ESSAYS IN LABOR ECONOMICS

A Dissertation

Submitted to the Faculty

of

Purdue University

by

Mary Kate Batistich

In Partial Fulfillment of the

Requirements for the Degree

of

Doctor of Philosophy

August 2020

Purdue University

West Lafayette, Indiana

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF DISSERTATION APPROVAL

Dr. Timothy N. Bond, Chair
Department of Economics
Dr. Kevin J. Mumford
Department of Economics
Dr. David Hummels
Department of Economics
Dr. Jack Barron
Department of Economics

Approved by:
Dr. Brian Roberson
Director of Economics Doctoral Program

ACKNOWLEDGMENTS

To my mom and dad Kathy Batistich and Nick Batistich, and my siblings Nicholas Batistich, Eileen Batistich, and Sam Batistich, thank you for many, many years of love and encouragement. To my advisor and coauthor Timothy Bond, thank you for entertaining thousands of ideas, staying patient, and giving me a solid foundation for pursuing research in economics. To my coauthor and committee member Kevin Mumford, thank you for dedicating so much of your time, for your clear thinking, and for always taking me seriously. To my other committee members Jack Barron and David Hummels, thank you for your wisdom and helpful insight. Thank you to Stanton Hudja for being there from the very beginning. To Kendall Kennedy, Clint Harris, Daniel Woods, my cohort and the many faculty, staff, and students I crossed paths with, thank you for your friendship and support during my time at Purdue. I am very lucky to be part of this community and to have had the opportunity to explore open-ended research for the past six years.

TABLE OF CONTENTS

Page

Page

Page

LIST OF TABLES

Table | Page

LIST OF FIGURES

ABSTRACT

Batistich, Mary Kate Ph.D., Purdue University, August 2020. Essays in Labor Economics. Major Professor: Timothy N. Bond.

My dissertation consists of three independent chapters in the field of labor economics. The first chapter studies the economic forces underlying employment declines and skill upgrading in the U.S. manufacturing sector around the turn of the 21st century. The second chapter assesses the role of Japanese import competition in explaining stalled racial progress in the U.S. during the 1970s and 1980s. The third chapter explores end-of-life medical spending for dogs who have been diagnosed with cancer.

In the first chapter, I propose a new method to decompose employment changes by skill type into changes caused by output, labor supply, production task concentration, and labor-augmenting technology, using market equilibrium conditions within a constant elasticity of substitution production framework. I apply this method to manufacturing industries between 1990 and 2007, a period of steep employment declines for non-college workers. I find that labor-augmenting technology, by reducing labor per unit of output, is the leading source of displacement overall. However, a shift toward high-skill tasks is even more important in displacing non-college workers, who represent a majority of employment. In contrast, output changes have little influence on upskilling or aggregate job loss. In applications, I explore the impacts of import penetration from China and susceptibility to automation and offshoring. Of these, only offshoring is associated with some task upgrading, suggesting these mechanisms are not the primary drivers of this source of employment loss.

The second chapter is written with Timothy N. Bond. We assess the impact of the rapid rise in imports from Japan in the 1970s and 1980s on domestic labor markets. We use commuting zone level variation in exposure and stratify our outcomes by

racial groups. We find it decreased black manufacturing employment, labor force participation, and median earnings, and increased public assistance recipiency. However these manufacturing losses for blacks were offset by increased white manufacturing employment. This compositional shift appears to have been caused by skill upgrading in the manufacturing sector. Losses were concentrated among black high school dropouts and gains among college educated whites. We also see a shifting of manufacturing employment towards professionals, engineers, and college educated production workers. We find no evidence the heterogeneous effects of import competition can be explained by unionization, prejudice, or changes in spatial mismatch. Our results can explain 66-86% of the relative decrease in black manufacturing employment, 17-23% of the relative rise in black non-labor force participation, and 34-44% of the relative decline in black median male earnings from 1970-1990.

The third chapter, written with Kevin Mumford, contributes to the literature on the causal effect of end-of-life medical spending by focusing on the pet health care industry. Using administrative records and an identification strategy based on the timing of pet health insurance benefit renewal, we create an environment in which arrival of insurance benefits is quasi-random. We focus on how the availability of health insurance reimbursement funds affects spending, veterinary visits, and mortality over a two-year period after a serious cancer diagnosis. Increases in the generosity of health insurance reimbursement causes increases in both spending and veterinary visits, but we do not find evidence of a causal effect on mortality.

1. SKILL-BIASED TECHNICAL CHANGE AND EMPLOYMENT IN U.S. MANUFACTURING

1.1 Introduction

Following its postwar peak, U.S. manufacturing employment remained flat for decades and then dropped by about one fifth between 1990 and 2007 (Figure 1.7). Non-college workers, who had dominated the sector and for whom manufacturing had been a key employment source, bore these job losses. It is broadly accepted that this decline had economy-wide consequences such as increased income inequality and labor market polarization.[1] However, a debate persists on its exact causes, with attention focused on the roles of globalization and computerization which both surged around the turn of the century.[2] Despite robust evidence on the importance of Chinese import penetration (Autor et al., 2013; Pierce and Schott, 2016) and industrial robot adoption (Acemoglu and Restrepo, 2020), we lack a unified framework to quantify the relative importance of these and other factors.[3]

I address this shortcoming by developing a flexible new method to decompose manufacturing employment loss into its broad underlying forces. In doing so, I show that changes to production technology, rather than decreased output, are most important in explaining job loss. In particular, there has been a broad shift in the mix of production tasks toward those in which high-skill workers have a comparative advantage. The development of workplace tasks in response to process improvements has been analyzed within plants and firms for select industries (e.g., Bartel et al., 2003, 2007) but has never been studied comprehensively nor has it been considered as a source of employment loss. In this paper I document the importance of changes

[1] See, e.g., Ebenstein et al. (2014) for evidence on the wage premium in manufacturing, and Charles et al. (2019) for evidence on the broad labor market impacts of declining manufacturing employment. See Autor et al. (2003) and related papers which document polarization in the U.S. and Goos and Manning (2007) and Goos et al. (2014) which define and document polarization in the United Kingdom and throughout Europe.

[2] See Fort et al. (2018) for a review.

[3] Theory developed in Acemoglu (2003) and elsewhere demonstrates that trade shocks may cause technological change, creating an additional challenge in cleanly identifying the relative contribution of these two forces. Batistich and Bond (2019) show that the Japanese import shock of the 1970s and 1980s led to employment changes through technical change.

in task mix for job loss, not only of low-skill jobs to the benefit of high-skill, but for the overall job count due to the lower labor requirements of high-skill tasks.

To conduct this analysis, I employ the canonical skill-biased technical change ("SBTC") model. I use the first order conditions to write equilibrium low- and high-skill labor in terms of output, concentration between low- and high-skill tasks, labor-augmenting technology, and relative labor supply. I then identify theoretical equilibria by changing one variable at a time. For example, I solve for the equilibrium labor levels implied by keeping tasks, technology, and labor supply fixed but changing the value of output to that from an earlier time period. The difference between this theoretical equilibrium and the original tells me the change in labor had only output changed over time. In this manner I trace out the effect of every channel, so that the four effects sum to the total observed change as an identity.

Imposing the structure of the SBTC model, I distinguish between technology that directly augments labor productivity (the "productivity channel") and technology that transfers tasks between low- and high-skill labor (the "task channel"). The productivity channel may cause employment loss by augmenting labor productivity within a task, reducing the labor required to meet product demand. The task channel may cause low-skill employment loss if tasks are transferred to high-skill production, and may cause overall employment loss due to the imperfect substitutability between processes as well as their differences in labor intensity. To separately identify task shifts, I extract information from the different production materials and intermediates used by industries over time. This strategy and the decomposition framework are the two main methodological contributions of this paper.

I implement these methods using industry level data for the manufacturing sector between 1990 and 2007. The results show that the productivity channel has the largest total impact, displacing over 3.7 million workers. Such effects would be expected for example from widespread adoption of industrial robots and other computer-assisted technologies which augment labor productivity. I later confirm this association between automation and productivity channel displacement for both skill types in an application of my decomposition. The productivity channel's strong impact is in line with the evidence and widely held perspective that manufacturing tasks are routine and codifiable, making them especially vulnerable to automation (Autor and Dorn, 2013; Frey and Osborne, 2017; Akst, 2013).

Along with these productivity improvements is a sweeping transition toward high-skill production tasks, reducing low-skill employment by over 4 million, while adding nearly 1.5 million high-skill jobs. For low-skill workers, task shifts have the strongest impact among the four channels, causing nearly twice as much job loss as the productivity channel. Because high-skill production is inherently less labor intensive, this channel is also responsible for about 40 percent of overall employment loss.

In contrast to these technology channels, the scale channel (governed by output changes) had a strong and positive impact on employment for both skill types. Scale increased employment by over 22 percent, relative to an overall observed decrease of 17.3 percent. This result means that jobs lost through output decreases–such as those associated with import competition–are more than offset by gains from other means, like increased product demand. Scale changes are also not responsible for a large degree of upskilling, which would occur if high-skill industries grew faster than low-skill.

The fourth channel is labor supply, captured by relative market wages which represent the outside option (non-manufacturing employment) for each skill type. The rising economy-wide skill premium directly reduced high-skill labor and increased low-skill labor for most manufacturing industries. This channel somewhat offsets the effects of task upgrading, but magnitudes are small relative to other channels.

In sum, this decomposition tells the story of a sector that has restructured rather than disappeared. Improvements in the productivity of labor paired with a shift toward high-skill tasks explain the majority of employment declines. This latter evidence for a substantial task upgrading draws attention to an understudied cause of job loss. Recent literature has focused on the importance of task allocation in shaping labor demand (e.g., Autor and Dorn, 2013). However, this literature generally emphasizes the reallocation of tasks away from manufacturing and toward other sectors such as services. The results here instead emphasize the reallocation of tasks *within* manufacturing, providing a new perspective on the sector's role in labor market polarization.

This evidence is consistent with, and provides explanation for, two key trends noted in the literature. First, manufacturing value added has continued to grow despite drops in employment, roughly keeping pace with non-manufacturing value added. Second, the educational attainment of manufacturing workers has been on the rise.[4] The share of workers with a college degree has grown from 7.7% in 1962 to over 30% in 2018.[5] My approach is the first to identify the underlying sources of this

[4]See Appendix A.2.1 for figures of these trends. Similar trends for value added are documented in Fort et al. (2018). See Charles et al. (2019) for a discussion of upskilling.

[5]Author's calculations from the Current Population Survey. See figure footnotes in Appendix A.2.1 for details.

upskilling and quantify the importance of task upgrading in explaining overall job loss.

My framework allows me to assess the impact of particular events, such as import competition or the arrival of a new technology. In theory these events affect employment through different market forces, which we can test empirically. To do this, I combine my decomposition with current approaches to assess the employment impacts of three recent economic shocks: Chinese import penetration (Acemoglu et al., 2016), automation, and offshoring (Autor and Dorn, 2013). For the import shock, I find that employment losses were triggered primarily through scale, consistent with the theoretical implication of price competition in the product market. Absent this import competition, employment gains through scale would have been even greater. I also find a lack of task upgrading, consistent with recent studies finding a lack of new capital investments and other signs of innovation in trade-exposed industries (Pierce and Schott, 2018; Autor et al., 2016).[6]

Automation susceptibility leads to job loss for both skill types by the productivity channel, paired with possibly slower shifts toward high-skill tasks. This suggests that task upgrading may be an alternative to automation, rather than a symptom of it. Finally, offshoring propensity is associated with sharp scale declines especially for low-skill workers. It also plays a role in task upgrading, but in smaller magnitudes. Such task shifts imply parts of the production process are moving offshore, especially those associated with low-skill tasks, while domestic production increases its focus on high-skill tasks.

These applications provide new insight into the channels through which recent structural shocks have affected the manufacturing sector. They also call attention to the substantial variation across industries in terms of which channels are most important. For instance, scale does not explain aggregate employment loss but was important for certain industries. About 35 percent experienced negative scale impacts between 1990 and 2007, and 47 percent between 2000 and 2007.[7] Many of these industries were subject to Chinese trade exposure, and others experienced declines in demand. These examples demonstrate that worker displacement has varied in underlying economic cause, which may shape the way we consider potential remedies.

[6]For different conclusions regarding innovation responses to trade in an analysis of European markets, see Bloom et al. (2016).

[7]When including the recession years, a decomposition between 2000 and 2010 shows aggregate decreases due to the scale channel, although magnitudes are small relative to the other channels. The results from a 2000-2010 decomposition are available upon request.

This paper builds on the long tradition in economics of using accounting exercises to deepen our understanding of structural changes (e.g., Berman et al., 1994; Juhn et al., 1991, 1993; Firpo et al., 2011). In my approach, I observe equilibrium labor levels in two different time periods and use theory to deconstruct the differences into underlying economic channels. A large literature focuses on a related question of the causes of the economy-wide skill premium, or the wedge between high- and low-skill wages. Originating with Katz and Murphy (1992), both reduced form and structural anlayses of skill-biased technical change have been developed (see, e.g., Feenstra and Hanson, 1999; Burstein et al., 2019; Krusell et al., 2000; Lindquist, 2005).[8] My framework departs from this literature in two important ways. First, I am focused on one particular sector rather than the aggregate labor market. Second, I am primarily interested in explaining employment loss rather than wage changes. While I focus here on manufacturing industries, this intuition could be applied in a variety of competitive settings using a broad class of production technologies.

The rest of the paper proceeds as follows. In Section 1.2, I describe my theoretical framework. In Section 1.3, I describe my data sources and treatment. In Section 1.4, I explain my estimation procedure for production parameters and derive each channel of the decomposition. In Section 1.5, I apply my framework to describe national trends between 1990 and 2007, and explore industry-level heterogeneity. In Section 1.6, I use my framework in applications to assess the importance of each channel in explaining employment declines from the China imports shock of the 1990s and 2000s, and automation and offshoring as predicted by the initial occupational mix of each industry. I conclude in Section 1.7.

1.2 Deriving Equilibrium Labor Levels From Firm Optimization Conditions

Each industry produces a single consumption good $Y_{i,t}$ by combining low- and high-skill processes through a constant elasticity of substitution ("CES") production function.[9] For industry i and time t, output is

[8]For a review and historical context of the relationship between technological change and the skill premium, see Acemoglu (2002).

[9]Alternatively one might consider a representative firm in the industry.

$$Y_{i,t} = \left[\alpha_{i,t} \left(a_{i,t} L_{i,t}\right)^{\rho} + \left[1-\alpha_{i,t}\right] \left(b_{i,t} H_{i,t}\right)^{\rho}\right]^{\frac{1}{\rho}} \tag{1.1}$$

where $L_{i,t}$ is low-skill (non-college) labor, $H_{i,t}$ is high-skill (college) labor, and $\rho < 1$ governs the elasticity of substitution σ between low- and high-skill processes ($\sigma = \frac{1}{1-\rho}$).[10] This production technology is widely used in the skill-biased technical change literature assessing the roles of supply and demand forces in explaining growth in the skill premium (see, e.g., Katz and Autor, 1999; Card and DiNardo, 2002; Autor et al., 2008). Parameters $a_{i,t}$ and $b_{i,t}$ represent unskilled and skilled labor augmenting technology, respectively, while $\alpha_{i,t}$ represents the allocation of tasks between low- and high-skill processes.[11]

Skill-neutral technological change occurs as $a_{i,t}$ and $b_{i,t}$ grow together.[12] Skill-biased technological change occurs through shifts in $a_{i,t}/b_{i,t}$ or $\alpha_{i,t}$. Shifts in $a_{i,t}/b_{i,t}$ can be thought of as "intensive" skill bias because it stems directly from unequal progression in the marginal productivity of labor (Johnson and Stafford, 1998; Katz and Autor, 1999). Changes in $\alpha_{i,t}$ in contrast can be thought of as "extensive" skill bias because it relates to shifts in concentration between low- and high-skill production tasks.[13] A newly adopted technology might influence $H_{i,t}/L_{i,t}$ by increasing the relative marginal productivity of $H_{i,t}$ but also by increasing the industry's emphasis on high-skill tasks.

I proceed by writing the equilibrium levels of $L_{i,t}$ and $H_{i,t}$ in terms of output, relative wages, and intensive and extensive technology parameters, so that changes in $L_{i,t}$ and $H_{i,t}$ can be interpreted as the combined effect of changes in these variables. Assuming a perfectly competitive market in which workers are paid the value of their marginal product, Equation 1.1 can be used to solve for the equilibrium ratio of high-

[10]This flexible framework could be adapted to analyze other worker categories, such as routine and non-routine, or those with and without high school degrees. Likewise it could be extended to consider more than two worker types. Capital could be included in a general way, such as Cobb-Douglas, which would not affect the results.

[11]Note that as $\rho \to 0$, the function approaches Cobb-Douglas where $\alpha_{i,t}$ is the share parameter. The canonical SBTC model can be nested as a special case of the task assignment model developed by Acemoglu and Autor (2011). I derive one case in Appendix A.3. See also Autor (2013).

[12]Alternative representations of this framework will pull out a common term and re-cast $a_{i,t}$ and $b_{i,t}$ to sum to one. Equation 1.1 is mathematically equivalent to $Y_{i,t} = A_{i,t}(\alpha_{i,t}(\omega_{i,t}L_{i,t})^{\rho} + (1-\alpha_{i,t})((1-\omega_{i,t})H_{i,t})^{\rho})^{1/\rho}$ where $A_{i,t} \equiv (a_{i,t}+b_{i,t})$ and $\omega_{i,t} \equiv a_{i,t}/(a_{i,t}+b_{i,t})$. I utilize this equivalence to estimate $a_{i,t}$ and $b_{i,t}$ in Section 1.4.1.

[13]Goldin and Katz (1998) for example document the shift in production tasks from low- to high-skill workers in the manufacturing sector between 1909 and 1929.

to low-skill labor:

$$\frac{w_{L,i,t}}{w_{H,i,t}} = \frac{\alpha_{i,t}}{1-\alpha_{i,t}} \left(\frac{a_{i,t}}{b_{i,t}}\right)^{\rho} \left(\frac{H^*_{i,t}}{L^*_{i,t}}\right)^{1-\rho} \tag{1.2}$$

where $w_{S,i,t}$ is the market wage for skill S.[14] Given this equilibrium, Equation 1.2 can be rearranged to express $H^*_{i,t}$ in terms of $L^*_{i,t}$ and substituted into Equation 1.1. Now $L^*_{i,t}$ is a function of output, relative wages, and technology parameters. Specifically,

$$L^*_{i,t} = Y^*_{i,t}\left(\alpha_{i,t}a^{\rho}_{i,t} + (1-\alpha_{i,t})\left[b_{i,t}\left(\frac{1-\alpha_{i,t}}{\alpha_{i,t}}\right)^{\frac{1}{1-\rho}}\left(\frac{b_{i,t}}{a_{i,t}}\right)^{\frac{\rho}{1-\rho}}\left(\frac{w_{L,i,t}}{w_{H,i,t}}\right)^{\frac{1}{1-\rho}}\right]^{\rho}\right)^{-\frac{1}{\rho}} \tag{1.3}$$

where $Y^*_{i,t}$ is equilibrium output. Likewise, equilibrium high-skill labor can be expressed as:

$$H^*_{i,t} = Y^*_{i,t}\left(\alpha_{i,t}\left[a_{i,t}\left(\frac{\alpha_{i,t}}{1-\alpha_{i,t}}\right)^{\frac{1}{1-\rho}}\left(\frac{a_{i,t}}{b_{i,t}}\right)^{\frac{\rho}{1-\rho}}\left(\frac{w_{H,i,t}}{w_{L,i,t}}\right)^{\frac{1}{1-\rho}}\right]^{\rho} + (1-\alpha_{i,t})b^{\rho}_{i,t}\right)^{-\frac{1}{\rho}} \tag{1.4}$$

This can be thought of in a cost minimization setting in which firms choose ($L^*_{i,t}$,$H^*_{i,t}$) to meet target output $Y^*_{i,t}$ given relative market wages, as depicted in Figure 1.7. A change over time from $(L^*_{i,t}, H^*_{i,t})$ to $(L^*_{i,t+1}, H^*_{i,t+1})$ must therefore be due to a finite set of causes. First is changes in $Y_{i,t}$, or movements to a new isoquant, as depicted in Figure 1.7. Second and third are changes in intensive and extensive technology through $a_{i,t}$, $b_{i,t}$, and $\alpha_{i,t}$. This corresponds to movements in the location and curvature of the isoquant, as depicted in Figure 1.7. Fourth is changes in relative wages, or a movement along the isoquant to a new optimal bundle of $(L_{i,t}, H_{i,t})$, as depicted in Figure 1.7. This set lends itself to four economically relevant channels.

The first channel is changes to output $Y_{i,t}$, such as in response to consumer demand shifts. I call this the scale channel. This channel will capture the extent that declining production is responsible for employment changes, whether from declining demand for manufactured goods, or because production has moved offshore. The scale channel

[14]In line with the data, these wages may vary by industry.

affects low- and high-skill workers equally within an industry, because it does not alter the technology of the firm. However there can be skill bias in the scale channel to the extent that skilled workers are concentrated in different industries.

The second channel is shifts between low and high-skill processes, represented by $\alpha_{i,t}$, such as in response to firms increasingly concentrating their resources on high-skill tasks. I call this the task channel. By construction, this channel will increase one type of employment while decreasing the other, with the overall effect on employment dependent on the magnitude of the shift, the labor intensity of each process, and the substitutability between processes.

The third channel comes from changes to the marginal productivity of labor within each process, represented by $a_{i,t}$ and $b_{i,t}$, such as in response to a new technology that impacts one or both types of workers. I call this the productivity channel. Unlike the task channel, the productivity channel might force equilibrium low- and high-skill labor in the same direction. However it may be biased to the extent that $a_{i,t}$ and $b_{i,t}$ grow at different rates.

The productivity and task channels identify distinct elements of technical change. The concept is captured in Bartel et al. (2007) who analyze the effects of new information technologies in the valve manufacturing industry. They find the new technologies increased the efficiency of existing processes while also causing a change in business practices to a different mix of processes, including increased production of customized valves. This change in process mix was associated with an increase in worker skill requirements.

Finally, there may be changes in relative wages $w_{H,i,t}/w_{L,i,t}$, such as increased relative high-skill wages in response to reduced supply. I call this the labor supply channel. Reduced supply to a manufacturing industry of a particular skill type may be in response to an increase in demand for that worker skill type in another sector. If for example there was an increase in demand for high-skill workers in health care and education, this would increase the market wage for this type of worker and reduce their supply to the manufacturing sector. Similar to the task channel, the labor supply channel will work in opposite directions for low- and high-skill labor within an industry.

My goal is to quantify the role of each of the four channels in explaining observed employment changes for each skill type. I do so in two broad steps. First, I develop and implement methods for estimating ρ, $\alpha_{i,t}$, $a_{i,t}$, and $b_{i,t}$ by industry and time period. Then, I use these estimates to calculate theoretical equilibria under different combinations of these parameters. For example, I can calculate the theoretical equilibrium implied by holding all parameters fixed in time t, but modifying the task parameter to equal that from time $t+1$. This tells me how much low- and high-skill labor would change if only task shifts had occurred, while output, productivity, and wages remained constant. A series of these theoretical exercises enables me to quantify the impact of each channel. I explain these exercises in more detail in Section 1.4.

1.3 Data

1.3.1 Industry Employment, Output, and Material Use

I use employment counts, the value of output, and spending on production materials and supplies by industry from the Census of Manufactures ("CoM"), which is published every 5 years for years ending in 2 or 7. The CoM is a component of the Economic Census and covers all establishments with one paid employee or more primarily engaged in manufacturing. It collects and reports a variety of statistics at various geographic levels, including number of establishments, employment, payroll, value added by manufacture, cost of materials consumed, capital expenditures, and product shipments. I take data beginning in 1987, the first year in which 1987 SIC codes are used, up through 2012.[15]

Total employment and output (for which I use value of shipments) are available for all industries in all years. Materials use comes from the "Materials Consumed by Kind" tables available by industry at the national level. These materials include all materials, ingredients, containers, and supplies used in production. They do not include any capital expenditures, such as rental payments or spending on new machinery, equipment, or computers.[16] I follow this definition of materials, which includes both raw materials and semifinished goods, throughout. I use these data to estimate the share of production belonging to low-skill and high-skill processes as described in Section 1.4.1. Details on treatment of output and materials data can be found in Appendix Section A.1.1.

While employment counts are available by industry in the CoM, total work hours and the share of work hours belonging to high-skill labor are not.[17] I calculate labor

[15]From the 1992 CoM, I take the years 1987 and 1992. From the 2002 CoM, I take the years 1997 and 2002. These are available at https://www2.census.gov. The years 2007 and 2012 are available separately, and I downloaded these from the American Fact Finder at https://factfinder.census.gov.
[16]They also do not include include resales, fuels, purchased electricity, or contract work.
[17]The CoM does provide an employment breakdown between production and non-production work-

hours and hourly wages by skill type by multiplying the total employment counts reported in the CoM by mean hours and wages by skill type and SIC industry as described in Section 1.3.2. In order to link these data, in the CoM I average 1987 and 1992 and call this 1990, average 1997 and 2002 and call this 2000, and average 2007 and 2012 and call this 2010.

1.3.2 Labor and Wages by Skill Type and Industry

To determine the skilled share of labor at the national level by SIC industry, I exploit geographic overlap at the commuting zone ("CZ") level between employment shares by SIC industry from the the County Business Patterns ("CBP") and skilled employment shares by Census industry from the Census of the Population (or American Community Survey, "ACS", for years after 2000).[18] In this procedure, I first calculate employment by SIC industry and CZ in the relevant year (1980, 1990, 2000, 2007, and 2010) from the CBP. I then connect these to calculations of the high-skill share of workers by Census industry and CZ in the same year.[19] I assume that, within a CZ and year, the skill share of Census industries is constant across all the SIC industries it maps to. The national level share of high-skill workers by SIC industry i in time t is calculated as follows:

$$\eta_{i,t} = \sum_{m} \frac{E_{m,i,t}}{E_{nat,i,t}} \eta_{m,n,t}$$

where $E_{m,i,t}$ is the total employment in commuting zone m and SIC industry i in time t, $E_{nat,i,t}$ is the industry's national employment in time t, and $\eta_{m,n,t}$ is the high-skill share in commuting zone m of Census industry n in time t, where Census industry n maps to SIC industry i. With these high-skill employment shares, I then calculate the number of high-skill workers as the total employment count reported in the CoM

ers. However this is an unsatisfactory proxy for the share of workers who are are high-skill, which I define as workers with at least four years of college education. For example, according to the 1992 CoM, cafeteria personnel and highway truckdrivers and their helpers are considered non-production workers. Further, production workers have become more educated over time.

[18] I use the 5 percent Census samples for 1980, 1990, and 2000. I use the 2005-2007 ACS for the year 2007 and the 2008-2012 ACS for 2010. I provide more details on how I define my Census/ACS samples and data treatment in Appendix Section A.1.2. Details on data handling of the CBP series can be found in Appendix Section A.1.2.

[19] This is a many-to-one mapping as multiple SIC industries may connect to the same Census industry.

multiplied by $\eta_{i,t}$, while the number of low-skill workers is total employment multiplied by $(1 - \eta_{i,t})$.

I also use these high-skill employment shares to determine the mean annual hours worked for high-skill workers in industry i in year t by

$$\mu_{i,t} = \sum_{m} \frac{E_{m,i,t}}{E_{nat,i,t}} \frac{\eta_{m,n,t}}{\eta_{i,t}} \mu_{m,n,t}$$

where $\mu_{m,n,t}$ is the mean annual hours worked for high-skill workers employed in Census industry n commuting zone m and time t, where again Census industry n maps to SIC industry i. In the same fashion I calculate mean annual hours worked for low-skill (non-college) workers, and mean hourly wages for each skill type. This requires the assumption that, within a CZ and year, the annual hours and wages of workers in a given Census industry are constant across all the SIC industries their Census industry maps to. Total annual labor hours by skill type, $L^*_{i,t}$ and $H^*_{i,t}$, are calculated as total employment of the skill type multiplied by mean annual hours of the skill type.

1.4 Empirical Methods

1.4.1 Estimating Production Parameters

Estimation of $\alpha_{i,t}$

An important component of my analysis is to separately identify $\alpha_{i,t}$, which represents the allocation of tasks between low- and high-skill processes, from $a_{i,t}$ and $b_{i,t}$, which represent the marginal productivity of low- and high-skill labor. To this end, I exploit the detailed information on materials, ingredients, containers, and supplies use by industry available in the CoM.[20] These data signal information about the underlying production processes of the firm. Certain materials, such as diagnostic substances (SIC product 2835) and other biological products (SIC product 2836) are predictive of a high share of skilled labor. Other materials, such as logging and lumber products (produced by SIC industries 2411 and 2421) are predictive of a low share of skilled labor. I interpret this to mean that given currently available technology, di-

[20] I give details on the CoM in Section 1.3.1 and on data treatment in Appendix Section A.1.1.

agnostic substances and other biological products are typically handled by high-skill workers in high-skill tasks while logging and lumber products are typically handled by low-skill workers in low-skill tasks.

