in regular recessions, the ratio of women's to men's labor supply increases. This dynamic reflects the greater job losses for men and, to a lesser extent, within-family insurance, i.e., married women increasing labor supply in response to their husband's job loss. In contrast, consistent with the evidence, in a pandemic recession women's labor supply declines sharply relative to men. This drop in women's relative labor supply is largely the flip side of the uneven burden of childcare across genders discussed in the previous section.

Table 3.7 breaks down changes in employment from pre-pandemic times to the second quarter of the recession by marital status, gender, and presence of children.[26] The table shows that during a pandemic recession parents reduce labor supply by much more than people without children, women more than men, and that the age of the children plays a large role, especially for mothers. Not surprisingly, employment declines for single mothers are by far the largest. These model implications can be easily tested once enough data is available and to a large extent have already been confirmed.

Our results also indicate that the ability to telecommute cushions the employment decline only by a small amount. The ability to telecommute primarily has a level effect on labor supply (see Figure C.2b in Appendix C.2): being able to telecommute leads mothers to supply more labor both in regular times and during a recession.

Whether a women is part of a modern or traditional couple also plays an important role. In regular times, the labor supply of traditional mothers is only slightly lower than that of modern mothers (see Figure C.2a in Appendix C.2). Indeed, with the relatively low childcare requirements in normal times, many traditional mothers are able to both work and provide the majority of childcare within the family. In a pandemic recession, in contrast, the traditional division of labor is reinforced, and traditional mothers reduce their labor supply more than modern mothers.

[26]In Figure C.1 in Appendix C.2.1 we contrast these results to a regular recession and we depict the evolution over time.

Table 3.7: Percent Decline in Hours Worked in Model, Q2 in Pandemic Recession relative to Normal.

	Decline in Hours Worked (%)
Men	12.82
Women	25.20
Fathers	20.07
Mothers	39.64
Single mothers	51.89
Married mothers	36.76
Mothers of small kids	47.34
Mothers of big kids	35.35
Fathers of small kids	21.86
Fathers of big kids	18.80
Parents of small kids	33.10
Parents of big kids	27.10
Non-parents	3.40
TC parents	29.26
non-TC parents	30.95

The available evidence to date lines up well with the evolution of labor supply in the model.[27] Bick and Blandin (2020) conduct an online survey to provide real-time evidence on the labor market impact of the current recession. The survey is designed to be comparable to the data typically provided by the Current Population Survey (CPS), and matches the CPS well for the period when the surveys overlap. At the trough of the recession, the decline in labor supply generated by the model roughly matches the 24 percent decline in hours reported by Bick and Blandin for the US economy from February 2020 to the May to June average. Moreover, Bick and Blandin show that women's employment rate (employed and at work)

[27]See Appendix C.3.3 for further evidence.

dropped by 17.8 percentage points from February to June 2020, compared to only 15.8 percentage points for men. The gender gap in hours worked is even larger: between February and May, women's average hours fell by 27 percent, versus a drop of only 20 percent for men.[28]

Other studies shed light on the role of the sector/occupation and childcare channels for the employment impact of the recession. Papanikolaou and Schmidt (2020) examine whether the ability to telecommute, based on ATUS data (as used by Alon et al. 2020a), actually predicts employment losses during the current recession. They find (using industry data from the Bureau of Labor Statistics) that, indeed, sectors with a lower ability to telecommute experienced larger declines in employment. Moreover, the employment of women with young children was particularly affected, underlining the importance of the childcare channel. Similarly, Collins et al. (2020) examine changes in work hours from February to April 2020 in the CPS data, and find that mothers with young children reduced their labor supply by four to five times as much as fathers.

Evidence on the impact of the Covid-19 recession on employment in other countries comes to similar conclusions. Both Adams-Prassl et al. (2020b) and Sevilla and Smith (2020) conducted real-time surveys in the UK and find that women were more likely to have reduced their labor supply during the pandemic than men. The studies show that occupa-

[28] Cajner et al. (2020) come to similar conclusions using data from a major payroll processing company, which show a 21.5 percent decrease in women's employment from February to April 2020, compared to a 17.8 percent decrease for men.

tion plays an important role but cannot explain the entire gender gap in employment rates. Rather, the presence of children and the division of childcare in the household is crucial. Farré et al. (2020) document that in Spain, women have been more likely than men to lose their jobs during the pandemic. Meanwhile, Lemieux et al. (2020) examine the labor market impact of the pandemic in Canada, and find that from February to April labor supply dropped by 30.1 percent for women compared to 27.7 percent for men. In Germany, the differential impact on women is small in comparison (Adams-Prassl et al. 2020b), which might be related to the policy instrument of *Kurzarbeit*, i.e., subsidized reduced employment without terminating the employment relationship. However, even in Germany, the increase in the unemployment rate from February to May has been higher for women (a rise of 19 percent) than for men (14 percent).[29]

3.6 The Transmission from Income to Consumption in Regular and Pandemic Recessions

So far, we have established that our model can explain the distinct impact of a pandemic recession on women's versus men's employment, as well as the central role that childcare obligations play in generating this outcome. The next question to address is whether the different impact of regular versus pandemic recessions on women and men matters at the aggregate

[29]See Table 1.1 in Bundesagentur für Arbeit (2020).

level and for the effects of possible policy interventions. In other words, is a pandemic shecession just like a regular mancession with the roles of women and men reversed, or are there important qualitative differences between these two types of economic downturns? We argue that a shecession is indeed qualitatively different, primarily because women's and men's labor supply respond differently to shocks.

3.6.1 Family Insurance

Family members provide one another with an important insurance mechanism against shocks to earnings and employment (see Attanasio, Low, and Sánchez-Marcos 2005; Blundell, Pistaferri, and Saporta-Eksten 2016, 2018; Ellieroth 2019). If a primary earner faces wage cuts or unemployment, the family's secondary earner can either enter the labor force or increase their hours to make up for the reduction in the family's income. This insurance mechanism is particularly relevant during regular recessions, when many men (who are often primary earners) lose their jobs, while women's employment prospects are less affected. Doepke and Tertilt (2016) argue that family insurance is a primary reason behind the low cyclical volatility of married women's labor supply (as documented in Section 3.2).

The family insurance mechanism is quantitatively important in our model. Figure 3.5 shows how labor supply changes over the course of recessions for married women who worked part time just before the recession while their husbands worked full time. This group of households

Figure 3.5: Spousal Insurance: Hours Worked for Married Women Who Worked Part-Time Before Recession while Husband Worked Full-Time

(a) All Women in this Group

(b) By Husband's Job Loss

generally displays the highest levels of family insurance because the secondary earner is already in the labor force, and is thus able to increase hours. The left panel of the figure shows that women in this group increase their labor supply during a regular recession. In the right panel, we further decompose labor supply in this group to compare women whose husband loses his job (i.e., is not working in the current period, even though he was working full time before the recession) versus those whose husband remains employed. We observe that the increase in hours in a regular recession is indeed driven by women whose husbands lost a job, as suggested by the family insurance mechanism. The effect is quantitatively large: conditional on the husband's job loss, labor supply during the recession increases by more than 50 percent for this group of women.

Figure 3.5 displays labor supply for these same groups during a pandemic recession. The left panel shows that the family insurance mechanism is no longer present in terms of total labor supply, which drops throughout the entire recession for this group of women. Again, the right panel decomposes the overall change in labor supply between women whose husbands lost their jobs and those whose husbands are still employed. Women whose husbands become unemployed still increase their labor supply in the initial period of the recession, though only by half as much as in a regular recession. However, this insurance effect becomes smaller in subsequent periods. As the pandemic regression progresses, many of the women who initially worked part time drop out of the labor force to meet childcare needs, which makes it more difficult to find a job and expand employment later on. Family insurance continues to exist in the sense that women whose husbands are unemployed work more than others, but this takes the form of not cutting hours rather than increasing hours. Families are able to soften the blow of falling earnings, but truly compensating for income losses by working more is not feasible for most couples during a pandemic recession.

3.6.2 Marginal Propensities to Consume

In addition to driving the labor supply response to a pandemic, family insurance plays a role in the transmission of income shocks to household spending and consumption. Households that lose access to insurance mechanisms are less able to compensate for income losses, resulting in a strong transmission from income shocks to reduced consumption. These changes are reflected in the distribution of marginal propensities to consume (MPCs) throughout the economy.

A recent body of macroeconomic work demonstrates the central role of MPC distributions for the transmission of macroeconomics shocks (e.g., Berger et al. 2017, Auclert 2019, Patterson 2019, Auclert, Bardóczy, and Rognlie 2020). If the average MPC is high, a negative shock to household income will result in a larger reduction in consumption demand. In models where demand shocks affect output (e.g., because of nominal frictions), a higher average MPC results in deeper recessions for a given initial shock. Thus, understanding the ways in which MPCs change over time during a pandemic recession is crucial to ascertain how the shock of a pandemic recession is transmitted throughout the economy and to assess the possibility of a highly persistent downturn driven in part by demand channels. Figure 3.6 addresses this question by illustrating how the average MPC evolves in the economy during both types of recessions.

Two important differences between regular and pandemic recessions stand out. First, on impact the pandemic recession raises MPCs by a greater amount than a regular recession, especially for single households. This initial difference arises primarily because a pandemic recession causes a bigger drop in earnings, which pushes households closer to financial constraints. Second, the rise in MPCs is more persistent during a pandemic recession than a regular recession for both single and married house-

Figure 3.6: Average Marginal Propensities to Consume

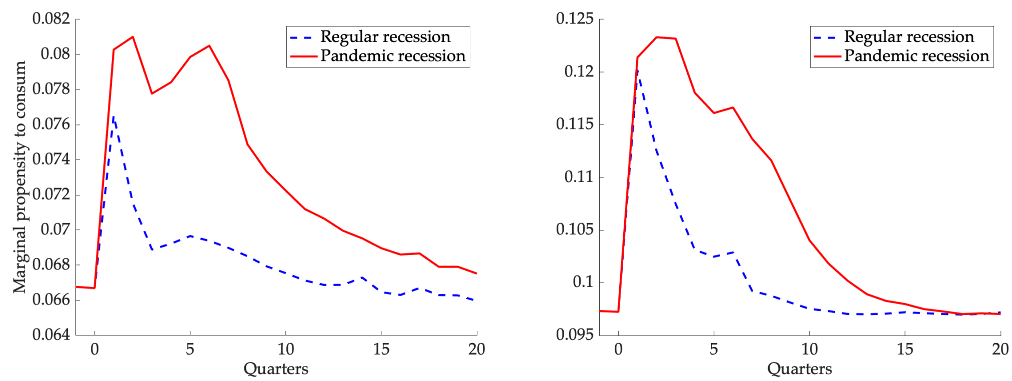

(a) Singles (b) Couples

holds. Two different mechanisms contribute to this persistence. For single households, the persistent increase in MPCs is primarily driven by single parents, a large number of whom drop out of the labor force for the entire pandemic recession. This persistent earnings loss drives assets down and leaves little room for self-insurance, even during the early years of the recovery. The same factor is at play for married households, but these households also suffer from the loss of family insurance as shown above. The loss of family insurance implies that married households are less able to compensate for earnings losses; they consequently draw down their assets and ultimately end up with a high MPC.

The persistent rise in MPCs during a pandemic recession and the subsequent recovery implies that the downturn can be amplified and the recovery delayed through demand-driven channels. Conversely, high MPCs also imply that economic stimulus measures are likely to be highly effective. Overall, these results highlight the important role of the dynamics of female labor supply and family decision-making in shaping the macroeconomic properties of recessions.

3.7 Implications for Gender Inequality

We now move on from the macroeconomic implications to focus on the repercussions of regular and pandemic recessions for gender inequality. We have already shown that unlike regular recessions, pandemic recessions reduce women's labor supply relative to men's, and that mothers' childcare responsibilities play an important role in this reduction. These shifts in labor supply have direct implications for gender inequality in the labor market through the accumulation of experience while working and skill loss while not employed. Regular recessions primarily lower men's employment and therefore result in a corresponding reduction in men's labor market experience that contributes to a narrowing of the gender wage gap. Conversely, a pandemic recession puts many women out of work and, at least initially, lowers women's relative wages.

We also consider the possibility that the experience of a pandemic recession can lead to changes in gender inequality that long outlast the pandemic itself. Gender inequality in the labor markets of advanced economies is linked, in large part, to childbearing and the unequal division of childcare responsibilities between women and men (Miller 2011; Adda, Dustmann, and Stevens 2017; Kleven, Landais, and Søgaard 2019; Kleven et al. 2019; Gallen 2018; Hannusch 2019; Xiao 2020). As we have documented, the current pandemic recession has led to massive changes in how families organize childcare: along with mothers, many fathers have also increased the time they spend caring for their children during the crisis, while numerous employers have reorganized work to enable their staff to continue working while caring for children at home. We argue that some of these changes are likely to persist, leading to long-term changes in gender inequality in the labor market.

3.7.1 The Evolution of the Gender Wage Gap During the Recovery

The link between job losses and persistent losses in earnings is well-documented in the literature (e.g., Stevens 1997), as is the fact that such losses are especially severe for layoffs that occur in recessions (Davis and von Wachter 2011). Laid-off workers forgo returns to experience, may face difficulty finding a new job in the same occupation or with the same level of responsibility, and are less likely to have secure employment in the future (Jarosch 2015). These consequences are not limited to workers who lose their jobs, but also affect those about to enter the labor market for the first time.[30]

[30] See, for example, Altonji, Kahn, and Speer (2016), Oreopoulos, von Wachter, and Heisz (2012), and Schwandt and von Wachter (2019).

Figure 3.7: Gender Inequality in the Labor Market during Recessions

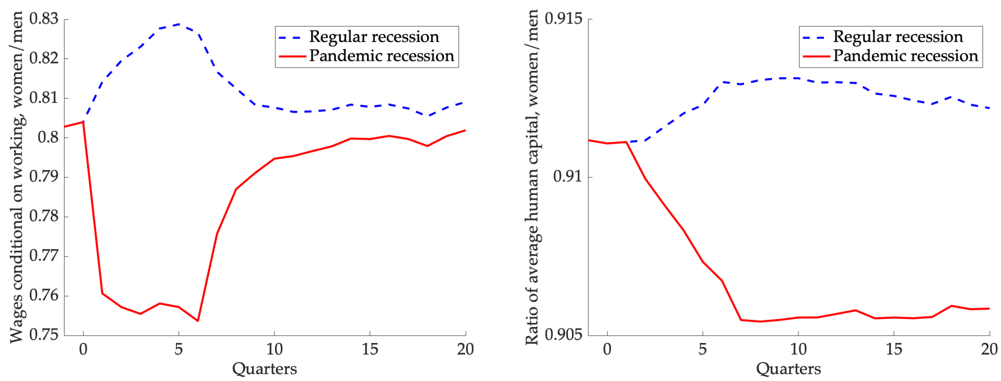

(a) Relative Wages (b) Relative Skills

We have documented that both in the data and in our model recessions affect women's and men's employment in different ways. These differences have consequences for the evolution of gender inequality in the labor market during and after recessions. Figure 3.7a shows that gender inequality shrinks during a regular recession, with women's wages increasing by close to two percent relative to those of men. This matches empirical evidence that gender wage gaps usually narrow during recessions (Solon, Barsky, and Parker 1994), an effect that was particularly pronounced in the Great Recession of 2007–2009 (Marchand and Olfert 2013; Chen and Kelly 2019). In contrast, we find that a pandemic recession leads to a widening of the gender gap by five percentage points, as it hits women's employment harder than men's.[31] Changes in relative wages during recessions

[31] We abstract from general equilibrium effects that could arise from limited substitutability between women's and men's labor. Such general equilibrium effects would dampen the increase in the gender wage gap during the pandemic but not after, because

do revert to some extent during the recovery, but the gap is persistent: even five years after a recession, the gender wage gap is smaller after a regular recession compared to a pandemic recession.

The changes in the observed gender wage gap are due both to skill accumulation and loss, and to selection effects. Figure 3.7b isolates the contribution of relative skill levels by displaying how the ratio of human capital (i.e., efficiency units of labor) between women and men changes during a recession. As expected, in regular recessions (when men face high unemployment) women's skills increase relative to men's, whereas in a pandemic recession (when many women stop working) women's relative skills drop sharply. Changes in skills are more persistent than changes in the wage gap, reflecting how some workers who face skill loss stop working permanently, and therefore no longer affect the measured gender gap among those in the labor force.[32] Figures 3.7a and 3.7b show that the initial changes in the gender wage gap during a recession are primarily due to selection, but the importance of skill accumulation increases over time.

A qualitative difference between a pandemic and a regular recession is that the movement in the the gender wage gap is more persistent after a pandemic recession. Most men who lose employment in a regular

women's relative labor supply actually increases in the recovery from the pandemic.

[32]These effects on the relative skills of women and men are similar to the finding by Heathcote, Perri, and Violante (2020) that if less-skilled workers lose their jobs in a recession, their attachment to the labor force tends to decrease.

recession ultimately return to full-time work and gradually regain labor market experience. In contrast, women's long-run labor supply is more responsive to lost human capital. Some women who worked full time before the pandemic but then lost employment either drop out of the labor force permanently or return only to part-time work, because the increased wage gap within the family (relative to the husband) induces more specialization.

3.7.2 The Long-Run Impact on the Gender Gap: Work Organization and Social Norms

The coronavirus pandemic has resulted in a historically unprecedented increase in the provision of childcare by working mothers and fathers, with many fathers becoming primary providers of childcare for the first time. The pandemic has also led to an equally unprecedented reorganization of the workplace, with a large fraction of the labor force working from home during the crisis and employers quickly adjusting to this new reality of pervasive remote work.[33]

Experience shows that such a temporary but profound shift in the division of labor between genders and the reorganization of the workplace can lead to permanent shifts in gender norms and economic outcomes. One example is the entry of millions of married women into the US labor force

[33] See Appendix C.3.1 for evidence supporting our assumption that the ability to work from home has increased permanently.

during World War II. Before the war, most women would stop working once they got married, a convention that was supported by social norms that favored the single-earner model and formal restrictions such as bans on the participation of married women in many occupations. The unparalleled rise in women's wartime labor force participation had a large and persistent effect on female employment.[34] The long-term impact of World War II on women's labor market participation was attributable in part to shifting social norms.[35] Similarly, Fernández (2013) and Fogli and Veldkamp (2011) argue that in the 1960s and 1970s observing working women in their families and neighborhoods created an awareness of the costs and benefits of employment and was a major engine behind the secular rise in married women's labor force participation from the 1950s to the 1990s.[36] This implies that temporary shocks can accelerate social change, in this case by providing additional learning opportunities.

Our model of a pandemic recession and the subsequent new normal incorporates the expectation that the substantial changes in childcare re-

[34] See Acemoglu, Autor, and Lyle (2004) and Goldin and Olivetti (2013). Doepke, Hazan, and Maoz (2015) argue that the persistent impact of World War II on the female labor market was also one of the root causes of the post-war baby boom.

[35] Fernández, Fogli, and Olivetti (2004) show that boys who grow up with a working mother are more likely to marry women who likewise continue to work when married. The example provided by their own parents arguably created a preference among these boys for a more equal division of labor in the family that was then reflected in their own choices as husbands and fathers. See Grosjean and Khattar (2018) for evidence on the persistence of gender norms over even longer periods.

[36] Along similar lines, Olivetti, Patacchini, and Zenou (2020), show that girls who are exposed to their peers' working mothers during their teenage years are more likely to end up working themselves.

sponsibilities and work organization during the crisis will have long-term effects. In particular, the pandemic recession has been marked from the outset by more couples switching from traditional to modern family roles, with modern couples especially prevalent among younger cohorts. While we do not model the exact nature of the adjustment process, we view this transformation as being driven by "learning by doing" as many fathers experience a major increase in childcare responsibilities, and by the role model effect produced by the increasing share of fathers who are the primary providers of childcare during the crisis.[37]

We also expect that the increased work flexibility that arises at the beginning of the pandemic, with a larger fraction of jobs done by telecommuting, will persist in the new normal. This change can once again be justified with learning by doing, in this case by both employers and employees. Furthermore, it is consistent with numerous news reports of employers planning to keep work-from-home arrangements in place after the pandemic. More flexible work arrangement can benefit women by lowering the overall burden of childcare and by increasing the childcare responsibilities of men who find telecommutable jobs. The notion that low workplace flexibility is a barrier for women's careers has been advanced by Goldin and Katz (2011), Goldin (2014), and Erosa et al. (2017), among others.[38]

[37] See Appendix C.3.2 for evidence from the context of parental leave policies that short term changes in the division of labor in the family can have lasting effects.

[38] See also Cubas, Juhn, and Silos (2019) and Iacopo and Moser (2020).