Once I know the degree to which each material is associated with each process, I calculate $\alpha_{i,t}$ as the share of total materials spending on the low-skill process, assuming the share of materials spending is proportional to the share of tasks. This is consistent for example with any production function in which, in equilibrium, materials are distributed uniformly across tasks.[21] In my data, I need to systematically allocate production materials between the low- and high-skill processes. I do so by estimating the impact of each material on the share of labor hours that are low-skill. For each material j, I estimate the linear regression

$$\frac{L_{i,t}}{H_{i,t}+L_{i,t}} = \lambda_{j,t} \times 1\left[\frac{z_{j,i,t}}{z_{tot,i,t}} > 0.01\right] + \epsilon_{j,i,t} \tag{1.5}$$

where the left-hand side variable is the share of labor hours that are low-skill in industry i and time t, $z_{j,i,t}$ is industry i's spending on material j in time t, $z_{tot,i,t}$ is industry i's total resources spending in time t, $1[\cdot]$ is an indicator function equal to one when the spending share is greater than one percent, and $\epsilon_{j,i,t}$ is an idiosyncratic error term.[22] I collect the coefficients $\widehat{\lambda}_{j,t}$.

Equation 1.5 is useful for materials used across many industries. For less commonly used materials, some of which may only appear in one or two industries in a given year, measurement error may be a concern. To address this, I exploit the empirical relationship between a material's prediction of skill share and its complexity, defined here as the skilled labor share in the SIC industry that is the primary producer of the material. I develop a "complexity index" for materials based on this definition, which I hold fixed across time.[23]

[21] See Acemoglu and Autor (2011) and Autor (2013) which describe how the canonical SBTC model can be derived as a special case of the task assignment model they develop. I derive one case in Appendix A.3. See Rosen (1978) for more details on the microfoundations of CES production functions.

[22] I use a low threshold of one percent rather than zero to avoid trace amounts of materials spending that appear as a result of imputations and bridging industry codes across classification systems.

[23] For manufacturing industries, I calculate skill share by industry as described in Subsection 1.3.2. For materials produced by non-manufacturing industries, I connect SIC products to Census industries according to Census Bureau Technical Paper #65, and use national level skill shares according to the 1980 Census of the Population.

A material's complexity score is highly correlated with its prediction of skill share.[24] I fit this relationship using a localized linear regression method with observations of above-median most commonly used materials in the relevant year.[25] This exercise reduces measurement error and allows for out-of-sample predictions for the less commonly used materials. For out-of-range predictions at the low and high end of the complexity index, I assign the maximum and minimum value from the within-sample predictions, respectively. The predicted values (in essence the predicted $\widehat{\lambda}_{j,t}$) represent the proportion of the material that is associated with the low-skill process. I call this proportion $\widetilde{\lambda}_{j,t}$. For each material j, I allocate the total amount an industry uses into the low-skill process according to $\widetilde{\lambda}_{j,t}$, while $1-\widetilde{\lambda}_{j,t}$ is allocated to the high-skill process.

It is quite possible that materials shift processes over time, especially if new technologies require materials to be increasingly handled in high-skill tasks. For this reason I repeat the estimation separately for each year in the data. In Appendix Section A.2.2, I provide a figure showing the estimated values for $\widetilde{\lambda}_{j,t}$ for 1990 and 2007. A level shift toward the high-skill process is apparent, but there is no dramatic change in curvature.

Maintaining my assumption on the relationship between materials and tasks, I estimate $\alpha_{i,t}$ as

$$\widehat{\alpha}_{i,t} = \sum_{j=1}^{N} \frac{z_{j,i,t}\widetilde{\lambda}_{j,t}}{z_{tot,i,t}}$$

where $z_{j,i,t}$ is industry i's spending on material j in time t for $j \in (1, ..., N)$, and $z_{tot,i,t}$ is industry i's total materials spending in time t. In Figure 1.7, I plot the distribution of $\widehat{\alpha}_{i,t}$ across my industry sample in 1990, 2000, and 2007. A clear shift leftward of the distribution reflects a transition away from low-skill tasks over time.

Estimation of ρ, $a_{i,t}$ and $b_{i,t}$

I can now include my estimates for $\widehat{\alpha}_{i,t}$ into Equation 1.2, take logs, and rearrange the parameters for an estimating equation. Specifically, I seek to estimate

[24] See Appendix A.2.2 for figures.

[25] Specifically I use the "lowess" command in Stata.

$$\ln\left(\frac{w_{L,i,t}/\widehat{\alpha}_{i,t}}{w_{H,i,t}/(1-\widehat{\alpha}_{i,t})}\right) = \rho\ln\left(\frac{a_{i,t}}{b_{i,t}}\right) + (1-\rho)\ln\left(\frac{H^*_{i,t}}{L^*_{i,t}}\right)$$

I do so by estimating the regression

$$\ln\left(\frac{w_{L,i,t}/\widehat{\alpha}_{i,t}}{w_{H,i,t}/(1-\widehat{\alpha}_{i,t})}\right) = \beta_0 + \beta_1\ln\left(\frac{H^*_{i,t}}{L^*_{i,t}}\right) + \epsilon_{i,t}$$

using all industries and time periods (379 industries by 4 years yielding 1,516 observations). The coefficient β_1 provides an estimate of $1-\hat{\rho}$. The industry and time specific estimates of $a_{i,t}/b_{i,t}$ are exactly identified and come from adding the constant and the error term, so that

$$\frac{\widehat{a}_{i,t}}{\widehat{b}_{i,t}} = e^{(\beta_0+\epsilon_{i,t})/\hat{\rho}}$$

Table 1.1 displays my estimates for $\widehat{\rho}$ and the elasticity of substitution $\widehat{\sigma}$. The estimate for $\widehat{\rho}$ is 0.651, implying an elasticity of substitution of 2.86 between low- and high-skill processes. For multiple reasons, this figure is somewhat high relative to those in the SBTC literature, which typically estimate a substitution elasticity between non-college and college workers around 1.5 or 2 (see, e.g., Katz and Murphy, 1992; Katz and Autor, 1999; Autor et al., 2008). First, my elasticity is within the manufacturing sector. Since manufacturing is understood to be a mid-skill sector, it is likely that the non-college workers in manufacturing are higher skilled than in the economy overall, while college workers in manufacturing might be lower skilled than college workers in the economy overall. Removing the tails of the skill distribution is likely to increase the substitutability between the two skill groups. Second, my unit of observation is an industry while typical estimates look across industries, using experience groups as the unit of observation. Within-industry estimates will be trivially higher where industry-specific capital more easily flows between low- and high-skill processes. However, to ensure that my results are not sensitive to my elasticity estimate, I estimate my decomposition using alternative elasticities of 1.5 and 2.0. The results, reported in Appendix Section A.2.4, are qualitatively similar to my main estimates.

The distribution of my estimates for the ratio of labor-augmenting technology parameters $\widehat{a}_{i,t}/\widehat{b}_{i,t}$ are shown in Figure 1.7. This ratio is always less than one throughout my sample, meaning high-skill labor is relatively more productive within its process.There is an apparent shift upward over time, implying faster growth in low-skill labor-augmenting technology relative to high-skill. I need one final step to separately identify $a_{i,t}$ and $b_{i,t}$ using my estimates of $\widehat{a}_{i,t}/\widehat{b}_{i,t}$. For this I combine these estimates with observations of equilibrium output. Equation 1.1 can be written as

$$Y_{i,t} = (a_{i,t} + b_{i,t}) \left[\alpha_{i,t} \left(\frac{a_{i,t}/b_{i,t}}{1 + a_{i,t}/b_{i,t}} L_{i,t} \right)^{\rho} + (1 - \alpha_{i,t}) \left(\frac{1}{1 + a_{i,t}/b_{i,t}} H_{i,t} \right)^{\rho} \right]^{1/\rho}$$

where I have simply pulled out a common term $(a_{i,t}+b_{i,t})$ to the front of the equation. Now I can calculate an estimate of $(\widehat{a}_{i,t} + \widehat{b}_{i,t})$ by

$$\widehat{a}_{i,t} + \widehat{b}_{i,t} = \frac{Y^*_{i,t}}{\left[\widehat{\alpha}_{i,t} \left(\frac{\widehat{a}_{i,t}/\widehat{b}_{i,t}}{1+\widehat{a}_{i,t}/\widehat{b}_{i,t}} L^*_{i,t} \right)^{\widehat{\rho}} + (1 - \widehat{\alpha}_{i,t}) \left(\frac{1}{1+\widehat{a}_{i,t}/\widehat{b}_{i,t}} H^*_{i,t} \right)^{\widehat{\rho}} \right]^{1/\widehat{\rho}}}$$

Combining these estimates of $(\widehat{a}_{i,t} + \widehat{b}_{i,t})$ with my estimates of $\widehat{a}_{i,t}/\widehat{b}_{i,t}$ allows me to separately identify $\widehat{a}_{i,t}$ and $\widehat{b}_{i,t}$. I report figures of the distributions of these parameters by year in Appendix Section A.2.3.

1.4.2 Calculating Decomposition Components

Using my parameter estimates and Equations 1.3 and 1.4, I can calculate the theoretical equilibrium low- and high-skill labor under any combination of parameters. I begin by considering a set of five equilibrium levels of low-skill labor, where I have temporarily dropped the industry subscripts. Equations 1.6 through 1.10 trace out the change in L from its equilibrium in time $t+1$ back to its initial value in time t. First I consider:

$$L^*_{t+1} = Y^*_{t+1} \left(\alpha_{t+1} a^{\rho}_{t+1} + (1 - \alpha_{t+1}) \left[b_{t+1} \left(\frac{1 - \alpha_{t+1}}{\alpha_{t+1}} \right)^{\frac{1}{1-\rho}} \left(\frac{b_{t+1}}{a_{t+1}} \right)^{\frac{\rho}{1-\rho}} \left(\frac{w_{L,t+1}}{w_{H,t+1}} \right)^{\frac{1}{1-\rho}} \right]^{\rho} \right)^{-\frac{1}{\rho}} \tag{1.6}$$

I now replace output with output from time t:

$$L^*_{scale} = Y^*_t \left(\alpha_{t+1} a^{\rho}_{t+1} + (1-\alpha_{t+1}) \left[b_{t+1} \left(\frac{1-\alpha_{t+1}}{\alpha_{t+1}} \right)^{\frac{1}{1-\rho}} \left(\frac{b_{t+1}}{a_{t+1}} \right)^{\frac{\rho}{1-\rho}} \left(\frac{w_{L,t+1}}{w_{H,t+1}} \right)^{\frac{1}{1-\rho}} \right]^{\rho} \right)^{-\frac{1}{\rho}} \tag{1.7}$$

Equation 1.7 represents the theoretical equilibrium low-skill labor implied by producing the original output with new technology and relative wages. Given the production structure, this is equivalent to a movement along the expansion path as depicted in Figure 1.7.[26] The difference between L^*_{t+1} and L^*_{scale} represents the scale effect. From Equation 1.7, I modify the task parameter to that from time t:

$$L^*_{task} = Y^*_t \left(\alpha_t a^{\rho}_{t+1} + (1-\alpha_t) \left[b_{t+1} \left(\frac{1-\alpha_t}{\alpha_t} \right)^{\frac{1}{1-\rho}} \left(\frac{b_{t+1}}{a_{t+1}} \right)^{\frac{\rho}{1-\rho}} \left(\frac{w_{L,t+1}}{w_{H,t+1}} \right)^{\frac{1}{1-\rho}} \right]^{\rho} \right)^{-\frac{1}{\rho}} \tag{1.8}$$

The difference between L^*_{scale} and L^*_{task} represents the task effect. Output and relative wages have remained the same, but the shape of the isoquant has changed. Next I modify the productivity parameters a and b:

$$L^*_{productivity} = Y^*_t \left(\alpha_t a^{\rho}_t + (1-\alpha_t) \left[b_t \left(\frac{1-\alpha_t}{\alpha_t} \right)^{\frac{1}{1-\rho}} \left(\frac{b_t}{a_t} \right)^{\frac{\rho}{1-\rho}} \left(\frac{w_{L,t+1}}{w_{H,t+1}} \right)^{\frac{1}{1-\rho}} \right]^{\rho} \right)^{-\frac{1}{\rho}} \tag{1.9}$$

Equation 1.9 represents the equilibrium labor implied by time t output and technology, with time $t+1$ relative wages. The difference between L^*_{task} and $L^*_{productivity}$ represents the productivity effect. Finally, I adjust relative wages:

$$L^*_{supply} = Y^*_t \left(\alpha_t a^{\rho}_t + (1-\alpha_t) \left[b_t \left(\frac{1-\alpha_t}{\alpha_t} \right)^{\frac{1}{1-\rho}} \left(\frac{b_t}{a_t} \right)^{\frac{\rho}{1-\rho}} \left(\frac{w_{L,t}}{w_{H,t}} \right)^{\frac{1}{1-\rho}} \right]^{\rho} \right)^{-\frac{1}{\rho}} \tag{1.10}$$

[26] Note that the illustrative Figures 1.7 through 1.7 represent the movement from time t to $t+1$ while here I decompose from time $t+1$ to time t.

Equation 1.10 differs from Equation 1.9 only in the wage ratio, which is now from time t. This represents a movement along the isoquant to its tangency point with the time t relative wages. The difference between $L^*_{productivity}$ and L^*_{supply} is the labor change caused by the supply channel.

Now all parameters have been replaced with the initial values, so that L^*_{supply} equals L^*_t. The entire change in labor is thus decomposed as an identity:

$$\begin{aligned} L_{i,t+1} - L_{i,t} = &L_{i,t+1} - L_{i,scale} \\ &+ L_{i,scale} - L_{i,task} \\ &+ L_{i,task} - L_{i,productivity} \\ &+ L_{i,productivity} - L_{i,t} \end{aligned} \tag{1.11}$$

where each line represents the scale, task, productivity, and supply channels, respectively. Common to decompositions in general, magnitudes may be sensitive to the ordering of the theoretical equilibria.[27] To address this, I calculate the effects under all 24 possible arrangements, and take the mean as my estimate. The identity still holds. I conduct a parallel decomposition for high-skill labor.

1.5 National Trends

1.5.1 Aggregating the Decomposition

The decomposition procedure described in Section 1.4.2 provides me with estimates of the effect of each channel by industry and skill type between any two time periods. I first use this to describe changes in the manufacturing sector between 1990 and 2007. To evaluate the overall effect of each channel, I simply sum up the industry-level effects as follows:

$$\Delta S_{2007,1990,c} = \sum_i \Delta S_{2007,1990,c,i} \tag{1.12}$$

where S represents labor hours for skill type $S \in L, H$, and $\Delta S_{2007,1990,c,j}$ represents the change in labor hours caused by channel c for skill type S in industry i between

[27] As one example from consumer theory, Eugen Slutsky and John Hicks each propose a method to decompose price change responses into income and substitution effects. Their methods differ based on which effect you calculate first.

2007 and 1990. To interpret the changes in terms of job counts, I divide the labor hours by the mean annual hours observed for each skill type for manufacturing workers in the 1980 Census.[28] In Section 1.5.2, I describe the overall results and then break down the analysis by ten broad industry groups. Appreciable advancement in computer technology during this time frame is associated with dramatic restructuring for computer producing industries.[29] While the results for these industries follow similar trends to the other industries, the impacts of each channel are of high magnitudes. For ease of interpreting the results, I omit computer industries from the main analysis but report them separately in Appendix Section A.2.5.[30]

1.5.2 Results

National Landscape

I show the decomposition results in Table 1.2.[31] Panel A displays overall employment changes by skill type. Total employment between 1990 and 2007 decreased from 15.00 million to 12.41 million, a loss of 2.6 million job equivalents. These losses were driven entirely by low-skill work, while the number of high-skill jobs *increased* by nearly 11 percent. This evidence of job growth for college workers simultaneous to steep job declines for non-college workers hints at the importance of skill-biased change in the restructuring of this sector. In Column (4) I report the share of employment belonging to high-skill workers, which increased by 4.9 percentage points, over one-third of its initial skill share.

Panel B breaks the observed employment changes into each channel, both in terms of levels and percent changes. Because the decomposition is based on an identity, each channel in terms of levels and percent changes will sum up to the overall changes reported in Panel A.[32] Beneath each level estimate, I report the 95 percent confidence

[28] These annual hours are 2141.77 for low-skill workers and 2224.04 for high-skill workers. Because average hours have increased since 1980, overall magnitudes are somewhat muted.

[29] See Houseman (2018) for an analysis of productivity improvements in computer industries relative to other manufacturing industries.

[30] Also omitted throughout are the few industries for which consistent materials use data is not available. I provide more details, including a list of omitted industries, in Appendix Section A.1.1.

[31] For a figure of these results, see Appendix Section A.2.5. For results using alternative elasticity estimates, see Appendix Section A.2.4.

[32] The implied percentage point change in skill-share reported in Column (4) does not necessarily sum to the overall percent change in skill-share.

interval based on 1000 bootstrapped industry samples. For each bootstrapped sample, I repeat the entire estimation as described in Section 1.4.

It is immediately apparent that scale forces have had a strong and positive impact on employment for both types of workers, adding nearly as many jobs as the overall observed loss for low-skill workers, and well over twice the total observed increase for high-skill workers. This is evidence that declining domestic production is not the key driver behind employment loss, and that these losses must be coming through other channels. The scale channel is also not especially important for upskilling, contributing only slightly to the increase in skill share. This implies that low- and high-skill industries had roughly even production growth overall.

The strong positive forces from scale are more than offset by job displacement from the two technology channels, in terms of overall employment. Together these two channels displaced over 6 million workers. For low-skill workers, the task channel is the most important contributor, costing 4.13 million jobs. Shifts toward high-skill production in turn created high-skill jobs. High-skill tasks are inherently less labor-intensive, so that each low-skill job lost translates to less than one high-skill job gained. The task channel is also a major contributor to sector-wide upskilling, with an implied increase in the skill share doubling the initial share in 1990.

The productivity channel, in contrast, displaced high-skill workers at a faster rate than low-skill workers, somewhat offsetting the skill-biased effect of the task channel. The productivity channel was also the largest source of job loss overall, leading to a decline of 3.77 million jobs. Larger losses in percent terms for high-skill workers are the result of heterogeneous growth in labor-augmenting technology and imperfect substitutability between low- and high-skill processes.[33]

Finally, the labor supply channel reflects changes due to relative wages facing the firm. These effects are small in light of the other channels. We see an increase in low-skill employment by 7.2 percent. This implies a general flattening of low-skill wages relative to high-skill caused firms to increase their share of low-skill workers. These changes lead to a lowering of the skill-share by 3.3 percent. A possible explanation for these results is that there was an increase in demand for high-skill workers in non-manufacturing industries, raising their market wage and reducing their residual supply to manufacturing firms.

[33] This is in spite of the fact that $a_{i,t}/b_{i,t}$ generally increased during this time frame, as shown in Figure 1.7.

Taken together, these results indicate that reduced domestic production, for example as firms move production offshore, or in response to consumers shifting from domestic to foreign-sourced goods, is not the driving force behind employment loss in manufacturing or its widespread upskilling. Instead, production increased during this time period, driving up employment for all workers. These gains are more than offset by investment in technology and shifting production processes away from labor-intensive low-skill processes. Task upgrading, an understudied factor driving employment loss, explains 64 percent of low-skill job loss and 41 percent of overall job loss. I explore heterogeneity in these trends by broad industry group in Section 1.5.2.

Breakdown by Broad Industry Groups

Figure 1.7 shows the decomposition by low- and high-skill employment for each of ten broad industry groups.[34] In each bar of the histogram, the effects of all four channels on employment (in terms of thousands of job equivalents) are stacked, so that overall employment change for the skill group is the sum of the above-zero changes net of the below-zero changes. A few interesting findings emerge.

First, there is simultaneously employment growth due to some channels and employment loss due to other channels *within* each industry group and skill type. These changes point to a significant restructuring, even in cases where the net employment effects are small. The largest evidence of restructuring in terms of levels is in Metal Products, Machinery and Equipment, and Transportation, which are also the three largest industry groups. Steep employment loss from productivity for both skill types is consistent with investments in automation as firms continue to increase production. In percentage terms, perhaps the most interesting evidence of restructuring is Chemicals and Petroleum, which was initially the highest skill group yet still saw the largest percent increase in skill share, from 28.6 percent in 1990 to 38.6 percent in 2007. This change is driven by an upskilling of production tasks paired with scale increases.

There are also several fairly consistent trends across all industry groups. For high-skill employment, the magnitude of job gains due to task shifts is similar to the magnitude of job losses due to the productivity channel. This correlation implies that

[34] I follow the industry groupings in Autor et al. (2014). See Appendix Figure A.2.5 for the same figure with an additional column for the omitted computer industry group.

as industries shift toward higher skill production processes, workers are becoming more productive within their process. This may be due to increased investment in high-skill augmenting technology. For low-skill workers, the relationship between task and productivity is not as strong. On the one hand, industries may lay off their lowest productivity workers as they task upgrade, which can appear as additional losses through the productivity channel. Industry groups such as Textiles and Apparel and Metal Products exhibit large job losses due to task upgrading paired with large job losses due to the productivity channel. On the other hand, industries that task upgrade may also reduce their investment in low-skill technology, so that productivity effects are smaller than in other industries that do not upgrade. Food and Tobacco, Paper and Printing, and Chemicals and Petroleum all exhibit large task shifts with relatively small effects from the productivity channel.

Another common trend is that supply shifts tend to transfer jobs from high-skill to low-skill workers, as high-skill workers have become more expensive to these industries. About 90 percent of the industry sample experiences a low-skill bias due to the supply channel. These effects are consistently small relative to the other channels. Finally, the scale channel is consistently positive across these groups. One particular exception is the Textiles and Apparel group, which experienced a scale-induced decline of over 1.2 million workers, heavily concentrated on low-skill jobs. One possible explanation is that much of this production has moved offshore: This group is known to have experienced a large influx of imports from China.[35]

While Textiles and Apparel is the only group with scale losses overall, select industries in the other groups also showed scale losses. In percentage terms, two of the highest are tobacco stemming and redrying (SIC code 2141) and manifold business forms (SIC code 2761), neither of which experienced competition from China.[36] Absent import competition, institutional or demand forces may have played a role in these declines. I explore the tobacco industries in more detail in Appendix Section A.2.6.

[35]Of the 351 industries, Textiles and Apparel industries make up over a third of the top 10 percent in terms of import penetration between 1990 and 2007. For details on my calculation of import penetration, see Section 1.6.1.

[36]Both industries fall in the bottom 10 percent in terms of import penetration between 1990 and 2007. See Section 1.6.1 for details on my calculation of import penetration.

1.6 Applications

1.6.1 China Shock Decomposition

Section 1.5 shows substantial evidence that scale was not the leading force behind employment declines between 1990 and 2007, nor was it an important driver of upskilling in manufacturing. Still, many industries, especially low-skill industries in Textiles and Apparel, experienced job losses predominantly through the scale channel.

A likely cause is China's export-oriented economic expansion in the 1990s and 2000s, which lead to unprecedented import penetration into the U.S. and other developed countries. The significant role of this shock in spurring the decline of manufacturing employment is well-documented (see Autor et al., 2013; Acemoglu et al., 2016; Pierce and Schott, 2016). In theory, this price competition should lead to job loss through scale decreases rather than technology changes. Indeed, the literature has found little evidence that trade-exposed industries responded by increasing capital investments or patent grants (Pierce and Schott, 2018; Autor et al., 2016), and while low-wage workers experienced greater earnings losses and lower ability to transition out of manufacturing, all worker skill types suffered employment loss (Autor et al., 2013, 2014).

In this section I explore the impact of the Chinese imports shock using established methods with my decomposed employment change variables. This way I can quantify the extent to which employment changes operated through scale or other channels.

Empirical Strategy

I generally follow the empirical approach described in Acemoglu et al. (2016). I provide details on data sources and treatment in Appendix Section A.2.7 and construction of the import exposure measure and its instrument in Appendix Section A.2.7. My main specification is

$$\Delta lnS_{i,t} = \alpha_t + \beta_1 \Delta IP_{i,t} + e_{i,t} \tag{1.13}$$

where $\Delta lnS_{i,t}$ is the change in log annual labor hours for skill type $S \in L, H$, whether overall or through a particular channel, α_t is an indicator for time period, $\Delta IP_{i,t}$ is import penetration for industry i, instrumented, and $e_{i,t}$ is an error term.[37] Regressions are weighted by start-of-period employment, and standard errors are clustered at the 3-digit SIC industry. I focus my analysis on the rise in imports from 1991-2007. In the main estimates I use stacked first differences in outcomes from 1990-2000 and 2000-2007, stopping in 2007 to avoid idiosyncracies resulting from the recession years. These specifications exhibit a strong first stage.

I show descriptive statistics for low-skill and high-skill samples in Appendix Section A.2.7. Import penetration variables have been annualized so that they are interpreted as 100 times the annual change in import penetration, following the literature. The outcome variables have likewise been annualized so that they can be interpreted as 100 times the annual change in log labor hours (in thousands). Import penetration variables will differ between low- and high-skill workers only to the extent that they are initially concentrated in different industries.

Results

I report the main estimates in Table 1.3.[38] Panels A and B show results from separate regressions on low-skill and high-skill employment, respectively. Column (1) displays the impact of import penetration on overall employment, and is generally comparable to the results in Column (4) in Table 2 of Acemoglu et al. (2016).[39]

The results in this column indicate that for every 1 percentage point increase in import penetration, there is a 1.4 log point decrease in low-skill employment and a 0.63 log point decrease in high-skill employment. Even though both skill types suffer losses, the marginal impact is much higher for low-skill workers. There are several possible explanations for this. It could be that industries in which high-skill workers are concentrated are better able to absorb imports shocks, so that fewer workers are laid off overall. It could instead be that industries are responding by shifting production toward high-skill processes, shielding high-skill workers from displacement. It could also be due to uneven labor-augmenting technology advancements, such as

[37]For this analysis, I modify the decomposition expressed in Equation 1.11 by using a multiplicative identity rather than additive. This way, log changes from each channel will sum to the total log change in employment. Details are in Appendix Section A.2.7.

[38]Estimates for alternative time horizons can be found in Appendix Section A.2.7.

[39]My results differ because I omit computer industries and because I split the sample by skill type. I also use data from the Census of Manufactures and Census of the Population to create my employment variables, while Acemoglu et al. (2016) uses County Business Patterns. Finally, my outcome variables are in terms of thousands of annual labor hours rather than job counts.

adoption of industrial robots that replace low-skill workers. To explore the possibilities, in Columns (2) through (5) I replace the total effect with the effect of each channel. The coefficients in Columns (2) through (5) sum to the total effect in Column (1), so that the decomposed elements explain the entire observed change as an identity.

Starting with Column (2), about 66 percent of low-skill employment loss comes through scale. For high-skill workers, the magnitude of scale is nearly the same as the overall observed impact. This reinforces the notion that reductions in the manufacturing workforce caused by Chinese import penetration have occurred through output declines. Since log scale effects for low-skill workers closely track that for high-skill workers within an industry, the difference in marginal effects suggests that industries in which high-skill workers were concentrated were better able to absorb the imports shock. These coefficients however are not statistically different at conventional levels, suggesting that any contribution of import penetration to upskilling on this margin are minor.[40]

Likewise, there was no impact of import penetration on the allocation of tasks within the industry, shown in Column (3). This is consistent with recent literature finding that Chinese import competition reduced spending on research and development spending and patent adoption (Autor et al., 2016).

There is however some evidence of job displacement due to productivity gains for both skill types in Column (4). These changes may be in part due to firms laying off their least productive workers, raising the average productivity of the remaining labor. The effect may also be related to the availability of capital. Analyses on the impacts of Chinese imports exposure during this time period have not found significant decreases in industry-level capital stock, despite decreased capital investment (Pierce and Schott, 2016, 2018). Given the employment losses, capital per worker still increases.[41] The relative sluggishness of capital may contribute to increased productivity for the remaining workers. In light of the evidence in other research of declining capital investments in exposed industries, automation is not a likely explanation for these productivity channel effects.

[40] A fully interacted stacked IV regression comparing the effect of low- and high-skill employment through the scale channel has a p-value of 0.104.

[41] Pierce and Schott (2016) finds increases in capital per worker at the industry and plant level.