Figure 3.8: Fraction of Married Couples with Children in which the Father is the Main Childcare Provider

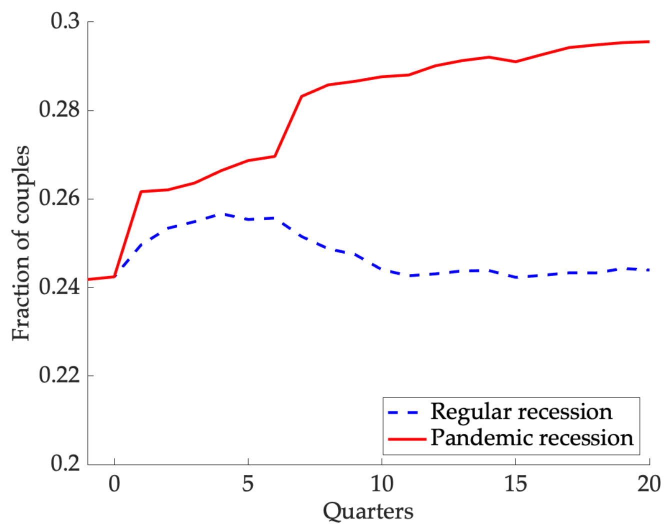

Notes: For telecommuters, childcare time in the model is computed as $t^g + 0.5\phi(k)\, n^g$*, that is, time that is spent on childcare and work simultaneously is counted as 50 percent childcare. Counting all of the combined time as childcare leads to similar results.*

The changing gender roles in the model can be seen in Figure 3.8, which shows that the fraction of couples where the father is the main childcare provider slowly increases to almost 30 percent. The initial jump from 24 to 26 percent is primarily due to telecommuting fathers. Later on, the gradual increase in the share of modern couples (i.e., a change in social norms) also plays a role.

Given these driving forces of long-run changes, Figure 3.9a shows how women's relative labor supply changes over the long term (40 years) in pandemic versus regular recessions. Despite the losses in employment and job skills that women face during a pandemic recession (see Figure 3.7b),

Figure 3.9: Gender Inequality in the Long Run with Decomposition of Channels

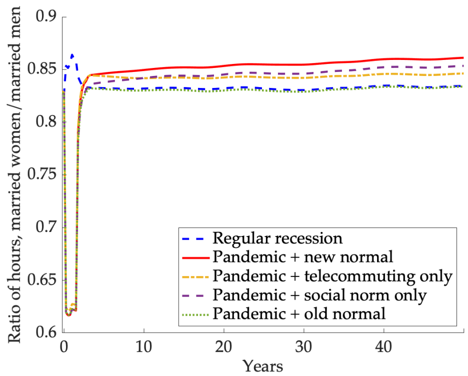

(a) Labor Supply, Married Women/Married Men

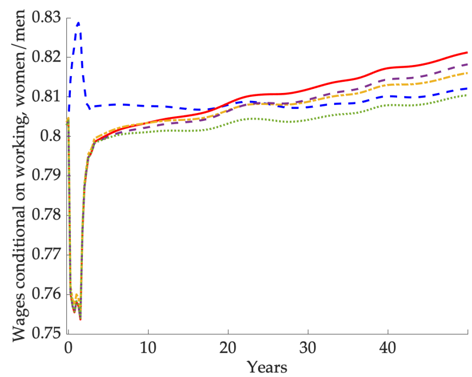

(b) Gender Wage Gap

female labor supply rises above its original level early in the recovery. Figure 3.9a also provides a decomposition that shows how the two long-run forces (changes in social norms and increase in TC jobs) contribute to this outcome. We observe that if, counterfactually, both social norms and the share of telecommutable jobs reverted to the pre-pandemic levels at the beginning of the recovery, women's labor supply would continue to be depressed and remain slightly below the level observed in the aftermath of a regular recession over the long term. Both long-run forces are crucial for raising women's labor supply.

Figure 3.9b shows the impact of this change in women's labor supply on the gender wage gap. As shown in Figure 3.7a, a pandemic recession increases the gender wage gap in the medium term through the depreciation of women's skills during the pandemic. However, the subsequent rise in female labor supply accelerates the accumulation of skills and gradually raises women's relative wages. After about 20 years, the gender wage gap returns to its original level, and continues to close in response to women's higher labor supply. The decomposition in Figure 3.9b shows that both the change in social norms and the increase in job flexibility play a quantitatively important role in narrowing the gender gap.

Two decades is a long time, and our long-run results do not contradict our basic finding that a pandemic recession is a setback for women's equality in the workplace. Nevertheless, the long-run results do provide a silver lining. A pandemic recession has the potential to be a watershed moment in terms of the division of labor in the family and in terms of a family-friendly organization of the workplace. Through these channels, the pandemic can contribute to reducing gender inequality over the long run.

3.8 Policy Implications for Pandemic Recessions

The severe impact of the current downturn on employment, earnings, and, ultimately, welfare raises the question of what public policy can do to offset some of the economic consequences of the pandemic. Our economic model can help inform this debate.

3.8.1 Fiscal Policy

Our findings on family insurance and MPCs suggest that fiscal policy, such as extended unemployment insurance and transfer payments to affected families, can be disproportionately effective during a pandemic recession in terms of stimulating aggregate demand. Our model focuses on the household sector of the economy and does not spell out an aggregate demand channel explicitly; however, it would be straightforward to add this along the lines of, e.g., Hagedorn, Manovski, and Mitman (2019). In such a model, fiscal policy would be unusually powerful as long as MPCs are elevated, and our model suggests that a pandemic recession is characterized by high MPCs that persist for about two years during the recovery (see Figure 3.6). Additional transfer payments during this phase would accelerate the recovery, especially so if the payments are targeted to households such as single parents and, more generally, to households with children.

For a full analysis of fiscal policy during a pandemic recession, additional channels that are not modeled here also need to be taken into account. In terms of the optimal provision of unemployment insurance, Mitman and Rabinovich (2020) argue that unemployment benefits should be higher as long as job finding rates are low during a lockdown, which adds another argument in favor of higher transfer payments. An argument against high transfer payments at the height of the pandemic is that during strict lockdowns consumption possibilities are reduced; for example, travel, indoor dining, and many entertainment options become unavailable, and shopping at physical stores is either impossible or comes with additional risks. In such a period, there may be less need for transfer payments. The evidence suggests so far that the lockdown-induced reduction in demand was relatively short-lived. At any rate, the force in favor of higher transfer payments spelled out in our analysis is likely to be especially relevant later during the recession and the recovery when the direct effect of lockdowns loses force.

3.8.2 School Openings

The policy issue most directly linked to our analysis is the role that school openings can play in accelerating the recovery from the crisis. A full analysis of this question would require an assessment of the health consequences of opening schools and daycare centers while the pandemic is still ongoing, an issue that we abstract from here.[39] Our analysis can, however, shed light on the repercussions of school openings for the labor market and the evolution of gender inequality during the recession and recovery.

In our setting, the primary effect of opening schools and daycare centers is to free up the labor supply of women and men who are currently

[39] A cautionary note is provided by Alon et al. (2020b), who argue that schools can be a major vector of disease transmission, particularly in developing countries due the high prevalence of multi-generation households, a feature that Bayer and Kuhn (2020) argue can contribute to high case-fatility rates. Baqaee et al. (2020) emphasize that measures such as reintroducing restrictions on social gatherings, wearing masks, and increasing testing and quarantine are necessary before wider re-openings are feasible. The effect of school closures on the US healthcare workforce specifically is analyzed in Bayham and Fenichel (2020).

not working because they need to look after and homeschool their children. Empirical estimates show that this effect may be especially important. Dingel, Patterson, and Vavra (2020) show that 32 percent of the US workforce has a child under the age of 14 in their household. Fuchs-Schündeln, Kuhn, and Tertilt (2020) report that the same is true for 26 percent of the workforce in low-fertility Germany, while this share is as high as 41 percent in other European countries.

Figure 3.10: Hours Worked and Aggregate Labor Earnings under School Reopenings

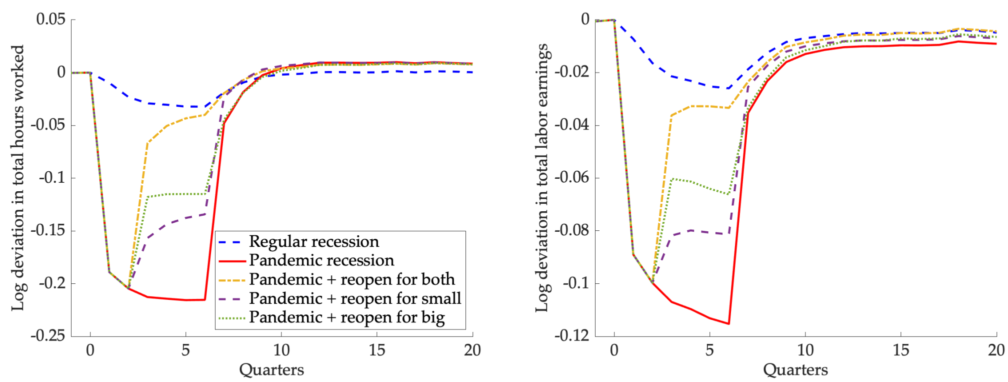

(a) Total Hours Worked (b) Aggregate Labor Earnings

Figure 3.10 shows how aggregate hours worked would change over time in a pandemic recession if schools were to reopen two quarters into the recession, rather than staying shut for the entire pandemic. Formally, opening schools would mean that childcare requirements return to the pre-recession level after two quarters, while job separation probabilities would continue to be elevated and the other aspects of the pandemic recession (changes to telecommuting and social norms) would remain in place. The figure also illustrates the results of returning only young children to school (i.e., by opening daycare centers and preschools) or only older children (opening K-12 schools). We observe that opening schools would immediately mitigate the economic impact of the pandemic by reversing more than half of the decline in labor supply brought about by the recession. The impact on labor earnings is even larger: losses in labor earnings are reduced by about two-thirds. This large economic impact underscores the key role of increased childcare requirements for the drop in economic activity during the pandemic, and shows that reopening schools is much more effective, in economic terms, than reopening specific sectors with small shares of aggregate employment shares (such as gyms, bars, and restaurants).

Figure 3.11: The Impact of School Reopenings on Gender Inequality

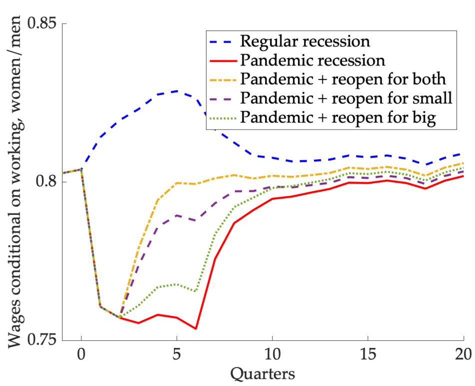

(a) Relative Wages

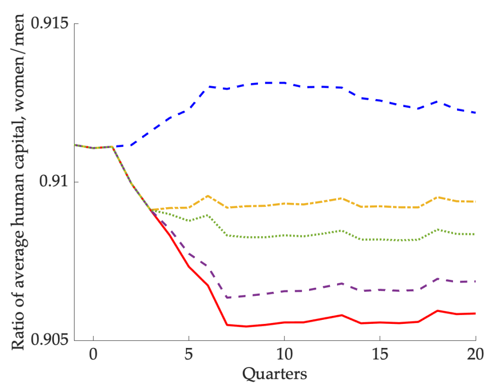

(b) Relative Skills

Figure 3.11 shows how opening schools early affects gender inequality in the labor market in terms of women's relative wages and relative skills. Since women bear the largest part of the extra childcare burden during the pandemic, relieving this burden will disproportionately benefit women. Opening schools early immediately reduces the gender wage gap, and avoids about half of women's recession-induced relative skill losses. Sending younger kids back to school does the most to reduce the gender wage gap. This result is primarily driven by a composition effect, with high-skill women more likely to return to work when childcare becomes available. As Panel (b) in Figure 3.11 shows, opening schools early for older kids reduces the widening skill gap by more than opening daycare centers for younger kids. In part, this is because there are more families with older than with younger kids, so more households are affected by opening schools. In addition, many women with young kids work part-time or not at all even in normal times, leaving a bit more room for dealing with extra childcare needs during the pandemic. Women also benefit disproportionately from school openings in terms of welfare (see Figure C.4 in Appendix C.2.3).

These results suggest that prioritizing school openings (relative to, say, opening bars and restaurants) can be an effective strategy for mitigating the economic impact of a pandemic recession. Of course, this policy implication comes with the caveat that the health consequences of opening schools must also be taken into account. Such a policy is a realistic option only if the pandemic is sufficiently controlled such that opening schools will not reignite or amplify the pandemic itself. Even when large-scale school re-openings are not feasible, our analysis suggests that similar policies which provide targeted childcare assistance can be helpful. Such limited policies have already been implemented by several countries, including Germany which provided emergency childcare assistance to single parents. These provisions allows those with large MPCs, namely single parents, to continue working, while minimizing the impact of infections.

3.9 Conclusions

As a result of the Covid-19 pandemic, countries around the world, including the United States, have entered the sharpest economic downturn since the Great Depression. In this paper, we argue that the central economic distinction between this downturn and other recent recessions, aside from its severity, lies in its impact on women's employment.

The lockdown measures accompanying a pandemic recession have a large effect on high-contact sectors such as hotels and restaurants, which have large shares of female employment. Thus, unlike in a regular recession, more women than men are directly affected by layoffs. In addition, daycare and school closures during the pandemic result in considerably higher childcare obligations. Women shoulder the majority of this additional responsibility, further decreasing their ability to work.

We develop a macroeconomic model that can account for the distinct features of regular and pandemic recessions. We use the model to examine the wider economic repercussions of the disproportionate impact of a pandemic recession on working women. In terms of macroeconomic implications, we find that the outsized impact of a pandemic recession on women's employment reduces the role of families as a shock absorber. Very few married workers are able to increase employment to make up for their spouse's lost earnings. As a result of this loss of insurance, earnings losses are strongly translated to lower consumption demand, and marginal propensities to consume increase by a greater amount than in regular recessions.

These findings have important policy implications. First, we show that reopening schools and daycare centers, if it can be safely done, have a first order effect on the speed of recovery. If policy-makers have to choose between reopening one or the other, we find that in terms of total economic impact, reopening schools is more important. The main reason is that there are more employees with school age children and they are more likely to work full time than those with smaller children. Second, our analysis suggests that fiscal policy is more effective during a pandemic than in usual recessions. The reason is that due to reduced possibilities for family insurance, marginal propensities to consumer are higher than in normal recessions, and particularly high for single parents. Third, going forward, our framework could be used for studying alternative policies such as emergency childcare for singles, or paid parental leave for school closures, including specific leave days ear-marked for fathers.

We also find that a pandemic recession has sizeable repercussions for gender inequality. In the short and medium term, a pandemic recession erodes women's position in the labor market, first through direct employment losses, and later through the loss in labor market experience brought about by low employment during the recession. These forces lead to a widening of the gender wage gap during a pandemic recession and in its immediate aftermath.

Nevertheless, we also argue that a pandemic recession can trigger changes that ultimately reduce gender inequality over the longer term. Specifically, the rise in work flexibility during a pandemic recession is likely to be persistent, and disproportionately benefits women who have major childcare responsibilities. We also note the possibility of shifting social norms towards a more equal division of childcare obligations between mothers and fathers, triggered by an increase in men's childcare provision and a rising fraction of men who are the main provider of childcare in their family. In our quantitative analysis, these changes imply that a pandemic recession ultimately reduces the gender wage gap, although it takes many years to fully make up for women's initial skill losses.

A more general lesson from our analysis is that accounting for family behavior and gender differences should be a central element of research on economic fluctuations. Authors such as Albanesi (2020), Doepke and Tertilt (2016), and Fukui, Nakamura, and Steinsson (2019) have already shown that the secular rise in female labor force participation in the twentieth century has changed the nature of aggregate labor supply and is the underlying cause behind recent changes in the nature of economic fluctuations. Our study adds to these arguments by accounting for the macroeconomic consequences of childcare responsibilities, skill accumulation, and work organization, factors that all play a central role in the current pandemic recession. A traditional, single-gender macroeconomic model would be unable to capture some of the most distinct characteristics of the economic environment brought about by the coronavirus pandemic.

Our work could be extended to consider the impact of the Covid-19 crisis on additional dimensions of gender equality, such as the rise in domestic violence that appears to have occurred during the crisis (see Leslie and Wilson 2020, Bullinger, Carr, and Packham 2020, and Rivera et al. 2020) or the impact on fertility (Wilde, Chen, and Lohmann (2020)). Moreover, our analysis has focused on advanced economies that are characterized by high income levels and high participation of women, including many mothers, in the formal labor market. As we have documented, the current pandemic recession has similar features in terms of the relative economic impact on women and men across countries in this group. An urgent challenge for future research is to assess the impact of pandemic recessions in middle-income and developing countries. The existing work on this issue (e.g., Alon et al. 2020b) has generally focused on issues other than gender or women's labor force participation. Yet, the pandemic is a global phenomenon, and policy measures such as school closings are being implemented around the world. At the same time, different economic conditions in terms of income levels, women's labor force participation, and the ability to work remotely suggest that the impacts of the pandemic recession and the resulting policy tradeoffs may be substantially different in developing economies. Given the severity of the ongoing health and economic crisis, research on the impact of the coronavirus epidemic on women's work and gender inequality in a wider range of countries is urgently needed.

Bibliography

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. "Innovation, Reallocation, and Growth." *American Economic Review* 108 (11): 3450–3491.

Acemoglu, Daron, David H. Autor, and David Lyle. 2004. "Women, War, and Wages: The Effect of Female Labor Supply on the Wage Structure at Midcentury." *Journal of Political Economy* 112 (3): 497–551.

Acemoglu, Daron, and Dan Cao. 2015. "Innovation by Entrants and Incumbents." *Journal of Economic Theory* 157:255–294.

Acharya, Viral V., Tim Eisert, Christian Eufinger, and Christian Hirsch. 2018. "Real effects of the sovereign debt crisis in Europe: Evidence from syndicated loans." *Review of Financial Studies* 31 (8): 2855–2896.

Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh. 2020a. "The Impact of the Coronavirus Lockdown on Mental Health: Evidence from the US." Human Capital and Economic Opportunity Working Group, Working Papers 2020-030.

———. 2020b. "Inequality in the Impact of the Coronavirus Shock: Evidence from Real Time Surveys." CEPR Discussion Paper 14665.

Adda, Jérôme, Christian Dustmann, and Katrien Stevens. 2017. "The Career Costs of Children." *Journal of Political Economy* 125 (2): 293–337.

Aghion, Philippe, Ufuk Akcigit, and Peter Howitt. 2014. "What Do We Learn From Schumpeterian Growth Theory?" Chapter 1 of *Handbook of Economic Growth*, Volume 2, 515–563. Elsevier B.V.

Aghion, Philippe, Stefan Bechtold, Lea Cassar, and Holger Herz. 2014. "The Causal Effects of Competition on Innovation : Experimental Evidence." pp. 1–31.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li. 2019a. "A Theory of Falling Growth and Rising Rents." *Federal Reserve Bank of San Francisco, Working Paper Series*, pp. 01–43.

Aghion, Philippe, Antonin Bergeaud, Mattheiu Lequien, and Marc J. Melitz. 2019b. "The Heterogenous Impact of Market Size on Innovation: Evidence from French Firm-Level Exports." *NBER Working Paper*.

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. 2005. "Competition and Innovation: An Inverted-U Relationship." *Quarterly Journal of Economics* 120 (2): 701–728.

Aghion, Philippe, Christopher Harris, Peter Howitt, and John Vickers.

2001. "Competition, Imitation and Growth with Step-by-Step Innovation." *Review of Economic Studies* 68 (3): 467–492.

Akcigit, Ufuk, and Sina Ates. 2019. "What Happened to U.S. Business Dynamism?"

———. 2020. "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory." *AEJ: Macroeconomics, Forthcoming*.

Akcigit, Ufuk, Sina Ates, and Giammario Impullitti. 2018. "Innovation and Trade Policy in a Globalized World." *NBER Working Paper*.

Akcigit, Ufuk, and William R. Kerr. 2018. "Growth Through Heterogeneous Innovations." *Journal of Political Economy* 126, no. 4.

Albanesi, Stefania. 2020. "Changing Business Cycles: The Role of Women's Employment." NBER Working Paper 25655.

Albanesi, Stefania, and Ayşegül Şahin. 2018. "The Gender Unemployment Gap." *Review of Economic Dynamics* 30:47–67.

Albanesi, Stefania, Rania Gihleb, Jialin Huo, and Jiyeon Kim. 2020. "Household Insurance and The Macroeconomic Impact of the Novel Coronavirus." Unpublished Manuscript, University of Pittsburgh.

Alipour, Jean-Victor, Oliver Falck, and Simone Schüller. 2020. "Germany's Capacities to Work from Home." CESifo Working Paper No. 8227.

Alon, Titan, Sena Coskun, and Matthias Doepke. 2020. "Trends in Work and Leisure: It's a Family Affair." Unpublished Manuscript, Northwestern University.

Alon, Titan, Matthias Doepke, Jane Olmstead-Rumsey, and Michèle Tertilt. 2020a. "The Impact of COVID-19 on Gender Equality." *Covid Economics: Vetted and Real-Time Papers* Issue 4:62–85.

Alon, Titan M., Minki Kim, David Lagakos, and Mitchell VanVuren. 2020b. "How Should Policy Responses to the COVID-19 Pandemic Differ in the Developing World?" NBER Working Paper 27273.

Alstadsæter, Annette, Bernt Bratsberg, Gaute Eielsen, Wojciech Kopczuk, Simen Markussen, Oddbjorn Raaum, and Knut Røed. 2020. "The First Weeks of the Coronavirus Crisis: Who Got Hit, When and Why? Evidence from Norway." NBER Working Paper 27131.

Altonji, Joseph G., Lisa B. Kahn, and Jamin D. Speer. 2016. "Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success." *Journal of Labor Economics* 34 (S1): S361–S401.

Amiti, Mary, and David E. Weinstein. 2018. "How Much Do Idiosyncratic Bank Shocks Affect Investment? Evidence from Matched Bank-Firm Loan Data." *Journal of Political Economy* 126 (2): 525–587.

Anderson, Gary, and Audrey Kindlon. 2019. "Indicators of R&D Performance by Size of Company." *National Science Foundation*, pp. 1–9.

Andrew, Alison, Sarah Cattan, Monica Costa Dias, Christine Farquharson, Lucy Kraftman, Sonya krutikova, Angus Phimister, and Almudena Sevilla. 2020. "The Gendered Division of Paid and Domestic Work under Lockdown." IZA Discussion Paper 13500.

Andrews, Dan, Chiara Criscuolo, and Peter Gal. 2016. "The Global Productivity Slowdown, Technology Divergence and Public Policy: A Firm Level Perspective." *The Future of Productivity: Main Background Papers*, pp. 1–50.

Andrews, Dan, Chiara Criscuolo, and Peter N Gal. 2015. "Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence From OECD Countries." *The Future of Productivity: Main Background Papers*, pp. 1–38.

Anzoategui, Diego, Diego Comin, Mark Gertler, and Joseba Martinez. 2019. "Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence." *AEJ: Macroeconomics* 11 (3): 67–110.

Arrow, Kenneth J. 1962. "The Economic Implications of Learning by Doing." *The Review of Economic Studies* 29 (3): 155–173.

Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos. 2005. "Female Labor Supply as Insurance Against Idiosyncratic Risk." *Journal of the European Economic Association* 3 (2/3): 755–64.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109 (6): 2333–67 (June).

Auclert, Adrien, Bence Bardóczy, and Matthew Rognlie. 2020. "MPCs, MPEs and Multipliers: A Trilemma for New Keynesian Models." Unpublished Manuscript, University of Arizona.

Aum, Sangmin, Sang Yoon (Tim) Lee, and Yongseok Shin. 2018. "Computerizing industries and routinizing jobs: Explaining trends in aggregate productivity." *Journal of Monetary Economics*.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms." *Quarterly Journal of Economics* 135 (2): 645–709.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John van Reenen. 2017. "The Fall of the Labor Share and the Rise of Superstar Firms." *NBER Working Paper Series*, no. 23396.

Baert, Stijn, Louis Lippens, Eline Moens, Philippe Sterkens, and Johannes Weytjens. 2020. "The COVID-19 Crisis and Telework: A Research Survey on Experiences, Expectations and Hopes." IZA DP No. 13229.

Baqaee, David, Emmanuel Farhi, Michael Mina, and James H. Stock. 2020. "Policies for a Second Wave." Forthcoming, Brookings Papers on Economic Activity.

Baqaee, David Rezza, and Emmanuel Farhi. 2020. "Productivity and Misallocation in General Equilibrium." *Quarterly Journal of Economics* 135 (1): 105–163.

Bardóczy, Bence. 2020. "Spousal Insurance and the Amplification of Business Cycles." Unpublished Manuscript, Northwestern University.

Barkai, Simcha, and Seth G Benzell. 2018. “70 Years of US Corporate Profits 70 Years of US Corporate Profits *.” no. 22.

Barrero, Jose Maria, Nick Bloom, and Steven J. Davis. 2020, October. “Why Working From Home Will Stick.” Unpublished Manuscript, Stanford University .

Bartelsman, Eric J., and Mark E. Doms. 2000. “Understanding Productivity: Lessons from Longitudinal Microdata.” *Journal of Economic Literature* XXXVIII (September): 569–594.

Bartik, Alexander W, Zoe B Cullen, Edward L Glaeser, Michael Luca, and Christopher T Stanton. 2020. “What jobs are being done at home during the COVID-19 crisis? Evidence from firm-level surveys.” NBER Working Paper 27422.

Bayer, Christian, and Moritz Kuhn. 2020. “Intergenerational Ties and Case Fatality Rates: A Cross-Country Analysis.” IZA Discussion Paper 13114.

Bayham, Jude, and Eli P Fenichel. 2020. “Impact of School Closures for COVID-19 on the US Health-Care Workforce and Net Mortality: A Modelling Study.” *The Lancet: Public Health* 5 (5): e271–e278.

Beauregard, Pierre-Loup, Marie Connolly, Catherine Haeck, and Tímea Laura Molnár. 2020, August. “Primary School Reopenings and Parental Work.” Research Group on Human Capital Working Paper No. 20-06.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2017. "House Prices and Consumer Spending." *The Review of Economic Studies* 85 (3): 1502–1542.

Berger, David, Kyle Herkenhoff, and Simon Mongey. 2020. "An SEIR Infectious Disease Model with Testing and Conditional Quarantine." NBER Working Paper 26901.

Berger, David W., Kyle F. Herkenhoff, and Simon Mongey. 2019. "Labor Market Power." *NBER Working Paper Series* 25719:1–80.

Bessen, James. 2017. "Information Technology and Industry Concentration." *Boston University School of Law, Law & Economics Paper* 17, no. 41.

Bessen, James, Erich Denk, Joowon Kim, and Cesare Righi. 2020. "Declining Industrial Disruption." *Boston University School of Law, Law & Economics Paper* 20, no. 28.

Bick, Alexander. 2016. "The Quantitative Role of Child Care for Female Labor Force Participation and Fertility." *Journal of the European Economic Association* 14 (3): 639–668.

Bick, Alexander, and Adam Blandin. 2020. "Real-Time Labor Market Estimates During the 2020 Coronavirus Outbreak." Unpublished Manuscript, University of Arizona.

Bick, Alexander, Adam Blandin, and Karel Mertens. 2020. "Work from Home After the COVID-19 Outbreak." Unpublished Manuscript, University of Arizona.

Birinci, Serdar. 2019. "Spousal Labor Supply Response to Job Displacement and Implications for Optimal Transfers." Federal Reserve Bank of St. Louis Working Paper 2019-020B.

Biroli, Pietro, Steven Bosworth, Marina Della Giusta, Amalia Di Girolamo, Sylvia Jaworska, and Jeremy Vollen. 2020. "Family Life in Lockdown." IZA Discussion Paper 13398.

Black, Iain S. 1995. "Money, information and space: banking in early-nineteenth-century England and Wales." *Journal of Historical Geography* 21 (4): 398–412.

Bloom, Nicholas, Charles Jones, John Van Reenen, and Michael Webb. 2020. "Are Ideas Getting Harder to Find?" *American Economic Review* 110 (4): 1104–1144.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. "Consumption Inequality and Family Labor Supply." *American Economic Review* 106 (2): 387–435.

———. 2018. "Children, Time Allocation, and Consumption Insurance." *Journal of Political Economy* 126 (S1): S73–S115.

Bocola, Luigi. 2016. "The Pass-Through of Sovereign Risk." *Journal of Political Economy* 124 (4): 879–926.

Bordo, Michael D. 1998. "Michael D. Bordo." *Federal Reserve Bank of St. Louis Review*, no. May/June:77–82.

Borella, Margherita, Mariacristina De Nardi, and Fang Yang. 2018. "The Aggregate Implications of Gender and Marriage." *Journal of the Economics of Ageing* 11:6–26.

Bornstein, Gideon. 2018. "Entry and Profits in an Aging Economy: The Role of Consumer Inertia." pp. 909–927.

Bottero, Margherita, Simone Lenzu, and Filippo Mezzanotti. 2017. "Sovereign Debt Exposure and the Bank Lending Channel: Impact on Credit Supply and the Real Economy."

Bresnahan, Timothy F., and M. Trajtenberg. 1995. "General Purpose Technologies 'Engines of growth'?" *Journal of Econometrics* 65 (1): 83–108.

Brotherhood, Luiz, Philipp Kircher, Cezar Santos, and Michèle Tertilt. 2020. "An Economic Model of the COVID-19 Epidemic: The Importance of Testing and Age-Specific Policies." CEPR Discussion Paper 14695.

Brunt, Liam. 2006. "Rediscovering Risk: Country Banks as Venture Capital Firms in the First Industrial Revolution." *The Journal of Economic History* 66 (1): 74–102.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2018. "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies." *NBER Working Paper*, no. 25148.

Bullinger, Lindsey Rose, Jilian B. Carr, and Analisa Packham. 2020.

"COVID-19 and Crime: Effects of Stay-at-Home Orders on Domestic Violence." Unpublished Manuscript, Vanderbilt University.

Bundesagentur für Arbeit. 2020. "Monatsbericht zum Arbeits- und Ausbildungsmarkt, Mai." Statistik der Bundesagentur für Arbeit Berichte: Blickpunkt Arbeitsmarkt.

Cajner, Tomaz, Ryan A. Crane, Leland D.and Decker, John Grigsby, Adrian Hamins-Puertolas, Erik Hurst, Christopher Kurz, and Ahu Yildirmaz. 2020. "The U.S. Labor Market during the Beginning of the Pandemic Recession." Unpublished Manuscript, University of Chicago.

Calomiris, Charles, and Stephen Haber. 2014. *Fragile By Design: The Political Origins of Banking Crises and Scarce Credit*. Princeton, N.J.: Princeton University Press.

Calomiris, Charles, and Joseph Mason. 2003. "Consequences of Bank Distress During the Great Depression." *American Economic Review* 93 (3): 937–947.

Calomiris, Charles W., Matthew Jaremski, and David C. Wheelock. 2019. "Interbank Connections, Contagion and Bank Distress in the Great Depression."

Carley, Michael, Deepak Hegde, and Alan Marco. 2015. "What is the Probability of Receiving a U.S. Patent?" *Yale Journal of Law and Techonology*, vol. 203.

Carlson, Daniel L., Richard J. Petts, and Joanna Pepin. 2020. "Men and Women Agree: During the COVID-19 Pandemic Men Are Doing More at Home." Council on Contemporary Families, Brief Reports, May 20.

Cavenaile, Laurent, Murat Alp Celik, and Xu Tian. 2020. "Are Markups Too High? Competition, Strategic Innovation, and Industry Dynamics."

Census of Great Britain. 1821. Observations, enumeration and parish register abstracts.

Chatterjee, Satyajit, and Burcu Eyigungor. 2019. "The Firm Size and Leverage Relationship and Its Implications for Entry and Concentration in a Low Interest Rate World." *Federal Reserve Bank of Philadelphia, Working Papers, Research Department* 19, no. 18.

Chen, Yanan, and Kyle A. Kelly. 2019. "The Gender Difference in Wages and the Returns to Schooling Over the Great Recession in the U.S." *Research in Economics* 73 (2): 190–198.

Chodorow-Reich, Gabriel. 2014. "The Employment Effects of Credit Market Disruptions: Firm-Level Evidence from the 2008–9 Financial Crisis." *Quarterly Journal of Economics* 129 (1): 735–774.

Chodorow-Reich, Gabriel, Gita Gopinath, Prachi Mishra, and Abhinav Narayanan. 2018. "Cash and the Economy: Evidence From India's Demonetization." *NBER Working Paper* 25370:1–61.

Collins, Caitlyn, Liana C. Landivar, Leah Ruppanner, and William J. Scarborough. 2020. “COVID-19 and the Gender Gap in Work Hours.” Forthcoming, Gender, Work, and Organization.

Corhay, Alexandre, Howard Kung, and Lukas Schmid. 2020. “Q: Risk, Rents, or Growth ?”

Coskun, Sena, and Husnu Dalgic. 2020. “The Emergence of Procyclical Fertility: The Role of Gender Differences in Employment Risk.” CRC TR 224 Discussion Paper Series No. 142.

Costoya, Victoria, Lucia Echeverria, Maria Edo, Ana Rocha, and Agustina Thailinger. 2020. “The Impact of Covid-19 on the Allocation of Time Within Couples: Evidence for Argentina.” Unpublished Manuscript, Universidad de San Andrés.

Council of Economic Advisers. 2016. “Benefits of Competition and Indicators.” *Issue Brief*, no. May.

Covarrubias, Matias, German Gutierrez, and Thomas Philippon. 2019. “From Good to Bad Concentration? U.S. Industries over the Past 30 Years.” *NBER Macro Annual*.

Cowan, Benjamin W. 2020. “Short-run Effects of COVID-19 on U.S. Worker Transitions.” NBER Working Paper 27315.

Crouzet, By Nicolas, and Janice Eberly. 2018. “Intangibles, Investment, and Efficiency.” *AEA Papers and Proceedings* 108:426–431.

Crouzet, Francois. 1972. *Capital Formation in the Industrial Revolution*. Edited by Francois Crouzet. Bungay, Suffolk: Metheun and Co., Ltd.

Cubas, German, Chinhui Juhn, and Pedro Silos. 2019. "Coordinated Work Schedules and the Gender Wage Gap." Unpublished Manuscript, University of Houston.

Dahl, Gordon B., Katrine Loken, and Magne Mogstad. 2014. "Peer Effects in Program Participation." *American Economic Review* 104 (7): 2049–2074.

Danieli, Ana, and Jane Olmstead-Rumsey. 2020. "Sector-Specific Shocks and the Expenditure Elasticity Channel During the COVID-19 Crisis." Unpublished Manuscript, Northwestern University.

Davis, Steven J., and Till von Wachter. 2011. "Recessions and the Costs of Job Loss." *Brookings Papers on Economic Activity*, no. 2:1–72.

Dawson, Frank Griffith. 1990. *The First Latin American Debt Crisis*. New Haven: Yale University Press.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2016. "Where Has All the Skewness Gone? The Decline in High-Growth (Young) Firms in the U.S." *European Economic Review* 86:4–23.

Del Boca, Daniela, Noemi Oggero, Paola Profeta, and Maria Cristina Rossi. 2020. "Women's Work, Housework and Childcare, before and during COVID-19." *Covid Economics: Vetted and Real-Time Papers* 28:70–90.

de Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *Quarterly Journal of Economics* 135 (2): 561–644.

de Loecker, Jan, and Frederic Warzynski. 2012. "Markups and Firm-Level Export Status: Appendix." *American Economic Review* 102 (6): 2437–2471.

de Ridder, Maarten. 2020. "Market Power and Innovation in the Intangible Economy."

Dias, Felipe A., Joseph Chance, and Arianna Buchanan. 2020. "The motherhood penalty and the fatherhood premium in employment during covid-19: Evidence from the united states." *Research in Social Stratification and Mobility*, vol. 69.

Dingel, Jonathan I., Christina Patterson, and Joseph Vavra. 2020. "Childcare Obligations Will Constrain Many Workers When Reopening the US Economy." BFI Working Paper 2020-46.

Doepke, Matthias, Moshe Hazan, and Yishay D. Maoz. 2015. "The Baby Boom and World War II: A Macroeconomic Analysis." *Review of Economic Studies* 82 (3): 1031–1073.

Doepke, Matthias, and Fabian Kindermann. 2019. "Bargaining over Babies: Theory, Evidence, and Policy Implications." *American Economic Review* 109 (9): 3264–3306.

Doepke, Matthias, and Michèle Tertilt. 2016. "Families in Macroeconomics." Chapter 23 of *Handbook of Macroeconomics, Vol. 2*. North Holland.

Duffy, Ian P.H. 1973. *Bankruptcy & Insolvency in London During the Industrial Revolution*. New York & London: Garland Publishing.

Dunbar, Geoffrey R. 2013, October. "American Dad: Families, Fathers and Childcare Shocks." Unpublished Manuscript, Simon Fraser University.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu. 2018. "How Costly Are Markups?"

Eggertsson, Gauti, Neil Mehrotra, and Jacob Robbins. 2019. "A Model of Secular Stagnation." *American Economic Journal: Macroeconomics* 11, no. 1.

Eggertsson, Gauti, Jacob Robbins, and Ella Getz Wold. 2018. "Kaldor and Piketty's Facts: The Rise of Monopoly Power in the United States." pp. 1–69.

Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt. 2020. "The Macroeconomics of Epidemics." NBER Working Paper 26882.

Ellieroth, Kathrin. 2019. "Spousal Insurance, Precautionary Labor Supply, and the Business Cycle." Unpublished Manuscript, Indiana University.

Elliott, Matthew, Benjamin Golub, and Matthew O Jackson. 2014. "Financial Networks and Contagion." *American Economic Review* 104 (10): 3115–3153.

Engbom, Niklas. 2019. "Firm and Worker Dynamics in an Aging Labor Market." *Working Paper*, pp. 1–115.

Erdsiek, Daniel. 2020. "Companies Plan to Keep Remote Work Arrangements After Crisis." ZEW News, September/October.

Erosa, Andrés, Luisa Fuster, Gueorgui Kambourov, and Richard Rogerson. 2017. "Hours, Occupations, and Gender Differences in Labor Market Outcomes." NBER Working Paper No. 20271.

Farré, Lídia, Yarine Fawaz, Libertad González, and Jennifer Graves. 2020, July. "How the COVID-19 Lockdown Affected Gender Inequality in Paid and Unpaid Work in Spain." IZA Discussion Paper No. 13434.

Farré, Lídia, and Libertad González. 2019. "Does Paternity Leave Reduce Fertility?" *Journal of Public Economics* 172:52–66.

Fernald, John. 2014. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." *Federal Reserve Bank of San Francisco*.

Fernández, Raquel. 2013. "Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century." *American Economic Review* 103 (1): 472–500.

Fernández, Raquel, Alessandra Fogli, and Claudia Olivetti. 2004. "Mothers and Sons: Preference Formation and Female Labor Force Dynamics." *Quarterly Journal of Economics* 119 (4): 1249–99.

Fernando, Chitru S., Anthony D. May, and William L. Megginson. 2012. "The value of investment banking relationships: Evidence from the collapse of Lehman brothers." *Journal of Finance* 67 (1): 235–270.

Flandreau, Marc, and Juan H. Flores. 2009. "Bonds and brands: Foundations of sovereign debt markets, 1820-1830." *Journal of Economic History* 69 (3): 646–684.

Fogli, Alessandra, and Laura Veldkamp. 2011. "Nature or Nurture? Learning and the Geography of Female Labor Force Participation." *Econometrica* 79 (4): 1103–38.

Friedman, Milton, and Anna Jacobson Schwartz. 1963. *A Monetary History of the United States, 1867-1960*. Princeton, N.J.: Princeton University Press.

Frydman, Carola, Eric Hilt, and Lily Y. Zhou. 2015. "Economic Effects of Runs on Early "Shadow Banks": Trust Companies and the Impact of the Panic of 1907." *Journal of Political Economy* 123 (4): 902–940.

Fuchs-Schündeln, Nicola, Moritz Kuhn, and Michèle Tertilt. 2020. "The Short-Run Macro Implications of School and Child-Care Closures." CEPR Discussion Paper 14882.

Fukui, Masao, Emi Nakamura, and Jón Steinsson. 2019. "Women, Wealth Effects, and Slow Recoveries." Unpublished Manuscript, University of Berkeley.

Gallen, Yana. 2018. “Motherhood and the Gender Productivity Gap.” Unpublished Manuscript, University of Chicago.

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow. 2019. “How Destructive Is Innovation?” *Econometrica* 87 (5): 1507–1541.

García-Pérez, J. Ignacio, and Sílvio Rendon. 2020. “Family Job Search and Wealth: The Added Worker Effect Revisited.” NBER Working Paper 26901.

Gayer, Arthur, W.W. Rostow, and Anna Jacobson Schwartz. 1975. *The Growth and Fluctuation of the British Economy 1790-1850, vol. 1 & 2*. Second. Brighton: The Harvester Press.

Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and Jose-Victor Rios Rull. 2020. “Health versus Wealth: On the Distributional Effects of Controlling a Pandemic.” CEPR Discussion Paper 14606.

Goldberg, Jonathan, David Lopez-Salido, and Craig A Chikis. 2020. “Interest Rates and the Innovation Multiplier.”

Goldin, Claudia. 2014. “A Grand Gender Convergence: Its Last Chapter.” *American Economic Review* 104 (4): 1091–1119.

Goldin, Claudia, and Lawrence F. Katz. 2011. “The Cost of Workplace Flexibility for High-Powered Professionals.” *The Annals of the American Academy of Political and Social Science* 638 (1): 45–67.

Goldin, Claudia, and Claudia Olivetti. 2013. “Shocking Labor Supply: A

Reassessment of the Role of World War II on Women's Labor Supply." *American Economic Review* 103 (3): 257–62.

Graham, Stuart, and Deepak Hegde. 2015. "U.S. Patents and Their Disclosure Status." *Science* 347, no. 6219.