Back-of-the-Envelope Job Losses

Using the coefficients from Table 1.3, I calculate back-of-the-envelope estimates for the effect of Chinese import penetration on employment losses by channel. Given the mean imports exposure for 1990-2000 and 2000-2007 shown in Appendix A.2.7, and a partial R-squared from the first stage regression 0.429 for low-skill and 0.498 for high-skill, I follow equation (4) of Acemoglu et al. (2016) by writing the difference between actual and counterfactual manufacturing employment in time $t+1$ as

$$\Delta S_{t+1,t}^{counterfactual} = \sum_i S_{i,t+1}(1 - e^{-\hat{\beta}_c \Delta \widetilde{IP}_{i,t+1,t}/100*\text{years}}) \tag{1.14}$$

where $\Delta\widetilde{IP}_{i,t}$ is the increase in import penetration from China that can be attributed to China's improving competitive position during the time period, meaning it is the observed change import penetration multiplied by the predictive power of the first stage. This number is then divided by 100 and multiplied by the number of years between t and $t+1$ to convert the annualized percentage point changes into overall effects. I then convert these effects from annual hours into job equivalents by dividing the effects by the mean annual hours of a manufacturing worker by skill type in 1980, as described in Section 1.5.

These estimates indicate that, between 1990 and 2007, Chinese import penetration cost a job loss of about 561,000 jobs, of which 386,000 were due to scale.[42] These results provide evidence that industries under pressure from Chinese competition did not respond primarily by adopting a more capital-intensive production method or otherwise innovating, but rather by scaling down production.

1.6.2 The Roles of Automation and Offshoring

Two additional mechanisms credited with the decline in production jobs are the falling costs of automation technologies and the movement of production activities offshore to countries with cheaper labor. Automation is understood to replace workers who carry out routine, codifiable tasks that can be programmed and accomplished by machines (Autor et al., 2003). Offshoring transfers tasks from domestic workers to

[42]This may be compared to the estimate in Acemoglu et al. (2016) which reports 853,000 jobs lost between 1990 and 2007. Given the smaller levels of manufacturing employment I observe in my industry sample, this smaller estimate is unsurprising.

workers abroad, meaning tasks which do not require physical proximity to customers or specific worksites are likely most susceptible (Autor and Dorn, 2013). While distinct, the occupations concentrated in offshorable tasks and automatable tasks are largely overlapping. I therefore consider both mechanisms together.

Data and Specification

I determine each industry's potential for automation and offshoring based on its mix of occupations in 1980. The variable I use to rate each occupation's potential for automation is based on its routine task share according to the 1977 Dictionary of Occupational Titles. For offshoring potential, I use a variable based on the occupation's requirements for face-to-face contact and physical presence on the job site, according to O*NET data. Both these variables are made available at the Census occupation level by Autor and Dorn (2013), and more details on their construction can be found in their paper.

To calculate the intensity of routine and offshorable tasks at the SIC industry level, I follow a procedure analogous to my calculations of SIC-level hours and wages by skill type described in Section 1.3.2. That is, I exploit the geographic overlap in 1980 between SIC industry locations in the CBP and workers' occupations by Census industry in the Census of the Population. Because these variables are correlated, I consider both together in one estimation. Specifically, I estimate

$$\Delta lnS_{i,2007,1990} = \beta_0 + \beta_1 \text{routine}_{i,1980} + \beta_2 \text{offshore}_{i,1980} + \gamma X_i + e_i \tag{1.15}$$

where $\Delta lnS_{i,2007,1990}$ is the annual log change in thousands of labor hours between 1990 and 2007 for skill type S in industry i, whether overall or through a particular channel. The coefficients on routine and offshore are the effects of interest, and X_i is the high-skill share of employment in 1990 which acts as a control. Therefore the comparison is between industries with the same initial share of high-skill employment, but which differ in their propensity to move parts of production offshore or to automate certain production tasks. Within each regression sample, I standardize the routine and offshore variables to be mean zero with a standard deviation of one. As in the China shock application, I weight the regressions by start-of-period labor hours of the relevant skill group, and use robust standard errors clustered at the 3-digit SIC level.

Results

I report my results by skill type in Table 1.4. Exposure to automation based on routine task share is associated with productivity-induced job losses for both low- and high-skill workers, paired with a downgrading of production tasks, to some benefit for low-skill work. This means industries which are initially concentrated in more automatable tasks subsequently adopt more labor-augmenting technologies which reduce employment, and also are slower to transition toward high-skill production tasks. While precisely estimated, the magnitudes are small, so that the overall effect on employment is muted.

Offshorability, in contrast, is associated with job loss overall for both types of workers, though the marginal impact is stronger for low-skill. The decomposition demonstrates that these job losses are predominantly through the scale channel, as production is moved offshore. There is also some evidence of task upgrading in these industries. This suggests that within an industry, low-skill tasks are offshored, and domestic activity shifts toward high-skill tasks. These skill biases are consistent with Hummels et al. (2014) who, looking at another high-income country during this time period, find that offshoring causes firms to reduce their workforce primarily through a reduction in low-skill workers. Interestingly, there is no impact on productivity, in contrast to automation and Chinese import competition.

Unlike the China shock application in which exposure is determined by volumes of imports, exposure to automation and offshoring is determined here by relative concentration of particular occupational tasks. Quantifying the impact in terms of job counts therefore comes from evaluating the effect of moving along the distribution of industries. For example, in Panel A of Table 1.4 we see that a one standard deviation increase in routine share leads to a 0.634 log point annual decrease in low-skill employment through the productivity channel. We can use this estimate to calculate the impact of moving from the 10th to 90th percentile among manufacturing industries in concentration of routine tasks, an increase of about 2.56 standard deviations.[43] This implies an annual log point decrease of 1.62, equivalent to about 40 jobs over the 17

[43]This is based on a z-score movement from -1.28 to 1.28, or the 10.03 percentile to the 89.97 percentile. I multiply the marginal effect by 2.56 to capture the annual effect of a movement from the 10th to 90th percentile.

year period. This number can be compared to a movement along the distribution of the low-skill productivity job loss from the 10th to the 90th percentile, which implies a 3.65 log point annual decrease, or 305 jobs over 17 years.[44] By this comparison, automation explains about 13 percent of the total low-skill job loss from the productivity channel. For high-skill workers, who experienced sharper job losses through the productivity channel, this estimate is just over 3 percent.

I likewise calculate back-of-the-envelope estimates for the impacts of offshoring. These estimates suggest a movement from the 10th to the 90th percent of the offshorability distribution is associated with a scale-induced low-skill job loss that is about 1,070 jobs, or 23 percent of the magnitude of the scale channel job loss associated with a movement from the 10th to the 90th percentile of the scale channel distribution. For high-skill workers, it is about 5 percent. With respect to task shifts, this same comparison is nearly 27 percent of task channel job losses for low-skill workers and about 13 percent of task channel job increases for high-skill.

Because of this relative interpretation, it is important to note that these effects are possibly attenuated in that they do not capture the average impact of offshoring and automation on manufacturing. Instead they capture cross-industry differences.

1.7 Conclusion

U.S. manufacturing employment dropped sharply during the 1990s and 2000s. At the same time, the use of technology such as computers and electronic networks increased dramatically, and volumes of imports into the U.S., especially from China, reached unprecedented levels. These simultaneous and in some ways intertwined events make it a challenge to separate the causes underlying the steep employment declines we observe.

This paper offers a new approach to understanding employment loss, by reinterpreting observed employment changes as the net effect of four distinct, and often opposing, forces. Combining market equilibrium conditions with equilibrium output, I separate employment changes by skill type into what is explained by scale (output), the mix of production tasks, labor productivity within those tasks, and labor supply. While applied here to manufacturing, this decomposition could be used in a variety of competitive settings using a broad class of production technologies. Adopting

[44] As in the corresponding regressions, I weight this distribution by 1990 low-skill hours in the industry when taking the 10th and 90th percentile.

this particular model, originally developed in Katz and Murphy (1992), I distinguish task shifts from labor productivity changes using the share of production materials allocated to each skill type.

The results indicate that a sweeping shift toward high-skill tasks explains 64 percent of employment loss for low-skill workers. Because low-skill workers predominate the manufacturing workforce, and high-skill tasks are inherently less labor-intensive, these task shifts are also responsible for 41 percent of overall employment loss. This evidence is surprising given the prevailing view that employment loss has been caused by some combination of jobs moving to low-income countries and jobs being replaced by machines. Instead, the results show that scale has worked to *increase* employment for both skill types, overcoming downward pressure from foreign competition and reflecting continued production in the U.S. Labor-augmenting technology, associated with automation, does cause displacement, but it is not nearly as important as task shifts in explaining low-skill job loss. In an application I find that automation is if anything associated with a slower transition to high-skill tasks, suggesting that this task upgrading is an alternative to automation, rather than a symptom of it.

This leaves open the question of what factors have led to the widespread task upgrading in manufacturing that I document. Consumer demand may have driven industries toward products that require more high-skill tasks to produce (Xiang, 2005). Another possibility is that increased government oversight raised the demand for high-skill tasks such as quality control and supervision. Changing management practices and other process improvements distinct from automation could also have given high-skill workers a comparative advantage (Bender et al., 2018). I leave further exploration of these and other factors for future research.

Irrespective of cause, these results point to a new era of manufacturing production, characterized by high-skill, high-productivity tasks. This insight informs our thinking about the future of work in manufacturing, as we consider ways to improve education and vocational training. While new trade policies and evolving consumer preferences will continue to drive the ebb and flow of demand for manufacturing workers, the evidence herein shows that the nature of the jobs at hand will be remarkably different from those of the past. Given the importance of manufacturing in the broader labor market, this knowledge can also further our understanding of recent economy-wide trends, for example in inequality. It likewise raises the question of whether these patterns exist in other sectors, where they would be harder to measure.

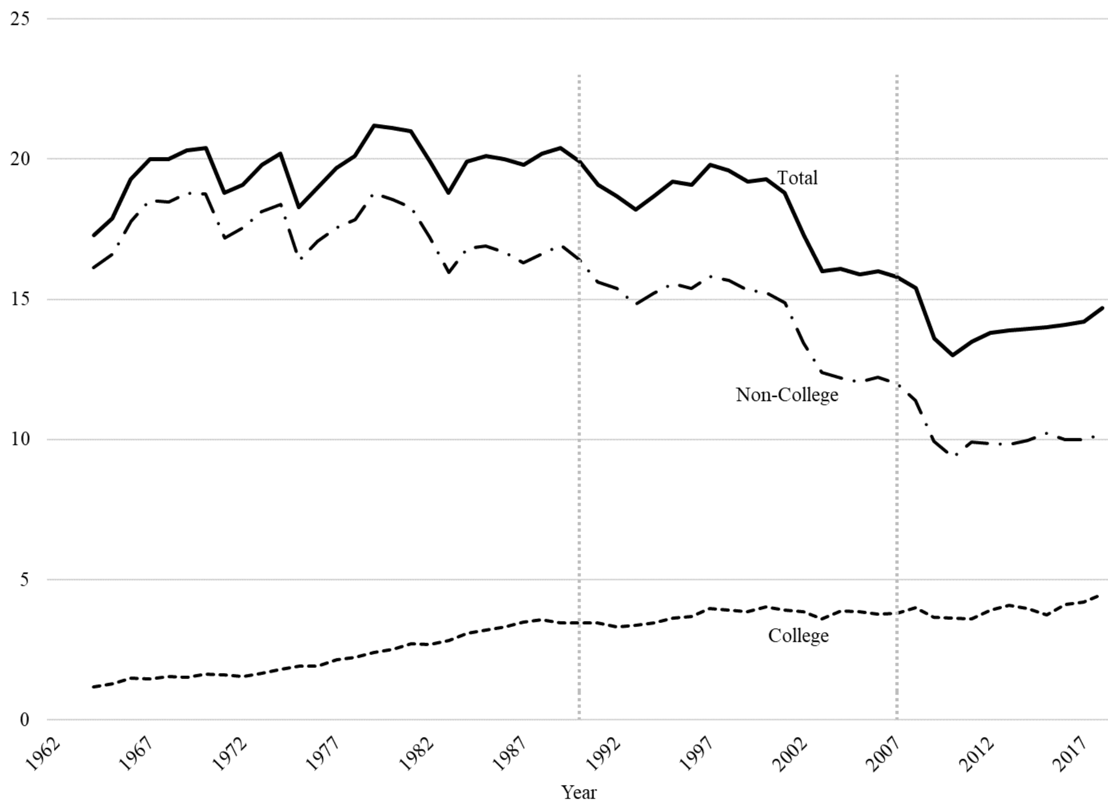

Figure 1.1.: U.S. Manufacturing Employment, Millions

Notes - Author's calculations from the CPS ASEC annual surveys 1962-2018. Sample is employed wage and salary workers ages 16-64, exclusive of self-employed, unpaid family workers, and military workers. Prior to 1992, individuals reporting at least 4 years of college are considered bachelor's degree holders.
Source - Flood et al. (2018)

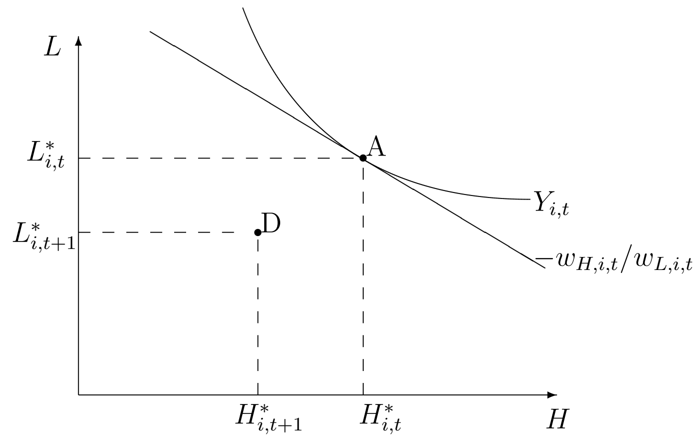

Figure 1.2.: Total Equilibrium Change: A $\rightarrow$ D

Notes - This figure illustrates the change in employment between time t and $t+1$ for industry i, where A and D are the equilibrium employment levels in times t and $t+1$, respectively, for low-skill labor L and high-skill labor H. See Section 1.2 for details.

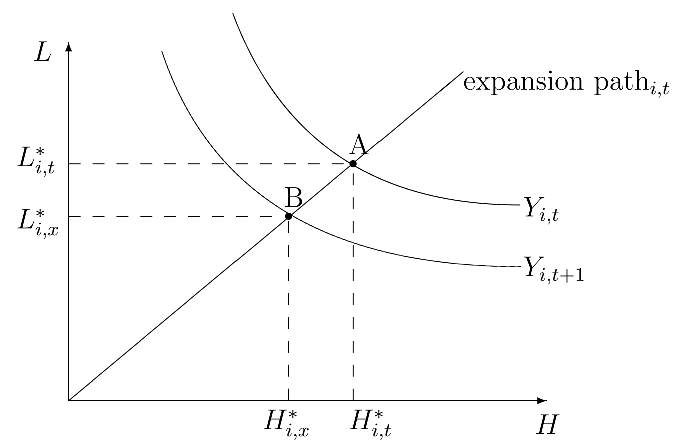

Figure 1.3.: Scale Channel: A $\rightarrow$ B

Notes - This figure illustrates the change in employment between time t and $t+1$ for industry i due to the scale channel. See Section 1.2 for details.

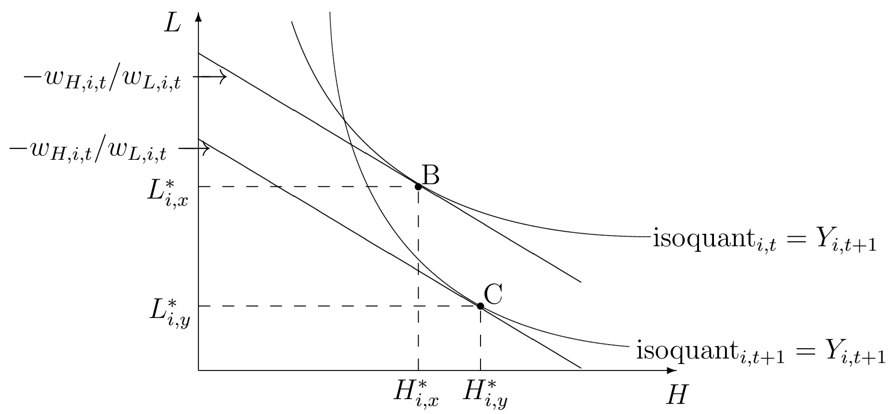

Figure 1.4.: Task and Productivity Channels: B $\rightarrow$ C

Notes - This figure illustrates the change in employment between time t and $t+1$ for industry i due to the task and productivity channels. See Section 1.2 for details.

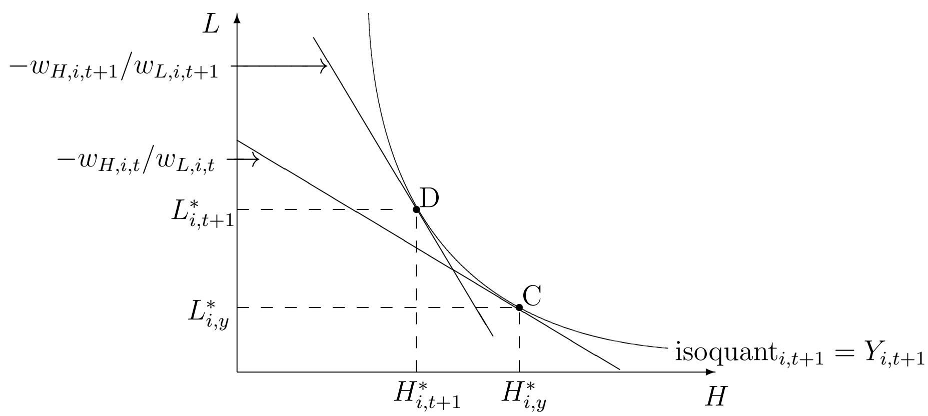

Figure 1.5.: Supply Channel: C $\rightarrow$ D

Notes - This figure illustrates the change in employment between time t and $t+1$ for industry i due to the labor supply channel. See Section 1.2 for details.

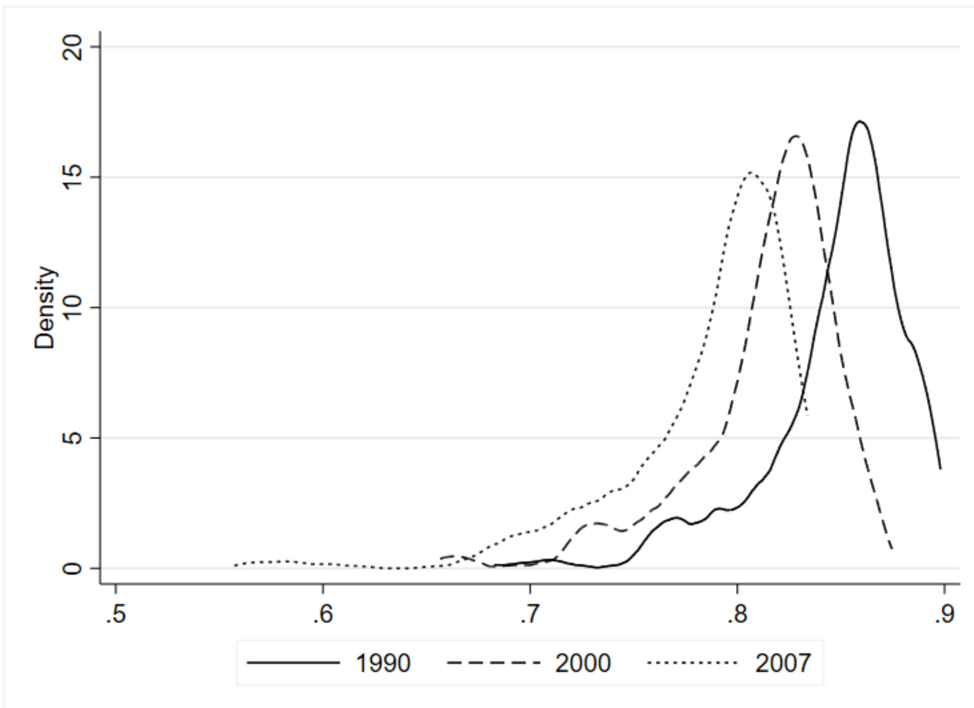

Figure 1.6.: Distribution of Task Parameter ($\alpha_{i,t}$) by Time Period

Notes - This figure shows the distribution of the task share parameter $\alpha_{i,t}$ across my industry sample for the years 1990, 2000, and 2007. A leftward shift of the distribution over time indicates a shift away from low-skill production tasks. See Section 1.4.1 for details.

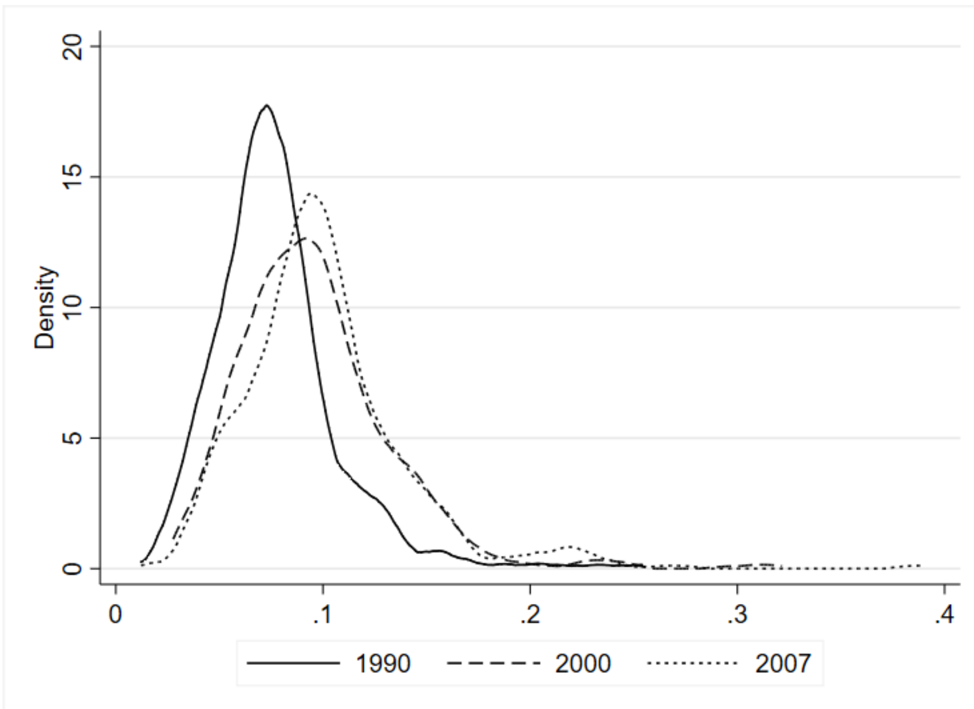

Figure 1.7.: Distribution of Ratio of Skill-Augmenting Productivity Parameters ($\frac{a_{i,t}}{b_{i},t}$) by Time Period

Notes - This figure shows the distribution of the ratio of low- to high-skill productivity parameters across my industry sample for the years 1990, 2000, and 2007. See Section 1.4.1 for details. For figures of the distributions of each parameter separately, see Appendix A.2.3.

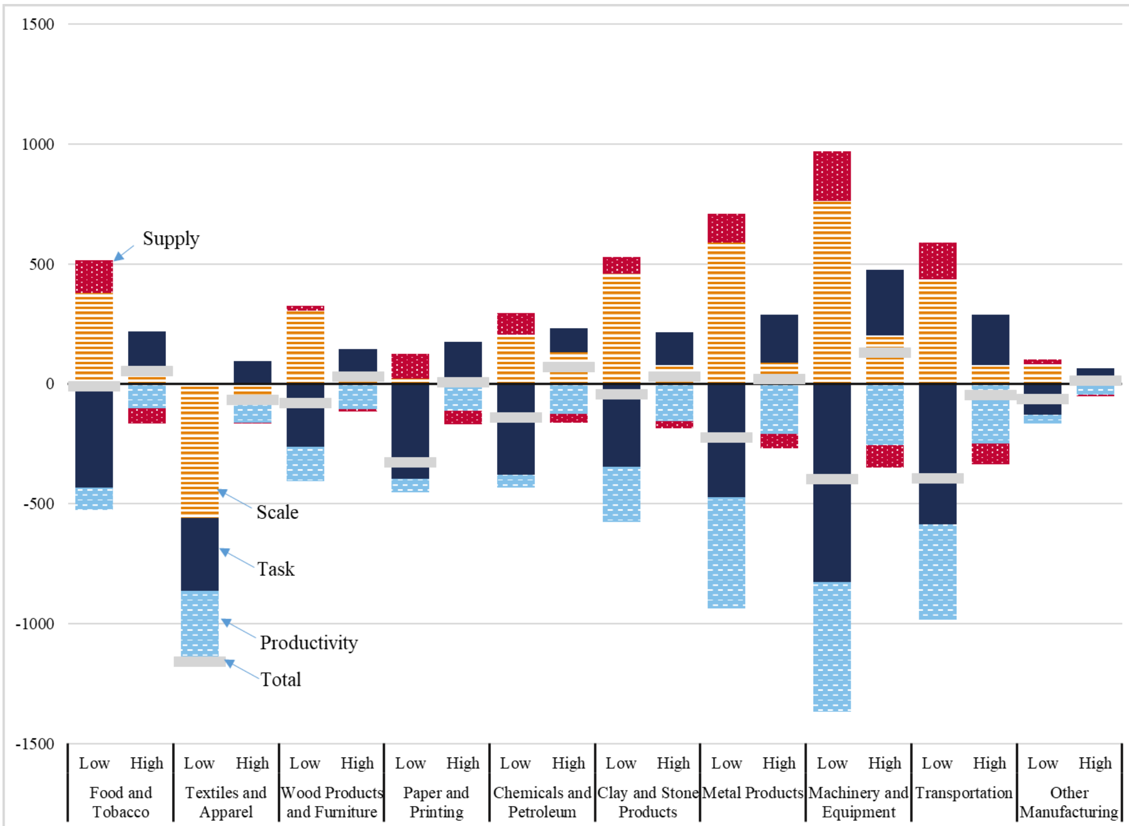

Figure 1.8.: Decomposition of Manufacturing Employment Changes, 1990-2007, by Industry Group, in Thousands of Job Equivalents

Notes - This figure shows the main decomposition results aggregated to 10 industry groups. See Section 1.5.2 for details.

Table 1.1.: Estimates of ρ and σ (Elasticity of Substitution)

	estimate (1)	standard error (2)
ρ	0.651	0.010
σ	2.863	0.086

Notes - This table reports the estimates of the elasticity of substitution between low- and high-skill production processes. See Section 1.4.1 for details.

Table 1.2.: Decomposition of Manufacturing Employment Changes 1990-2007 Into Four Channels

	Jobs (Millions)			
	Low-Skill (1)	High-Skill (2)	Total (3)	High-Skill Share (4)
		Panel A. Overall Change		
1990	12.83	2.17	15.00	14.5%
2007	10.00	2.41	12.41	19.4%
Δ	−2.84	+0.24	−2.60	+4.9pp
$\%\Delta$	−22.1%	+10.9%	−17.3%	
		Panel B. Decomposition		
Scale	+2.68	+0.66	+3.34	+0.9pp
	(1.75, 3.69)	(0.46, 0.84)	(2.21, 4.53)	
	+20.9%	+30.2%	+22.2%	
Task	−4.13	+1.48	−2.65	+15.1pp
	(−5.02, −3.21)	(1.11, 1.88)	(−3.91, −1.33)	
	−32.2%	+68.2%	−17.7%	
Productivity	−2.31	−1.46	−3.77	−8.1pp
	(−3.37, −1.54)	(−1.81, −1.09)	(−5.18, −2.63)	
	−18.0%	−67.2%	−25.1%	
Supply	+0.92	−0.44	+0.48	−3.3pp
	(0.71, 1.20)	(−0.58, −0.34)	(0.13, 0.86)	
	+7.2%	−20.3%	+3.2%	

Notes - Decomposition is calculated for 351 manufacturing industries and then summed to national totals. Decomposed changes in Columns (1) through (3) may not exactly sum to total due to rounding. Employment results are converted from annual hours to millions of estimated jobs based on mean annual hours of employed manufacturing workers of the same skill type in the 1980 Census. In each cell of Columns (1) through (3), I report the level estimate, a 95 percent confidence interval based on 1000 bootstrapped samples, and the estimate in percentage terms. Column (4) reports the implied change in high-skill share of employment.