Great Britain. 1832. Parliament. House of Commons. Select Committee on Bank of England Charter. Report from the Committee of Secrecy on the Bank of England Charter: with the minutes of evidence, appendix and index. London: J. & L. G. Hansard & Sons, Printers to the House of Commons. Accessed via Hathi Trust on 14 March 2018.

Greenwood, Jeremy, Ananth Seshadri, and Mehmet Yorukoglu. 2005. "Engines of Liberation." *Review of Economic Studies* 72 (1): 109–33.

Gregory, Victoria, Guido Menzio, and David Wiczer. 2020. "Pandemic Recession: L- or V-Shaped?" *Federal Reserve Bank of Minneapolis Quarterly Review* 40 (1): 1–28.

Griliches, Zvi. 1998. "Patent Statistics as Economic Indicators: A Survey." Chapter 13 of *R&D and Productivity: The Econometric Evidence*, 287–343. University of Chicago Press.

———. 2001. *R&D, Education, and Productivity: A Retrospective*. Cambridge, Mass.: Harvard University Press.

Grosjean, Pauline, and Rose Khattar. 2018. "It's Raining Men! Hallelujah? The Long-Run Consequences of Male-Biased Sex Ratios." *Review of Economic Studies* 86 (2): 723–754.

Grullon, Gustavo, Yelena Larkin, and Roni Michaely. 2019. "Are US Industries Becoming More Concentrated?" *Review of Finance* 23 (4): 697–743.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning. 2020. "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" NBER Working Paper 26918.

Guler, Bulent, Fatih Guvenen, and Gianluca Violante. 2012. "Joint-Search Theory: New Opportunities and New Frictions." *Journal of Monetary Economics* 59 (4): 352–369.

Guner, Nezih, Remzi Kaygusuz, and Gustavo Ventura. 2012. "Taxation and Household Labour Supply." *Review of Economic Studies* 79:1113–1149.

———. 2020. "Child-Related Transfers, Household Labor Supply and Welfare." *Review of Economic Studies* 87 (5): 2290–2321 (October).

Guner, Nezih, Yuliya Kulikova, and Arnau Valladares-Esteban. 2020. "Does the Added Worker Effect Matter?" Unpublished Manuscript, CEMFI.

Hagedorn, Marcus, Iourii Manovski, and Kurt Mitman. 2019. "The Fiscal Multiplier." NBER Working Paper 25571.

Hall, Robert E. 2015. "Quantifying the Lasting Harm to the US Economy from the Financial Crisis." *NBER Macroeconomics Annual* 29 (1): 71–128.

Hannusch, Anne. 2019, January. “Taxing Families: The Impact of Child-related Transfers on Maternal Labor Supply.” CRC TR 224 Discussion Paper Series, No. 67.

Hansen, Mary Eschelbach, and Nicolas L Ziebarth. 2017. “Credit Relationships and Business Bankruptcy during the Great Depression.” *American Economic Journal: Macroeconomics* 9 (2): 228–255.

Harris, Christopher, and John Vickers. 1987. “Racing with Uncertainty.” *Review of Economic Studies* 54 (1): 1–21.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante. 2020. “The Rise of US Earnings Inequality: Does the Cycle Drive the Trend?” NBER Working Paper 27345.

Heblich, Stephan, and Alex Trew. 2019. “Banking and Industrialization.” *Journal of the European Economic Association* 00 (0): 1–44.

Heggeness, Misty L. 2020, June. “Why Is Mommy So Stressed? Estimating the Immediate Impact of the COVID-19 Shock on Parental Attachment to the Labor Market and the Double Bind of Mothers.” Opportunity and Inclusive Growth Institute, Federal Reserve Bank of Minneapolis, Working Paper No. 33.

Helpman, Elhanan. 1998. *General Purpose Technologies and Economic Growth*. Edited by Elhanan Helpman. Cambridge: MIT Press.

Hereford Journal. 1825. Friday’s Post. 14 December 1825. Accessed via Gale NewsVault on 14 March 2018.

Hoberg, Gerard, and Gordon Phillips. 2010. “Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis.” *Review of Financial Studies* 23 (10): 3773–3811.

Hobjin, Bart, and Fernanda Nechio. 2019. “Sticker Shocks: Using VAT Changes to Estimate Upper-Level Elasticities of Substitution.” *Journal of the European Economic Association* 17 (3): 799–833.

Hopenhayn, Hugo, Julian Neira, and Rish Singhania. 2018. “From Population Growth to Firm Demographics: Implications for Concentration, Entrepreneurship and the Labor Share.” *National Bureau of Economic Research Working Paper Series*, pp. 1–55.

Hsieh, Chang-Tai, and Esteban Rossi-Hansberg. 2020. “The Industrial Revolution in Services.” pp. 139–158.

Huber, Kilian. 2018. “Disentangling the Effects of a Banking Crisis: Evidence from German Firms and Counties.” *American Economic Review* 108 (3): 868–898.

Hupkau, Claudia, and Barbara Petrongolo. 2020, October. “Work, Care and Gender during the Covid-19 Crisis.” CEPR Discussion Paper No. 15358.

Iacopo, Morchio, and Christian Moser. 2020. “The Gender Pay Gap: Micro Sources and Macro Consequences.” Unpublished Manuscript, Columbia University.

Indarte, Sasha. 2016. "Contagion via Financial Intermediaries in Pre-1914 Sovereign Debt Markets." *Working Paper*, pp. 1–36.

Iyer, Rajkamal, and Jose-Luis Peydro. 2015. "Interbank Contagion at Work: Evidence from a Natural Experiment." *Review of Financial Studies* 24 (4): 1337–1377.

James, John A. 2012. "Panics, payments disruptions and the Bank of England before 1826." *Financial History Review* 19 (3): 289–309.

James, John A., James McAndrews, and David F. Weiman. 2013. "Wall Street and Main Street: The macroeconomic consequences of New York bank suspensions, 1866-1914." *Cliometrica* 7 (2): 99–130.

Jarosch, Gregor. 2015. "Searching for Job Security and the Consequences of Job Loss." Unpublished Manuscript, Princeton University.

Jaume, David, and Alexander Willén. 2018, December. "The Neglected Impact of School Disruptions." NHH Discussion Paper, SAM 30-2018.

Jones, Callum, and Thomas Philippon. 2016. "The Secular Stagnation of Investment?" *Working Paper*, no. December:1–25.

Jovanovic, Boyan, and Peter L Rousseau. 2005. *General Purpose Technologies*. Volume 1. Edited by Philippe Aghion and Steven N. Durlauf. Elsevier B.V.

Kahn, Lisa B., Fabian Lange, and David G. Wiczer. 2020. "Labor Demand in the Time of COVID-19: Evidence from Vacancy Postings and UI Claims." NBER Working Paper 27061.

Kaplan, Greg, and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82 (4): 1199–1239.

Karahan, Fatih, Benjamin W. Pugsley, and Aysegul Sahin. 2019. "Demographic Origins of the Startup Deficit." *NBER Working Paper*, no. 25874.

Kehrig, Matthias, and Nicolas Vincent. 2020. "The Micro-Level Anatomy of the Decline of the Labor Share." *NBER Working Paper*, no. 25275:1689–1699.

Kelly, Bryan T., Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2018. "Measuring Technological Innovation over the Long Run." *SSRN Electronic Journal*.

Kenny, Sean, Jason Lennard, and John D. Turner. 2017. "The Macroeconomic Effects of Banking Crises: Evidence from the United Kingdom , 1750-1938."

Kenny, Seán, and John D. Turner. 2019. "Wildcat Bankers or Political Failure ? The Irish Financial Pantomime, 1797-1826."

Kikuchi, Shinnosuke, Sagiri Kitao, and Minamo Mikoshiba. 2020, July. "Who Suffers from the COVID-19 Shocks? Labor Market Heterogeneity and Welfare Consequences in Japan." Research Group on Human Capital Working Paper No. 20-06.

Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *Journal of Political Economy* 112 (5): 986–1018.

Kleven, Henrik, Camille Landais, Johanna Posch, Andreas Steinhauer, and Josef Zweimüller. 2019. "Child Penalties across Countries: Evidence and Explanations." *AEA Papers and Proceedings* 109:122–26.

Kleven, Henrik, Camille Landais, and Jakob Egholt Søgaard. 2019. "Children and Gender Inequality: Evidence from Denmark." *American Economic Journal: Applied Economics* 11 (4): 181–209.

Knott, Anne Marie, and Carl Vieregger. 2016. "Reconciling the Firm Size and Innovation Puzzle." *CES Working Paper*.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth." *Quarterly Journal of Economics*, pp. 665–712.

Kotsadam, Andreas, and Henning Finseraas. 2011. "The state intervenes in the battle of the sexes: Causal effects of paternity leave." *Social Science Research* 40:1611–1622.

———. 2013. "Causal Effects of Parental Leave on Adolescents' Household Work." *Social Forces* 92:329–351.

Kristal, Tali, and Meir Yaish. 2020. "Does the coronavirus pandemic level the gender inequality curve? (It doesn't)." *Research in Social Stratification and Mobility*, vol. 68.

Krueger, Dirk, and Chunzan Wu. 2019. "Consumption Insurance Against Wage Risk: Family Labor Supply and Optimal Progressive Income Taxation." Forthcoming, American Economic Journal: Macroeconomics.

Lemieux, Thomas, Kevin Milligan, Tammy Schirle, Wilfrid Laurier, and Mikal Skuterud. 2020. "Initial Impacts of the COVID-19 Pandemic on the Canadian Labour Market." Forthcoming, Canadian Public Policy.

Lentz, Rasmus, and Dale Mortensen. 2008. "An Empircal Model of Growth Through Product Innovation." *Econometrica* 76 (6): 1317–1373.

Leslie, Emily, and Riley Wilson. 2020. "Sheltering in Place and Domestic Violence: Evidence from Calls for Service during COVID-19." Forthcoming, Journal of Public Economics.

Liu, Ernest, Atif R. Mian, and Amir Sufi. 2019. "Low Interest Rates, Market Power, and Productivity Growth."

Long, J. Scott, and Jeremy Freese. 2014. *Regression models for categorical dependent variables using Stata.* 3rd. College Station, Texas: Stata Press.

Lucas, Robert E. 1977. "Understanding Business Cycles." *Carnegie-Rochester Conference Series on Public Policy* 5:7–29.

Lundberg, Shelly. 1985. "The Added Worker Effect." *Journal of Labor Economics* 3 (1): 11–37.

Luttmer, Erzo G.J. 2007. "Selection, Growth, and the Size Distribution of Firms." *Quarterly Journal of Economics*, no. August:1103–1145.

Ma, Sen, Zhengyun Sun, and Hao Xue. 2020, June. "Childcare Needs and Parents' Labor Supply: Evidence from the COVID-19 Lockdown." Unpublished Manuscript, Harvard University.

Malkov, Egor. 2020. "Nature of Work and Distribution of Risk: Evidence from Occupational Sorting, Skills, and Tasks." Unpublished Manuscript, University of Minnesota.

Mankart, Jochen, and Rigas Oikonomou. 2017. "Household Search and the Aggregate Labor Market." *Journal of Economic Perspectives* 84 (4): 1735–1788.

Marchand, Joseph, and Sara Olfert. 2013. "The US Gender Gap Through the Great Recession Using an Alternative Approach to Cyclicality." *Applied Economics Letters* 20 (3): 276–281.

Marriner, Sheila. 1980. "English Bankruptcy Records and Statistics before 1850." *The Economic History Review* 33 (3): 351–366.

Mathias, Peter. 1973. "Capital, Credit and Enterprise in the Industrial Revolution." *Journal of European Economic History* 2 (1): 121.

Mennuni, Alessandro. 2019. "The Aggregate Implications of Changes in the Labour Force Composition." *European Economic Review* 116:83–106.

Mian, Atif, and Asim Sufi. 2008. "Tracing the Impact of Bank Liquidity Shocks : Evidence from an Emerging Market." *American Economic Review* 98 (4): 1413–1442.

Mian, Atif R., and Amir Sufi. 2014. "What Explains the 2007-2009 Drop in Employment?" *Econometrica* 82 (6): 2197–2223.

Miller, Amalia R. 2011. "The Effects of Motherhood Timing on Career Path." *Journal of Population Economics* 24 (3): 1071–1100.

Mitman, Kurt, and Stanislav Rabinovich. 2020. "Optimal Unemployment Benefits in the Pandemic." Unpublished Manuscript, IIES Stockholm.

Mongey, Simon. 2017. "Market Structure and Monetary Non-neutrality."

Mongey, Simon, Laura Pilossoph, and Alex Weinberg. 2020. "Which Workers Bear the Burden of Social Distancing Policies?" NBER Working Paper 27085.

Montenovo, Laura, Xuan Jiang, Felipe Lozano Rojas, Ian M. Schmutte, Kosali I. Simon, Bruce A. Weinberg, and Coady Wing. 2020. "Determinants of Disparities in Covid-19 Job Losses." NBER Working Paper 27132.

Morgan, Don, and James Narron. 2015. Crisis Chronicles: The Panic of 1825 and the Most Fantastic Financial Swindle of All Time. Liberty Street Economics Blog. Accessed on 14 March 2018.

Murphy, Anne. 2014. "The financial revolution and its consequences." Chapter 11 of *The Cambridge Economic History of Modern Britain, Volume 1*, edited by Roderick Floud, Jane Humphries, and Paul Johnson, 321–342. Cambridge: Cambridge University Press.

Möhring, Katja, Elias Naumann, Maximiliane Reifenscheid, Annelies G. Blom, Alexander Wenz, Tobias Rettig, Roni Lehrer, Ulrich Krieger, Sebastian Juhl, Sabine Friedel, Marina Fikel, and Carina Cornesse. 2020. "Mannheimer Corona-Studie: Schwerpunktbericht zu Erwerbst ätigkeit und Kinderbetreuung." Unpublished Manuscript, University of Mannheim.

Neal, Larry. 1998. "The Financial Crisis of 1825 and the Restructuring of the British Financial System." *Federal Reserve Bank of St. Louis Review*, no. May/June:53–76.

Northampton Mercury. 1826, jan. Friday and Saturday's Post. 7 January 1826. Accessed via Gale NewsVault on 16 March 2019.

Olivetti, Claudia, Eleonora Patacchini, and Yves Zenou. 2020. "Mothers, Peers, and Gender-Role Identity." *Journal of the European Economic Association* 18 (1): 266–301.

Olley, G. Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 1263–1297.

Olsson, Jonna. 2019. "Structural Transformation of the Labor Market and the Aggregate Economy." Unpublished Manuscript, University of Amsterdam.

Ongena, Steven, José Luis Peydró, and Neeltje Van Horen. 2015. "Shocks Abroad, Pain at Home? Bank-Firm-Level Evidence on the International Transmission of Financial Shocks." *IMF Economic Review* 63 (4): 698–750.

Oreffice, Sonia, and Climent Quintana-Domeque. 2020. "Gender Inequality in COVID-19 Times: Evidence from UK Prolific Participants." Unpublished Manuscript, University of Exeter.

Oreopoulos, Philip, Till von Wachter, and Andrew Heisz. 2012. "The Short- and Long-Term Career Effects of Graduating in a Recession." *American Economic Journal: Applied Economics* 4 (1): 1–29.

Ortigueira, Salvador, and Nawid Siassi. 2013. "How Important is Intrahousehold Risk Sharing for Savings and Labor Supply?" *Journal of Monetary Economics* 60 (6): 650–666.

Papanikolaou, Dimitris, and Lawrence D. W. Schmidt. 2020. "Working Remotely and the Supply-side Impact of Covid-19." Unpublished Manuscript, Northwestern University.

Patnaik, Ankita. 2019. "Causal Effects of Parental Leave on Adolescents' Household Work." *Journal of Labor Economics* 37 (4): 1009–1059.

Patterson, Christina. 2019. "The Matching Multiplier and the Amplification of Recessions." Unpublished Manuscript, Northwestern University.

Pease, Joseph. 1848. "Secret Committee of the House of Commons on Commercial Distress, 1847-1848. Testimony of Joseph Pease." Technical Report, House of Commons Reports and Papers, London.

Pellegrino, Bruno. 2020. "Product Differentiation, Oligopoly, and Resource Allocation."

Peters, Michael, and Conor Walsh. 2019. "Declining Dynamism, Increasing Markups and Missing Growth: The Role of the Labor Force." *SSRN Electronic Journal*, pp. 1–49.

Peters, Ryan H., and Lucian A. Taylor. 2017. "Intangible Capital and the Investment-q Relation." *Journal of Financial Economics* 123 (2): 251–272.

Pilossoph, Laura, and Shu Lin Wee. 2020. "Household Search and the Marital Wage Premium." Forthcoming, American Economic Journal: Macroeconomics.

Pollard, Sidney. 1964. "Fixed Capital in the Industrial Revolution in Britain." *Journal of Economic History* 24 (3): 299–314.

Pressnell, L.S. 1956. *Country Banking in the Industrial Revolution*. Oxford: Clarendon Press.

Price, Frederick George Hilton. 1890. *A handbook of London bankers*. London: Simpkin, Marshall, Hamilton, Kent and co.

Prickett, Kate C., Alexa Martin-Storey, and Robert Crosnoe. 2015. "A Research Note on Time With Children in Different- and Same-Sex Two-Parent Families." *Demography* 52:905–18.

Pruitt, Seth, and Nicholas Turner. 2020. "Earnings Risk in the Household: Evidence from Millions of US Tax Returns." *American Economic Review: Insights* 2 (2): 237–54.

Pugsley, Benjamin W., Petr Sedlacek, and Vincent Sterk. 2020. "The Nature of Firm Growth."

Qian, Yue, and Sylvia Fuller. 2020. "COVID-19 and the Gender Employment Gap among Parents of Young Children." *Canadian Public Policy*, August, S89–S101.

Rege, Mari, and Ingeborg F Solli. 2013. "The impact of paternity leave on fathers' future earnings." *Demography* 50 (6): 2255–2277.

Rivera, Carolina, Yu-Chieh Hsu, Fernanda Pavez Esbry, and Esuna Dugarova. 2020. "Gender Inequality and the Covid-19 Crisis: A Human Development Perspective." UNDP Human Development Working Paper.

Rosen, Richard J. 1991. "Research and Development with Asymmetric Firm Sizes." *RAND Journal of Economics* 22 (3): 411–429.

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Nicholas Trachter. 2020. "Diverging Trends in National and Local Concentration." *NBER Macroeconomics Annual* 35:1–39.

Sargent, Thomas, and Francois Velde. 2002. *The Big Problem of Small Change*. Princeton, N.J.: Princeton University Press.

Schoonbroodt, Alice. 2018. "Parental Child Care During and Outside of Typical Work Hours." *Review of the Economics of the Household* 16:453–476.

Schwandt, Hannes, and Till von Wachter. 2019. "Unlucky Cohorts: Estimating the Long-Term Effects of Entering the Labor Market in a Recession in Large Cross-Sectional Data Sets." *Journal of Labor Economics* 37 (S1): S161–S198.

Sevilla, Almudena, and Sarah Smith. 2020. "Baby Steps: The Gender Division of Childcare During the COVID19 Pandemic." CEPR Discussion Paper 14804.

Shannon, H.A. 1934. "Bricks: A Trade Index, 1785-1849." *Economica* 1 (3): 300–318.

Silberling, Norman. 1919. "British Financial Experience, 1790-1830." *Review of Economics and Statistics* 1 (4): 282–297.

Solon, Gary, Robert Barsky, and Jonathan A. Parker. 1994. "Measuring the Cyclicality of Real Wages: How Important is Composition Bias." *Quarterly Journal of Economics* 109 (1): 1–25.

Stevens, Ann Huff. 1997. "Persistent Effects of Job Displacement: The Importance of Multiple Job Losses." *Journal of Labor Economics* 15 (1): 165–188.

Syverson, Chad. 2017. "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown." *Journal of Economic Perspectives*.

Tamm, Marcus. 2019. "Fathers' Parental Leave-Taking, Childcare Involvement and Labor Market Participation." *Labour Economics* 59:184–197.

Temin, Peter, and Hans Joachim Voth. 2013. *Prometheus Shackled: Goldsmith Banks and England's Financial Revolution After 1700*. Oxford: Oxford University Press.

The Examiner. 1825a, dec. State of the Money Market. 25 December 1825. Accessed via Gale NewsVault on 14 March 2018.

———. 1825b, nov. The Examiner. 20 November 1825. Accessed via Gale NewsVault on 14 March 2018.

The Morning Chronicle. 1825, dec. The Lord Mayor's Court. 8 December 1825. Accessed via Gale NewsVault on 14 March 2018.

Turner, John D. 2014. *Banking in Crisis: The Rise and Fall of British Banking Instability, 1800 to Present*. Cambridge: Cambridge University Press.

van Reenen, John. 2018. "Increasing Differences Between Firms: Market Power and the Macro-Economy." *CEP Discussion Paper*, no. 1576.

von Gaudecker, Hans-Martin, Radost Holler, Lena Janys, Bettina Siflinger, and Christian Zimpelmann. 2020. "The Early Impact of the COVID-19 Pandemic on Gender Inequality: Working Hours and Childcare Arrangements." Unpublished Manuscript, University of Bonn.