Table 1.3.: Effects of Direct Exposure to Chinese Imports on Employment 1990-2007: 2SLS Estimates

	Total	Scale	Task	Productivity	Supply
	(1)	(2)	(3)	(4)	(5)
	Panel A: Low-Skill Employment				
100 × annual Δ Chinese	-1.443**	-0.949**	0.008	-0.411**	-0.092*
import penetration	(0.591)	(0.420)	(0.099)	(0.194)	(0.049)
	Panel B: High-Skill Employment				
100 × annual Δ Chinese	-0.629**	-0.604**	0.185	-0.330*	0.120
import penetration	(0.304)	(0.266)	(0.119)	(0.184)	(0.092)
Observations	702	702	702	702	702

Notes - Regressions are weighted by start-of-period labor hours of the relevant skill group. Also included is an indicator for time period. Robust standard errors are clustered at the 3-digit SIC. * p< .1, ** p< .05, *** p< .01

Table 1.4.: Effects of Industry Routine Task Intensity and Offshorability on Employment 1990-2007

	Total	Scale	Task	Productivity	Supply
	(1)	(2)	(3)	(4)	(5)
	Panel A: Low-Skill Employment				
Routine Task Share	-0.499	0.050	0.116**	-0.634***	-0.031
	(0.339)	(0.299)	(0.050)	(0.134)	(0.048)
Offshorability	-1.995***	-1.915***	-0.371***	0.302	-0.011
	(0.307)	(0.377)	(0.054)	(0.240)	(0.034)
	Panel B: High-Skill Employment				
Routine Task Share	-0.474	-0.089	0.064	-0.432**	-0.016
	(0.294)	(0.232)	(0.094)	(0.180)	(0.110)
Offshorability	-0.851***	-1.153***	0.131*	0.057	0.113
	(0.275)	(0.223)	(0.076)	(0.244)	(0.087)
Observations	351	351	351	351	351

Notes - Routine Task Share and Offshorability are calculated by industry based on the occupations of its workers in the 1980 Census. Data for these variables at the Census Occupation level are made available by Autor and Dorn (2013). Both variables are standardized within each regression sample to have mean zero and standard deviation of one. Also included as a control is 1990 skill share of industry. Robust standard errors are clustered at the 3-digit SIC. * p< .1, ** p< .05, *** p< .01

2. STALLED RACIAL PROGRESS AND JAPANESE TRADE IN THE 1970S AND 1980S

with Timothy N. Bond

2.1 Introduction

The mid-1970s through the mid-1980s saw a striking reversal of the economic gains made by black men in the Civil Rights era. From 1962 to 1976, the black/white median earnings ratio rose from 52% to 70%.[1] By 1984, it had fallen to 61%, roughly the same level as it was in 1968 (Figure 2.5). There was a similar erosion in labor force participation and employment (Figure 2.5). Blacks were hit especially hard in areas that experienced manufacturing declines (Gould, 2018), after having made rapid gains in this sector during the 1960s, even surpassing whites (as a fraction of employment; Figure 2.5). These losses are even more surprising given that the black workforce was gaining ground in both quality and quantity of education in this time period (e.g, Card and Krueger, 1992; Neal, 2006). While some important racial inequality indicators would stabilize in the late 1980s, these economic losses continue to be felt today.[2] The causes of this change in fortune remain an open question.

Also during this time period the United States experienced an unprecedented increase in import competition from a rapidly growing East Asian economy: Japan.[3] From 1975 to 1986 American imports of Japanese manufactured goods would grow by an average of $8.5 billion dollars per year, representing an increase from 1.1% to

[1] These figures are constructed from the Current Population Survey (CPS). See Appendix B.1.1 for details of data construction. Note that unlike with our main empirical analysis, these figures include Hispanic whites, as the CPS does not track Hispanic ethnicity in the earlier years. For a more comprehensive review of trends in racial differences in this era, see Smith and Welch (1989), Bound and Freeman (1992), and Lang and Lehmann (2012).

[2] For example, after taking into account the continued declines in labor force participation, the racial gap in median earnings today is at 1950 levels, substantially larger than it was in 1970 (Bayer and Charles, 2018).

[3] Japanese GDP grew by 240% during the 1960s as part of the "Japanese Economic Miracle." It would grow by another 50% during the 1970s due to both capital accumulation and improvements in technology (Boskin and Lau, 1990). This growth was contemporaneous with declining barriers to trade, and a strong U.S. dollar which made U.S. industries particularly vulnerable to rising international competition. See Irwin (2017) for a comprehensive review of U.S. trade in this era.

3.5% of total U.S. spending in this sector (Figure 2.5). This surge in imports would cease in the late 1980s in part due to U.S. trade restraints, a devaluation of the dollar, and a shift of Japanese firms towards foreign direct investment in the United States (Irwin, 2017).

In this paper, we assess the extent to which the Japanese trade boom can explain the deterioration of black economic well-being. We use geographic variation in imports exposure, following the identification and instrumental variable approach introduced in the "China shock" literature by Autor et al. (2013), to look at differences in changes in racial disparities across local labor markets. We find a substantial negative impact of this import competition on black employment outcomes. A $1,000 increase in Japanese imports per worker led to a 0.59 percentage point decrease in a commuting zone's black manufacturing employment rate. However, we find no impact on manufacturing employment in the *aggregate*. Instead, we find *higher* manufacturing employment for whites, offsetting the effect on blacks.

Our results suggest that this disparate impact was a consequence of trade-induced skill upgrading in the manufacturing sector. Job losses were concentrated among black high school dropouts, who found at most limited re-employment in non-manufacturing, while gains in manufacturing employment centered on the primarily white college educated. Likewise, we find a growth in professional occupations within manufacturing, particularly for engineers, and a shift to higher-educated production workers.

Black manufacturing workers in this time period were particularly vulnerable to changes in the relative demand for skill. In 1970, 60% of black manufacturing workers had less than a high school degree compared to 38% of whites, and blacks occupied 15% of manufacturing jobs for those with less than a high school degree (compared to 10% of manufacturing jobs overall).[4] Further these education figures will understate true skill differences given racial disparities in school quality (Smith and Welch, 1989; Card and Krueger, 1992; Neal and Johnson, 1996). 84% of black manufacturing workers were working in production jobs, compared to 66% of whites, and among production workers blacks had on average .5 years less formal education. In the North, where the manufacturing sector was largest, more than half of black workers were recent migrants who were educated in segregated schools during the Jim Crow Era South.[5] In fact, we see the strongest negative effects for Southern-born blacks, who

[4]Unless otherwise noted, all figures in this paragraph are authors' calculations from the CPS.
[5]Figure from authors' calculations from U.S. Census Integrated Public Use Microdata Series samples.

had the lowest quality of education.

This cross-race redistribution of jobs had important consequences for labor market disparities. Nearly all black workers displaced by trade left the labor force altogether rather than finding reemployment, leading to a 0.54 percentage point increase in the labor force non-participation gap for every $1,000 increase in import competition. This same increase led to a 3.6 log point widening of the median male earnings gap, 2.6 log point widening of the household income gap, and 0.6 percentage point widening of the welfare recipiency gap. Given that the average black worker faced a $1,413 increase in exposure to Japanese imports, these effects are substantial, accounting for 17-23% of the decline in relative labor force participation and 34-44% of the decline in relative earnings during this time period.

We explore several alternative mechanisms for this disparate impact. While we find evidence that Japanese trade hastened the "white flight" of residents from central cities, we find no evidence for a suburbanization of manufacturing jobs themselves. We also find little evidence that unionization or racial prejudice can explain the differing impact of trade on employment outcomes. Further, black workers neither lived in areas that were more exposed to imports than whites, nor worked in industries that received a higher degree of import competition.[6] The evidence we present is for disparate responses to exposure, not disparate exposure itself.

A large literature has focused on the negative impacts of the recent growth in Chinese import competition on the American manufacturing sector. The "China shock", an average annual increase of $14.6 billion imports per year from 1991-2007, negatively impacted employment, unemployment, earnings, and job growth; and spurred the decline in manufacturing (Autor et al., 2013; Acemoglu et al., 2016; Pierce and Schott, 2016).[7] However, we know little about whether the American economy's response to China is typical or atypical of trade shocks. Previous studies of trade in the 1980s have focused on the role of exchange rates and trade deficits generally (e.g., Katz and Revenga, 1989; Revenga, 1992), which were influenced by Japan as well as traditional Western trading partners and developing countries such as the Asian

This was a consequence of the 1940-1970 "Great Migration" which saw 4 million blacks move from the rural South to the industrial North. See, for example, Boustan (2009) for a comprehensive review.

[6]The average black worker in 1970 lived in a CZ which experienced an increase in imports per worker of $1,413 and worked in an industry that experienced a .020 increase in the Japanese import penetration ratio. This compares to $1,601 and .024 for whites. Note, however, these industry figures can only be calculated at fairly aggregated level compared to the data we use for CZ-level exposure. See the discussion in section 2.3.1.

[7]Figures are in 1999$ and taken from the U.S. Census Bureau.

tigers; or been limited to specific industries (e.g., Grossman, 1986). We are the first to provide a comprehensive look at the impact of Japan's rapid export expansion on U.S. labor markets.

Our findings suggest the economic consequences of Japanese imports on black Americans were similar to the consequences China has had on the overall labor market. Yet, there are several important differences worth highlighting between both our results and the nature of the trade shock. First, Japan was already a highly developed country when the import expansion began, trailing only the United States and the Soviet Union in GDP in 1972. Second, while China's story has focused on its abundance of cheap labor, much of Japan's success was attributed to innovative management practices. Many would later be copied to mixed success by American firms (Powell, 1995; Ichniowski and Shaw, 1999). Finally, our evidence suggests that Japanese competition led to a change in the skill composition of manufacturing; for China the negative effects have been felt at all skill levels (Autor et al., 2013).

Our identification strategy is based on Autor et al. (2013), which uses variation in import exposure across local labor markets due to differences in the product composition of their manufacturing industries. We account for the endogeneity of trade by instrumenting with the exposure predicted by historical local industry shares and the products exported by Japan to six other highly developed countries. While there has been a recent debate about the validity of such approaches (e.g., Goldsmith-Pinkham et al., 2018; Borusyak et al., 2018), we show that our instrument performs well with respect to several different robustness and validation exercises proposed in the literature.

What caused the reversal of black economic progress is still not well understood. Wilson (1987) and other supporters of "demand-side" explanations proposed this was a symptom of the de-industrializing economy, trade being one of its causes.[8] In support of this theory, several empirical studies have found black workers were disproportionately negatively affected by decreases in labor demand (proxied by changes in national employment by industry) in the 1970s and 80s (e.g, Acs and Danziger, 1993; Bound and Holzer, 1993, 2000). However such studies are unable to disentangle demand decreases caused by foreign competition from other important factors of the time period, such as skill-biased technical change.[9] Murphy and Welch (1991) exam-

[8]See also Kasarda (1989). This was in contrast to "supply-side" explanations, advanced by, among others, Mead (1986), that centered around a decreased willingness of black workers to accept low wage work.

[9]For example, Reardon (1997) finds that blacks were more affected by within-industry skill compo-

ine the susceptibility of various race, gender, and skill groups to trade deficits based on their distribution of employment across four broadly defined industry categories. From this they calculate that the 1980s trade deficits should have increased the black-white wage gap, but their projection is much smaller than the actualized growth, and their model does not allow for differential effects within industry or account for the endogeneity of trade exposure. We provide the first direct evidence, using credible exogenous geographic variation, that increased foreign competition was responsible for a large portion of the decreased labor demand for and subsequent economic malaise of black workers.

The 1980s especially was a time of broad manufacturing declines and increased economic hardship for low-skill workers, and previous work has found some evidence that import competition played a role in these changes (Borjas et al., 1992; Borjas and Ramey, 1995). However, the consensus view is that these structural shifts were primarily driven by other factors, especially skill-biased technical change (e.g., Berman et al., 1994; Feenstra and Hanson, 1999; Katz and Autor, 1999; Autor et al., 2008). Our results are consistent with this view. While we find large aggregate decreases in manufacturing in commuting zones whose pre-existing industrial composition made them vulnerable to Japanese competition, all of this effect can be explained by differences in workforce composition, particularly worker education levels, and the size and occupation mix of the manufacturing sector. Further, because blacks made up a small portion of the labor force, and because black high school dropouts especially were overwhelmingly located at the lowest tail of the skill distribution, any aggregate changes in inequality were small.

Theoretical models of trade generally predict the most disruptive labor market effects occur when import increases come from low wage countries (Krugman, 2000, 2008). While Japan had lower wages than the United States throughout this time period, it was already an OECD member by 1970. Still, several recent theoretical papers have demonstrated that trade can lead to increases in inequality even when both partners are similarly developed. For example, trade can trigger technological advancement within firms as they preempt competitive threats through skill-biased innovations (e.g., Neary, 2002; Thoenig and Verdier, 2003). Alternatively, trade can cause factors to reallocate across firms toward those of higher productivity and skill

sition changes in the 1980s than cross-industry changes in demand, and concludes that technological change is responsible for widening racial disparities. However trade can also cause changes in the relative demand for skilled workers, as domestic firms adopt more competitive practices and close unproductive factories (Bernard et al., 2006).

intensity (Monte, 2011; Burstein and Vogel, 2017). Epifani and Gancia (2008) develop a model where the increased market size caused by reductions in trade barriers can increase demand for skilled workers because of stronger returns to scale in the skill-intensive sector. Our results lend support to these theories. More generally, it is by now well recognized that trade has winners and losers. In this instance, it appears the losses were concentrated on black workers.

The remainder of this paper is organized as follows: Section 2 describes our data sources and treatment; Section 3 explains our identification strategy; Section 4 discusses our results; and Section 5 concludes and conjectures on the significance of our results for the persistence of economic and racial inequalities.

2.2 Data

2.2.1 Labor Market Data

Our primary sources for labor market data are the 1960 5%, 1970 1% form 1 and form 2 metro, and 1990 5% Integrated Public Use Microdata Series (IPUMS) samples of the United States Decennial Census (Ruggles et al., 2015). We also use the 1980 5% state sample in a robustness exercise. As in Autor et al. (2013), we define local labor markets by commuting zones (CZs) using the definitions created by Tolbert and Sizer (1996). We match workers to CZs using Public Use Micro Areas (PUMAs) in 1990 and 1960, and Census County Groups in 1980 and 1970 following crosswalks provided in Autor and Dorn (2013) and Rose (2018). Unless otherwise stated, we restrict our attention throughout to working age males due to concerns about changes in female labor force participation across time. This is particularly important given the racial differences in selection of women into the labor market (Neal, 2004). To ensure an adequate sample for calculating race-specific statistics, we restrict attention to CZs in the continental U.S. which had a black male working age population of at least 500 in both 1970 and 1990. This restriction primarily affects rural Western commuting zones (see Appendix B.2.4) and results in a sample of 358 CZs. We winsorize all CZ-level control and outcome variables by race/year to the 2nd and 98th percentiles to further account for measurement error due to sample size. See Appendix B.1.2 for more details on data construction.

We present descriptive statistics for our sample in Table 3.4.[10] Perhaps most im-

[10]In this table and throughout, we weight by the race-specific 1970 commuting zone working age

mediately striking is the 6.8 percentage point relative decrease in black employment in non-manufacturing, a sector that cannot be directly affected by foreign competition. However, as we show in Appendix B.2.3, this appears to be drive by a continuation of pre-existing trends in the agricultural sector.[11] In the 1960s such workers were likely absorbed into the manufacturing sector as discussed before. The 1970-1990 period instead saw a relative decline in black manufacturing employment, paired with relative increases in black unemployment (3.9 percentage points) and labor force non-participation rate (3.6 percentage points). While blacks saw small gains among those with positive earnings, once including non-earners the median earnings gap for working age males widened by 23 log points. We also see a relative decrease in household income and increase in welfare recipiency.

2.2.2 Import Competition

To calculate industry-level exposure to Japanese imports, we begin with bilateral trade data in SITC Revision 1 from UN Comtrade.[12] Autor et al. (2013) provide a crosswalk from 1992 Harmonized System (HS) product-level codes to SIC 87 industry codes. However, the HS system was not introduced until 1988 and is not consistently available for our countries of interest until the early 1990s. We therefore constructed a new crosswalk from SITC Revision 1 to HS, which we describe in Appendix B.1.3. We then utilize the Autor et al. (2013) crosswalk to bridge our trade data to industries.

Following Autor et al. (2013), we measure import competition through changes in imports per worker (IPW).[13] For each CZ i we calculate

$$\Delta IPW_{uit} = \sum_j \frac{L_{ijt}}{L_{ujt}} \frac{\Delta M_{ujt}}{L_{it}} \tag{2.1}$$

where L_{ijt} is the number of workers in commuting zone i in industry j at the beginning

male population.

[11]These factors including improvements in education, mechanization of Southern agriculture, agricultural mechanization advancements and industrialization in the South, and continued urbanization of the black population. See, for example, Cogan (1982) and Smith and Welch (1989).

[12]These data are available at https://comtrade.un.org.

[13]We prefer this measure as it has been used by the vast majority of research studying the impact of import competition on local labor markets (e.g., Greenland and Lopresti, 2016; Feler and Sense, 2017). It thus allows us to easily compare our estimates with those found in the China Shock literature. Our results are largely robust to instead using the import penetration ratio as is more common in papers studying industry-level effects (e.g., Acemoglu et al., 2016) as well as the recent study of local marriage markets by Autor et al. (2018). See Appendix B.2.2.

of period t, L_{ujt} is that same value for the United States, L_{it} is the total number of workers in these industries in commuting zone i at the beginning of period t, and ΔM_{ujt} is the change in imports in that industry's product space (in \$1000s) during the time period. As in Autor et al. (2013), we restrict IPW to include only manufacturing imports. We explore the geographic dispersion of IPW in more depth in Appendix B.2.4. In general, we find that the most exposed areas were in the Midwest and Northeast, and the least were the inland West and South.

We calculate 1970 CZ-level industry employment using the County Business Patterns (CBP). The CBP is an annual series that provides county-level economic data by industry, including the number of establishments, employment during the week of March 12, and payroll information extracted from the U.S. Census Bureau's Business Register. The 1970 series is reported in SIC 1967 codes, which we convert to SIC 1987 codes.[14] The CBP suppresses the employment counts for some counties to avoid identifying individual employers. As detailed in Appendix B.1.4, we impute employment in these instances based on establishment counts following Autor and Dorn (2013).

As nationwide CBP data is not available prior to 1970, for our instrument we calculate CZ-level industry employment using the 1960 5% IPUMS sample from the Decennial Census. We disaggregate the 1960 Census industry codes into SIC 1987 according to each CZ's industry composition in the 1970 CBP.[15] See Appendix B.1.5 for more details.

2.2.3 Geographic Outcomes

We also explore the impact of import competition on the distribution of workers and jobs between cities and suburbs. Here, we utilize the 1970 Census definition of central cities, under which 167 commuting zones include a central city.[16] We calculate residential populations from place- and county-level tabulations of the 1970 and 1990 Census of the Population available from IPUMS National Historical Geographic

[14] We use an employment-based weighted crosswalk from the NBER-CES Manufacturing Industry Database to convert SIC 1972 to SIC 1987, and construct a parallel crosswalk using the 1972 Census of Manufactures to convert SIC 1967 to SIC 1972.

[15] Our findings are broadly robust to using an instrument constructed directly from Census industry codes, or an instrument based only on the 1970 CBP. These results are available upon request.

[16] We omit four cities (and their associated CZs) which consolidated with their county between 1970 and 1990: Lexington, KY; Indianapolis, IN; Jacksonville, FL; and Columbus, GA. This is to avoid conflating changes in population and employment with the substantial geographical changes they experienced.

Information System (NHGIS; Manson et al., 2017).

For the location of jobs, we use tabulations from the 1972 and 1987 Census of Manufactures (CoM).[17] These data include counts of manufacturing establishments and employees by state, county, and place. Occasionally, employee counts are suppressed. For counties, we impute missing observations by utilizing the state-level counts of establishments and employees net of the counts in non-suppressed counties. This generates a residual employment count that we distribute to the suppressed counties according to their number of establishments, which is always available. For suppressed cities an analogous calculation is not possible, as the CoM provides information only for places with at least 450 manufacturing employees. We therefore impute missing employment counts by multiplying the number of establishments in the city by the average establishment size in the state.

2.3 Empirical Strategy

2.3.1 Specification

We adopt a similar approach to Autor et al. (2013). In our preferred specification, for outcome Y of race $k \in \{w, b\}$ in commuting zone i we estimate

$$\Delta(Y_{ik,1990-1970}) = \alpha_k + \beta_k \Delta IPW_{ui,1990-1970} + \gamma_k X_{i,1960} + \varepsilon_{ik}, \tag{2.2}$$

where $X_{i,1960}$ is a vector of commuting zone characteristics measured in 1960. In words, we estimate a fully-interacted regression that allows for local labor market conditions to affect blacks and whites in different ways. Our main interest is the disparate impact of import exposure on blacks, $\beta_b - \beta_w$, which is most easily displayed as the coefficient on the interaction between $\Delta IPW_{ui,1990-1970}$ and a black indicator.

We prefer using the long difference approach over stacked first differences including 1980 data for several reasons. First, 1980 was a recession year, while 1970 and 1990 were relatively normal economic times.[18] The recession was caused in large part by sudden, steep increases in interest rates by the Federal Reserve, and was thus felt almost exclusively by consumer durables typically purchased on credit, especially

[17] The CoM is part of the Economic Census, and is conducted in regular five year intervals during off-years of the Census of the Population.

[18] According to the NBER, the United States entered recession in June of 1990. However, the 1990 census was taken April 1st, and the income data reflect 1989 outcomes.

automobiles (Westcott and Bednarzik, 1981). Because Japanese import competition was also highest in these industries, we are concerned about conflating the effects of trade with the peculiarities of this recession. Second, as illustrated in Figure 2.5, Japanese imports peaked in 1986 before receding in the latter half of the 1980s. We are thus concerned that the 1990-1980 difference may not accurately reflect the effects of Japanese import competition since all of the change in this decade came four years before the measurement of the economic outcomes. The long difference will be less sensitive to this issue, since it encompasses the totality of the Japanese trade influx.[19] In Appendix B.2.5 we provide estimates for each decade separately, and using a stacked first differences approach as in Autor et al. (2013). We find stronger evidence for negative effects on black employment in the 1970s, and positive effects on white employment in the 1980s, but it is unclear if this is due to the patterns of trade adjustment, or the reasons discussed above. We also perform a validation exercise for our IV by estimating a placebo regression of 1970-1990 import increases on the 1960-1970 change in manufacturing employment, and find no evidence of any effects.

An alternative approach would be to exploit variation in exposure by industry of employment, such as in Acemoglu et al. (2016). Using geographic variation offers several advantages for our context. First, we can measure whether job losses led to re-employment or changes in labor force participation, the latter of which saw very important changes for black men in this era. Second, the industry-based approach would not allow us to measure indirect effects that could be particularly important for understanding racial differences. For example, whites who experience job losses in a trade-affected manufacturing industry may displace black workers in an industry which received little exposure (due to prejudice or otherwise). Finally, industry of employment variation requires information on industrial employment counts by race, while our strategy requires only (non race-specific) geographic employment counts and race-specific population counts. The latter is readily found in publicly available data sets, while the former is available only for a highly aggregated set of industries.[20]

[19]This also partially addresses concerns recently raised by Jaeger et al. (2018) that shift-share instruments can conflate the short- and long-run effects of economic shocks when these shocks are ongoing and correlated across time.

[20]Our import exposure is based on 380 different SIC87 manufacturing industries as reported by the CBP. The only publicly available race-specific employment data come from the CPS or the Census, which have just 59 manufacturing industries in 1960. Census/CPS data also include a non-trivial number of workers who are simply classified as working in manufacturing without a specific industry to which we can map trade. While we do find evidence of reduced employment when implementing the Acemoglu et al. (2016) approach on these limited data, the estimates are much too imprecise to identify whether these effects differed by race. These results are available upon request.

2.3.2 Instrumental Variable and Identification

Because Japanese imports may be driven by domestic changes in American industries, we adopt the strategy implemented by Autor et al. (2013) for China, and instrument with the observed change in Japanese import penetration in other highly developed economies.[21] Specifically, our instrument is defined as

$$\Delta IPW_{oi,1990-1970} = \sum_j \frac{L_{ij1960}}{L_{uj,1960}} \frac{\Delta M_{oj,1990-1970}}{L_{i1960}}, \tag{2.3}$$

where the subscript o indicates the sum across these other countries.[22] In words, our instrument is the change in import exposure faced by the average worker that would have been predicted from (1) the commuting zone's industrial composition in 1960 (i.e., before Japanese import competition began), and (2) the ability of Japan to penetrate these industries in other countries during our time period. The variation in the exposure each CZ receives can be further subdivided into two avenues: the manufacturing share of the local economy and the composition of products they manufacture. Our preferred specification will control for initial manufacturing share, and thus isolate the latter variation.

Our instrument is a "shift-share" instrument that combines local industry employment shares and national industry-level "shifts" (trade shocks). There has been a recent debate on the sources of identification for such instruments (e.g., Goldmsith-Pinkham et al., 2018). Borusyak et al. (2018) develop a quasi-experimental framework that views the trade shocks (i.e., the industry-level exports from Japan to other countries) as "as-if" randomly assigned across industries. They then prove shift-share estimators are consistent under two conditions: (1) industry trade shocks are orthogonal to the unobservable factors in the CZs in which they are located, and (2) shocks across industries are sufficiently independent.[23]

[21] Hummels et al. (2014) use a similar instrument to predict the offshoring and export behavior of Danish firms.

[22] We use a similar set of countries as Autor et al. (2013): Australia, Denmark, Finland, New Zealand, Spain, and Switzerland. Unlike them, we exclude Germany because of complications arising with reunification, and Japan for obvious reasons.

[23] Goldmsith-Pinkham et al. (2018) argue instead that shift-share estimators are consistent only if the industry employment shares in each CZ are orthogonal to the CZ-level unobservables. The key difference between their result and Borusyak et al. (2018) is that the latter relies on large sample asymptotics for both CZs and industries, while the former assumes only a large sample of CZs. In our setting, we have 358 CZs and 380 industries, with little correlation in shocks across industries outside of 135 3-digit industry codes. See Appendix B.2.7.

Our instrument will satisfy the orthogonality condition provided that exports that are common across countries are driven by changes within Japan (e.g. productivity shocks) rather than forces in the United States. Specifically, we assume that any demand increases or negative productivity shocks for U.S. industries are uncorrelated with similar changes in our IV countries, and that changes in other countries' Japanese imports are not driven by negative productivity shocks to U.S. exporting industries. These assumptions are similar to those outlined in Autor et al. (2013).

The most obvious concerns for these assumptions center around the computer and automobile industries. Advances in computer technology during this era may represent a worldwide positive demand shock. In general, this should bias us away from finding negative impacts of trade, because U.S. firms also experienced this shock, but how this would bias the effect on racial disparities is unclear. Likewise any bias caused by the automobile industry, which faced the largest increase in import competition in absolute terms, is also uncertain. While much of this growth was due to improvements in Japanese manufacturing technology, the 1970s oil shocks caused a worldwide shift in demand from large cars (a specialty of American firms) to the smaller, more fuel efficient cars already preferred in Japan (Crandall, 1984; Ohta and Griliches, 1986).[24] Note that none of the countries we use in constructing our instrument were major importers of U.S. automobiles in our time period, which should minimize the impact of any global drop in demand for these products on our estimates.[25]

In Appendix B.2.6 we perform a numerically equivalent transformation of our main specification developed by Borusyak et al. (2018) that isolates variation caused by each industry, and show that our results are robust to excluding automobiles and computers. We also perform a series of additional robustness checks recommended by Borusyak et al. (2018). These include excluding industries with outlying instrument exposure, including 1-digit and 2-digit CZ-level industrial classification employment shares (thus allowing that the expected trade shock from Japan may have varied at these levels), and using the individual highly developed countries' imports as separate instruments. Our results are robust to all of these exercises, and only the 2-digit industry employment shares meaningfully reduce the magnitudes of our estimates.

[24]Further complicating this industry is the Voluntary Export Restraint (VER) Japan implemented under U.S. pressure which led to a strategic shift by Japanese manufacturers to higher quality automobile exports in the 1980s (Feenstra, 1984, 1988).

[25]In 1970, Switzerland, the largest importer of U.S. automobiles in our set of other developed countries, accounted for just over 1% of American automobile exports. The largest customer, Canada, accounted for almost 75%.

Further, we fail to reject the overidentifying restrictions in the multiple instruments case.

We also provide a test for the second condition (dispersion) in Appendix B.2.7. Following Borusyak et al. (2018), we exploit the hierarchical design of the SIC system, and estimate intraclass correlation coefficients for clusters of similar industries. We find little correlation in the trade shock within two- and one-digit industry clusters, consistent with a high level of independence in the distribution of shocks.

We show the time variation in imports from Japan for the United States and our six other developed countries in Table 2.2. From 1970 to 1990, U.S. imports from Japan rose by \$94.5 billion (in 1999\$), a 374% increase. In the same time period, the six other countries saw an even larger increase in percentage terms of 389%. The United States also saw an increase in exports to Japan, but not nearly at the same rate, resulting in a trade deficit of \$57.8 billion by 1990. We also see in column (3) that this period was one of a general increase in globalization. But, the pace of import increases from Japan outstripped that from the rest of the world, both in the United States and the other developed countries we study.

In Table 2.3 we estimate our first-stage regression. Unsurprisingly our instrument is very strong, with an F-statistic over 80.