Wang, Haomin. 2019. "Intra-household Risk Sharing and Job Search Over the Business Cycle." *Review of Economic Dynamics* 34:165–182.

Wilde, Joshua, Wei Chen, and Sophie Lohmann. 2020. "COVID-19 and the Future of US Fertility: What Can We Learn from Google?" *Covid Economics: Vetted and Real-Time Papers* Issue 54:158–190.

Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross-Section and Panel Data*. MIT Press.

Wozniak, Abigail. 2020. "Disparities and Mitigation Behavior during COVID-19." Opportunity and Inclusive Growth Institute Working Paper 32.

Xiao, Pengpeng. 2020. "Wage and Employment Discrimination by Gender in Labor Market Equilibrium." Unpublished Manuscript, Yale University.

Zachariadis, Marios. 2003. "R&D, Innovation, and Technological Progress: A Test of the Schumpeterian Framework Without Scale Effects." *Canadian Journal of Economics/Revue Canadienne d'Economique* 36 (3): 566–586.

Zhang, Linyi. 2018. "Escaping Chinese Import Competition? Evidence from U.S. Firm Innovation." *Working Paper*, pp. 1–69.

Appendix A

Appendix: Market Concentration and the Productivity Slowdown

A.1 Data Appendix

The main source of data for the paper is the Compustat Fundamentals Annual database, 1962-2017 (though most analysis focuses on the post-1980 period). I restrict attention to firms incorporated in the U.S. reporting in U.S. dollars. I further restrict attention to non-financial, non-agricultural, non-utilities firms.

A.1.1 Data Sources and Moment Computations

Table A.1 lists the source and, where necessary, computation method for each target moment from the data.

Moment	Source	Computation/Series Name
TFP growth	Fernald (2014)	Utilization-adjusted annual total factor productivity growth
Leader market share	Compustat	Average of sales share (SALE) of largest firm in each 4-digit SIC industry (weighted by industry size)
Patent quality $\equiv$ patent stock growth per patent (psgpp)	Kogan et al. (2017)	$rTsm_{it} = \frac{Tsm_{it}}{GDPdefl_t}$ is the real value of firm i's patents issued in year t. $psgpp_{it} = \frac{\frac{rTsm_{it}}{fNpats_{it}}}{\sum_{s=1}^{t-1} rTsm_{is}}$; s =first year in Compustat. Citation-based version substitutes Tcw (not deflated).
R&D share of GDP	OECD Main Science and Technology Indicators	Business Expense R&D (private)/GDP
R&D intensity	Compustat	XRD/SALE, mean across all firms with real sales over 1 million in 2012 USD, assuming 0 if XRD missing.
Profit share of GDP	Bureau of Economic Analysis/FRED	Profits after tax with inventory valuation and capital consumption adjustments/Gross domestic income
Leader's share of R&D	Compustat	Average sales leader share of total R&D in 4-digit sector (weighted by industry size)
Leadership turnover	Compustat	Share of 4-digit SIC industries with new sales leader per year

Table A.1: Data sources and computation method for each moment used in the text.

A.1.2 Additional Patent Quality Figures

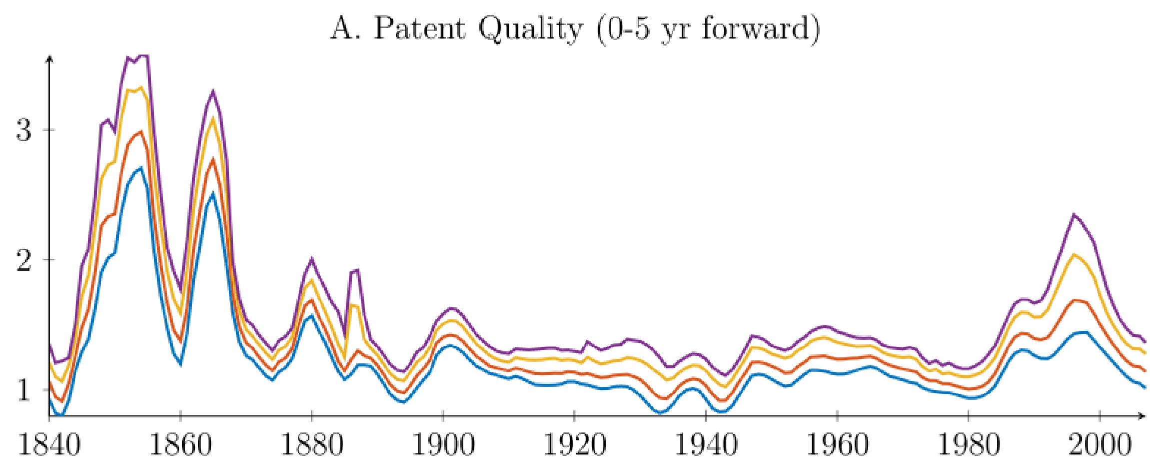

Figure A.1: Percentiles of text-based patent quality distribution over time. Blue = P50, Red = P75, Yellow = P90, Purple = P95. Source: Kelly et al. (2018) Figure 3a.

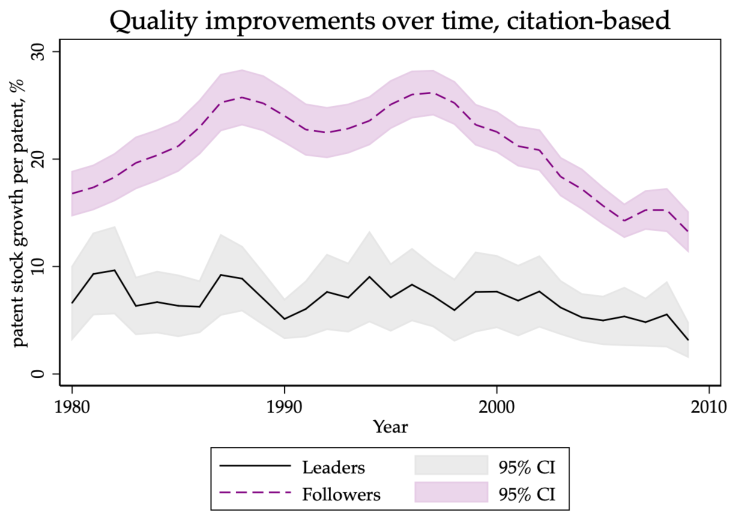

Figure A.2: Contribution of average new patent to firm's existing stock of patents, substituting forward citations counts for dollar value, from Kogan et al. (2017). Leader indicates sales leaders in 4-digit SIC industries and followers are all other firms.

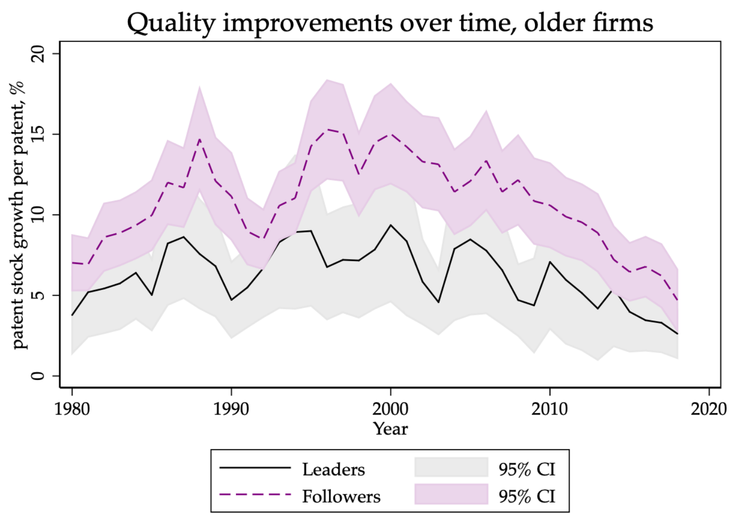

Figure A.3: Contribution of average new patent to value of firm's existing stock of patents, using estimated patent values from Kogan et al. (2017). Leader indicates sales leaders in 4-digit SIC industries and followers are all other firms, restricting attention to firms that have been public at least 20 years in the year patent is issued.

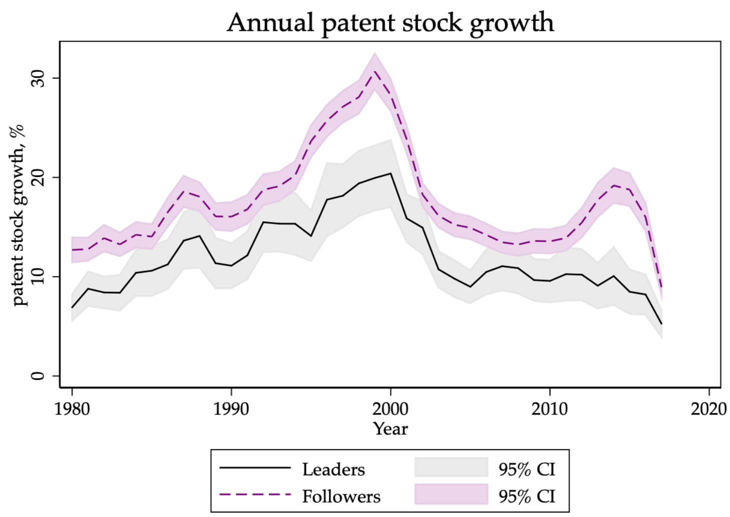

Figure A.4: Average annual growth of firm's patent stock conditional on patenting at least once in that year, using estimated patent values from Kogan et al. (2017). Leader indicates sales leaders in 4-digit SIC industries and followers are all other firms.

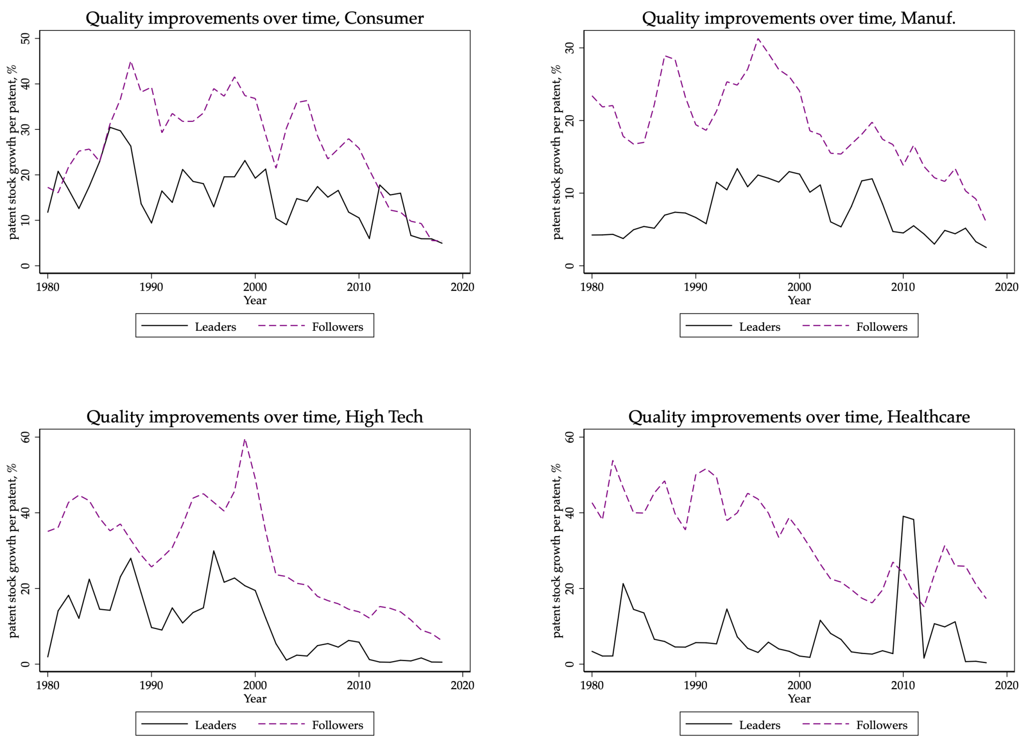

Figure A.5: Average patent quality differences between leaders and followers in Fama-French 5 broad industry categories (excluding "Other" category), using estimated patent values from Kogan et al. (2017). Leader indicates sales leaders in 4-digit SIC industries and followers are all other firms.

A.1.3 TFP and Markup Estimation

I use Compustat data on U.S. public firms from 1962-2017 to estimate revenue-based total factor productivity (TFPR) and markups at the firm level. I focus on the non-farm, non-financial sector and exclude utilities and firms without an industry classification. I keep only those companies that are incorporated in the U.S. The sample includes around 3,000 firms per year, though this number varies over time.

I construct each firm's capital stock $K_{i,t}$ by initializing the capital stock as PPEGT (total gross property, plant, and equipment) for the first year the firm appears. I then construct $K_{i,t+1}$ recursively:

$$K_{i,t+1} = K_{i,t} + I_{i,t+1} - \delta K_{i,t}$$

where PPENT (total net property, plant, and equipment) is used to capture the last two terms (net investment). I deflate the nominal capital stock using the Bureau of Economic Analysis (BEA) deflator for non-residential fixed investment.

In de Loecker and Warzynski (2012) the authors show that under a variety of pricing models firm i's markup at time t, μ_{it}, can be computed as a function of the output elasticity θ_{it}^{V} of any variable input and the variable input's cost share of revenue[1] :

$$\mu_{it} = \theta_{it}^{V} \frac{P_{it} Q_{it}}{P_t^V V_{it}} \tag{A.1}$$

where P_{it} is the output price of firm i's good at time t, Q_{it} its output, P_t^V the price of the variable input and V_{it} the amount of the input used.

Following de Loecker, Eeckhout, and Unger (2020) I use COGS (cost of goods sold) deflated by the BEA's GDP deflator series as the real variable input cost $M_{i,t}$ of the firm. While the number of employees is well measured in Compustat and would be sufficient to estimate productivity, the wage bill is usually not available and would be needed to compute the labor cost share needed to compute the markup simultaneously with productivity.

For the results presented in this paper, I assume a Cobb-Douglas production function[2] for firm i in 2-digit SIC sector s in year t so that factor shares may vary across sectors but not over time:

$$Y_{i,s,t} = A_{i,s,t} M_{i,s,t}^{\beta_{M,s}} K_{i,s,t}^{\beta_{K,s}}$$

I use the variable SALE to measure firm output $Y_{i,s,t}$. I deflate SALE using the GDP deflator series to obtain real revenue at the firm level. I include

[1]This approach requires several assumptions. First, the production technology must be continuous and twice differentiable in its arguments. Second, firms must minimize costs. Third, prices are set period by period. Fourth, the variable input has no adjustment costs. No particular form of competition among firms need be assumed.

[2]Alternative estimation of a translog production function yielded similar estimates.

firm and time fixed effects and obtain revenue-based TFP in logs (lower case variables denote variables in logs) by computing the residual (including fixed effects) of the following regressions for each 2-digit sector:

$$y_{i,t} = \alpha + \eta_t + \delta_i + \beta_{M,s} m_{i,t} + \beta_{K,s} k_{i,t-1} + \varepsilon_{i,t}.$$

In the above equation, $\beta_{M,s}$ captures the sector specific variable output elasticity, so I use equation A.1 to obtain the markup from the estimated $\hat{\beta}_{M,s}$ and the inverse cost share $\frac{SALE}{COGS}$.

A.1.4 Industry Profit Shares

The competitive fringe assumption generates empirically plausible predictions about profit shares: the largest U.S. public firms (by sales) capture by far the largest share of industry profits (see Figure A.6).[3]

A.1.5 Additional Model Validation Figures

An empirical exploration of the causal relationships among productivity growth, productivity gaps, and concentration is beyond the scope of this paper. However, especially given the sectoral heterogeneity in the decline in laggards patent quality in Figure A.5 which suggests that the extent of this phenomena differs across industries, we might expect rising con-

[3]TFP and sales share are correlated, and the figure looks similar if one uses a productivity ranking instead of sales-share based ranks.

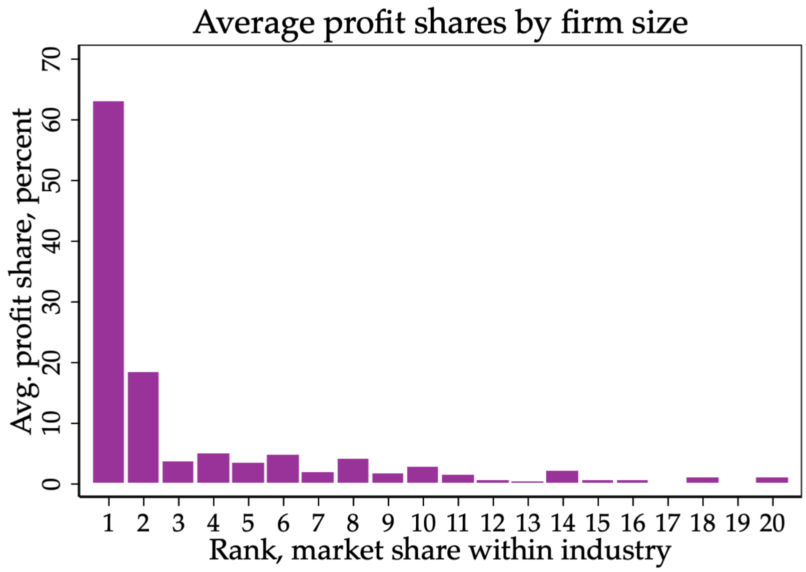

Figure A.6: Source: Compustat, 1975-2015. Firms are ranked by market share (sales) within 4-digit SIC industries, and these ranks are compared to profit shares (firm's own operating income (OIDBP) as a share of industry-total operating income). The Figure averages across 4-digit sectors.

centration and the productivity slowdown to be correlated at the sector level. I use data from Bureau of Economic Analysis estimates of multi-factor productivity[4] at the 3-digit NAICS level and data from Compustat to check the association between the change in the leader's market share in Compustat and the change in the sector's average productivity growth rate from 1994-2003 to 2004-2017 at the sector level. Sectors experiencing greater slowdowns in average productivity growth rates between 1994-2003 and 2004-2017 also saw greater increases in concentration, measured

[4]https://www.bea.gov/data/special-topics/integrated-industry-level-production-account-klems

as the market leader's share of total industry sales, on average (Figure A.7).

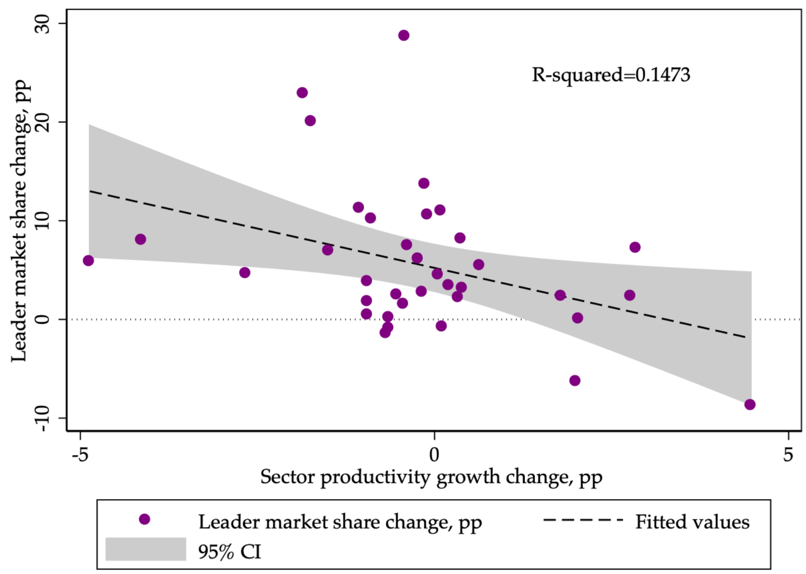

Figure A.7: Author's calculations from Compustat and BEA Integrated Industry-Level Production Accounts. 3-digit NAICS sectors, comparing 1994-2003 average to 2004-2017 average.

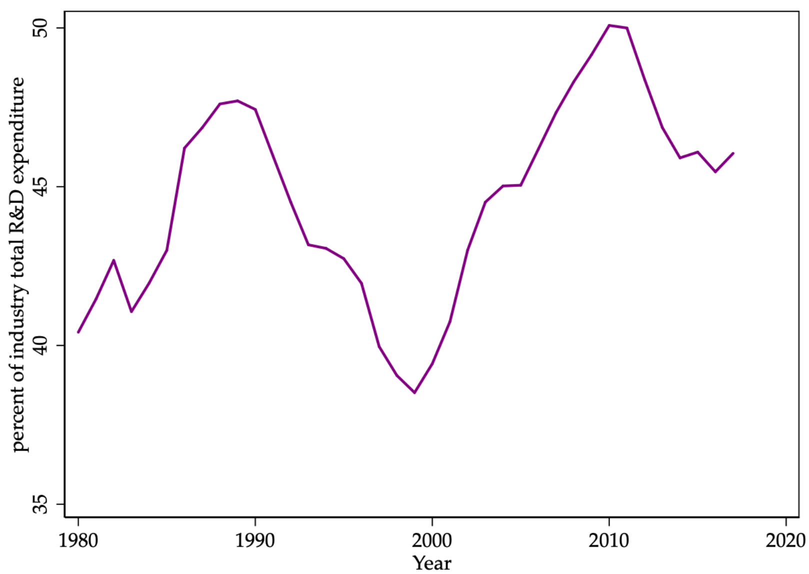

Figure A.8: Research and development expenditures (XRD) of sales leaders in 4-digit SIC industries in Compustat as a share of total R&D expenditures of all firms in that sector. Average across industries, sale-weighted by industry size.