2.4 Results

2.4.1 The Impact of Japanese Trade on Employment Disparities

In Table 2.4 we perform the 2SLS estimation of equation (2.2) on the manufacturing employment share of the male working age population.[26] All percentage variables are scaled in percentage points. Column (1) is the standard regression in the literature that does not allow for racially heterogeneous effects. Without any additional controls, we find a large, negative and statistically significant effect of Japanese imports on commuting zone manufacturing share, with a point estimate nearly double that found by Autor et al. (2013). However, once accounting for the pre-existing size of the manufacturing sector in column (2), and characteristics of the CZ's workforce in column (3) any potential effects are reduced to 0. Instead, we find strong evidence for

[26] Standard errors are clustered at the state level. Borusyak et al. (2018) derive an alternative set of standard errors which are asymptotically equivalent to those derived by Adão et al. (2018) and allow for correlations within similar industries across CZs. We find in practice that these standard errors are smaller than the state-clustered errors. See Appendix B.2.8.

a secular decline in manufacturing. Also of importance appear to be the pre-existing stock of college educated workers and the occupational mix of the local manufacturing sector.[27] This is consistent with evidence presented in, for example, Autor et al. (2003) and Autor et al. (2008), that skill-biased technical change was the dominant driver of changes in the wage structure during this time period.

This specification, however, masks substantial heterogeneity by race. In columns (4) and (5) we estimate equation (2.2) separately for whites and blacks, respectively. The results are striking. While a \$1,000 increase in Japanese import competition led to a .59 percentage point decline in black manufacturing share, it also led to a .19 percentage point *increase* in white manufacturing share. In other words, columns (3)-(5) suggest that, rather than eliminating manufacturing jobs, Japanese competition led to a shifting of employment from blacks to whites.

For our remaining results, we use the full set of CZ-level controls and estimate the fully-interacted version of (2.2), reporting $\beta_b - \beta_w$ as the interaction between $\Delta IPW_{ui,1990-1970}$ and a black indicator.[28] Column (1) of Table 2.5 repeats the estimates from columns (4) and (5) of Table 2.4 using this approach. Columns (2)-(4) provide the same estimation for non-manufacturing employment share, the unemployment rate, and non-labor force participation rate, respectively.[29] For whites, we see evidence of a movement of workers from the non-manufacturing sector and from out of the labor force into manufacturing, although neither of these effects is statistically significant. In contrast, we find at best weak evidence that black workers who moved out of manufacturing found re-employment in non-manufacturing. We also see no increase in black unemployment. Instead the vast majority of displaced black manufacturing workers drop out of the labor force. We estimate a \$1,000 increase in Japanese import competition led to a .45 percentage point increase in the black labor force non-participation rate, or a .54 percentage point widening of the racial labor force non-participation gap.

We provide a series of robustness checks of our main results on manufacturing employment in Table 2.6. In column (1) we replace our CZ-level controls with race-

[27] Given that many of the information technology advancements that enabled certain types of jobs to be offshored were yet to occur, the offshorability index is best viewed as an additional measure, beyond routine-intensity, of manufacturing task composition. See Appendix B.1.6 for details on the construction of these indices.

[28] We now add our instrument interacted with a black indicator to the first stage.

[29] While by definition all individuals must at any given time be either employed in manufacturing, employed in non-manufacturing, unemployed, or out of the labor force, our coefficients do not add up exactly to 0 because of winsorization.

specific CZ workforce characteristics. That is, we define the college educated percentage of the population as the fraction of white working age males with a college education, and similarly for blacks.[30] While this provides the benefit of, for example, better measuring how one racial subgroup may have been more vulnerable to skill-biased technical change, it provides a drawback in that it implicitly assumes that when these factors impact the white labor force there are no spillovers onto the black. Making this change has a negligible impact on our point estimates. In column (2), we include 10 1-digit manufacturing sector employment shares and their interactions with race, following the classification system in Autor et al. (2014). This allows us to better account for secular trends within manufacturing that may be correlated across similar industries. However it also reduces some good variation given that similar industries tend to co-locate, and measurement error in our mapping from products to industries will likely misclassify trade within these large sectors.[31] These controls slightly reduce the point estimate of our interaction term, but the magnitude remains large and statistically significant.

In column (3) we re-estimate column (1) of Table 2.5 using OLS. Similar to Autor et al. (2013), we find that OLS biases our estimates of trade on manufacturing employment upwards for both black and white workers. In column (4), we estimate the OLS specification measuring exposure as net imports rather than imports, and our results are essentially unchanged. In column (5), we adopt a 2SLS strategy for net imports. Following Autor et al. (2013) our first stage in this specification includes an analogous second instrument reflecting the change in exports to Japan from the same set of high-income countries. An important caveat is that our exports instrument is not statistically significant in the first stage once controlling for our main imports instrument. Nonetheless, our results are virtually unchanged from column (1) of Table 2.5. In column (6), we again follow Autor et al. (2013) and construct a measure of imports that isolates final goods from intermediates, exploiting 1972 input-ouput data from the Bureau of Economic Analysis (BEA; see Appendix B.1.7 for more details). If anything here we find a stronger effect on disparities.

[30]Note that we are unable to compute a race-specific $\Delta IPW_{ui,1990-1970}$ as we lack a sufficient sample of black manufacturing workers in 1960, or any race-specific employment data in the CBP. However, the argument against using a race-specific measure in this case is particularly strong. It rules out, for instance, that white workers who are displaced by trade do not in turn displace black workers in unaffected industries.

[31]This is especially a concern given the imputation method used in constructing the 1960 CZ-level industry distribution.

2.4.2 Understanding the Mechanisms of the Disparate Impact

Skill Upgrading in Manufacturing

The previous subsection established that Japanese trade caused an influx of white workers into manufacturing, replacing black workers who dropped out of the labor force. We now analyze the mechanisms that caused this disparate impact.

We first explore heterogeneous effects by skill group in Table 2.7. We divide our sample by race and education: high school dropouts, high school graduates, and college educated.[32] Due to the small number of college educated black workers, particularly in 1970, we are unable to explore effects for this subgroup. We then estimate equation (2.2) separately for each group.

First, in Panel A we find high school dropouts moved out of manufacturing and into non-manufacturing employment. However, underlying this is substantial heterogeneity. Black high school dropouts saw a large decrease in manufacturing employment, roughly half of which manifests itself in higher non-labor force participation. The drop in manufacturing employment for white non-manufacturing workers is statistically insignificant. Instead, they see large gains in non-manufacturing employment fueled by higher labor force participation.

We see little effect of Japanese import competition on the labor market outcomes of high school graduates in Panel B. Black workers saw decreases in labor force participation, but it is unclear to what extent this was due to lower manufacturing employment or a shifting of the unemployed out of the labor force, both of which have non-trivial but imprecisely estimated effects. White high school graduates saw a small, though statistically significant increase in unemployment, but no substantial effects on any other outcome. In Panel C, we see that all of the gains in manufacturing employment came from college educated workers, particularly among whites. This was fueled by a corresponding drop in non-manufacturing employment.

The results in Table 2.7 are strongly suggestive of blacks being disparately affected by a change in the demand for skill within manufacturing. While we cannot directly rule out all other factors for the racial differences in outcomes among high school dropouts, we note the substantial differences in skills within the same quantity of education because of historical differences in school quality.[33] We find only mild evi-

[32] High school graduates have exactly 12 years of education, while we define college educated as those with more than 12.

[33] For example, in the National Longitudinal Survey of Youth 1979 cohort, which due to their later

dence for negative effects on higher skill blacks, and the positive employment effects accrued entirely to our highest measurable skill group, college educated whites.

While quality of education data is not itself readily available in the Census, the "Great Migration" presents the opportunity to look for heterogeneous effects within black workers who plausibly differed in schooling quality. From 1940 to 1970, 4 million blacks moved out of the rural South. Due to both differences in resources and as a consequence of segregation, we would expect these workers to have lower quality formal education than their Northern born counterparts.[34] To explore how the effect of Japanese trade differed among blacks born in and outside of the South, we restrict attention to CZs which had a substantial population of Southern and non-Southern born blacks.[35] We then calculate employment outcomes in each CZ separately for these groups, and estimate an analogous set of regressions to Table 2.5.

There are some important caveats to this exercise. First, the skill levels of Southern born blacks in 1990 will look much different than in 1970. Even before desegregation in the 1960s, black Southern schools were seeing improvements on many measurable dimensions (Card and Krueger, 1992). Further, given that the Migration ended by 1970, many Southern born blacks in the North will be children of migrants that were educated primarily in higher quality Northern schools. Our census division fixed effects (and their interactions with the Southern born indicator) should alleviate some of these concerns. It is not obvious why import exposure (or our instrument) would be correlated with changes in the relative skill-level of Southern born blacks beyond these regional differences, though it cannot be ruled out.

We display the results of this exercise in Table 2.8. Remarkably, within this sample of CZs we see only mild evidence for negative effects for black workers born outside of the South, concentrated on labor force participation. In contrast, we see strong evidence for negative impacts on Southern born black employment outcomes.

birth year would be a cohort with a lower school quality gap than the majority of the black working age population in both 1970 and 1990, white high school drop outs actually scored in higher percentiles of the Armed Forces Qualifying Test than black high school *graduates*. See also Lang and Manove (2011), who show that blacks are incentivized to receive higher levels of formal education relative to their skill level in the presence of statistical discrimination.

[34]For example, in 1940 Southern blacks attended schools which with 25% higher pupil-teacher ratios and 10 percent shorter terms than Southern whites (Card and Krueger, 1992). Northern black newspapers expressed concern that these new Southern migrants would hurt the reputation of the local black workforce (Grossman, 1991). See Boustan (2009) for more discussion of skill differences between Great Migrants and Northern-born blacks.

[35]We make a similar restriction to that in our main results, requiring at least 500 working age Southern and non-Southern born black males in both 1970 and 1990. This restriction leaves us with a sample of 185 CZs.

In Table 2.9 we look for direct evidence of skill upgrading in manufacturing. The first two columns look at the education composition of manufacturing jobs. We find a \$1,000 increase in Japanese import competition led to .89 percentage point increase in the share of manufacturing jobs held by college educated workers. We also estimate a decrease in the share held by high school dropouts, though this is not statistically significant. The final four columns look at the occupation composition of manufacturing. We find that a \$1,000 import increase led to a .14 percentage point increase in the share of manufacturing employment held by managers and professionals (column 3), all of which is accounted for by an increase in engineers (a subcategory of professionals, column 4). In contrast, we see no change in the share of employment to production workers (column 5). However, in column (6) we see an increase in the skill level of production workers. The share of manufacturing employment belonging to college educated production workers rose by .65 percentage points for every \$1,000 increase in Japanese imports.

Japanese Trade and the Geography of Employment

In Table 2.10 we explore the impact of trade on the distribution of workers and jobs across geographies using Census population tabulations from the NHGIS. In column (1) we first look at changes in CZ population in response to Japanese import competition. The empirical treatment of CZs as separate labor markets relies in part on the idea that workers are slow to migrate across CZs in response to changes in economic conditions. Consistent with several previous studies (e.g., Bound and Holzer, 2000; Autor and Dorn, 2013), and despite the long time horizon we look at, we do not find any evidence of aggregate out-migration from CZs in response to Japanese trade. However, when we instead look at the share of the commuting zone population that is black in column (2), we find that imports exposure caused CZs to become *blacker*.[36] In other words, despite blacks bearing the negative economic effects of trade, CZs which faced a high degree of import competition experienced increases in the black population, offsetting any out-migration from whites or non-black minority groups. While surprising on its face, this is consistent with work by Glaeser and Gyourko (2005) that shows that weak labor demand causes increases in the population of low-skill workers who are attracted by the now lower prices of housing.

[36]Consistent with this, we also find a positive effect of Japanese import competition in a regression on log black population. These results are available upon request.

In columns (3) and (4) we instead look at the distribution of workers within a CZ between the central city and the suburbs.[37] While we find little evidence that central city populations declined, we find strong evidence that the black share of their population increased; a $1,000 increase in Japanese import competition increased the black share of the central city population by 1.6 percentage points. Thus, white residents left the inner cities in response to trade, and were replaced by an inflow of new black residents, which is again consistent with the work of Glaeser and Gyourko (2005).

A popular explanation for black-white employment differences is the "spatial-mismatch" hypothesis originally advanced by Kain (1968). That is, jobs are located in areas where black workers do not live and are difficult for them to reach. The previous set of results are suggestive of this mechanism if jobs followed white workers to the suburbs. In column (5) of Table 2.10 we find little evidence that manufacturing jobs shifted from residents of central cities to residents of suburbs. In column (6) we use data from the 1987 and 1972 CoM to look at the location of jobs themselves, and likewise find no evidence they moved out of central cities.[38]

Other Explanations: Prejudice and Unions

Unionization in the United States remained relatively high in 1970, particularly for manufacturing workers. Another hypothesis is that whites were insulated from this trade shock due to better union protections. This seems unlikely given that blacks actually had higher unionization rates than whites throughout the 70s and 80s (Farber et al., 2018). Further, testing this hypothesis is difficult, given that unionization data for this time period is notoriously poor. For example, the CPS does not begin tracking unionization rates until 1973, and even these are only available at the level of often arbitrary state groupings. Nonetheless, we followed recent work by Farber et al. (2018) and constructed state-level estimates of unionization rates for the 1967-1972 time period using data from Gallup surveys. As unionization rates are largely driven by industrial composition, we took the residual of this variable from a regression on state-

[37] We remind the reader that we have a reduced sample size here as only 167 CZs have a Census-defined central city in 1970. All of our main results from Section 2.4.1 are robust to using just these 167 CZs (results available upon request).

[38] As discussed before, the CoM is not conducted simultaneously with the Census, but instead at a different set of five year intervals. Consistent with Table 2.4, we find no effect on CZ-level manufacturing employment from trade in the 1972-1987 period. However the CoM does not track employment by race, so we cannot replicate our main results using these data.

level manufacturing share, and then matched it to the state of the largest city in each CZ. Column (1) of Table 2.11 includes this variable along with its interaction with CZ-level import competition in our main manufacturing specification, while column (2) adds interactions with race.[39] While our results suggest that unionization may have shielded manufacturing jobs from import competition, we find no evidence that whites received greater protections.

Another alternative explanation for our findings is that, when forced to lay off workers due to increased Japanese competition, managers chose to only lay off blacks due to racial prejudice. If this were the case, it is not clear how such managers were then able to gather resources to hire high-skill whites. That notwithstanding, we tested this hypothesis using county-level voting data from the Atlas of U.S. Presidential Elections for the 1968 presidential election, which included George Wallace, a serious pro-segregation third party candidate.[40] In column (3) of Table 2.11 we include an indicator for whether the CZ was at or above the national median in Wallace vote share, along with its interactions with race and import penetration.[41] If anything, blacks appear to have seen less negative effects from trade in highly prejudiced areas. Column (4) instead includes an indicator for whether the CZ was at or above the median of its census division, with similar results.

2.4.3 The Impact of Japanese Trade on Earnings

In the previous sections we established that Japanese trade led to a displacement of the relatively low-skill black population from manufacturing and a replacement with the relatively higher-skill white population. We now explore how these structural changes influenced black financial outcomes.

In the first three columns of Table 2.12, we estimate the impact of Japanese competition on median male wages and earnings.[42] These results must be taken with

[39]We see a small reduction in sample size here as not all states were surveyed by Gallup.

[40]Wallace received 13.5% of the popular vote and won five states. Since 1948, Wallace is the only third party candidate to have won a state, and only H. Ross Perot in 1992 received a higher vote share.

[41]We exclude CZs in Louisiana, as parish-level voting data is not available.

[42]We prefer working with medians for several reasons. First, earnings data is topcoded, and the topcode varies across censuses. Second, medians will be less sensitive to outliers, which is relevant particularly for smaller CZs that contain few black workers. Finally, we cannot calculate a mean log earnings inclusive of non-earnings as in column (3) of Table 2.12 as the log of 0 is undefined. In Appendix B.2.9, we report results using means and generally find results which are less negative but consistent with those reported here.

caution because of the effect of Japanese import competition on the composition of employment. While we saw in Table 2.7 trade caused a movement of high-skill whites into manufacturing, these workers were primarily drawn from employment in non-manufacturing. The strongest *net* employment effects were an increase in the labor force participation of the lowest skill whites and a decrease in the labor force participation of the lowest skill blacks. It is therefore not surprising that we see no disparate impact on weekly wages of those with earnings in column (1). In column (2), we find a negative, but statistically insignificant effect on the black-white annual earnings gap for these workers.

However, as is well-known in the literature (e.g., Butler and Heckman, 1977; Brown, 1984; Chandra, 2003) and can be seen in Table 2.5, estimates of changes in the earnings gap in this time period can understate the true magnitude of the changes in relative black financial circumstances due to the large decrease in labor force participation by black men. In column (3), we estimate the impact on median male earnings inclusive of individuals who report zero income.[43] Once we allow for non-earners we see a large and statistically significant negative impact of trade on black workers, with little impact on whites. Our estimate suggests a $1,000 increase in Japanese import competition led to a 3.6 log point increase in the black-white median earnings gap.

The final three columns of Table 2.12 look at household finances.[44] The impact of reduced black male employment may have been partially offset if other household members, including black women, found employment opportunities in response. The ability to do so is hampered by the fact that black women's labor force participation has historically been higher than whites' (Neal, 2004). We see this was not the case for earnings in column (4). While the economic standing of white families did not change in response to import competition, the median black-white family earnings gap rose by 2.6 log points for every $1,000 of exposure. We see a smaller effect when we look at household income rather than earnings in column (5), which appears to be due to a relative increase in welfare recipiency (column 6).

Interestingly, we also find increased welfare recipiency for white families. Because

[43]While the log of 0 is undefined, this does not cause problems for calculating median earnings as we simply assume these earnings are below median. After performing the winsorization, there is no commuting zone in which the median worker of any race reported 0 earnings.

[44]We calculate the sum of all income of individuals in the household ages 16-64 and divide by the total number of 16-64 year olds in the household. The race of the household is determined by the race of the household head.

welfare is an outcome that specifically measures the economic health of those near the bottom of the income distribution, this further suggests the importance of the demand for skill. While white high school dropouts do not appear to be affected on the aggregate, perhaps due to historical differences in school quality, within the left tail of this group should be a set of workers more comparable in skill to black high school dropouts. Those workers should reasonably have felt similar impacts of trade as low-skill blacks. The evidence of increased welfare recipiency supports this story.[45]

2.4.4 Quantifying the Impact of Japanese Trade

Our previous results have shown that Japanese import competition exacerbated racial differences in employment, earnings, and the financial standing of households. In Table 2.13 we perform two back of the envelope calculations to quantify these impacts. First in column (1) we calculate the national change in the racial gap from 1970-1990 across several economic variables using the full IPUMS samples of the respective Decennial Censuses. Note that this includes geographies excluded from our regression analysis due to the small number of black workers living in these communities. We show descriptive statistics for this sample in Appendix B.2.10. In general, we see that blacks appear to perform slightly better in the national sample relative to what we see in our regression sample in Table 3.4, though the trends are similar.[46]

The average black worker lived in a commuting zone which was exposed to \$1,413 (in 1999 dollars) worth of new Japanese imports, while the average white worker was exposed to \$1,601. In column (2) we use these values, as well as our estimates from Tables 2.5 and 2.12 to estimate the change in national disparities were Japanese imports to have remained at 1970 levels for both white and black workers. This will overstate the explanatory power of trade if part of the import increase was due to domestic demand increases, and demand-induced imports have a smaller impact on racial disparities than imports induced by exogenous factors. In column (3), we

[45]As an additional test of this hypothesis, we calculated what percentile in each CZ the median black earner would have been in the 1970 white distribution. Similar to Bayer and Charles (2018), we then compared the change in black median male earnings to the earnings of this percentile white in response to Japanese trade. When we considered only those with positive earnings, we found a strong negative impact on whites near the median of the black distribution, despite no negative impact on median black earnings. However, once including those without earnings, we found a large negative impact on median blacks and no evidence of negative effects on comparable whites, similar to those reported in Table 2.12. These results are available upon request.

[46]This is possibly due to the CZs outside of our regression sample being exposed to less Japanese trade, which we have shown negatively impacted black workers.

follow Autor et al. (2013) and Acemoglu et al. (2016) and obtain a more conservative estimate using just the exogenous increase in imports determined by our instrument. We first multiply the realized per worker import increases by the partial R^2 from the first stage regression (.764), and then compute our counterfactuals as before using these values.

There are some important caveats to this exercise. Because our identification is entirely from cross-commuting zone exposure, these estimates are best viewed as accounting for only the direct effects of foreign competition. They will not take into account, for example, a common national effect on racial disparities caused by access to lower prices, higher quality, or increased variety of consumer goods. They will also not take into account changes caused by movements of capital out of highly affected CZs and redistributed in a way orthogonal to trade exposure.[47] Nonetheless, we believe these estimates are informative on the importance of trade in explaining changes in disparities.

The impact is substantial. We can explain 66-86% of the relative decrease in black manufacturing employment, and 17-23% of the relative rise in black non-labor force participation due to Japanese import competition. In the absence of this trade influx, the median earnings gap for working age males would have seen a 34-44% slower divergence, while household earnings would have converged at 1.7-2.0 times the rate. Welfare recipiency would have grown by 37-48% less.

2.5 Conclusion

Much of the popular press has focused on the effects of Chinese import competition on white working class communities. But many of the identified impacts, including declines in manufacturing employment, labor force participation, and earnings, are reminiscent of the economic hardships experienced by black communities in the 1970s and 1980s. Using modern methods, we find strong evidence that import competition in this time period from Japan played a sizable role in these hardships. Every $1,000 increase in imports exposure per worker resulted in a decrease in black manufacturing

[47]For example, these calculations may overstate the overall negative impact of Japanese trade if it caused a relocation of manufacturing firms from the North to the South, and these new Southern factories employed blacks at high rates. While our primary identification strategy is not able to account for such changes, in unreported results we found, using cross-industry differences in Japanese import exposure, that trade competition had no impact on the propensity of industries to increase Southern employment.

employment by .59 percentage points, a rise in labor force non-participation of .46 points, and a decline in median household earnings by 2.8 log points.

However, we do not see evidence for aggregate losses for the American manufacturing sector. Instead we find a shifting of employment, particularly from low-skill blacks to high-skill whites. Thus the net effect of this period of globalization was a redistribution of welfare from a disadvantaged community to an advantaged one. Our results suggest that the costs of foreign competition in the 1970s-1980s were obscured by disproportionately loading onto black Americans. They also provide a wealth of evidence that increased import exposure was instrumental in the stalling of black economic progress during this time period, mirroring the effects widely acknowledged for white working class communities in the 2000s.

To the extent that these disparities were caused by changes in the demand for low-skill manufacturing workers, one natural question is to ask whether this reversal was inevitable. The subsequent national declines in American manufacturing have been accompanied by changes in technology which have made the remaining sector more high-skill (Charles et al., 2018). However, the timing of the Japanese trade shock may have made it particularly damaging. Black workers had only recently made advances in manufacturing. The inability to sustain this success may have played a role in the failure to close gaps for longer-term progress markers, such as the home ownership (Collins and Margo, 2001). It likewise conceivable that this economic disruption reduced the ability of parents to invest in human capital for the next generation. Indeed, progress on education and test gaps would begin to stall and reverse at the end of the 1980s, concurrent with a rise in youth incarceration and drug violence (Neal, 2006; Evans et al., 2016). Had blacks continued to make economic progress during these decades, they may have been less vulnerable to skilled biased technical change and mechanization during the 1990s and 2000s.

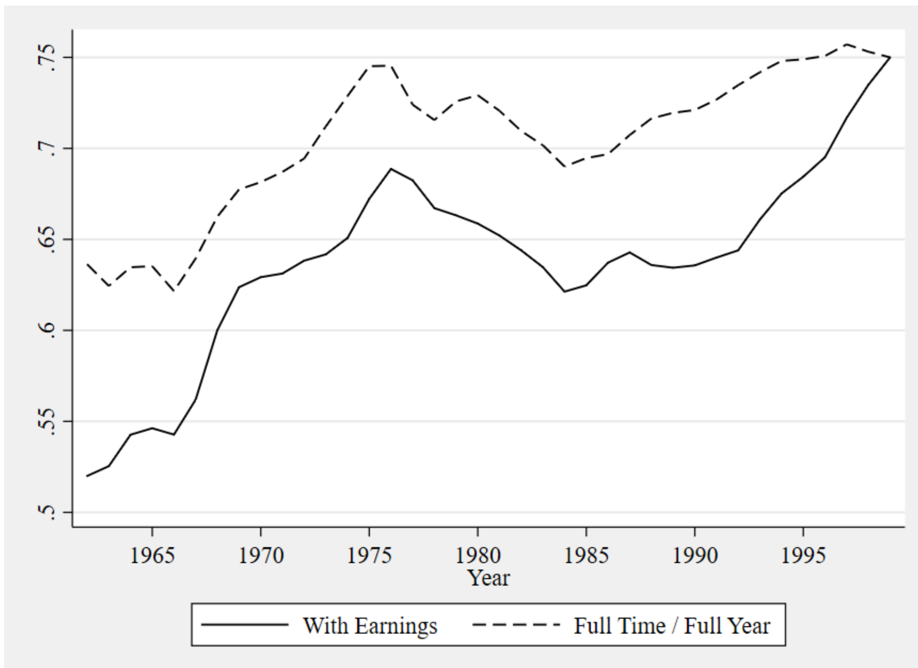

Figure 2.1.: Ratio of Median Earnings for Working Age Population: Black Men/White Men, 1962-1999

Notes -

Source - Current Population Survey (1962-1999).

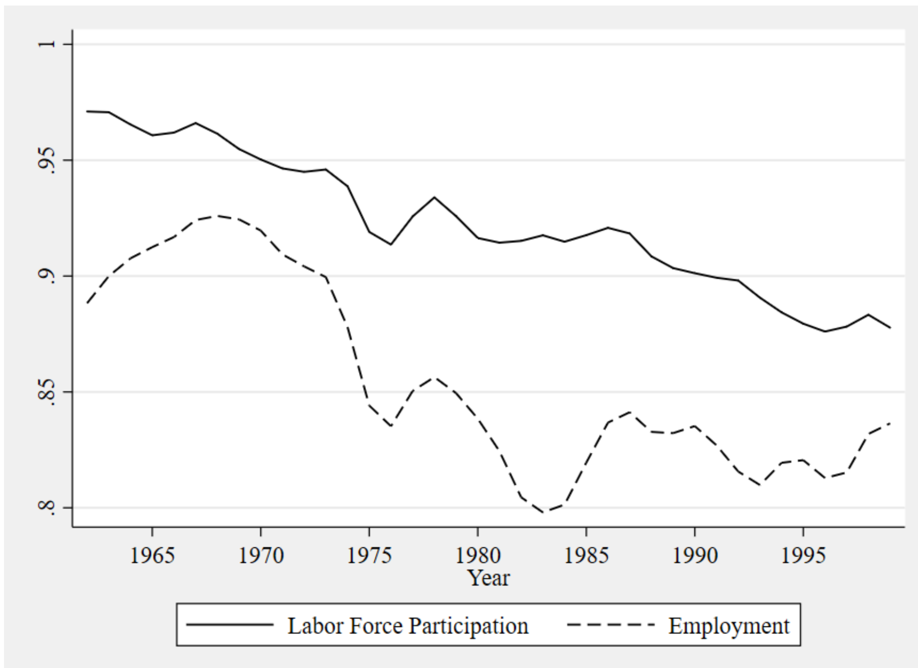

Figure 2.2.: Ratio of Employment Rates for Working Age Population: Black Men/White Men, 1962-1999

Notes - Yearly scatterplot data smoothed using LOWESS with bandwidth=0.15

Source - Current Population Survey (1962-1999).

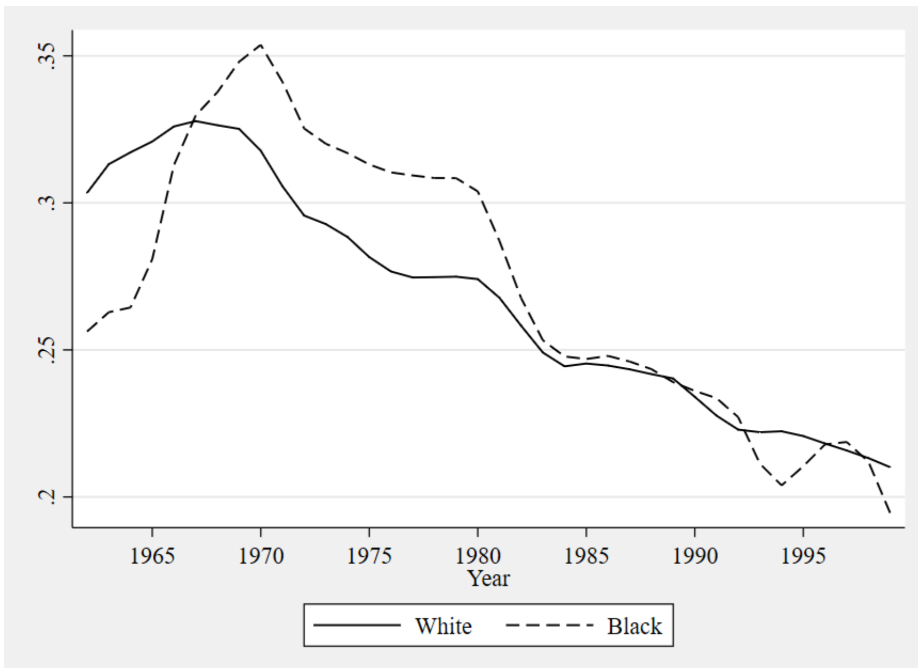

Figure 2.3.: Fraction of Employment in Manufacturing: Working Age Men, 1962-1999

Notes - Yearly scatterplot data smoothed using LOWESS with bandwidth=0.15

Source - Current Population Survey (1962-1999).