A.2 Model Appendix

A.2.1 Proof Prices Depend on Relative Quality

Relative quality refers to the ratio of qualities of the two incumbent firms in a sector (dropping the sector notation j) $\frac{q_1}{q_2}$ for firm 1 and $\frac{q_2}{q_1}$ for firm 2. Below I show that the firms' pricing strategies depend only on relative quality, not the level of their own or their rival's quality.

First, this is clearly satisfied for the technology follower ($m_i < 0$) who sets price equal to marginal cost η regardless of absolute quality, and for sectors where $m_1 = m_2 = 0$, that is, when firms are neck-and-neck, because of the presence of the competitive fringe.

For the leader ($m_i > 0$), plugging the final good firm's demand for good i into the definition of the market share and using the definition of the price index yields:

$$\begin{aligned} s_i &= q_i^{\epsilon-1}\left(\frac{p_i}{P_j}\right)^{1-\epsilon} \\ &= \frac{q_i^{\epsilon-1}p_i^{1-\epsilon}}{q_i^{\epsilon-1}p_i^{1-\epsilon} + q_{-i}^{\epsilon-1}\eta^{1-\epsilon}} \\ &= \frac{1}{1+\left(\frac{q_{-i}}{q_i}\right)^{\epsilon-1}\left(\frac{p_i}{\eta}\right)^{\epsilon-1}}, \end{aligned}$$

where $-i$ denotes the follower. Now using the pricing decision of the leader:

$$s_i = \frac{1}{1 + \left(\frac{q_{-i}}{q_i}\right)^{\epsilon-1} \left(\frac{\epsilon-(\epsilon-\frac{1}{\beta})s_i}{\epsilon-(\epsilon-\frac{1}{\beta})s_i-1}\right)^{\epsilon-1}}.$$

Thus there is a mapping from technology gaps to market shares and prices that is independent of quality levels. ■

A.2.2 Value Function Boundary Equations

For the firm that's furthest behind (at gap $-\bar{m}$ with quality q_t):

$$\begin{aligned}
r_t V_{-\bar{m},t}(q_t) - \dot{V}_{-\bar{m},t}(q_t) = \max_{x_{-\bar{m},t}} \{ & 0 - \alpha \frac{(x_{-\bar{m},t})^{\gamma}}{\gamma} q_t^{\frac{1}{\beta}-1} \\
& + x_{-\bar{m},t} \sum_{n_t=-\bar{m}+1}^{\bar{m}} \mathbb{F}_m(n_t)[V_{nt}(\lambda^{n_t-(-\bar{m})} q_t) - V_{-\bar{m},t}(q_t)] \\
& + x_{\bar{m},t}(V_{-\bar{m},t}(\lambda q_t) - V_{-\bar{m},t}(q_t)) \\
& + \delta_e(0 - V_{-\bar{m},t}(q_t) \}.
\end{aligned}$$

The difference between this and equation 1.5 is in the third line, where if the firm's competitor innovates, there is a spillover that causes the firm at gap $-\bar{m}$ to improve its quality by λ.

For a firm at gap $\bar{m}$ the value function is:

$$\begin{aligned}
r_t V_{\bar{m},t}(q_t) - \dot{V}_{\bar{m},t}(q_t) = \max_{x_{\bar{m},t}} \{ & \pi(\bar{m}, q_t) - \alpha \frac{(x_{\bar{m},t})^{\gamma}}{\gamma} q_t^{\frac{1}{\beta}-1} \\
& + x_{\bar{m},t}(V_{\bar{m},t}(\lambda q_t) - V_{\bar{m},t}(q_t)) \\
& + x_{-\bar{m},t} \sum_{n_t=-\bar{m}+1}^{\bar{m}} \mathbb{F}_{-\bar{m}}(n_t)[V_{nt}(q_t) - V_{\bar{m},t}(q_t)] \\
& + \delta_e(0 - V_{\bar{m},t}(q_t) \},
\end{aligned}$$

where:

$$\pi(m,q_t)=\begin{cases}0 & \text{if } m\le 0\\ q_t^{\frac{1}{\beta}-1}(p(m)-\eta)p(m)^{-\epsilon}(p(m)^{1-\epsilon}+(\lambda^{-m})^{\epsilon-1}\eta^{1-\epsilon})^{\frac{\epsilon-\frac{1}{\beta}}{1-\epsilon}} & \text{for } m\in\{1,\ldots,\bar{m}\}\end{cases}.$$

A.2.3 Derivation of Final Output

Dropping the time subscript t, plugging the pricing strategies in equation 1.4 and $p_i=\eta$ for firms with $m\le 0$ into the demand curve 1.3 to obtain the output of each incumbent and plugging these outputs into equation 1.2 and equation 1.2 into equation 1.1 simplifies as:

$$\begin{aligned}
Y &= \frac{1}{1-\beta}\left(\int_0^1 K_j^{1-\beta}dj\right)L^\beta\\
&= \frac{1}{1-\beta}\left(\int_0^1\left(\sum_{i=1}^2 q_i^{\frac{\epsilon-1}{\epsilon}}(q_i^{\epsilon-1}\left(\frac{p_i}{P_j}\right)^{-\epsilon}\left(\frac{P_j}{P}\right)^{-\frac{1}{\beta}}L)^{\frac{\epsilon-1}{\epsilon}}\right)^{\frac{\epsilon}{\epsilon-1}1-\beta}dj\right)L^\beta\\
&= \frac{L}{1-\beta}P^{\frac{1-\beta}{\beta}}\left(\int_0^1 P_j^{\epsilon(1-\beta)-\frac{1-\beta}{\beta}}\left(\sum_{i=1}^2 q_i^{\epsilon-1}p_i^{1-\epsilon}\right)^{\frac{\epsilon(1-\beta)}{\epsilon-1}}dj\right)\\
&= \frac{L}{1-\beta}P^{\frac{1-\beta}{\beta}}\left(\int_0^1 P_j^{-\frac{1-\beta}{\beta}}dj\right).
\end{aligned}$$

The demand shifter $P^{\frac{1}{\beta}}L$ index is common to all firms and can be taken out entirely (and normalized to one since I assume zero population growth). The quality-adjusted price index P_j of each sector falls as the qualities of the two firms in the sector grow, and the exponent is negative for all $\beta\in$

$(0, 1)$ so Y is increasing in firms' qualities.

Common to all firms with a particular technology gap m are the prices $p(m)$ of the firm at gap m and its competitor at $-m$, $p(-m)$. At time t, therefore, Y can be expressed as:

$$Y_t = \frac{1}{2}\frac{L}{1-\beta}P^{\frac{1-\beta}{\beta}} \sum_{m=-\bar{m}}^{\bar{m}} \left(\int_0^1 \left(q_{it}^{\epsilon-1} p_i(m)^{1-\epsilon} + q_{-it}^{\epsilon-1} p_{-i}(-m)^{1-\epsilon}\right)^{-\frac{(1-\beta)}{\beta(1-\epsilon)}} \mathbb{1}_{\{i\in\mu_{mt}\}} di \right)$$

where μ_{mt} is the measure of firms at technology gap m at time t and the above integration is taken over firms rather than sectors. More simply:

$$Y_t \equiv \frac{1}{2}\frac{L}{1-\beta}P^{\frac{1-\beta}{\beta}} \sum_{m=-\bar{m}}^{\bar{m}} Q_{mt},$$

where Q_{mt} is defined as:

$$\begin{aligned} Q_{m,t} &= \int_0^1 \left(q_{it}^{\epsilon-1} p(m)^{1-\epsilon} + q_{-it}^{\epsilon-1} p(-m)^{1-\epsilon}\right)^{-\frac{(1-\beta)}{\beta(1-\epsilon)}} \mathbb{1}_{\{i\in\mu_{mt}\}} di \\ &= \left(p(m)^{1-\epsilon} + (\lambda^{-m})^{\epsilon-1} p(-m)^{1-\epsilon}\right)^{\frac{1-\beta}{\beta(\epsilon-1)}} \int_0^1 q_{i,t}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i\in\mu_{mt}\}} di. \end{aligned}$$

A.2.4 Output Growth on Balanced Growth Path

To understand how aggregate output evolves, this section studies the evolution of $\tilde{Q}_{m,t}$ (defined in equation 1.11) between t and $t + dt$ for all m. These expressions are similar to those for the stationary distribution (equations 1.8-1.10) because they are based on the movement of firms to different technology gaps from their rival, but account for the quality improvements that occur because of innovation.

Assuming fixed distribution $\mu_{mt} = \mu_m$ for all m, t:

$$\dot{\tilde{Q}}_{mt} = \int_0^1 q_{m,t+dt,i}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i \in \mu_m\}} di - \int_0^1 q_{m,t,i}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i \in \mu_m\}} di.$$

that is, quality growth at gap m is due to the change an index of the qualities of all the firms with technology gap m. Consider an arbitrary $m \in (-\bar{m}, \bar{m})$ ($-\bar{m}$ and $\bar{m}$ are special cases because of spillovers). A portion of firms at m at t innovate to a different gap, and another portion leave gap m because their competitor innovates. Because all firms at gap m choose the same arrival rate x_m, these are a random sample of the firms at gap m at time t. The outflows from $\dot{\tilde{Q}}_m$ are:

$$-(x_m + x_{-m}) \int_0^1 q_{m,t,i}^{\frac{1-\beta}{\beta}} \mathbb{1}_{\{i \in \mu_m\}} di = -(x_m + x_{-m}) \tilde{Q}_m.$$

The inflows to m's quality index come from two sources. First, some firms innovate into position m from a lower position n, improving their quality by λ^{m-n}. The probability they innovate and reach gap m is given by $x_n F_n(m)$. Some firms fall back to m from a higher gap n because their competitor innovates to $-m$. The probability their competitor reaches $-m$

is given by $x_{-n}F_{-n}(-m)$. So cumulative inflows are:

$$\sum_{n=-\bar{m}}^{m-1} x_n F_n(m)(\lambda^{(m-n)})^{\frac{1-\beta}{\beta}}\tilde{Q}_n + \sum_{n=m+1}^{\bar{m}} x_{-n}F_{-n}(-m)\tilde{Q}_n.$$

Putting it together:

$$\dot{\tilde{Q}}_{mt} = \sum_{n=-\bar{m}}^{m-1} x_n F_n(m)(\lambda^{(m-n)})^{\frac{1-\beta}{\beta}}\tilde{Q}_n + \sum_{n=m+1}^{\bar{m}} x_{-n}F_{-n}(-m)\tilde{Q}_n - (x_m + x_{-m})\tilde{Q}_m. \tag{A.2}$$

For lowest gap there are spillovers when competitor innovates:

$$\dot{\tilde{Q}}_{-\bar{m}t} = \sum_{n=-\bar{m}+1}^{\bar{m}} x_{-n}F_{-n}(\bar{m})\tilde{Q}_n + x_{\bar{m}}(\lambda^{\frac{1-\beta}{\beta}} - 1)\tilde{Q}_{-\bar{m}} - x_{-\bar{m}}\tilde{Q}_{-\bar{m}}. \tag{A.3}$$

For highest gap the firm does not exit that gap when they innovate:

$$\dot{\tilde{Q}}_{\bar{m}t} = \sum_{n=-\bar{m}}^{\bar{m}-1} x_n F_n(\bar{m})(\lambda^{(m-n)})^{\frac{1-\beta}{\beta}}\tilde{Q}_n + x_{\bar{m}}(\lambda^{\frac{1-\beta}{\beta}} - 1)\tilde{Q}_{\bar{m}} - x_{-\bar{m}}\tilde{Q}_{\bar{m}}. \tag{A.4}$$

Given equations A.2, A.3, and A.4, on a balanced growth path where $\frac{\dot{\tilde{Q}}_{mt}}{Y_t}$ is constant, it's sufficient to assume $\frac{\tilde{Q}_{mt}}{Y_t}$ is constant over time for all $m \in [-\bar{m}, \bar{m}]$. Differentiating $\frac{\tilde{Q}_{mt}}{Y_t}$ with respect to time yields:

$$\begin{aligned}
\dot{(\frac{\tilde{Q}_m}{Y})} &= \frac{\dot{\tilde{Q}}_m}{Y} - \frac{\tilde{Q}_m}{Y}\frac{\dot{Y}}{Y} \\
&= \frac{\dot{\tilde{Q}}_m}{Y} - g\frac{\tilde{Q}_m}{Y}.
\end{aligned}$$

Imposing that the left hand side is zero implies:

$$\frac{\dot{\tilde{Q}}_m}{Y} = g\frac{\tilde{Q}_m}{Y}.$$

The vector on the left hand side is defined above by the flow equations (A.2), (A.3), and (A.4) divided by GDP. Use those equations to form a matrix A that captures the flow equations:

$$\frac{\dot{\tilde{Q}}_m}{Y} = A\frac{\tilde{Q}_m}{Y} = g\frac{\tilde{Q}_m}{Y}.$$

The values in A depend on λ, ϕ, and x_m. The above equation means that the growth rate g is an eigenvalue of the matrix A and $\frac{\tilde{Q}_m}{Y}$ is the corresponding eigenvector of A. If there is only one positive, real eigenvalue there is only one such balanced growth path where the contribution of the growth of the quality index of each technology gap to the total growth rate is constant and the growth rate of the economy is constant.

A.2.5 Alternate Model With No Competitive Fringe

It is also possible to solve the full dynamic model without the presence of the competitive fringe imitating the follower's variety so that both firms exercise market power over their variety i of sector j's good. In this alternative model, only the incumbent firms compete a la Bertrand. There is still the possibility of exogenous entry/exit, though this assumption can be relaxed as well. The analogy from the model to the data becomes less obvious under this assumption, since the laggard firm can no longer be thought of representing many firms producing generic products that are perfectly substitutable with other generic products but imperfectly substitutable with the brand produced by the leader. In this setup the quality leader always has at least 50% market share, unlike in the data. This assumption also gives empirically counterfactual predictions that the profit shares of total industry profits of the market leader and the other firm in the industry are relatively similar, contradicting the pattern shown in Figure A.6.

Nonetheless, many of the main results carry through under this alternate assumption. Before describing these alternate results, I return to the pricing problem of the firms assuming the follower can now choose its optimal markup. Using the same derivation as in section 1.3.3 it can be shown that both firms follow the pricing policy the leader follows in the baseline model:

$$p_i = \frac{\epsilon - (\epsilon - \frac{1}{\beta})s_i}{\epsilon - (\epsilon - \frac{1}{\beta})s_i - 1}\eta,$$

where

$$s_i = q_i^{\epsilon-1}\left(\frac{p_i}{P_j}\right)^{1-\epsilon}.$$

I look for a Markov perfect equilibrium with balanced growth where each firm's price is the best response to its competitor's price at time t. The algorithm for finding the steady state remains the same, plugging in the pricing functions of the firms, illustrated in Figure A.9.

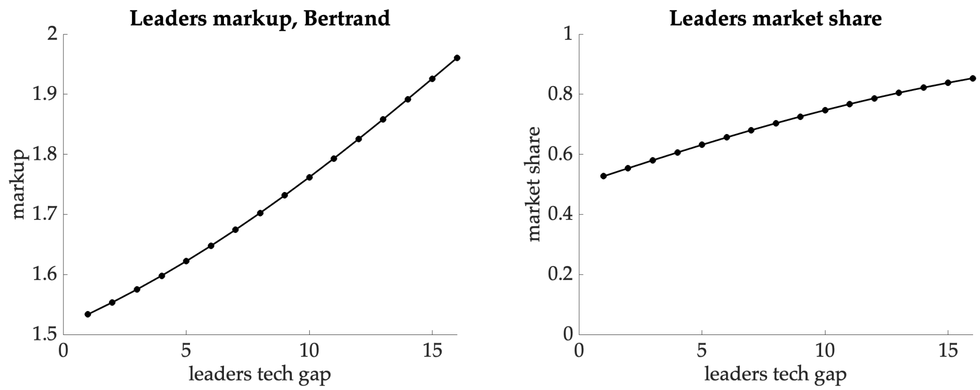

Figure A.9: Markups and resulting market shares as a function of the technology gap (ratio of firm qualities), Bertrand pricing.

Table A.2 gives the results of the same experiment as in section 1.4.3 under the alternate pricing strategies with the same parameters as in table 1.1 and Figure A.10 shows the policy functions and stationary distributions. Note that the escape competition motive around the neck and neck state disappears in the version without the competitive fringe. As before, changing ϕ has a level effect on total innovation effort but also changes the location of R&D from laggard firms to leading firms.

The level of the growth rates and the change in the growth rate from one steady state to the other under Bertrand pricing due to a change in ϕ are very similar to the baseline model with marginal cost pricing of the follower. The increase in concentration is smaller since the change in technology gaps is not as dramatic as in the main case (Figure A.10), though technology gaps do increase modestly. As for the growth decomposition, the effects of the firms' innovation responses is smaller, and the first order effect of lowering the probability of radical innovations is a bit larger than in the baseline model with the competitive fringe.

Moment	Data			Model		
	1990s	2000s	Chg. (pp)	1990s	2000s	Chg. (pp)
TFP growth, %	1.74	0.49	-1.25	1.82	0.27	-1.55
Leader market share, avg, %	43.34	48.12	4.78	62.44	63.52	1.08
R&D share of GDP, %	1.8	1.89	0.09	2.29	1.72	-0.57
Profit share of GDP, %	5.24	6.61	1.37	14.58	14.45	-0.13
R&D intensity, avg, %	2.56	3.8	1.24	8.13	5.33	-2.8
Pat stock growth/patent, avg, %	23.52	11.81	-11.71	21.44	11.10	-10.44
Leadership turnover, %	13.74	9.27	-4.47	13.49	10.43	-3.06

Table A.2: Model and data comparison, role of ϕ, model with no competitive fringe.

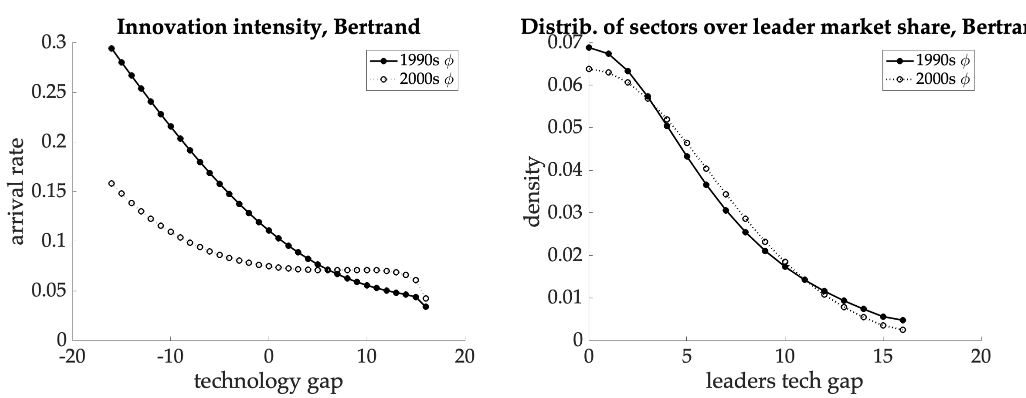

Figure A.10: Firm policy functions depending on technology gap (a) and stationary distribution of firms over technology gaps (b), model with no competitive fringe.

Decomposition	% of slowdown explained
Role of effort (ϕ fixed, x changes)	49.9
First order effect (x fixed, ϕ changes)	82.37

Table A.3: Growth decompositions, model without competitive fringe.

A.3 Numerical Appendix

A.3.1 Solution Algorithm

For a given set of parameter values, the solution algorithm involves first guessing a steady state interest rate. Given this interest rate, solve the value functions for each technology gap by policy function iteration using the fact that $\dot{v}_{mt} = 0$ on a balanced growth path. This process yields the optimal innovation policies of firms at each technology gap. Given the policy functions the stationary distribution of firms over technology gaps can be obtained by solving the system of equations described in section 1.3.5. To obtain the growth rate of GDP, solve the system described in appendix A.2.4. Check whether this growth rate is consistent with the interest rate guess using the household's Euler equation: $r = g\psi + \rho$. Update the guess of the interest rate and repeat until the interest rate guess and the interest rate implied by the resulting growth rate and the Euler equation are consistent.

To obtain micro-level moments, I simulate a discrete time version of the model with ten subperiods per year for a panel of 3000 firms for 400 years after the model reaches the steady state distribution over technology gaps.

A.3.2 Simulated Method of Moments Estimation

Let $M_j(\boldsymbol{\theta})$ denote the steady state value of moment j in the model as a function of the model parameters in vector $\boldsymbol{\theta}$. Let D_j denote the same moment in the data. The simulated method of moments estimation procedure seeks to find the vector of parameters $\boldsymbol{\theta}^*$ that solves:

$$\min_{\boldsymbol{\theta}} \sum_{j=1}^{J} \left| \frac{M_j(\boldsymbol{\theta}) - D_j}{\frac{1}{2}M_j(\boldsymbol{\theta}) + \frac{1}{2}D_j} \right|$$

for J moments. The moments are weighted equally.

A.3.3 Decomposition Table

Moment	Model		Data		Effect of each parameter						
	1990s	2000s	1990s	2000s	ρ	ϵ	η	δ_e	λ	α	ϕ
TFP	1.75	0.74	1.74	0.49	1.76	1.73	1.87	1.81	1.94	2.06	0.47
Concent.	44.62	48.89	43.34	48.12	44.66	45.01	44.67	44.68	45.61	44.7	46.8
R&D/GDP	1.91	1.32	1.8	1.89	1.95	1.88	2.33	2.06	2.06	2.13	0.72
Profits/GDP	6.02	6.71	5.24	6.61	6.02	6.01	6.01	6.0	6.23	6.02	6.34
R&D/Sales	5.18	3.54	2.56	3.8	5.22	5.09	6.23	5.5	5.61	5.7	1.93
Pat. qual.	22.26	11.88	23.52	11.71	22.26	22.28	22.27	22.27	24.06	22.28	10.98
New leader	13.26	9.27	13.74	9.27	13.28	13.21	13.53	12.6	13.4	13.91	9.85

Table A.4: Effect of each estimated parameter change in Table 1.5 on the model steady state, holding other parameters fixed at estimated 1990s values.

A.3.4 Transition Dynamics

This appendix details the computational approach to solving the model's transition dynamics from one steady state to another and then presents the results. I assume the economy begins in the initial (1990s) steady state in period $t = 1$ and arrives at the new steady state by T, where T is large. I consider a discrete version of the model with small time increments dt. Given a conjecture for the interest rate path over the transition, I solve the value and policy functions backwards from T. Then, given the sequence of innovation policies, I use the flow equations A.2, A.3, A.4 and 1.8, 1.9, 1.10 to solve for aggregate output and the distribution of sectors over technology gaps on the transition path respectively. Then I use the household's Euler equation to check the consistency of the growth rate over the transition with the interest rate guess, and update the guess until the conjectured and implied interest rate paths are within some minimum distance from one another. Formally:

1. Guess an interest rate path $\mathbf{r} = \{r_1, r_{1+dt}, r_{1+2*dt}, \ldots, r_T\}$.

2. Given the steady state values $v_{m,T}$ assumed at T, solve for innovation policies at $T - dt$ as:

$$x_{m,T-dt} = \begin{cases} \left(e^{-r_T dt} \frac{\sum_{n=m+1}^{\bar{m}} \mathbb{F}_{m,T-dt}(n)[(\lambda^{n-m})^{\frac{1}{\beta}-1} v_{n,T} - v_{m,T}]}{\alpha} \right)^{\frac{1}{\gamma-1}} & \text{for } m < \bar{m} \\ \left[e^{-r_T dt} \frac{1}{\alpha} (\lambda^{\frac{1}{\beta}-1} - 1) v_{\bar{m},T} \right]^{\frac{1}{\gamma-1}} & \text{for } m = \bar{m} \end{cases}.$$

3. Given the policy functions at $T-dt$ and the interest rate guess, solve for the value functions $v_{m,T-dt}$:

$$
\begin{aligned}
v_{m,T-dt} = &\left(\pi(m) - \alpha\frac{x^{\gamma}_{m,T-dt}}{\gamma}\right) dt \\
&+ e^{-r_T dt}(x_{m,T-dt}dt \sum_{n=m+1}^{\bar{m}} \mathbb{F}_{m,T-dt}(n)[v_{n,T}(\lambda^{n-m})^{\frac{1}{\beta}-1} - v_{m,T}] \\
&+ x_{-m,T-dt}dt \sum_{-m+1}^{\bar{m}} \mathbb{F}_{-m,T-dt}(n)[v_{-n,T} - v_{m,T}] \\
&+ \delta_e dt(0 - v_{m,t}) + v_{m,T}).
\end{aligned}
$$

4. Repeat this backwards iteration for $x_{m,t}$ and $v_{m,t}$ until $t = 1$.

5. Beginning at $t = 1$, initialize $Y = 1$ and use the flow equations A.2-A.4 and 1.8-1.10 and the sequence of innovation policies $\{x_{m,t}\}_{t=1}^{T-dt}$ to obtain the distribution of firms over technology gaps and the sequence of growth rates g over the transition.

6. Check if the interest rate sequence $\mathbf{r}$ is consistent with the resulting sequence of growth rates g using the households' Euler equation.

7. Update the guess of $\mathbf{r}$ to $\mathbf{r}_{new}$ using the implied sequence of interest rates from the Euler equation.

8. Repeat until $|\mathbf{r}_{new} - \mathbf{r}| < \epsilon_{tol}$ for some small tolerance value ϵ_{tol}.

A.3.5 Transition Results

The experiment is to consider a surprise, permanent decrease in laggards' patent quality consistent with the increase in ϕ estimated in section 1.5 and with the pattern of declining patent quality in Figure 1.2 in section 1.2.2.

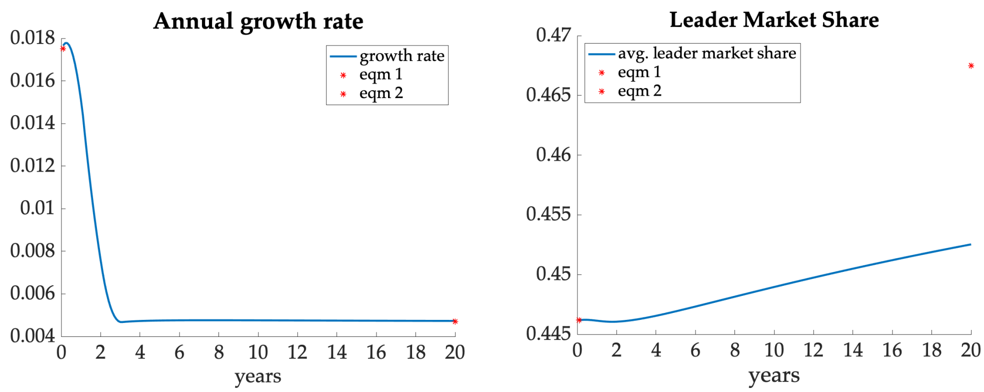

Figure A.11: Productivity growth and concentration over the transition to a lower level of patent quality.

For illustrative purposes, I assume the transition takes 20 years, with ϕ increasing smoothly for the first two years and then remaining permanently higher (that is, innovations becoming more incremental). Figure A.11 illustrates the evolution of productivity growth and concentration over the transition. Productivity growth closely tracks the decline in patent quality, while concentration rises more slowly and after 20 years has not fully reached its new steady state value. Because innovations are somewhat infrequent, it takes a long time for market leaders to pull ahead and laggards to fall as far behind as they are in the new steady state on average. These patterns are fairly consistent with Figure 1.1 which shows the

TFP growth rate declining quickly in the early 2000s and then remaining roughly flat while concentration has continued to rise through 2017. As discussed in section 1.5.3, an increase in product substitutability immediately increases concentration, and fitting the model over the transition would help inform the relative roles of innovation technology and the superstar firm hypothesis, and might also imply an even sharper decline in laggards patent quality in order to arrive at the high level of concentration observed recently by 2017.

Appendix B

Appendix: Country Banks and the Panic of 1825

B.1 Data Appendix

B.1.1 Banking Network

As described in Section 2.3, the banking network data comes from the Post Office London Directories for five years: 1820, 1823, 1825, 1827, 1830. Historical records suggest these directories were published annually, but I was only able to locate and access these five years. The directories were published under the patronage of the postmaster general, and were meant to provide the names and addresses of government officials, merchants, and public companies. The relevant section in each year is called "A List of Country Bankers", which contains the name (almost always the partner names), town, and London agent of all country bank branches that drew on a London bank.

I transcribed these records by hand since the number of entries was manageable, around 600 country bank branches per year. After transcribing each year separately I matched bank branches over time using their name and location. Country banks were limited to six partners, and partnerships changed somewhat frequently over time according to (Pressnell 1956), so I chose to consider two branch-year pairs a match as long as the location and at least one partner name was the same. Within each year I also tried to match country banks branches across different locations to identify branches of the same banks. Here I used partner names and London agent, assuming that different branches would have the same London agent. Around 2% of country bank branches had two London agents rather than one and I recorded each of these connections separately.

Identifying Agent and Bank Failures

As discussed in Section 2.3, my identification of branch failures is based on branches that disappeared from the Post-Office Directories between 1825 and 1827 rather than on direct evidence. London agent failures identified using the same method match narrative evidence from (James 2012) and (Dawson 1990) closely (both sources list examples of London agents that failed during the crisis). I further use (Price 1890)'s Handbook of London

Bankers, which attempts to provide a comprehensive guide to all London banks in the 19th century, to verify when agents failed. In any case, agent failures are less likely to be mismeasured by my strategy since each agent shows up in the Directories as many times each year as they have clients, which is nine on average. Country bank branches, which only appear once, are much more likely to be mismeasured.

B.1.2 Firm Bankruptcy Data

The bankruptcy data comes from the Edinburgh Gazette. There is a gap in the Edinburgh Gazette publications available online from July to December of 1826, so I end the analysis with June 30, 1826. The crisis period is thus December 1, 1825-June 30, 1826. To have a comparable sample and to help control for seasonal variation in bankruptcies, I begin collecting data December 1, 1824. Excluding bankruptcies in the urban areas of London and Middlesex, I am left with 1,440 individual bankruptcies over the study period (excluding bankruptcies from July 1, 1825 to November 30, 1825). I exclude that period because of uncertainty about whether it falls into the pre- or post-crisis period.

When an individual declared bankruptcy or was sued for bankruptcy by a creditor, a notice was required to be posted in the London, Edinburgh, and Dublin Gazettes so other creditors were aware of the proceedings. (Marriner 1980), an authority on bankruptcy statistics from this period, argues that Gazette records provide an accurate picture of bankruptcy statistics because the government-appointed Commission of Bankrupt had to certify that creditors had a legitimate claim before notices were posted in the Gazettes. (Duffy 1973) analyzed 50 bankruptcies from 1810-1811 and found that for 41 out of 50 cases of payment stoppage, bankruptcy proceedings began within one month, so the dating is fairly accurate.

By searching all banker names in the Post-Office Directories that I identify in those databases as bankrupt, I find that some bankers' occupations are not listed as banker in their bankruptcy notices in the Gazette. This misclassification could introduce bias by understating banker failures and counting them as non-financial firms, so I drop bankruptcies of individuals with the same name as the bankers who failed, but I may miss some since not all partner names are listed in the London Directories. For banks with two or more partners listed, however, I actually find few cases where more than one partner went bankrupt.

Occupation Classifications

Occupation titles vary widely, so after transcribing the data I classify each occupation into eight broad categories: bankers, other financial occupations, trade, manufacturing, retail, food, clothing, and construction. In cases where multiple occupations are listed I make the classification based only on the first occupation listed. I also separately classify occupations according to the four categories (tradables, non-tradables, construction, and other) used in Mian and Sufi (2014), matching occupation titles in the

Gazette to the 4-digit NAICS industries listed in the appendix of that paper as closely as possible, with the exception of food- and beverage-related occupations all of which I classify as non-tradable. I also collapse construction occupations into non-tradables unlike Mian and Sufi (2014) because the exogenous shock in 1825 was unrelated to house prices.

Focus on England

As discussed in Section 2.3, I focus exclusively on English bankruptcies rather than Irish or Scottish bankruptcies. This is acceptable because both countries differed from England in important ways. In Scotland, for example, joint-stock banks were not prohibited, banks had limited liability, and most banks had many branches,[1] more closely resembling modern day financial systems and allowing Scottish banks to better weather crises like the 1825 panic.[2] Both Scotland and Ireland had their own (proto-)central banks, adding an additional layer to the institutional differences between these national contexts. It is not clear how and where Welsh bankruptcies were reported, but I find very few Welsh bankruptcies reported in the Edinburgh Gazette despite Wales having a large number of country banks, so I also exclude Wales.

[1] Black (1995) provides a useful description of how the Scottish system differed from the English system.

[2] The *Edinburgh Courant*, quoted in the December 25th, 1825 edition of The Examiner (1825a), wrote "The consideration of these circumstances forces upon our notice the superior security which our Scottish banking establishments afford...the alarm of the money-market in London has scarcely been at all felt."

B.2 First Stage Regressions for Main IV Results

Table B.1: First Stages, Including for Main IV Results in Table 2.8

	(1)	(2)	(3)	(4)
Bank failures				
Exposed banks	0.393**	0.178+	0.173	0.130
	[0.084]	[0.101]	[0.111]	[0.120]
Has bank		16.095**	17.418**	18.604**
		[0.107]	[0.110]	[0.144]
Population, 1821, thousands		-0.001	-0.002	0.007
		[0.003]	[0.008]	[0.007]
Firm bankruptcies, pre-period			0.005	-0.004
			[0.035]	[0.045]
County FE				Yes
Observations	616	616	616	616

Source: Post-Office London Directories, 1820-1830; Edinburgh Gazette; (Census of Great Britain 1821). Pre-period for firm bankruptcies is Dec. 1 1824-Jun. 30 1825. Firm bankruptcies exclude banks. Each regression includes a log(total banks+1) control with coefficient constrained to be one. Columns 2-4 correspond to columns 1-3 in Table 2.8. Reported coefficients are the estimated coefficients in model 1.3. Robust standard errors in parentheses. + $p<0.1$, * $p<0.05$, ** $p<0.01$.

Appendix C

Appendix: This Time It's Different: The Role of Women's Employment in a Pandemic Recession

with Titan Alon, Matthias Doepke & Michèle Tertilt

C.1 Additional Tables and Details on the Calibration

C.1.1 Overview of Calibration Data Sources

The calibration targets draw on data from several different sources. Data on childcare hours by gender and marital status come from the American Time Use Survey (ATUS). The telecommuting status of different occupations is derived from the Leave Module of the American Time Use Survey(2017-2018) and is then merged into the Current Population Survey (CPS) to calculate the aggregate occupation shares. All data on employment status, household composition, and the presence of children is likewise taken from the CPS or related Census data sources. Labor market flows are calculated using the CPS matched basic monthly files from 2000–2020. Data on the share of households with traditional or modern social views is derived from questions in the General Social Survey (GSS). Finally, auxiliary data used to calculate average child rearing duration comes from the National Survey of Family Growth (NSFG) and data on the returns to (broad) labor market experience is estimated using the National Longitudinal Survey of Youth 1997 (NLSY97).

C.1.2 Further Details on the Calibration Procedure

Moments on the gender wage gap, labor supply, and labor market flows are calculated from the Current Population Survey. The primary sample includes all households ages 25 to 55 with non-missing entries for marital status, gender, and employment status. The age limit of 55 is chosen to be consistent with our focus on prime-age workers below an age when early retirement becomes common. Unless otherwise stated, the sample period spans the years 2017 to 2018. Individuals are grouped by gender (male, female), marital status (single, married), type of children (none, younger, older), employment status (not employed, part-time, full-time), and occupation type (telecommuting, non-telecommuting).

Child groups correspond to the age of the parents' youngest child in a household, with younger kids corresponding to ages 0–5 and older kids corresponding to ages 6–16. Employment groups are identified using labor force status and usual hours worked. The non-employed includes those who are either unemployed or not in the labor force, part-time includes all those who are employed and usually work fewer than 35 hours per week, and full-time includes all those who usually work more than 35 hours per week.

Telecommuting status is assigned using Census occupational codes following the classification procedure in Alon et al. (2020a). Subsequent labor market flows between telecommuting and non-telecommuting jobs are calculated to match the employment shares of each type during the period immediately preceding and during the pandemic, as documented in Bick and Blandin (2020).

The gender wage gap is calculated as the average hourly wage of employed women relative to employed men, where wages are derived from CPS data on total annual income, weeks worked, and usual weekly hours.

Table C.1: Job Flows during Regular Recessions, by Gender and Employment Status

Recession Labor Market Flows	Data	Model
men E-to-E	0.93	0.91
men U-to-U	0.64	0.67
women E-to-E	0.91	0.91
women U-to-U	0.76	0.72

Moments on labor market flows by gender, marital status, employment status, and aggregate state of the economy are calculated using the matched CPS Basic Monthly Files from 2000 to 2020. Recessions are identified using the NBER's business cycle dates. Monthly flows are then converted to the quarterly frequency so as to conform to the timing convention in our model. The flows during normal times are included as targets in the model's joint internal calibration. Flow parameters during recessions are fit separately in an auxiliary calibration to reflect their typical cyclical variation. Table C.1 summarizes the data and model fit for labor market flows during recessions; flow targets for normal times are included in Table 3.3 of the main text.

Data on childcare requirements by gender, telecommuting status, and

Table C.2: Parameters Governing Child-Rearing Dynamics

Parameter	Value	Target	Data	Model
$\bar{\pi}_f$	0.1500	Share single females have first child by age 25	0.15	0.15
$\pi^f(s\|0)$	0.00467	Single women, share with children	0.35	0.35
$\pi^f(b\|s)$	0.02604	Single moms, ratio older-to-younger children	1.67	1.67
$\pi^f(0\|b)$	0.00002	Single moms, avg. duration of child-rearing in quarters	88.61	81.36
$\bar{\pi}_m$	0.0850	Share single men have first child by age 25	0.085	0.085
$\pi^m(s\|0)$	0.00133	Single men, share with children	0.15	0.15
$\pi^m(b\|s)$	0.02083	Single dads, ratio older-to-younger children	1.30	1.32
$\pi^m(0\|b)$	0.00003	Single dads, avg. duration of child-rearing in quarters	83.23	83.92
$\bar{\pi}_c$	0.5280	Share married couples have first child before age 25	0.528	0.528
$\Pi(s\|0)$	0.05429	Couples, share with children	0.69	0.69
$\Pi(b\|s)$	0.05952	Couples, ratio older-to-younger children	1.17	1.18
$\Pi(0\|b)$	0.04167	Couples, avg. duration of child-rearing in quarters	88.89	82.59

age of child are calculated using the American Time Use Survey. Childcare time includes all time diary entries related to (1) caring for and helping household children [030100], (2) activities related to household children's education [030200], and (3) activities related to household children's health [030300]. Time use variables are converted to average weekly levels by collapsing across household types using the ATUS supplied weights. The resulting childcare variables are then re-normalized to be consistent with the time endowment of the model, which sets full-time work equal to unity.

The initial shares of households with traditional versus modern social norms are derived from the General Social Survey. Specifically, we consider the survey question "Preschool kids suffer if their mothers work (agree/disagree)" and calculate the share of modern married couples as the fraction answering either disagree or strongly disagree in the 2018 wave of the GSS. The procedure yields a 30 percent share of couples with traditional social norms.

C.1.3 Calibrating Child Dynamics

The parameters governing the arrival and aging of children are set to jointly match targets on the life cycle of child-rearing by gender and marital status. The share of households initialized with children ($\bar{\pi}$) is calculated to match the share of each gender and marital status group with children by age 25, the model's first period. These shares are taken from Table 1 in the 2018 Census Fertility Report and Table 2 in the Census Fatherhood Report.

The remaining moments governing the arrival rate of younger children (after age 25), the aging of younger children into older children, and the aging of older children into adults are chosen to jointly match (1) the share of households with children, (2) the ratio of older to younger children, and (3) the average duration of child-rearing. Targets (1) and (2) are calculated from our primary CPS dataset so as to be consistent with our other targets. The average duration of child rearing is calculated by summing the duration of childhood in quarters (16×4) with the median inter-pregnancy interval (measured in quarters) multiplied by the average number of children minus one. The inter-pregnancy interval value is taken from the National Survey of Family Growth. The resulting parameters, data targets, and model fit are summarized in Table C.2.

Table C.3: Initial Distribution of Human Capital by Gender and Marital Status

Couples					
Husband \ Wife	(1)	(2)	(3)	(4)	(5)
(1)	0.652	0.094	0.003	0.000	0.000
(2)	0.155	0.089	0.002	0.000	0.000
(3)	0.003	0.002	0.000	0.000	0.000
(4)	0.000	0.000	0.000	0.000	0.000
(5)	0.000	0.000	0.000	0.000	0.000
Singles					
Men	0.825	0.170	0.005	0.000	0.000
Women	0.856	0.140	0.004	0.000	0.000

C.1.4 Calibrating Skill Formation

The human capital grid consists of five grid points with a constant ratio of 1.4 between adjacent points (i.e., moving one step up the ladder increases full-time earnings by 40 percent). The constant ratio of grid points implies that returns to experience and the impact of skill loss are equalized along the grid. We identify the initial position of individuals in the human capital grid using their hourly wage in the CPS. The grid values are initialized so that the boundary between the first and second skill regions equals the average wage of the employed population. The initial distribution of individuals on the grid is chosen to match the (joint) distribution of wages by gender and marital status for those aged 25 to 30.[1] Specifically, we assign to the first grid point the share of people with incomes below the first grid point, to the second grid point we assign the share of all those between

[1]Couples are included in the sample based on the age of the husband.

the first and second grid points, and so on. Couples are initialized on a two-dimensional grid to capture the assortativeness of marriage markets. Table C.3 summarizes the initial distribution of human capital for single men, single women, and the joint distribution for couples.