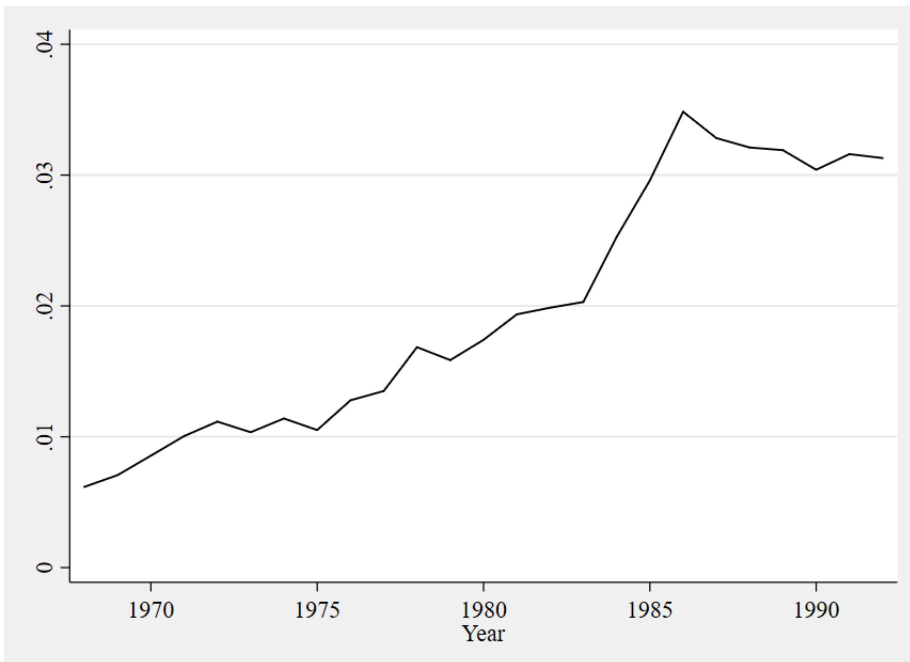

Figure 2.4.: U.S. Import Penetration Ratio in Manufactured Goods for Japan, 1968-1992

Source - Authors' calculations using trade data from UN Comtrade and domestic output data from the BEA.

Table 2.1.: Descriptive Statistics: Regression Sample

	Black		White		ΔGap
	1970	1990	1970	1990	
	(1)	(2)	(3)	(4)	(5)
Percentage of population employed in manufacturing	19.432 (8.68)	13.372 (6.27)	23.620 (8.73)	18.147 (6.29)	-0.587
Percentage of population employed in non-manufacturing	51.588 (10.18)	48.695 (9.69)	59.078 (8.34)	62.997 (7.20)	-6.811
Unemployed share of population	4.380 (1.85)	9.808 (2.83)	2.748 (0.99)	4.311 (1.14)	3.866
Labor force non-participation rate	24.531 (5.16)	28.084 (6.27)	14.530 (2.99)	14.447 (3.02)	3.637
Median log weekly wage, male earners	613.088 (27.25)	613.290 (18.38)	652.737 (14.86)	650.471 (16.06)	2.467
Median log annual earned income, male earners	996.847 (29.87)	995.625 (20.49)	1043.513 (16.15)	1040.489 (17.19)	1.803
Median log annual earned income, all working-age males	977.015 (37.81)	948.890 (34.56)	1035.449 (18.09)	1030.179 (21.93)	-22.855
Median log HH earned income	933.542 (37.27)	947.747 (30.98)	989.553 (16.15)	1003.265 (21.18)	0.493 (0.00)
Median log HH total income	939.089 (35.47)	958.054 (26.96)	990.343 (15.93)	1010.350 (19.35)	-1.041 (0.00)
HH welfare rate	14.142 (4.55)	18.064 (5.55)	2.842 (1.23)	4.374 (1.78)	2.389 (0.00)
Observations	358	358	358	358	

Notes - Standard deviations in parentheses. Percentage and rate variables are scaled in percentage points, while earnings and income variables are scaled in log points.

Source - 1970 form 1 and 2 1% metro and 1990 5% IPUMS samples of the United States Decennial Census.

Table 2.2.: Value of Trade with Japan for the U.S. and Other Selected High-Income Countries and Value of Imports from all Other Source Countries, 1970-1990

	Imports from Japan	Exports to Japan	Imports from rest of world
	(1)	(2)	(3)
Panel A: United States			
1970	25.2	19.8	146.3
1990	119.7	61.9	538.6
Growth 1970-1990	374%	213%	268%
Panel B: Six other developed countries			
1970	4.8	6.8	97.9
1990	23.7	20.1	311.8
Growth 1970-1990	389%	194%	219%

Notes - Values are in billions of 1999 U.S. Dollars.

Source - UN Comtrade

Table 2.3.: Japanese imports to the U.S. and to Other Countries: First Stage Estimates

	(ΔImports to US)/worker	
	(1)	(2)
(Δ Imports from Japan to OTH)/worker	4.694***	4.949***
	(0.395)	(0.537)
Controls	No	Yes
Observations	358	358
R^2	0.764	0.786

Notes - Robust standard errors clustered by state in parentheses. Models are weighted by 1970 population. Regression in column (2) includes census division fixed effects and commuting zone-level controls for black percentage of population, foreign-born percentage of population, percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, and percentage of employment in routine occupations in 1960; $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 2.4.: Japanese Imports on Change in Manufacturing Employment/ Working Age Population in CZs, 1990-1970 Long Difference: 2SLS Estimates

	All			White	Black
	(1)	(2)	(3)	(4)	(5)
(Δ Imports from Japan to US)/worker	-1.264***	-0.096	0.034	0.193*	-0.592***
	(0.427)	(0.199)	(0.111)	(0.117)	(0.137)
Percentage of employment in manufacturing$_{1960}$		-0.295***	-0.222***	-0.238***	-0.187**
		(0.032)	(0.044)	(0.046)	(0.078)
Black percentage of population$_{1960}$			-0.053**	-0.050**	-0.054
			(0.025)	(0.025)	(0.039)
College percentage of population$_{1960}$			-0.204***	-0.220***	-0.272***
			(0.067)	(0.067)	(0.093)
Foreign-born percentage of population$_{1960}$			0.003	0.017	0.039
			(0.048)	(0.045)	(0.081)
Average offshorability index of occupations$_{1960}$			0.079**	0.062	0.194***
			(0.035)	(0.038)	(0.053)
Percentage of employment in routine occupations$_{1960}$			-0.108	-0.087	-0.163*
			(0.067)	(0.072)	(0.098)
Census Division FE	No	No	Yes	Yes	Yes
Observations	358	358	358	358	358

Notes - Robust standard errors clustered at the state-level in parentheses. Models are weighted by 1970 population. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 2.5.: Japanese Imports and Change in Racial Employment Status Gap, 1990-1970 Long Difference: 2SLS Estimates

	Mfg emp	Non-mfg emp	Unemp	NILF
	(1)	(2)	(3)	(4)
(Δ Imports from Japan to US)/worker	0.193*	-0.097	-0.010	-0.086
	(0.117)	(0.096)	(0.034)	(0.054)
(Δ Imports from Japan to US)/worker $\times$ Black	-0.785***	0.228	-0.071	0.542***
	(0.173)	(0.217)	(0.127)	(0.121)
Observations	716	716	716	716

Notes - Robust standard errors clustered at the state-level in parentheses. Models are weighted by race-specific 1970 population. Each regression includes census division fixed effects; commuting zone-level controls for percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960; a black indicator; and interactions of the black indicator with all of these variables. $p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 2.6.: Japanese Imports on Change in Manufacturing Employment/ Working Age Population in CZs, 1990-1970 Long Difference: Robustness Exercises

	Race-specific	1-dig shares	Gross imports	Net imports		Final goods
	2SLS	2SLS	OLS	OLS	2SLS	2SLS
	(1)	(2)	(3)	(4)	(5)	(6)
(ΔImports from Japan to US)/worker	0.187	0.268	0.211**	0.166*	0.213*	0.227
	(0.157)	(0.189)	(0.122)	(0.139)	(0.109)	(0.146)
(ΔImports from Japan to US)/worker × Black	-0.638***	-0.509**	-0.595***	-0.564***	-0.830***	-1.035***
	(0.159)	(0.211)	(0.135)	(0.140)	(0.187)	(0.233)
Observations	716	716	716	716	716	716

Notes - Robust standard errors clustered at the state-level in parentheses. Models are weighted by race-specific 1970 population. Each regression includes census division fixed effects; commuting zone-level controls for percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960; a black indicator; and interactions of the black indicator with all of these variables. $p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 2.7.: Japanese Imports and Change in Employment Status by Race and Skill Group, 1990-1970 Long Difference: 2SLS Estimates

	Mfg emp	Non-mfg emp	Unemp	NILF
	(1)	(2)	(3)	(4)
	Panel A: HS Dropouts			
All Workers				
(Δ Imports from Japan to US)/worker	-0.271**	0.353***	0.044	-0.123
	(0.109)	(0.089)	(0.058)	(0.116)
Black Workers				
(Δ Imports from Japan to US)/worker	-0.877***	0.295	0.112	0.406*
	(0.109)	(0.204)	(0.120)	(0.216)
White Workers				
(Δ Imports from Japan to US)/worker	-0.090	0.402***	0.050	-0.345***
	(0.119)	(0.089)	(0.045)	(0.129)
	Panel B: HS Grads			
All Workers				
(Δ Imports from Japan to US)/worker	-0.202	0.105	0.059	0.036
	(0.155)	(0.112)	(0.042)	(0.049)
Black Workers				
(Δ Imports from Japan to US)/worker	-0.263	0.012	-0.228	0.350*
	(0.283)	(0.344)	(0.168)	(0.193)
White Workers				
(Δ Imports from Japan to US)/worker	-0.079	0.043	0.092**	-0.035
	(0.160)	(0.129)	(0.038)	(0.057)
	Panel C: College Educated			
All Workers				
(Δ Imports from Japan to US)/worker	0.264*	-0.321**	0.044	0.016
	(0.158)	(0.126)	(0.029)	(0.090)
White Workers				
(Δ Imports from Japan to US)/worker	0.300**	-0.312**	0.038	-0.028
	(0.150)	(0.121)	(0.029)	(0.090)
Observations	358	358	358	358

Notes - Robust standard errors clustered at the state-level in parentheses. Models for all workers are weighted by 1970 population. Models for racial subgroups are weighted by race-specific 1970 population. Each entry represents a separate regression for that race and/or skill group. Each regression includes census division fixed effects; commuting zone-level controls for percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 2.8.: Japanese Imports and Change in Employment Status for Southern versus Non-Southern Born Blacks, 1990-1970 Long Difference: 2SLS Estimates

	Mfg emp	Non-mfg emp	Unemp	NILF
	(1)	(2)	(3)	(4)
(Δ Imports from Japan to US)/worker	0.016	-0.028	-0.244*	0.237*
	(0.181)	(0.284)	(0.130)	(0.132)
(Δ Imports from Japan to US)/worker $\times$ Southern Born	-0.586*	-0.084	0.323***	0.264
	(0.347)	(0.476)	(0.073)	(0.193)
Observations	370	370	370	370

Notes - Robust standard errors clustered at the state-level in parentheses. Models are weighted by group-specific 1970 population. Each regression includes census division fixed effects; commuting zone-level controls for percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960; a Southern-born indicator; and interactions of the Southern-born indicator with all of these variables. The sample includes CZs with a population of at least 500 Southern born and 500 non-Southern born working age black males in 1970 and 1990. $p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 2.9.: Japanese Imports and Change in Skill Composition of Manufacturing, 1990-1970 Long Difference: 2SLS Estimates

	Share of Manufacturing Employment					
	College	HS dropout	Prof	Eng	Prd wrk	College prd wrk
	(1)	(2)	(3)	(4)	(5)	(6)
(Δ Imports from Japan to US)/worker	0.892***	-0.248	0.143*	0.154***	-0.056	0.648***
	(0.221)	(0.204)	(0.087)	(0.035)	(0.103)	(0.123)
Observations	358	358	358	358	358	358
1970 mean of DV	19.7	43.2	13.0	3.5	65.9	5.0

Notes - Robust standard errors clustered at the state-level in parentheses. Models are weighted by 1970 population. Left-hand side variable in column (1) is the share of manufacturing employment belonging to college educated workers. Left-hand side variable in column (2) is the share of manufacturing employment with less than a high school degree. Left-hand side variable in column (3) is the share of manufacturing employment in management and professional occupations. Left-hand side variable in column (4) is the share of manufacturing employment in engineering occupations. Left-hand side variable in column (5) is the share of manufacturing employment in production occupations. Left-hand side variable in column (6) is the share of manufacturing employment belonging to college educated workers in production occupations. Each regression includes census division fixed effects; and commuting zone-level controls for black percentage of population, foreign-born percentage of population, percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, and percentage of employment in routine occupations in 1960. $p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 2.10.: Japanese Imports and Changes in the Geography of Employment, 1990-1970 Long Difference: 2SLS Estimates

	CZ Population		Central City Population		Central City Man Share	
	Log pop (1)	Share black (2)	Log pop (3)	Share black (4)	Share reside (5)	Share jobs (6)
(Δ Imports from Japan to US)/worker	0.982 (0.894)	0.183*** (0.059)	-0.176 (0.235)	1.552*** (0.235)	-0.415 (0.320)	0.479 (0.538)
Observations	358	358	167	167	167	167
1970 mean of DV	1386.6	12.3	38.9	21.2	37.5	45.6

Notes - Robust standard errors clustered at the state-level in parentheses. Left-hand side variable in column (1) is the change in log population of the commuting zone. Left-hand side variable in column (2) is the change in the share of population that is black in the commuting zone. Left-hand side variable in column (3) is the change in log population of the central cities within the commuting zone. Left-hand side variable in column (4) is the change in the share of population that is black in the central cities within the commuting zone. Left-hand side variable in column (5) is the change in the share of manufacturing workers living in central cities within the commuting zone. Left-hand side variable in column (6) is the change in the share of manufacturing jobs located in central cities within the commuting zone. Long difference in columns (1)-(5) is 1990-1970. Long difference in column (6) is 1987-1972. Each regression includes census division fixed effects; commuting zone-level controls for percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 2.11.: Heterogeneous Effect of Japanese Imports on Change in Manufacturing Employment/ Working Age Population in CZs by Unionization and Prejudice, 1990-1970 Long Difference: 2SLS Estimates

	Unionization		Prejudice	
	Resid	Resid	Nat	Div
	(1)	(2)	(3)	(4)
(Δ Imports from Japan to US)/worker	-0.299	-0.228	0.367**	0.617***
	(0.395)	(0.416)	(0.152)	(0.238)
(Δ Imports from Japan to US)/worker $\times$ Black	-0.764***	-0.838	-1.407***	-1.646***
	(0.178)	(0.666)	(0.255)	(0.495)
(Δ Imports from Japan to US)/worker $\times$ Union Residual	0.055**	0.044		
	(0.027)	(0.027)		
(Δ Imports from Japan to US)/worker $\times$ Union Res. $\times$ Black		0.017		
		(0.049)		
(ΔImports from Japan to US)/worker $\times$ Wallace			-0.258	-0.467**
			(0.177)	(0.219)
(Δ Imports from Japan to US)/worker $\times$ Wallace $\times$ Black			0.773***	0.879**
			(0.297)	(0.447)
Observations	708	708	688	688

Notes - Robust standard errors clustered at the state-level in parentheses. Models are weighted by race-specific 1970 population. Each regression includes census division fixed effects; and commuting zone-level controls for black percentage of population, foreign-born percentage of population, percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, and percentage of employment in routine occupations in 1960. Columns (1) and (2) include the residual of a regression of 1967-1972 state-level unionization rates on 1970 state manufacturing share. Wallace indicator in column (3) equals one if the CZ was at or above the national median in Wallace vote share in 1968 presidential election. Wallace indicator in column (4) equals one if the CZ was at or above census division median in Wallace vote share in 1968 presidential election. $p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 2.12.: Japanese Imports and Changes in Financial Well-being, 1990-1970 Long Difference: 2SLS Estimates

	Log Weekly Wage	Log Annual Earnings			Log Ann Inc	% Welf Recp
	Earners	Earners	All	HH	HH	HH
	(1)	(2)	(3)	(4)	(5)	(6)
(Δ Imports from Japan to US)/worker	0.265	0.517	-0.108	-0.242	-0.121	0.137**
	(0.415)	(0.340)	(0.520)	(0.348)	(0.307)	(0.064)
(Δ Imports from Japan to US)/worker × Black	-0.067	-0.305	-3.570***	-2.607***	-1.861***	0.609***
	(0.263)	(0.431)	(0.830)	(0.491)	(0.394)	(0.106)
Observations	716	716	716	716	716	716

Notes - Robust standard errors clustered at the state-level in parentheses. Models are weighted by race-specific 1970 population. Wage, earnings and income variables are measured as CZ medians. Each regression includes census division fixed effects; commuting zone-level controls for percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960; a black indicator; and interactions of the black indicator with all of these variables. $p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 2.13.: Japanese Imports and Change in Racial Disparities, 1990-1970: Back of the Envelope Calculations

	Realized Change	Counterfactual Change	
		All	Exog
	(1)	(2)	(3)
	Panel A: Males, 16-64		
Manufacturing Employment	-1.34	-0.19	-0.46
NILF Rate	3.50	2.71	2.90
Log Median Earnings, All Males	-11.38	-6.36	-7.54
	Panel B: Households		
Log Median Earnings	3.71	7.35	6.49
Welfare Recipiency Rate	1.73	0.90	1.09

Notes - Realized changes calculated from 1970 1% form 1 and form 2 metro and 1990 5% IPUMS samples of the United States Decennial Censuses, and include individuals living in commuting zones that were not used in regression analysis due to the sample size of black workers. Counterfactual change calculations in column (2) based on regressions results from Tables 2.5 and 2.12, given that from 1970-1990, the average black worker was exposed to a $1,413 increase in Japanese trade competition, while the average white worker was exposed to a $1,601 increase. Counterfactual change calculations in column (3) instead use vales of $1079.53 and $1,223.16, which reflects the exogenous trade increase per worker based on a partial R^2 of 0.764 in the first-stage regression.

3. END-OF-LIFE MEDICAL SPENDING: EVIDENCE FROM PET INSURANCE

3.1 Introduction

In the United States, a large fraction of health spending happens in the last year of life.[1] The important question for policy makers is to what extent does increased medical care spending cause an increase in health for patients diagnosed with a high-mortality-rate illness. In both the RAND Health Insurance Experiment (Newhouse, 1993) and the Oregon Experiment (Finkelstein et al. 2012), there was no clear improvement in health as a result of the increase in health spending. Long lifespans make death rates a noisy outcome when working with small samples or short time periods. One problem is that both studies focused on health in general, rather than on increased health spending for those at an elevated risk of dying.

Unfortunately, the literature does not have a clear answer to the question, even when focused on those with an elevated risk of dying. Identifying a causal effect of additional health care spending is difficult because health insurance benefits are generally not random. In the United States, people with better jobs tend to have better health insurance. Using quasi-experimental methods has not produced consistent results. For example, using geographic variation in health care spending for serious illnesses, Skinner et al. (2005) find no decrease in the mortality rate. However, using differences in hospital quality for patients with a life-threatening illness who were randomly allocated between hospitals, Joseph J. Doyle (2011) finds that increased spending was associated with lower mortality. Melberg (2018) provides a helpful review of this literature.

In this paper we contribute to this literature by focusing on the pet health care industry. Using a creative identification strategy based on the timing of benefit renewal, we create an environment in which arrival of benefits is quasi-random. And because pet lifespan is much shorter, we are able to explore death rates as a reliable measure of health outcomes.

[1] Aldridge and Kelley (2015) estimate that 13 percent of health care costs in the United States was for care in the last year of life.

While pets are not people, we believe there is much to be learned from the pet health care industry. Researchers have noted striking similarities between human and pet health care spending patterns. Using a small extract of billing data from a pet hospital in California, Einav et al. (2017) document a large end-of-life spike in spending for dogs diagnosed with lymphoma. They compare this spending spike to a similar increase for Medicare patients diagnosed with lymphoma. They also note that "most dogs die cheaply" because there is no sharp increase at the end of life for the median dog, as opposed to the median Medicare patient. Instead, a smaller group of dogs drive the sharp end-of-life spending increase. As noted by these authors, the pet owners' financial status is likely an important component of pet health care decisions and ultimate outcomes.

In this paper we focus specifically on how the availability of health insurance reimbursement funds affects pet health after a serious cancer diagnosis. We analyze a sample of dogs who have been diagnosed with very serious cancer, and who have a health insurance plan. We ask whether the availability of insurance benefits affects the amount of treatment the dog undergoes. Availability of benefits varies exogenously by the point in time during the policy term that the dog is diagnosed. A policyholder whose dog is diagnosed late in the policy term will have the option to "double up" on benefits, by using the benefits available to them this term now, and the benefits available to them next term in the near future. Policyholders whose dogs are diagnosed early in the term, in contrast, will have to wait nearly a year for their benefits to renew. We ask whether this difference in timing has any effect on the policyholder's pet care decisions.

Our analysis is guided by a two-period model in which pet owners make cancer treatment decisions by responding to current and future benefit availability, the cost of treatment, and the probability that the treatment will be successful. Given that the treatment is expensive enough to prohibit at least some pet owners from treating, the decision to treat is increasing in the probability that the treatment will be successful. We demonstrate that the probability of successful treatment is conditional on the age of pet, where older dogs are less likely to survive the initial months of treatment. Because of this, pet owners tend to spend less on older dogs with the same diagnosis. In the next stage of analysis, we hold the success probability fixed by analyzing each age group separately, and including additional age controls.

We then turn to the impact of availability of insurance benefits on health care decisions and outcomes. The model predicts that a pet owner whose benefits will renew in the next period will be more likely to treat. We find that the availability of insurance benefits does affect spending behavior. Young dogs, ages 2 to 5, receive nearly one extra vet visit on average over the next 6 months if they are diagnosed late in their policy term. They also receive more medical care as indicated by higher spending over the next year, concentrated in the first 3 months. These differences are apparent only for young dogs, however. Old dogs, ages 6 to 9, do not experience any spending boost from a late-in-term diagnosis. Very old dogs, ages 10 to 12, actually experience a dip in spending over subsequent months if they are diagnosed late in the policy term.

This finding tells us that pet health care decisions are largely dependent on the age of the dog. Pet owners are more likely to ramp up spending if their dog is still young. We attribute this difference to the likelihood that treatment will be successful, since older dogs have higher death rates even conditional on amount of treatment. While we see effects on health care decisions, these effects do not translate to different outcomes for the pet in any age group. We find no significant difference in death rates for any age group in the months following diagnosis.

Finally, we explore heterogeneous effects by cost of treatment for young dogs. The model predicts that more expensive treatments will have stronger effects. Consistent with this prediction, we find that the effect of a late diagnosis on spending is increasing in cost of treatment. The effect of a late diagnosis on pet death in the next 12 and 24 months is decreasing in the cost of treatment. This suggests that while we do not find an impact of a late diagnosis on pet outcomes overall, the most expensive treatments are indeed sensitive to the timing of benefits.

Apart from any extrapolation made between the results here and human health care, these pet health care results are interesting for what we learn about pets. About half of all US households (64 million) have at least one pet.[2] Dogs are the most common pet in the U.S. with an estimated 77 million dogs in total.[3] Given the importance of dogs in the modern lifestyle, our investments in their health care are likely of interest to many.

[2]The U.S. Census Bureau's American Housing Survey (2017) of 30,000 randomly selected households implies that 49 percent of households have at least one pet. The Simmons National Consumer Study (2018) survey of 25,000 randomly selected households suggests that 53 percent of households have a pet. Results from this survey are not released publicly, but the pet data was reported on by the Washington Post (Jan 31, 2019).

[3]Both the Simmons National Consumer Study (2018) and the American Veterinary Medical Association Survey (2016) estimate that 38 percent of U.S. households have one or more dogs.

3.2 Model

We propose a simple model of the behavior of pet owners who own a dog who becomes ill with cancer in period 1. The pet owners in our model each have a pet health insurance policy which will pay up to benefit level B of the cost of cancer treatment. Each pet owners' utility is H in periods in which the dog is healthy, $-C$ in periods in which the dog is sick, and 0 after the dog dies. The pet owner has three choices: (1) veterinary treatment which, with probability p, will result in the dog being healthy, (2) euthanization which will result in a utility level of 0 for the current and all future periods, and (3) do nothing. All dogs in our model, who have not been euthanized, die at the end of period 2.

The cancer treatment costs B in each period in which it is provided. The full amount of treatment is covered by the pet health insurance, but this is the maximum amount that insurance will cover until the term ends and the new term begins. For type 1 pet owners, the term began in period 1 and the new B in benefits will not be available until after period 2. For type 2 pet owners, a new B in benefits becomes available at the beginning of period 2. Clearly, it is better to be a type 2 pet owner because insurance will cover B cost of cancer treatment in period 1 and B cost of cancer treatment in period 2 if needed. For the type 1 pet owners, pet insurance will cover B cost of cancer treatment in either period 1 or period 2, not both. The timing of the new term is the only difference between type 1 and type 2 pet owners and we indicate this by $i = 1$ or $i = 2$. We denote the expected continuation utility as U_i and assume that the pet owners do not discount.

To solve the model, we start in period 2. If the dog was successfully treated in period 1, the continuation utility is H. If the dog was euthanized in period 1, the continuation utility is 0. If the dog was either not treated in period 1 or the treatment was unsuccessful, then the pet owner will need to choose between euthanization or treatment in period 2. We can rule out doing nothing as a choice because this would guarantee a period 2 continuation utility of $-C$ which is less than zero. The expected continuation utility depends on if the pet owner has insurance benefits B available. If the owner has insurance benefits, expected utility is given by:

$$E\left[U_i\right] = \begin{cases} 0 & \text{if euthanize} \\ pH + (1-p)(-C) & \text{if treat} \end{cases}$$

If the owner has no insurance benefits available (the policy began in period 1 and the owner exhausted the benefits due to an unsuccessful treatment in period 1), the expected utility is given by:

$$E\left[U_i\right] = \begin{cases} 0 & \text{if euthanize} \\ pH + (1-p)(-C) - B & \text{if treat} \end{cases}$$

If the owner has insurance benefits B available, the owner will choose to treat the pet as long as the probability p of successful treatment is greater than $\frac{C}{H+C}$, where C and H are both positive. If the owner has no insurance benefits available, p must be greater than $\frac{C+B}{H+C}$ for the owner to choose treatment. The interesting case is when p takes an intermediate value:

$$\frac{C}{H+C} < p < \frac{C+B}{H+C}.$$

which implies that the owner will only choose the treatment in period 2 if there are insurance benefits available.

Now consider the owner's choice in period 1. All pets in this period have cancer and all owners have insurance benefits B available. Owners of type 1 could choose to do nothing in period 1, saving their insurance benefits for treatment in period 2, but this provides lower expected utility than treating in period 1 because there is no advantage to waiting to treat. If we are in the interesting case where p takes on an intermediate value as described above, owners of type 1 would choose to euthanize in period 2 if the treatment is not successful, so we can derive that they will choose to treat in period 1 only if $p > \frac{C}{2H+C}$.

Pet owners of type 2 can use their insurance benefits to pay for treatment in period 1 and after the policy term ends, they will have an additional B in benefits to pay for treatment in period 2 if needed. We can show that type 2 pet owners will always choose treatment in period 1 if p is greater than the cutoff for treatment in period 2.[4] If p is greater than $\frac{C}{2H+C}$, but less than $\frac{C}{H+C}$ the type 2 pet owners will choose to treat in period 1 and euthanize in period 2. If p is greater than $\frac{C}{H+C}$ then type 2 pet owners will choose to treat in period 1 and if that treatment is unsuccessful, they will choose to treat again in period 2.

To summarize, if p is greater than $\frac{C+B}{H+C}$, both types will choose to treat in period

[4]Type 2 pet owners will choose to treat in period 1 if $p > \frac{3C+3H+\sqrt{C^2+10CH+9H^2}}{2(C+H)}$. For $H > 0$ and $C > 0$, $\frac{C}{H+C} > \frac{3C+3H+\sqrt{C^2+10CH+9H^2}}{2(C+H)} > 0$.