The parameters that govern human capital dynamics on the grid are δ and η. Both parameters map analytically into observable data moments. Specifically, the expected wage growth amongst employed individuals will equal ηh_{step}. We therefore set η to match a 1.1 percent average quarterly return to labor market experience that we estimate from longitudinal micro-data in the NLSY97 controlling for individual and year fixed effects. Similarly, the expected wage loss from a quarter of unemployment is equal to δh_{step}. We therefore choose δ to match an average quarterly wage loss of 2.5 percent during non-employment, consistent with the annual estimates of lost earnings one year after job displacement in Davis and von Wachter (2011).

C.1.5 Details on Computing the Model

The model is solved via value function iteration with discrete grids for all state variables. The grid for human capital is described above. The asset grid has 25 equally spaced grid points from 0 to 2.5 times maximum individual earnings. This maximum asset level is set such that very few households have maximum assets in steady state. Dynamic simulations are carried out by simulating 250,000 individuals over many periods, so that an initial N steady state is reached before the recession shock takes place. For both regular recessions R and pandemic recessions P, the probability that the recession will end in every period is set to $1/6$, that is, $\rho_R = \rho_P = 5/6$.

C.2 Additional Model Results

C.2.1 Labor Supply by Types of Families

Figures C.1 and C.2 show the evolution of labor supply by types of families. We scale these figures so that 40 hours corresponds to the full-time employment of a single worker (80 hours for a couple in which both work full time). The left panels in Figures C.1a and C.1b show that for both singles and couples without children, the impact of a regular versus a pandemic recession is similar, whereas if kids are present (right panels) a pandemic recession leads to a much larger reduction in labor supply. Increased childcare obligations due to school closures affects all parents' ability to work during the pandemic. The impact on single parents is particularly large, because they lack certain margins of adjustment from which couples instead may benefit (i.e., couples where both parents can telecommute or where one was not working before the pandemic).

Within couples, mothers reduce working hours considerably more than fathers (Figure C.1d), which again coincides with the empirical observations discussed in Section 3.2. Among single parents, fathers' labor supply drops more than that of mothers. This difference is primarily due to the fact that single fathers (who make up a small share of the population) start out with a higher labor supply. For single parents of either gender with small children, working full time at a job that does not allow telecommuting is infeasible during the pandemic, necessitating a large drop in labor supply.

Figure C.2a highlights the role of traditional versus modern social norms for couples' labor supply. In regular times, the labor supply of traditional mothers is only slightly lower than that of modern mothers. Indeed, with the relatively low childcare requirements in normal times, many traditional mothers are able to both work and provide the majority of childcare within the family. In a pandemic recession, in contrast, the traditional division of labor is reinforced, and traditional mothers reduce their labor supply more than modern mothers.

Figure C.2b shows that occupation (TC vs NT) primarily has a level effect on labor supply. Being able to telecommute leads mothers to supply more labor both in regular times and during a recession. The reduction in labor supply in a pandemic recession is similar across occupations.

Another notable finding depicted in both panels of Figure C.2 is that during a normal recession, average hours worked by mothers are roughly constant. While some mothers lose their jobs, others are entering the labor force just as their husbands lose their jobs.

Figure C.1: Labor Supply by Types of Family

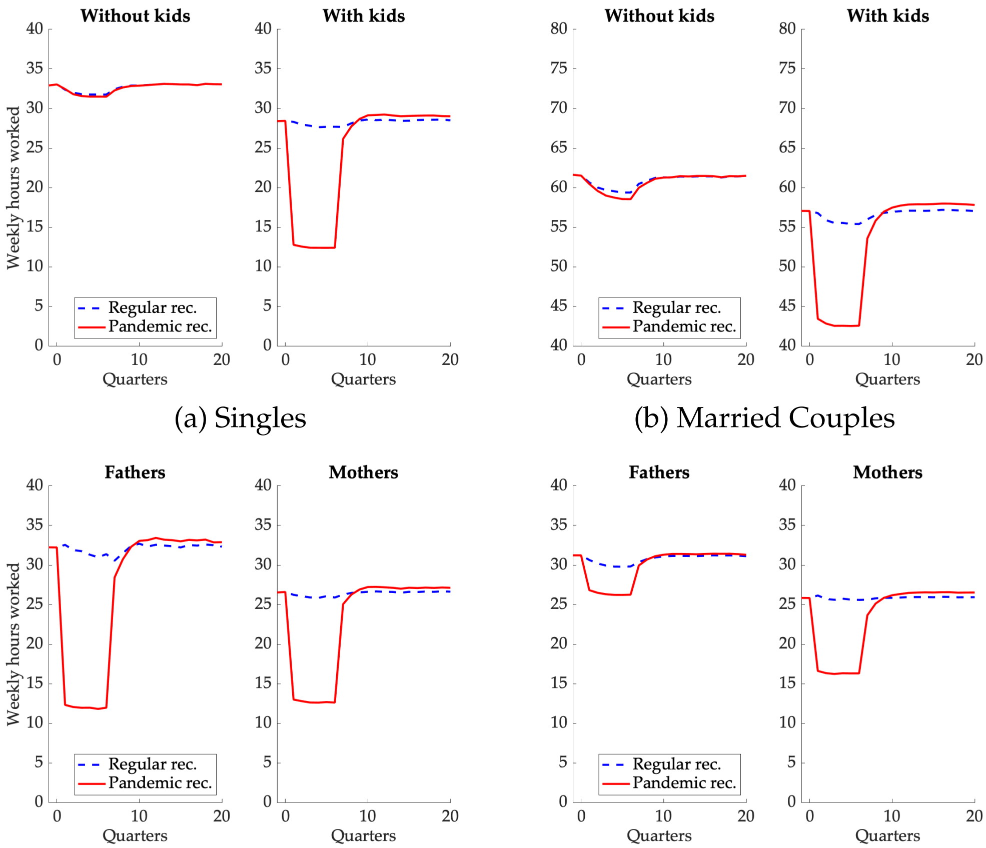

(a) Singles

(b) Married Couples

(c) Singles, Fathers vs. Mothers

(d) Couples, Fathers vs. Mothers

C.2.2 Leisure

Figure C.3 displays changes in leisure for single and married parents.

C.2.3 Welfare Implications of School Openings

Figure C.4 provides details on how welfare changes over time for singles, married women, and married men under different policy scenarios for

Figure C.2: Mothers' Labor Supply by Social Norm and Occupation

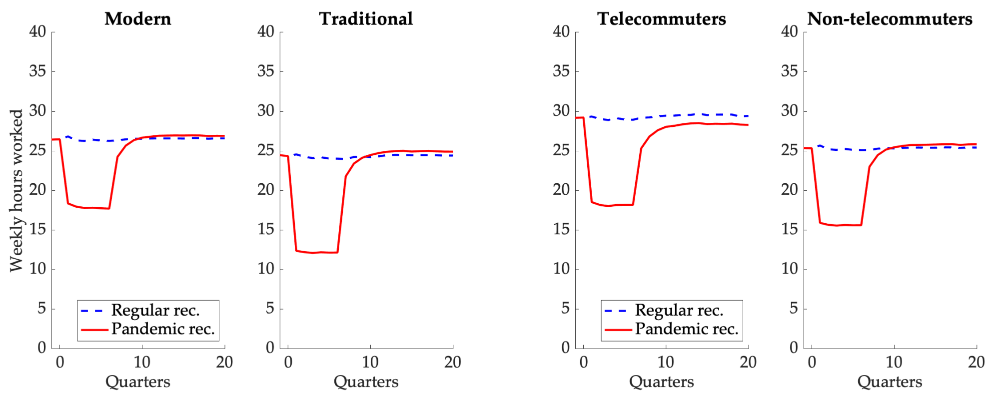

(a) in Modern vs. Traditional Couples

(b) By Occupation

Figure C.3: Leisure for Single and Married Parents

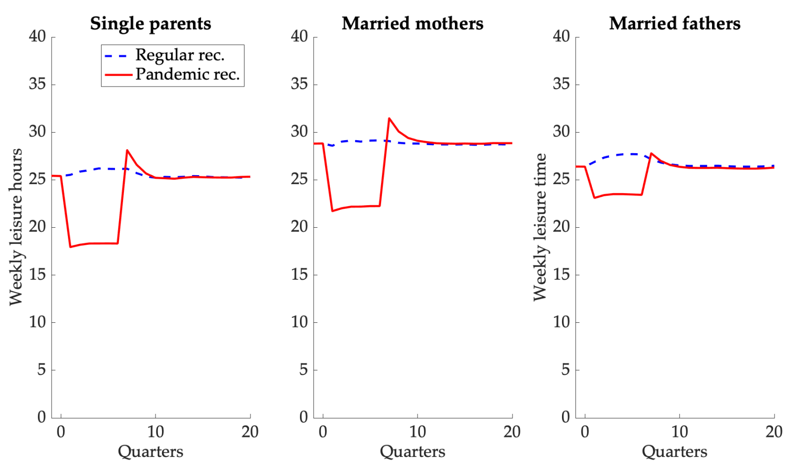

school openings. School openings occur either after the recession (quarter 6, panel a) or after two quarters of recession (in quarter 3, panels b–d).

Figure C.4: Welfare Implications of School Reopenings

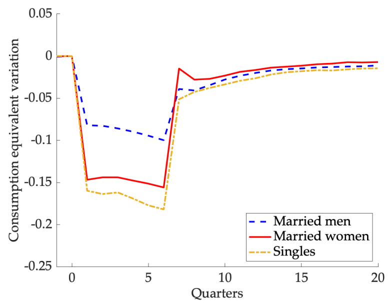

(a) Welfare during the pandemic

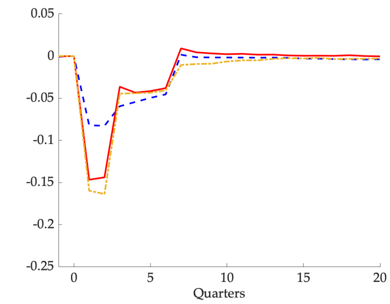

(b) Welfare with school reopenings

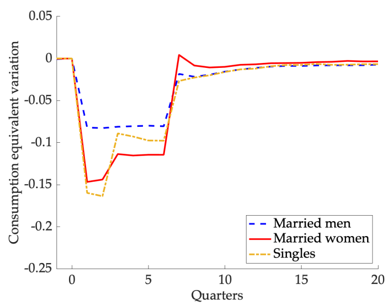

(c) Welfare with school reopening, big kids only

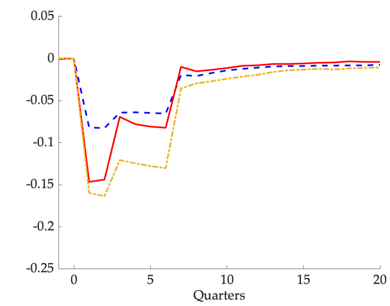

(d) Welfare with school reopening, small kids only

C.3 Additional Evidence

C.3.1 Evidence of a Permanent Shift in Telecommuting

There is much discussion in the media that working from home (WFH) is here to stay. Twitter famously announced that all employees could work from home permanently and Facebook CEO Mark Zuckerberg said that he expects as much as 50 percent of the company's workforce to work remotely in the long run.[2] The media is full of reports that demand for office space has plummeted. For example new lease signings in the first eight months of 2020 in New York City were only half of those in 2019.[3] This decline in demand has already led to falling rental prices of commercial property.[4] Since these commercial property leases are typically long-term contracts, these changes likely signal employer expectations of lasting changes in remote working

Are these newspaper reports exceptions hyped in the media or is there real evidence that something in the organization of work has permanently changed? And if so, what is the magnitude? There is some evidence by now from employee and firm surveys that point to a sizeable shift towards more WFH in the post-COVID world. Barrero, Bloom, and Davis (2020) conducted a survey of 12,500 US workers over the summer (May through September) asking about employee desires to work from home and about their employer planned post-COVID WFH days. They estimate that 20 percent of all full work days will be supplied from home after the pandemic ends, compared to just 5 percent before, i.e. a quadrupling. This

[2]"Coronavirus: Twitter allows staff to work from home 'forever"' BBC News, May 13, 2020 and "Half Of Facebook's Employees May Permanently Work From Home By 2030, Zuckerberg Says," Forbes, May 21, 2020.

[3]"Manhattan's Office Buildings Are Empty. But for How Long?," New York Times, September 8, 2020.

[4]"San Francisco Office Rents Tumble and Show No Sign of Bottoming," WSJ, October 6, 2020.

aligns well with evidence from two surveys (of small business owners and managers of large companies) conducted in March and April by Bartik et al. (2020) – where a large share of employers expects a sizeable increase in remote working after the Covid-19 crisis has passed.

Similar evidence is found in other countries. For example Baert et al. (2020) conducted a web survey among Flemish employees and find that the majority of surveyed employees believe that teleworking and digital conferencing is here to stay – 85 percent of respondents stated that they believe in overall more teleworking in the country in the future and 81 percent believe in more digital meetings in the country in the future. More than 50 percent of German firms indicated an increased importance of WFH after the crisis in the ifo business survey (IBS) conducted in May (Alipour, Falck, and Schüller (2020)). Similarly an employer survey conducted in June by the ZEW found that more than half of large manufacturing companies expect a permanent increase in remote work, which increases to three quarters for companies in the information industry (Erdsiek 2020).

Barrero, Bloom, and Davis (2020) lay out several mechanisms for why WFH will stay and provide some evidence for each channel: WFH stigma has diminished, WFH productivity is higher than was previously thought, large investments enabling WFH have been made (both in equipment but also time learning to use new technology), and finally because people expect the need for social distancing to stay for a long time even beyond when a vaccine is available.

C.3.2 Evidence from Paternity Leave Reforms

Rege and Solli (2013) use a paternity leave reform in Norway in 1993 to estimate the lasting effect of a short-term change in the division of labor in the household. The reform increased the length of subsidized parental leave by a month conditional on fathers taking at least one month. The paper finds a significant negative effect on fathers' earnings (1-3 percent lower for those treated with the reform), which persisted up to the last point of observation when the child is 5 years old. Kotsadam and Finseraas (2011) analyze the impact of the same reform of the division of household labor. They find that couples with children born after the reform have fewer conflicts about household work and that they share household tasks more equally 13 years later. Kotsadam and Finseraas (2013) find that the effect extends to the next generation – when fathers do more adolescent girls (not boys) do less domestic chores. Thus, gender norms seem to be permanently changed. Dahl, Loken, and Mogstad (2014) document sizeable peer effects (in coworkers and brothers) in the uptake of paternity leave in the context of the same Norwegian reform. The estimated effect snowballs over time, leading to a long-run uptake rate that is substantially higher than without the peer effects. They provide some suggestive evidence that the mechanism is likely related to information transmission about the costs and benefits of taking paternity leave.

Patnaik (2019) analyzes a reform in Quebec from 2006 and combines it with time diary data. Fathers exposed to the reform spend more time on housework and childcare activities and mothers spend more time working in the market even four years after the reform. Similar evidence is found in Tamm (2019) for Germany and Farré and González (2019) for Spain.

C.3.3 Evidence on Employment Effects of COVID-19 so far

Since we hypothesized likely effects based on pre-COVID data in March (Alon et al. 2020a), a sizeable literature has documented the actual effects since the beginning of the pandemic across the world. By and large, in most countries, female employment is suffering more than male employment. We briefly review the existing literature to date here.

Dias, Chance, and Buchanan (2020) use CPS data between December 2019 and May 2020 to assess the employment impact of the COVID-19 crisis by gender and household composition in the United States. The main finding is that employment was reduced much more for mothers than fathers. Interestingly, there is a fatherhood premium in the layoff rate – between March and April the layoff rate increased by 10.1 percentage points for mothers, by even more for non-parents, but only 6.8 percent for fathers. Similarly, Cowan (2020) finds a large gender gap in the employment declines in CPS data between February and April, especially when children are present. Controlling for many observables (such as age, race, education and industry/occupation), being female with children has a significant impact on the transition from at work to unemployment and a significant impact in the transition from full-time to part-time work. Montenovo et al. (2020) extend the analysis into May and also look at singles specifically. They find that women were substantially more likely to transition into unemployment between February and May as well and that women with young children experience higher rates of absence from work. They further find that single parents were particularly likely to have lost their jobs. Heggeness (2020) uses the differential timing of school closures across states to assess the impact on parental employment. Employees living in early closure states were 20 percent more likely to take temporary leave. This effect is almost entirely driven by women, who were 32 percent more likely to stop working. Even mothers who maintained their jobs in early closure states were 53 percent more likely to not be at work, compared to mothers in late closure states.

Andrew et al. (2020) provide evidence from the UK. They analyze the labor market outcomes for mothers and fathers in two-parent families based on an online survey during the first half of May. The decline in hours of paid work between February and May was dramatic in the UK. Proportionally, hours of paid work have shrunk more for mothers than fathers. To gain insights into productivity of parents during the pandemic, they also used a measure of uninterrupted work. While prior to the crisis, fathers and mothers used to be interrupted both proportionally to their work hours, now mothers are interrupted about 50 percent more often.

This may have implication for human capital accumulation on the job and future career prospects. The paper also documents large gender differences in domestic work, with mothers doing about 4 hours more per day. Large gender differences remain even when conditioning on parental work status. Their empirical findings cannot be explained by comparative advantage alone and thus seem to suggest that social norms play a role. At the same time, the average time fathers are involved in childcare doubled compared to pre-pandemic levels. Large gender differences in the provision of childcare in the UK are also documented in Hupkau and Petrongolo (2020). Interestingly, they do not confirm the gendered employment impact. Using the official UK labor market survey, they find evidence that men and women were equally impacted on the extensive margin and that at the intensive margin, women's hours fell by slightly less (comparing January/February with data from late April and May). At the same time, the reduction in hours was larger for parents with small children. The result that parents were impacted, but not specifically mothers, might be related to a sizeable fraction of fathers becoming the main childcare provider. In fact, the paper finds that the absolute and proportional increase in housework time in 2-adult households was larger for men than women, leading to a *reduction* in the gender gap in housework hours from 7.6 to 6 weekly hours – contrary to evidence from other countries.

Qian and Fuller (2020) provide evidence from Canada: They find sizeable gender gaps in employment declines for parents between February and May, using Canada's official monthly Labor Force Survey (LFS). The paper shows that these gaps grew by more for parents of elementary age children (6-12 years) compared to those with pre-K children. The findings that effects are larger for somewhat older children are in line with our model that shows that closing schools has a larger effect on the gender employment gap than closing daycare centers. They further find that gaps are particularly large for parents with only a high school (or less) degree. The gender gaps are even larger when "being employed and at work" is used as an outcome variable. Beauregard et al. (2020) analyze data from Quebec, Canada, based on primary school re-openings which started in May in some regions, but not others. Using a triple-difference-strategy, they find a positive effect of re-openings on parental work, a more pronounced effect on singles and a stronger impact when the job cannot be done from home.

Alstadsæter et al. (2020) provide evidence from Norway – using administrative data from the early crisis period (March and April). They find that women were more affected than men by layoffs (temporary and permanent combined) in March and April and so were parents. The effect that having young children has on layoffs remains largely unaffected once firm and job fixed effects are controlled for. Once firm and occupational differences are accounted for, the gender effect in layoffs is only associated with the presence of young children: women with young children are more likely to be laid off and this is a within-firm and within-occupation effect. Kristal and Yaish (2020) show in Israeli survey data that between early March and early May women's employment and income declined by more than that of men. Kikuchi, Kitao, and Mikoshiba (2020) conclude, based on official survey data, that female employment declined by more than male employment in Japan from January to April and May.

Ma, Sun, and Xue (2020) provide evidence of the effect of school closures due to Covid-19 in China. They find that school closures were an important reason for mothers (not fathers) not to be returning to work after the "economy reopened" in March/April and especially so if prior to the pandemic children were in boarding school. The effect was particularly large for migrant workers (who usually work away from home), and somewhat smaller if grandmother care was available. A few papers provide evidence of the effect of school closures on parental labor supply from other contexts. Dunbar (2013) provides evidence on the impact of school closures due to teachers' strikes on parental labor income in the US. They find a sizeable effect in families with school-aged children (6-12 years) in which both parents work. There is no significant effect for families with older children (12-18 years). Jaume and Willén (2018) analyze teachers' strikes in Argentina and find a sizeable impact on labor market participation and labor earnings of mothers. The effect is most pronounced for lower-skilled mothers and mothers in dual-income households. There is no effect on the average father; however for fathers with lower earnings than their wives, a small reduction in participation is found as well.