1 and both types would choose treat in period 2 if the treatment in period 1 is unsuccessful. If p is less than $\frac{C+B}{H+C}$ but greater than $\frac{C}{H+C}$, both types will choose to treat in period 1, but only type 2 pet owners would choose to treat in period 2.[5] If p is less than $\frac{C}{H+C}$ but greater than $\frac{C}{2H+C}$ then both types will choose to treat in period 1, but neither will choose to treat in period 2. Finally, if p is less than $\frac{C}{2H+C}$ neither type will choose to treat in period 1.

The implication of this model is that the timing of the pet health insurance term can have an important impact on the treatment decision. Those who have a pet diagnosed with a serious, but treatable, disease a short time before the end of the insurance term can "double up" by having B of veterinary expenses covered in the current term and another B covered in the next term. However, those who have a pet diagnosed with a serious disease soon after the beginning of the insurance term, will only be able to receive B of covered veterinary expenses in total.

3.3 Data

We use claims-level administrative data from Nationwide Pet Insurance, the largest provider of pet health insurance policies in the United State.[6] It is important to note that only a small share of dogs in the U.S., about 2 percent in 2018, have a health insurance plan.[7] We observe claims from January 2009 through December 2019 for the universe of pet policy holders. For each claim we observe the date of treatment, a description of each treatment provided, the cost of the treatment as indicated on the veterinary bill, and the amount reimbursed by the insurance company.

For this study, we select a sample of dogs who were diagnosed with serious cancers between January 1, 2009 and July 31, 2017, so that we can observe a full 24-month period from date of diagnosis. We define serious cancers as cancers which are associated with a death rate 12 months after diagnosis that is greater than 70 percent.[8] We

[5]A larger B implies a greater range of p over which type 2 owners will choose to treat in period 2 but type 1 owners will not.

[6]Nationwide's market share is 35 percent. Pet health insurance does not include insurance policies that cover livestock, horses, or other farm animals. Market share is the percentage of gross written premiums as reported by the North American Pet Health Insurance Association State of the Industry Report (2018).

[7]North American Pet Health Insurance Association State of the Industry Report (2018) reports 1,538,000 active health insurance policies for dogs.

[8]These include: heart/pericardium neoplasia; thorax neoplasia; metastatic or infiltrative neoplasia; brain or spinal cord neoplasia; peritoneal neoplasia; osteogenic sarcoma; stomach neoplasia; hepatic neoplasia; lymphosarcoma; urethral neoplasia; small intestine neoplasia; peripheral vessels neoplasia;

remove dogs who were diagnosed with a less serious cancer within the 2 years leading up to their serious cancer diagnosis, since these dogs were already sick. We further restrict the sample to dogs who are at least 2 years of age but no older than 12 at the time of diagnosis, and who have had an insurance policy for at least 1 year. This leaves us with a sample of 33,899 dogs.

We identify the date of diagnosis as the first treatment date in which a cancer-related medical claim was made. Determining if and when a dog dies from administrative insurance data involves some imputation. If there is a medical claim identifying pet death (i.e. claim description mentions death, euthanasia, and/or remains care), we use the date of this claim. If there is a cancellation of the policy, we use the date of cancellation, assuming that a cancellation when the pet has a serious medical condition indicates medical care ceased and the pet died. Likewise, if the policy is not renewed after the end of the policy term, and the policy term expired July 31, 2019 or earlier, we use the date that the term expired. Finally, if there is a denial of benefits that indicates pet death, we use the date of the medical claim associated with that denial.

Some imputations are also required for the cost of veterinary treatments. Claims where the cost of treatment is recorded as $0 are replaced with the median value by claim code and breed size.[9] We also replace cost of treatment values below the 10th percentile with the 10th percentile and values above the 90th percentile with the 90th percentile.

We also observe a variety of characteristics of the dog and the policy which we use as controls. These include indicators for female, 10 breed size categories (mixed and pure; great, large, medium, small, and toy), age at diagnosis, Census region, and plan type (including 26 different base plans, 10 wellness riders, and a cancer rider). We also use as controls total spending in the 12 months leading up to diagnosis, and month and year of diagnosis indicators. We show summary statistics for our full sample in Table 3.1.

spleen neoplasia; leukemia; prostate neoplasia; and islet cell tumor.

[9] For a small number of cases where that claim code and breed size cell is empty, we replace it with the median value for the claim code.

End-of-Life Spending

In the human health care system, an important debate is the extent to which end-of-life spending is a large and growing share of overall medical spending (Melberg, 2018). Among other trends, researchers have noted that end-of-life medical care declines with age (e.g., Lubitz et al., 1995; Kwok et al., 2011).

We report patterns in end-of-life spending for our sample in Figures 3.6 and 3.6.[10] Of the 33,899 in our sample, 31,484 die before the end of the sample period, so we consider end-of-life spending for these decedents. Overall, nearly $\$4,000$ or about 40 percent of lifetime medical spending takes place in the last 12 months of life. Pet owners spend less during the last 12 months of life if their pet is older, and spend less upon diagnosis for their older pets.

End-of-life spending also increases over time in our sample, from under $\$3,300$ in 2009 to over $\$4,500$ in 2017. While this may be connected to an overall increase in the cost of medical care, this spending as a share of lifetime spending also increases from 36 to 42 percent during these years.

The Role of Age in Probability of Successful Treatment

In the model, given that the treatment is expensive enough to prohibit at least some pet owners from treating, the decision to treat is increasing in the probability that the treatment will be successful. We first demonstrate that the probability of successful treatment is conditional on age. We estimate the following:

$$Death_{i,t} = \alpha_0 + \alpha_1 Old_i + \alpha_2 VeryOld_i + \alpha_3 Spending_{i,t} + \gamma X_i + \epsilon_i \tag{3.1}$$

where $Death_{i,t}$ is an indicator for death of dog i during time horizon t. Dogs are divided into three age groups: Young (ages 2-5), Old (ages 6-9), and Very Old (ages 10-12). We control for a variety of factors in the matrix X_i, including plan type, breed type and size, Census region, cumulative spending in the 12 months before diagnosis, month and year of diagnosis, and term week of diagnosis. The effect of treatment is captured by the coefficient on log spending, α_3.

We show the results in Table 3.2. While higher spending is associated with a higher death rate in the very short term (Column 1), for most time horizons higher

[10] For tables showing these trends, see Section C.1.

spending is associated with lower death rates. In Column 2, we see that a 100 percent increase in spending in the first 3 months is associated with a 7.5 percentage point reduced likelihood of death in those months.

However, even conditional on treatment amount, the age group of the dog plays an important role in the likelihood of survival. In Column 2, we see that old dogs have a 9 percentage point higher death rate in the first 3 months, and very old dogs have a 13.2 percentage point higher death rate. This result establishes that age affects the probability of successful treatment.

Next, we show that pet owners respond by treating older dogs less. We estimate the following:

$$Spending_{i,t} = \alpha_0 + \alpha_1 Old_i + \alpha_2 VeryOld_i + \gamma X_i + \epsilon_i \tag{3.2}$$

where $Spending_{i,t}$ is log spending for dog i during time horizon t. We show the results in Table 3.3. In Column 4, we see that in the 12 months following diagnosis, old dogs receive 12.1 percent less medical treatment than young dogs, and very old dogs receive 26.8 percent less.

The conclusion from this analysis is that (1) age affects the probability that treatment will be successful, and (2) the probability of success plays a role in the decision to treat the pet. For the remainder of the analysis, where we explore the impact of the timing of benefit renewal, we hold this probability fixed by analyzing each age group separately. Given the importance of age, we add indicators for age within each age group, and age by breed size category interaction terms as controls.

3.4 Estimation Strategy

All policies in our sample have a term length of 12 months. Unlike human health care policies, which are highly seasonal, the Nationwide pet policies have start dates throughout the calendar year.[11]

We rely on the assumption that cancer is a random shock that arrives independently of the month of term. Some dogs will be diagnosed toward the beginning of their policy term, so that there are many months until their benefits renew. Other

[11] A small group of policyholders sign up through a group program, indicated as "group program payroll" in the "origination" variable. This group, representing about 20 percent of the policy-term observations, is disproportionately more likely to have policy start dates in December and January and are removed from the data.

dogs will be diagnosed toward the end of their policy term, giving their owners the option to spend their cancer benefits for the current term now, and the benefits for the next term in the near future. This means dogs diagnosed toward the end of the term have the possibility of "doubling-up" on benefits over the course of a few months.

We compare dogs who are diagnosed in weeks 5-16 (months 2-4) of their policy term compared to those who are diagnosed in weeks 41-53 (months 10-12) of their policy term. We do not include dogs diagnosed in month 1 of their policy term because the high numbers of vet visits during this month may be associated with a biased sample of diagnoses uncovered during routine checkups. As shown previously, pet owners make health care decisions differently depending on pet age therefore we stratify our analysis to different age groups. Descriptive statistics for this sample of early-term and late-term diagnosis dogs are reported in Table 3.4.

We are interested in the effect of a late diagnosis on spending, number of vet visits, and pet death over the subsequent months. Specifically, we estimate the following:

$$Y_i = \beta_0 + \beta_1 LateDiagnosis_i + \gamma X_i + \epsilon_i \tag{3.3}$$

where Y_i refers to outcome Y for pet i, β_0 is a constant, γ is a vector of coefficients on control matrix X_i, and ϵ_i is an error term. β_1 is the coefficient of interest, representing the causal impact of a late term diagnosis on outcome Y.

3.5 Results: Effects of Benefit Renewal

Young Dogs

We report the main set of results for young dogs in Table 3.5. In each regression, we include the full set of controls. Each column shows the effect of a late diagnosis on the outcome over a different time horizon, ranging from 1 to 24 months. Starting with Panel A, we see that dogs diagnosed late in the term do experience an increase in spending over subsequent months. This suggests that owners are "doubling-up" on benefits by spending both current and next term benefits. In Panel B, we also see an increase in the number of vet visits by nearly one third of a visit during the first month, and nearly a full extra visit over the next six months. Consistent with the model, this is evidence that the presence of future benefits results in greater health care spending.

Despite evidence for increased medical care for these pets, we do not find any evidence that these pets are better off in terms of mortality. In Panel C, we see that later diagnosed pets are not less likely to die in the subsequent months. This result is reminiscent of findings elsewhere that the very high spending on human health care in the U.S. is not associated with better health outcomes.

Old Dogs

In Table 3.6, we repeat the analysis on older dogs, ages 6 through 9 at time of diagnosis. Here we see that these pets do not benefit from increased medical spending by their owner when they are diagnosed late in the policy term. This is consistent with the model predicting that pet owners will not be sensitive to benefit renewal if the likelihood of success is too low.

Curiously, late diagnosed dogs in this age group are less likely to die during the 12 and 24 month time horizons, despite no evidence of changes to medical care responses by the owner.

Very Old Dogs

For very old dogs, ages 10 to 12 at time of diagnosis, we actually find that late in term diagnoses are associated with less spending over the subsequent months. We report these results in Panel A of Table 3.7. It is possible that these pet owners are discouraged by the diagnosis and choose not to renew their policies, resulting in less spending in subsequent months. This is not associated with fewer vet visits, however, or any differences in ultimate outcomes, shown in Panels B and C of Table 3.7. It is possible that there is a divide between pet owners who have zero visits because they do not renew their policies and those who increase their number of visits, and these opposing effects cancel each other out.

Common across all three age groups is that later diagnosed pets do not have any differences in the likelihood of death. This is evidence that increased medical spending is not associated with a longer lifespan for the dog.

Expensive Treatments

Another prediction from the model is that more expensive treatments will have larger effects. In this final section, we interact the effect of a late diagnosis with the average cost of treatment.[12] We focus on young dogs since this group is associated with the strongest responses by pet owners, as shown above.

We report the results in Tables 3.8 through 3.10. Treatment cost is normalized to mean 0 with a standard deviation of 1. In Table 3.8, we see that more expensive treatment is unsurprisingly associated with higher spending across all time horizons. We also see that spending is even higher if the dog is diagnosed late in the policy term. In Column 4, we see that a 1 standard deviation increase in treatment cost is associated with a 16.7 percent increase in spending over the 12 months following diagnosis for dogs diagnosed early in the term, but a 20.9 percent increase in spending for late diagnosed dogs. Likewise, in Table 3.9, we see that cost of treatment is also associated with more vet visits for late diagnosed dogs.

Finally, in Table 3.10, we explore the impact of treatment cost on death. While more expensive treatments are associated with lower initial death rates (Column 1), these treatments are associated with higher likelihood of death over the next year. These effects are attenuated, however, if the dog is diagnosed late in the policy term. Treatment cost is less prohibitive in cases where the dog is diagnosed late, since the pet owner is able to double-up on the benefits from the following term. As a result, these dogs are more likely to survive.

3.6 Discussion and Conclusion

This paper provides evidence that pet owners are sensitive to changes in insurance benefits when making health care decisions for their pets. Furthermore, pet owners take into account the age of the pet when responding to a serious cancer diagnosis. Owners spend less on older dogs with the same cancer diagnosis, as these dogs are less likely to survive. Finally, we do not find any evidence that these health care decisions have a significant impact on the longevity of the pet. The exception to this is for the most expensive cancer treatments, where the marginal impact of treatment cost on death is attenuated in the 12 and 24 month time horizons when dogs are diagnosed late in their policy term.

While there are many differences between pet health insurance and human health insurance, there is something to be learned from pet health trends.

[12] We calculate average cost of treatment by diagnosis as median spending during the 12 months following diagnosis.

Figures and Tables

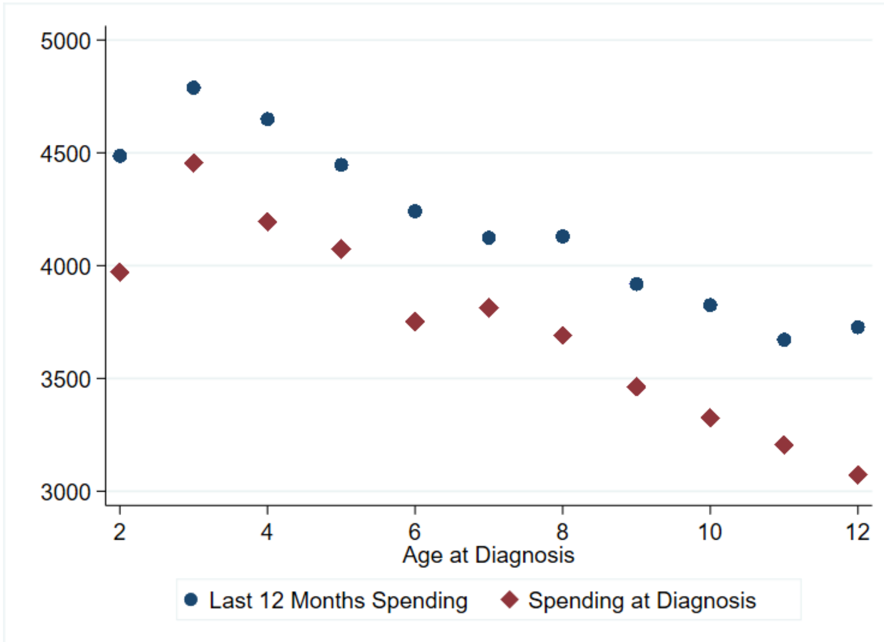

Figure 3.1.: End-of-Life and Post-Diagnosis Spending by Age at Diagnosis

Notes - Last 12 Months Spending is mean end-of-life spending for the 31,484 dogs in our sample who died during the sample period. Spending at Diagnosis is mean spending during 12 months following diagnosis for the 33,899 dogs in our sample. Y-axis is US Dollars.

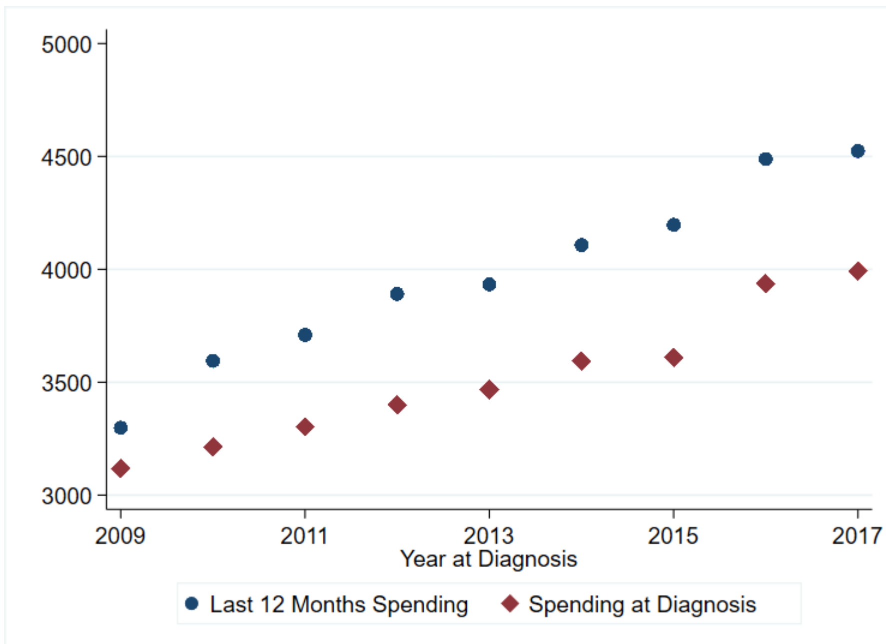

Figure 3.2.: End-of-Life and Post-Diagnosis Spending by Year of Diagnosis

Notes - Last 12 Months Spending is mean end-of-life spending for the 31,484 dogs in our sample who died during the sample period. Spending at Diagnosis is mean spending during 12 months following diagnosis for the 33,899 dogs in our sample. Y-axis is US Dollars.

Table 3.1.: Descriptive Statistics, Full Sample

	Mean (1)	St. Dev. (2)	Min. (3)	Max. (4)
Early Diagnosis Share	0.23	0.42	0.00	1.00
Late Diagnosis Share	0.22	0.41	0.00	1.00
Mean Log Spending Next 3 Months	7.58	0.81	3.69	10.60
Mean Total Visits Next 3 Months	3.32	2.97	1.00	26.00
Death Rate Next 3 Months	0.56	0.50	0.00	1.00
Mean Age at Diagnosis	8.69	2.29	2.00	12.00
Share Female	0.45	0.50	0.00	1.00
Mean Total Spending 12 Months Prior to Diagnosis	1013.37	1382.90	0.00	38974.34
Mean Year of Diagnosis	2013.14	2.39	2009.00	2017.00
Mean Month of Diagnosis	6.40	3.43	1.00	12.00
Observations	33,899	33,899	33,899	33,899

Table 3.2.: Death Rates by Age Group, Conditional on Treatment

	(1) 1 Month	(2) 3 Months	(3) 6 Months	(4) 12 Months	(5) 24 Months
Old Dogs (Ages 6-9)	0.095***	0.093***	0.078***	0.042***	0.041***
	(0.009)	(0.009)	(0.009)	(0.008)	(0.007)
Very Old Dogs (Ages 10-12)	0.149***	0.132***	0.111***	0.075***	0.085***
	(0.010)	(0.010)	(0.009)	(0.008)	(0.007)
Log 1-Month Spending	0.017***				
	(0.003)				
Log 3-Month Spending		-0.075***			
		(0.003)			
Log 6-Month Spending			-0.113***		
			(0.003)		
Log 12-Month Spending				-0.110***	
				(0.003)	
Log 24-Month Spending					-0.088***
					(0.002)
Observations	33,899	33,899	33,899	33,899	33,899

Notes - Each column refers to 1, 3, 6, 12, and 24 month time horizons from date of diagnosis. Omitted Category is Young Dogs, Ages 2-5. Controls include 26 plan indicators, 10 wellness rider indicators, a cancer rider indicator, 10 breed size category indicators (pure and mixed; great, large, medium, small, and toy), 4 Census region indicators, cumulative spending during 12 months prior to diagnosis, calendar year of diagnosis indicators, calendar month of diagnosis indicators, and a term week of diagnosis control. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 3.3.: Effect of Age on Spending Behavior After Cancer Diagnosis

	(1) 1 Month	(2) 3 Months	(3) 6 Months	(4) 12 Months	(5) 24 Months
Old Dogs (Ages 6-9)	-0.000	-0.058***	-0.097***	-0.121***	-0.131***
	(0.016)	(0.015)	(0.015)	(0.016)	(0.017)
Very Old Dogs (Ages 10-12)	-0.096***	-0.196***	-0.241***	-0.268***	-0.294***
	(0.016)	(0.016)	(0.016)	(0.017)	(0.018)
Observations	33,899	33,899	33,899	33,899	33,899

Notes - Each column refers to 1, 3, 6, 12, and 24 month time horizons from date of diagnosis. Omitted category is Young Dogs, Ages 2-5. Controls include 26 plan indicators, 10 wellness rider indicators, a cancer rider indicator, 10 breed size category indicators (pure and mixed; great, large, medium, small, and toy), 4 Census region indicators, cumulative spending during 12 months prior to diagnosis, calendar year of diagnosis indicators, calendar month of diagnosis indicators, and a term week of diagnosis control. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 3.4.: Descriptive Statistics, Regression Sample

	Young (2-5 yrs) (1)	Old (6-9 yrs) (2)	Very Old (10-12 yrs) (3)
Late Diagnosis Share	0.525 (0.50)	0.510 (0.50)	0.435 (0.50)
Mean Log Spending Next 3 Months	7.721 (0.85)	7.640 (0.80)	7.501 (0.79)
Mean Total Visits Next 3 Months	4.383 (3.79)	3.419 (3.03)	3.034 (2.72)
Death Rate Next 3 Months	0.453 (0.50)	0.565 (0.50)	0.596 (0.49)
Mean Age at Diagnosis	4.156 (0.95)	7.802 (1.06)	10.847 (0.80)
Share Female	0.415 (0.49)	0.437 (0.50)	0.484 (0.50)
Mean Total Spending 12 Months Prior to Diagnosis, USD	862.731 (1259.27)	974.703 (1322.47)	1113.827 (1427.75)
Mean Year of Diagnosis	2012.868 (2.34)	2013.163 (2.37)	2013.385 (2.34)
Mean Month of Diagnosis	6.427 (3.43)	6.345 (3.41)	6.423 (3.42)
Observations	1,524	7,618	6,171

Table 3.5.: Effect of Late-in-Term Cancer Diagnosis on Spending, Vet Visits, and Likelihood of Death, Young Dogs (Ages 2-5)

	(1) 1 Month	(2) 3 Months	(3) 6 Months	(4) 12 Months	(5) 24 Months
			Panel A. Spending		
Late Diagnosis	0.080*	0.124***	0.127***	0.092*	0.057
	(0.047)	(0.047)	(0.049)	(0.048)	(0.049)
			Panel B. Vet Visits		
Late Diagnosis	0.314***	0.631***	0.945***	0.914**	0.779
	(0.099)	(0.211)	(0.320)	(0.414)	(0.539)
			Panel C. Death		
Late Diagnosis	0.016	0.032	0.027	0.015	0.028
	(0.026)	(0.027)	(0.027)	(0.025)	(0.023)
Observations	1,524	1,524	1,524	1,524	1,524

Notes - Each column refers to 1, 3, 6, 12, and 24 month time horizons from date of diagnosis. Controls include 26 plan indicators, 10 wellness rider indicators, a cancer rider indicator, pet age at diagnosis indicators, 10 breed size category indicators (pure and mixed; great, large, medium, small, and toy), pet age by breed size category interaction terms, 4 Census region indicators, cumulative spending during 12 months prior to diagnosis, calendar year of diagnosis indicators, and calendar month of diagnosis indicators. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 3.6.: Effect of Late-in-Term Cancer Diagnosis on Spending, Vet Visits, and Likelihood of Death, Old Dogs (Ages 6-9)

	(1) 1 Month	(2) 3 Months	(3) 6 Months	(4) 12 Months	(5) 24 Months
			Panel A. Spending		
Late Diagnosis	0.017	0.004	-0.002	-0.015	-0.018
	(0.019)	(0.019)	(0.019)	(0.020)	(0.021)
			Panel B. Vet Visits		
Late Diagnosis	0.030	0.018	0.058	-0.078	-0.104
	(0.036)	(0.071)	(0.103)	(0.140)	(0.196)
			Panel C. Death		
Late Diagnosis	-0.001	0.001	-0.008	-0.028***	-0.024***
	(0.012)	(0.012)	(0.011)	(0.010)	(0.009)
Observations	7,618	7,618	7,618	7,618	7,618

Notes - Each column refers to 1, 3, 6, 12, and 24 month time horizons from date of diagnosis. Controls include 26 plan indicators, 10 wellness rider indicators, a cancer rider indicator, pet age at diagnosis indicators, 10 breed size category indicators (pure and mixed; great, large, medium, small, and toy), pet age by breed size category interaction terms, 4 Census region indicators, cumulative spending during 12 months prior to diagnosis, calendar year of diagnosis indicators, and calendar month of diagnosis indicators. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 3.7.: Effect of Late-in-Term Cancer Diagnosis on Spending, Vet Visits, and Likelihood of Death, Very Old Dogs (Ages 10-12)

	(1) 1 Month	(2) 3 Months	(3) 6 Months	(4) 12 Months	(5) 24 Months
			Panel A. Spending		
Late Diagnosis	-0.029	-0.035*	-0.033	-0.047**	-0.058**
	(0.021)	(0.020)	(0.021)	(0.021)	(0.023)
			Panel B. Vet Visits		
Late Diagnosis	0.060	0.092	0.146	0.045	-0.083
	(0.037)	(0.070)	(0.104)	(0.146)	(0.202)
			Panel C. Death		
Late Diagnosis	-0.010	0.007	0.013	-0.002	0.000
	(0.013)	(0.013)	(0.012)	(0.011)	(0.009)
Observations	6,171	6,171	6,171	6,171	6,171

Notes - Each column refers to 1, 3, 6, 12, and 24 month time horizons from date of diagnosis. Controls include 26 plan indicators, 10 wellness rider indicators, a cancer rider indicator, pet age at diagnosis indicators, 10 breed size category indicators (pure and mixed; great, large, medium, small, and toy), pet age by breed size category interaction terms, 4 Census region indicators, cumulative spending during 12 months prior to diagnosis, calendar year of diagnosis indicators, and calendar month of diagnosis indicators. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 3.8.: Effect of Late-in-Term Cancer Diagnosis on Spending, Young Dogs (Ages 2-5), Heterogeneous Effects by Cost of Treatment

	(1) 1 Month	(2) 3 Months	(3) 6 Months	(4) 12 Months	(5) 24 Months
Late Diagnosis	0.083*	0.128***	0.132***	0.097**	0.062
	(0.047)	(0.046)	(0.048)	(0.047)	(0.048)
Treatment Cost	0.087***	0.147***	0.173***	0.167***	0.155***
	(0.032)	(0.032)	(0.033)	(0.033)	(0.033)
Late Diagnosis	0.011	0.022	0.030	0.042	0.052
× Treatment Cost	(0.045)	(0.044)	(0.045)	(0.045)	(0.045)
Observations	1,524	1,524	1,524	1,524	1,524

Notes - Each column refers to 1, 3, 6, 12, and 24 month time horizons from date of diagnosis. Interaction term is median 12 month spending on treatment for cancer diagnosis, normalized to mean 0 standard deviation 1. Controls include 26 plan indicators, 10 wellness rider indicators, a cancer rider indicator, pet age at diagnosis indicators, 10 breed size category indicators (pure and mixed; great, large, medium, small, and toy), pet age by breed size category interaction terms, 4 Census region indicators, cumulative spending during 12 months prior to diagnosis, calendar year of diagnosis indicators, and calendar month of diagnosis indicators. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 3.9.: Effect of Late-in-Term Cancer Diagnosis on Vet Visits, Young Dogs (Ages 2-5), Heterogeneous Effects by Cost of Treatment

	(1) 1 Month	(2) 3 Months	(3) 6 Months	(4) 12 Months	(5) 24 Months
Late Diagnosis	0.323***	0.654***	0.977***	0.951**	0.817
	(0.097)	(0.206)	(0.313)	(0.406)	(0.532)
Treatment Cost	0.333***	0.822***	1.133***	1.301***	1.296***
	(0.067)	(0.143)	(0.217)	(0.282)	(0.369)
Late Diagnosis	-0.017	0.041	0.205	0.363	0.561
× Treatment Cost	(0.092)	(0.197)	(0.299)	(0.387)	(0.508)
Observations	1,524	1,524	1,524	1,524	1,524

Notes - Each column refers to 1, 3, 6, 12, and 24 month time horizons from date of diagnosis. Interaction term is median 12 month spending on treatment for cancer diagnosis, normalized to mean 0 standard deviation 1. Controls include 26 plan indicators, 10 wellness rider indicators, a cancer rider indicator, pet age at diagnosis indicators, 10 breed size category indicators (pure and mixed; great, large, medium, small, and toy), pet age by breed size category interaction terms, 4 Census region indicators, cumulative spending during 12 months prior to diagnosis, calendar year of diagnosis indicators, and calendar month of diagnosis indicators. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table 3.10.: Effect of Late-in-Term Cancer Diagnosis on Death, Young Dogs (Ages 2-5), Heterogeneous Effects by Treatment Cost

	(1) 1 Month	(2) 3 Months	(3) 6 Months	(4) 12 Months	(5) 24 Months
Late Diagnosis	0.015	0.032	0.027	0.016	0.029
	(0.026)	(0.027)	(0.027)	(0.025)	(0.023)
Treatment Cost	-0.046***	0.006	0.022	0.047***	0.055***
	(0.018)	(0.019)	(0.019)	(0.018)	(0.016)
Late Diagnosis	0.006	-0.051*	-0.062**	-0.042*	-0.037*
× Treatment Cost	(0.024)	(0.026)	(0.026)	(0.024)	(0.022)
Observations	1,524	1,524	1,524	1,524	1,524

Notes - Each column refers to 1, 3, 6, 12, and 24 month time horizons from date of diagnosis. Interaction term is median 12 month spending on treatment for cancer diagnosis, normalized to mean 0 standard deviation 1. Controls include 26 plan indicators, 10 wellness rider indicators, a cancer rider indicator, pet age at diagnosis indicators, 10 breed size category indicators (pure and mixed; great, large, medium, small, and toy), pet age by breed size category interaction terms, 4 Census region indicators, cumulative spending during 12 months prior to diagnosis, calendar year of diagnosis indicators, and calendar month of diagnosis indicators. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

BIBLIOGRAPHY

Acemoglu, D. (2002). Technical change, inequality, and the labor market. *The Journal of Economic Literature*, 40(1):7–72.

Acemoglu, D. (2003). Patterns of skill premia. *The Review of Economic Studies*, 70(2):199–230.

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. 4:1043–1171.

Acemoglu, D., Autor, D., Dorn, D., Hanson, G. H., and Price, B. (2016). Import competition and the great us employment sag of the 2000s. *Journal of Labor Economics*, 34(S1):S141–S198.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):000–000.

Acs, G. and Danziger, S. (1993). Educational attainment, industrial structure, and male earnings through the 1980s. *Journal of Human Resources*, 28(3):618–648.

Adão, R., Kolesár, M., and Morales, E. (2018). Shift-share designs: Theory and inference. NBER Working Paper No. 24944.

Akst, D. (2013). Automation anxiety. *The Wilson Quarterly*, 37(3):65.

Aldridge, M. D. and Kelley, A. S. (2015). The myth regarding the high cost of end-of-life care. *American Journal of Public Health*, 105(12):2411–2414.

Apollonio, D. E. and Glantz, S. A. (2016). Minimum ages of legal access for tobacco in the united states from 1863 to 2015. *American Journal of Public Health*, 106(7):1200–1207.

Autor, D., Dorn, D., Hanson, G., et al. (2018). When work disappears: Manufacturing decline and the falling marriage-market value of young men. *American Economic Review: Insights*.

Autor, D., Dorn, D., Hanson, G. H., Pisano, G., and Shu, P. (2016). Foreign competition and domestic innovation: Evidence from us patents. Technical report, National Bureau of Economic Research.

Autor, D. H. (2013). The “task approach” to labor markets: an overview.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review*, 103(5):1553–97.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, 103(6):2121–68.

Autor, D. H., Dorn, D., Hanson, G. H., and Song, J. (2014). Trade adjustment: Worker-level evidence. *The Quarterly Journal of Economics*, 129(4):1799–1860.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2008). Trends in us wage inequality: Revising the revisionists. *The Review of Economics and Statistics*, 90(2):300–323.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.

Bartel, A., Ichniowski, C., and Shaw, K. (2007). How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills. *The Quarterly Journal of Economics*, 122(4):1721–1758.

Bartel, A. P., Ichniowski, C., and Shaw, K. (2003). “new technology” and its impact on the jobs of high school educated workers: A look deep inside three manufacturing industries. *Low-Wage America*, pages 155–94.

Batistich, M. K. and Bond, T. N. (2019). Stalled racial progress and japanese trade in the 1970s and 1980s.

Bayer, P. and Charles, K. K. (2018). Divergent paths: A new perspective on earnings differences between black and white men since 1940. *Quarterly Journal of Economics*, 133(3):1459–1501.

Becker, R., Gray, W., and Marvakov, J. (2013). Nber-ces manufacturing industry database: Technical notes. *NBER Working Paper*, 5809.

Bender, S., Bloom, N., Card, D., Van Reenen, J., and Wolter, S. (2018). Management practices, workforce selection, and productivity. *Journal of Labor Economics*, 36(S1):S371–S409.

Berman, E., Bound, J., and Griliches, Z. (1994). Changes in the demand for skilled labor within us manufacturing: evidence from the annual survey of manufactures. *The Quarterly Journal of Economics*, 109(2):367–397.

Bernard, A. B., Jensen, J. B., and Schott, P. K. (2006). Survival of the best fit: Exposure to low-wage countries and the (uneven) growth of U.S. manufacturing plants. *Journal of International Economics*, 68(1):219–237.

Bloom, N., Draca, M., and Van Reenen, J. (2016). Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *The Review of Economic Studies*, 83(1):87–117.

Borjas, G. J., Freeman, R. B., and Katz, L. F. (1992). On the labor market effects of immigration and trade. In Borjas, G. J. and Freeman, R. B., editors, *Immigration and the Workforce: Economic Consequences for the United States and Source Areas*, pages 213–244. University of Chicago Press, Chicago.

Borjas, G. J. and Ramey, V. A. (1995). Foreign competition, market power, and wage inequality. *Quarterly Journal of Economics*, 110(4):1075–1110.

Borusyak, K., Hull, P., and Jaravel, X. (2018). Quasi-experimental shift-share research designs. NBER Working Paper No. 24997.

Boskin, M. J. and Lau, L. J. (1990). Post-war economic growth in the Group-of-Five countries: A new analysis. NBER Working Paper No. 3521.

Bound, J. and Freeman, R. B. (1992). What went wrong? the erosion of relative earnings and employment among young black men in the 1980s. *Quarterly Journal of Economics*, 107(1):201–232.

Bound, J. and Holzer, H. J. (1993). Industrial shifts, skills levels, and the labor market for white and black males. *Review of Economics and Statistics*, pages 387–396.

Bound, J. and Holzer, H. J. (2000). Demand shifts, population adjustments, and labor market outcomes during the 1980s. *Journal of Labor Economics*, 18(1):20–54.

Boustan, L. P. (2009). Competition in the Promised Land: Black migration and racial wage gap convergance in the North: 1940-1970. *Journal of Economic History*, 69(3):756–783.

Brown, C. (1984). Black-white earnings ratios since the Civil Rights Act of 1964: The importance of labor market dropouts. *Quarterly Journal of Economics*, 99(1):31–44.

Burstein, A., Morales, E., and Vogel, J. (2019). Changes in between-group inequality: computers, occupations, and international trade. *American Economic Journal: Macroeconomics*, 11(2):348–400.

Burstein, A. and Vogel, J. (2017). International trade, technology, and the skill premium. *Journal of Political Economy*, 125(5):1356–1412.

Butler, R. and Heckman, J. J. (1977). The government's impact on the labor market status of black Americans: A critical review. In Hausman, L. J., Ashenfelter, O., Rustin, B., Schubert, R. F., and Slaiman, D., editors, *Equal Rights and Industrial Relations*, pages 235–281. Industrial Relations Research Association, Madison, WI.

Card, D. and DiNardo, J. E. (2002). Skill-biased technological change and rising wage inequality: Some problems and puzzles. *Journal of Labor Economics*, 20(4):733–783.

Card, D. and Krueger, A. B. (1992). School quality and black-white relative earnings: A direct assessment. *Quarterly Journal of Economics*, 107(1):151–200.

Chandra, A. (2003). Is the convergance of the racial wage gap illusory? NBER Working Paper No. 9476.

Charles, K. K., Hurst, E., and Schwartz, M. (2018). The transformation of manufacturing and the decline in U.S. employment. In Eichenbaum, M. and Parker, J. A., editors, *NBER Macroeconomics Annual 2018*, volume 33. University of Chicago Press, Chicago.

Charles, K. K., Hurst, E., and Schwartz, M. (2019). The transformation of manufacturing and the decline in us employment. *NBER Macroeconomics Annual*, 33(1):307–372.

Cogan, J. F. (1982). The decline in black teenage employment: 1960-1970. *American Economic Review*, 72(4):621–638.

Collins, W. J. and Margo, R. A. (2001). Race and home ownership: A century-long view. *Explorations in Economic History*, 38(1):68–92.

Crandall, R. W. (1984). Import quotas and the automobile industry: The costs of protectionism. *Brookings Review*, 2(4):8–16.

Cross, G. S., Proctor, R. N., and Proctor, R. (2014). *Packaged Pleasures: How Technology and Marketing Revolutionized Desire*. University of Chicago Press.

Delnevo, C. D. (2006). Smokers' choice: What explains the steady growth of cigar use in the us? *Public Health Reports*, 121.

Drope, J. and Schluger, N. W. (2018). *The Tobacco Atlas*. American Cancer Society.

Ebenstein, A., Harrison, A. E., McMillan, M. S., and Phillips, S. (2014). Estimating the impact of trade and offshoring on american workers using the current population surveys. *The Review of Economics and Statistics*, 96(4):581.

Einav, L., Finkelstein, A., and Gupta, A. (2017). Is american pet health care (also) uniquely inefficient? *American Economic Review*, 107(5):491–95.

Epifani, P. and Gancia, G. (2008). The skill bias of world trade. *Economic Journal*, 118(530):927–960.

Evans, W. N., Garthwaite, C., and Moore, T. J. (2016). The white/black educational gap, stalled progress, and the long-term consequences of the emergence of crack cocaine markets. *Review of Economics and Statistics*, 98(5):832–847.

Farber, H. S., Herbst, D., Kuziemko, I., and Naidu, S. (2018). Unions and inequality over the Twentieth Century: New evidence from survey data. NBER Working Paper No. 24587.

Federal Register (1996). Regulations restricting the sale and distribution of cigarettes and smokeless tobacco to protect children and adolescents: Final rule. Department of Health and Human Services, Food and Drug Administration. Fed. Reg. Vol. 61, No. 168. August 28, 1996. Available at https://www.govinfo.gov/content/pkg/FR-1996-08-28/pdf/X96-10828.pdf.

Feenstra, R. C. (1984). Voluntary export restraint in U.S. autos, 1980-81: Quality, employment, and welfare effects. In Baldwin, R. E. and Krueger, A. O., editors, *The Structure and Evolution of Recent U.S. Trade Policy*, pages 35–66. University of Chicago Press, Chicago.

Feenstra, R. C. (1988). Quality change under trade restraints in Japanese autos. *Quarterly Journal of Economics*, 103(1):131–146.

Feenstra, R. C. (1996). Us imports, 1972-1994: Data and concordances. Technical report, National Bureau of Economic Research.

Feenstra, R. C. (1997). Us exports, 1972-1994: With state exports and other us data. Technical report, National Bureau of Economic Research.

Feenstra, R. C. and Hanson, G. H. (1999). The impact of outsourcing and high-technology capital on wages: estimates for the united states, 1979–1990. *The Quarterly Journal of Economics*, 114(3):907–940.

Feler, L. and Sense, M. Z. (2017). Trade shocks and the provision of local public goods. *American Economic Journal: Economic Policy*, 9(4):101–143.

Finkelstein, A., Taubman, S., Wright, B., Bernstein, M., Gruber, J., Newhouse, J. P., Allen, H., Baicker, K., and Oregon Health Study Group (2012). The oregon health insurance experiment: Evidence from the first year. *Quarterly Journal of Economics*, 127(3):1057–1106.

Firpo, S., Fortin, N. M., and Lemieux, T. (2011). Occupational tasks and changes in the wage structure.

Flood, S., King, M., Rodgers, R., Ruggles, S., and Warren, J. R. (2018). Integrated public use microdata series, current population survey: Version 6.0 [dataset]. University of Minnesota.

Fort, T. C., Pierce, J. R., and Schott, P. K. (2018). New perspectives on the decline of u.s. manufacturing employment. *Journal of Economic Perspectives*, 32(2):47–72.

Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280.

Glaeser, E. L. and Gyourko, J. (2005). Urban decline and durable housing. *Journal of Political Economy*, 113(2):345–375.

Goldin, C. and Katz, L. F. (1998). The origins of technology-skill complementarity. *The Quarterly Journal of Economics*, 113(3):693–732.

Goldmsith-Pinkham, P., Sorkin, I., and Swift, H. (2018). Bartik instruments: What, when, why, and how. NBER Working Paper No. 24408.

Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in britain. *The Review of Economics and Statistics*, 89(1):118–133.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–26.

Gould, E. D. (2018). Torn apart? the impact of manufacturing employment on black and white Americans. IZA Discussion Paper No. 11614.

Greenland, A. and Lopresti, J. (2016). Import exposure and human capital adjustment: Evidence from the U.S. *Journal of International Economics*, 100:50–60.

Grossman, G. M. (1986). Imports as a cause of injury: The case of the U.S. steel industry. *Journal of International Economics*, 20(3–4):201–223.

Grossman, J. R. (1991). *Land of Hope: Chicago, Black Southerners and the Great Migration*. University of Chicago Press, Chicago.

Houseman, S. N. (2018). Understanding the decline of us manufacturing employment. *Upjohn Institute working paper, 18-287*.

Hummels, D., Jørgensen, R., Munch, J., and Xiang, C. (2014). The wage effects of offshoring: Evidence from Danish matched worker-firm data. *American Economic Review*, 104(6):1597–1629.

Ichniowski, C. and Shaw, K. (1999). The effects of human resource management systems on economic performance: An international comparison of U.S. and Japanese plants. *Management Science*, 45(5):704–721.

Irwin, D. A. (2017). *Clashing Over Commerce: A History of US Trade Policy*. University of Chicago Press, Chicago.

Jaeger, D. A., Ruist, J., and Stuhler, J. (2018). Shift-share instruments and the impact of immigration. NBER Working Paper No. 24285.

Johnson, G. and Stafford, F. (1998). Technology regimes and the distribution of real wages. *Microfoundations of Economic Growth: A Schumpeterian Perspective, University of Michigan Press, Ann Arbor.*

Joseph J. Doyle, J. (2011). Returns to local-area healthcare spending: Evidence from health shocks to patients far from home. *American Economic Journal. Applied Economics*, 3(3):221–243.

Juhn, C., Murphy, K. M., and Pierce, B. (1991). Accounting for the slowdown in black-white wage convergence. *Workers and Their Wages: Changing Patterns in the United States*, pages 107–43.

Juhn, C., Murphy, K. M., and Pierce, B. (1993). Wage inequality and the rise in returns to skill. *Journal of Political Economy*, 101(3):410–442.

Kain, J. F. (1968). Housing segregation, Negro employment, and metropolitan decentralization. *Quarterly Journal of Economics*, 82(2):175–197.

Kasarda, J. D. (1989). Urban industrial transition and the underclass. *Annals of the American Academy of Political and Social Science*, 501(1):26–47.

Katz, L. and Autor, D. (1999). Changes in the wage structure and earnings inequality. ashenfelter, o. and d. card (eds), handbook of labor economics., 3a: 1463–1555.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Katz, L. F. and Revenga, A. L. (1989). Changes in the structure of wages: The United States vs Japan. *Journal of the Japanese and International Economies*, 3(4):522–553.

Krugman, P. R. (2000). Technology, trade, and factor prices. *Journal of International Economics*, 50(1):51–71.

Krugman, P. R. (2008). Trade and wages, reconsidered. *Brookings Papers on Economic Activity*, pages 103–154.

Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V., and Violante, G. L. (2000). Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica*, 68(5):1029–1053.

Kwok, A. C., Semel, M. E., Lipsitz, S. R., Bader, A. M., Barnato, A. E., Gawande, A. A., and Jha, A. K. (2011). The intensity and variation of surgical care at the end of life: a retrospective cohort study. *The Lancet*, 378(9800):1408–1413.

Lang, K. and Lehmann, J.-Y. K. (2012). Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature*, 50(4):596–1006.

Lang, K. and Manove, M. (2011). Education and labor market discrimination. *American Economic Review*, 101(4):1467–96.

Lindquist, M. J. (2005). Capital–skill complementarity and inequality in sweden. *Scandinavian Journal of Economics*, 107(4):711–735.

Lubitz, J., Beebe, J., and Baker, C. (1995). Longevity and medicare expenditures. *New England Journal of Medicine*, 332(15):999–1003.

Manson, S., Schroeder, J., Van Riper, D., and Ruggles, S. (2017). Ipums national historical geographic information system: Version 12.0 [database]. Technical report, University of Minnesota, Minneapolis.

Mead, L. M. (1986). *Beyond Entitlement: The Social Obligations of Citizenship*. Free Press, New York.

Meier, B. (1998). Court rules fda lacks authority to limit tobacco. *The New York Times*. August 15, 1998.

Melberg, H. O. (2018). The economics of end-of-life spending. *Economics and Finance*.

Monte, F. (2011). Skill bias, trade, and wage dispersion. *Journal of International Economics*, 83(2):202–218.

Murphy, K. M. and Welch, F. (1991). The role of international trade in wage differentials. In Kosters, M. H., editor, *Workers and Their Wages: Changing Patterns in the United States*, pages 39–69. American Enterprise Institute, Washington, D.C.

Neal, D. (2004). The measured black-white wage gap among women is too small. *Joural of Political Economy*, 112(S1):S1–S28.

Neal, D. (2006). Why has black–white skill convergence stopped? In Hanushek, E. A. and Welch, F., editors, *Handbook of the Economics of Education*, volume 1, pages 511–576. Elsevier North-Holland, Oxford.

Neal, D. A. and Johnson, W. R. (1996). The role of premarket factors in black-white wage differences. *Journal of Political Economy*, 104(5):869–895.

Neary, J. P. (2002). Foreign competition and wage inequality. *Review of International Economics*, 10(4):680–693.

Newhouse, J. P. (1993). *Free for all? Lessons from the RAND health insurance experiment*. Cambridge, MA: Harvard University Press.

Ohta, M. and Griliches, Z. (1986). Automobile prices and quality: Did the gasoline price increases change consumer tastes in the U.S.? *Journal of Business & Economic Statistics*, 4(2):187–198.

Pierce, J. R. and Schott, P. K. (2016). The surprisingly swift decline of us manufacturing employment. *American Economic Review*, 106(7):1632–62.

Pierce, J. R. and Schott, P. K. (2018). Investment responses to trade liberalization: Evidence from us industries and establishments. *Journal of International Economics*.

Powell, T. C. (1995). Total quality management as competitive advantage: A review and empirical study. *Strategic Management Journal*, 16(1):15–37.

Reardon, E. (1997). Demand-side changes and the relative economic progress of black men: 1940-90. *Journal of Human Resources*, 32(1):69–97.

Revenga, A. L. (1992). Exporting jobs? the impact of import competition on employment and wages in U.S. manufacturing. *Quarterly Journal of Economics*, 107:255–284.

Rose, E. K. (2018). The rise and fall of female labor force participation during World War II in the United States. *Journal of Economic History*, 78(3):673–711.

Rosen, S. (1978). Substitution and division of labour. *Economica*, 45(179):235–250.

Ruggles, S., Genadek, K., Goeken, R., Grover, J., and Sobek, M. (2015). Integrated public use microdata series: Version 6.0 [dataset]. Technical report, University of Minnesota.

Schott, P. K. (2008). The relative sophistication of chinese exports. *Economic policy*, 23(53):6–49.

Scott, A. E. (1999). The continuing tobacco war: State and local tobacco control in washington. *Seattle UL Rev.*, 23:1097.

Skinner, J. S., Fisher, E. S., and Wennberg, J. (2005). The efficiency of medicare. In Wise, D. A., editor, *Analyses in the Economics of Aging*, chapter 4, pages 129–160. Chicago: University of Chicago Press.

Smith, J. P. and Welch, F. R. (1989). Black economic progress after Myrdal. *Journal of Economic Literature*, 27(2):519–564.

Sperry, R. G., Wilhoit, J. H., and Duncan, G. A. (2013). Development of a semi-automated tobacco stripping machine utilizing string trimmers. *Applied Engineering in Agriculture*, 29(2):171–178.

Thoenig, M. and Verdier, T. (2003). A theory of defensive skill-biased innovation and globalization. *American Economic Review*, 93(3):709–728.

Tolbert, C. M. and Sizer, M. (1996). U.S. commuting zones and labor market areas: A 1990 update. Economic Research Service Staff Paper 9614.

US Department of Health and Human Services (1998). Cigars: Health effects and trends. Smoking and Tobacco Control Monograph 9, NIH Publication No. 98-4302.

US Department of Health and Human Services (2014). The health consequences of smoking: 50 years of progress. a report of the surgeon general.

Westcott, D. N. and Bednarzik, R. W. (1981). Employment and unemployment: A report on 1980. *Monthly Labor Review*, 104:4.

Wilhoit, J. H., Duncan, G. A., and Sperry, R. G. (2013). Labor productivity with various innovations for stripping burley tobacco. *Tobacco Science*, pages 1–10.

Wilson, W. J. (1987). *The Truly Disadvantaged*. University of Chicago Press, Chicago.

Xiang, C. (2005). New goods and the relative demand for skilled labor. *Review of Economics and Statistics*, 87(2):285–298.

A. APPENDIX FOR: SKILL-BIASED TECHNICAL CHANGE AND EMPLOYMENT IN U.S. MANUFACTURING

A.1 Data Appendix

A.1.1 Census of Manufactures

Consistent Industry Sample

Data are reported by SIC 1987 codes for 1987 and 1992 and by NAICS codes beginning in 1997. I convert data to 1987 SIC codes as follows. First, I convert from NAICS 2012 to NAICS 2007 to NAICS 2002 to NAICS 1997 to SIC 1987 as necessary, using weights based on value of shipments for output and materials, and based on employment for employment, from Census Bridge publications.

The sample of SIC industries is determined by industries which have consistently available data on material use. The industries that need to be dropped because of this restriction are 2097, 2813, 3295, 3398, 2371, 2395, 2397, 2999, and 3399. Together these industries represent less than 0.5 percent of manufacturing employment in 1987. I also drop the six industries that exit manufacturing upon transition into the NAICS system, and adjust output, employment, and material use data for partially exiting industry 3732 by multiplying 1997 and later values by 1/(1-0.127) following Becker et al. (2013).

The sample of SIC industries is then slightly aggregated according to the industry aggregations in Autor et al. (2013). This is to facilitate analysis of the China shock in Section 1.6.1. In addition to these aggregations, I also combine 2067 with 2064 (chewing gum with other confectionery products), and 3292 with 3299 (asbestos products with nonmetallic mineral products, not elsewhere classified), as the Census of Manufactures does not separately report these industries beginning in the 1990s. The resulting sample of industries is 379. When computer-producing industries are dropped, the final sample is 351. In this paper I define computer industries following Acemoglu et al. (2016), which generally captures industries associated with NAICS code 334.

Consistent Materials Sample

SIC materials are aggregated following the same scheme as the industries, described in Appendix Section A.1.1, with the exception that adjustments are not made for exiting industries because both manufacturing and non-manufacturing materials may be used in production. Some materials, with codes starting with "19", do not directly map to industries and were hand-matched to SIC codes by product description.[1] My final sample of materials includes 415 unique product codes. After converting both industries and their materials use information, the resulting data set is a panel of SIC manufacturing industries and their expenditures on SIC products, which are largely manufacturing but also include agriculture, forestry, and other industry products.

Output

I interpret value of shipments as output. Output is deflated to real values using industry-specific deflators from the NBER-CES Manufacturing Database (Becker et al., 2013). As this database only goes through 2011, I estimate 2012 prices as the 2007 price plus 1.25 times the difference between 2011 and 2007 prices.

Materials Use Imputations

According to the 1992 CoM, information was collected from surveyed establishments for those materials which were important parts of the cost of production in a particular industry and for which cost information was available from manufacturers' records. Material expenditures are reported by detailed industry code (SIC based prior to 1997 and NAICS based for 1997 and later years), for those materials which reporting establishments consumed at or above a specialized threshold, usually $25,000. For expenditures falling below that threshold, the materials are not separately classified. Also, the cost of materials for certain small establishments are not separately specified by product. For these reasons, and due to occasional grouping together of material expenditures to avoid disclosure of information on particular companies, some imputations are required. I describe these imputations here.

[1] These matches are available on request. Likewise, a small number of products are reported in very general 2-digit codes. I also hand-matched these to more detailed industries according to product description. A list of these assignments is also available upon request.

I use data from the 1992 CoM, which includes materials use for 1987 and 1992; the 2002 CoM, which includes materials use for 1997 and 2002; and the 2007 and 2012 CoM series. I begin by preserving the detailed industry codes as reported, combining some codes that only appear once or a few times and for which the expenditure value is hidden.[2] If the cost for a particular material is reported for one year but not the other, between 1987 and 1992, or between 1997 and 2002, I impute the missing value as the share of total materials expenditure reported in the non-missing year multiplied by the materials expenditure of the missing year.[3] If this exercise leaves only one missing value for a particular industry, I can now impute that missing value as the residual of the industry's total materials expenditure in that year. Next, I impute the remaining missing values by assigning the value of the average expenditure share of that material by the other industries in that year multiplied by the industry's total materials expenditure. Finally, I readjust imputed values so that the total cost of the materials sum to the reported total.

The imputed data still generally includes two non-specified categories, 970099 and 971000, representing materials that did not meet the minimum expenditure threshold or materials from small establishments not reporting this data. After converting both the industries and the materials into time-consistent SIC sample, described in Appendix Sections A.1.1 and A.1.1, I distribute the values in the non-specified categories across the specified materials, assuming the same distribution of expenditures within the industry.

A.1.2 Labor Hours and Wages By Skill Type

Employment, Hours, and Wages by CZ and Census Industry

In Section 1.3.2 I describe my procedure for calculating high- and low-skill hours by industry, and mean wages by skill type by industry. This procedure requires measures of skilled share of employment by Census industry by commuting zone, and mean hours and hourly wages by skill type by Census industry by commuting zone. I describe these measures here.

I use the 5 percent Census samples for 1980, 1990, and 2000. I use the 2005-2007 ACS for the year 2007 and the 2008-2012 ACS for 2010. I connect these to commuting

[2] These aggregations are available on request.

[3] As 2007 and 2012 are reported separately, and there are some changes to industry codes between these series, I cannot cross-impute in this way, so I skip this step.

zones by Census County Groups for 1980 and 1990 following Autor and Dorn (2013) and using Census Public Use Micro Areas for 2000 and the ACS samples following Autor and Dorn (2013) and Autor et al. (2018). Following these papers, I omit Hawaii and Alaska, leaving 722 CZs which comprehensively cover the continuous 48 states. I slightly aggregate 1990 Census industry codes so that they are a balanced panel across all years of data.[4]

For each commuting zone and Census industry code, I determine high-skill share of employment, mean hours worked by skill type, and mean hourly wages by skill type.[5] My sample includes employed private wage and salary workers (class of worker codes 22 and 23) ages 16 to 64 who report working at least 50 weeks in the previous year and who report non-zero usual hours worked per week. I replace topcoded annual wage and salary income with 1.5 times the topcoded value. I define high-skill as workers with at least four years of college education.

Annual hours are defined as weeks worked last year multiplied by usual hours worked per week. Since weeks worked last year are only available in interval categories for the 2007-2012 ACS, I impute weeks worked as the mean weeks worked for the same Census industry and skill type observed in the 2005-2007 ACS. Hourly wages are defined as annual wage and salary income divided by the product of weeks worked last year and usual hours worked per week. For both annual hours and hourly wages, I replace values falling below the first percentile in that year with the value at the first percentile in that year, and values falling above the ninety-ninth percentile in that year with the value at the ninety-ninth percentile in that year.

In the event that, in a CZ, the CBP reports presence of an industry but the Census does not, I substitute state or region level variables as needed.

Employment by CZ and SIC Industry

I calculate employment shares by CZ and SIC industry in the 1980, 1990, 2000, 2007, and 2010 CBP. The CBP is an annual series that provides county-level economic data by industry, including the number of establishments, employment during the week of March 12, and payroll information extracted from the U.S. Census Bureau's Business Register. The 1980 series is reported in SIC 1972 codes, which I convert to SIC 1987 codes using an employment-based weighted crosswalk from the NBER-CES

[4]These aggregations are available upon request.

[5]Means by CZ are weighted by Census sample weights multiplied by commuting zone weights.

Manufacturing Industry Database. The CBP suppresses the employment counts for some counties to avoid identifying individual employers. I impute employment in these instances based on establishment counts following Autor et al. (2013), since establishment counts are always available even when employment counts are suppressed. In this procedure I multiply the number of establishments in each bracket by the average firm size in that bracket that can be observed in the CBP.

A.2 Empirical Appendix

A.2.1 Related Manufacturing Trends

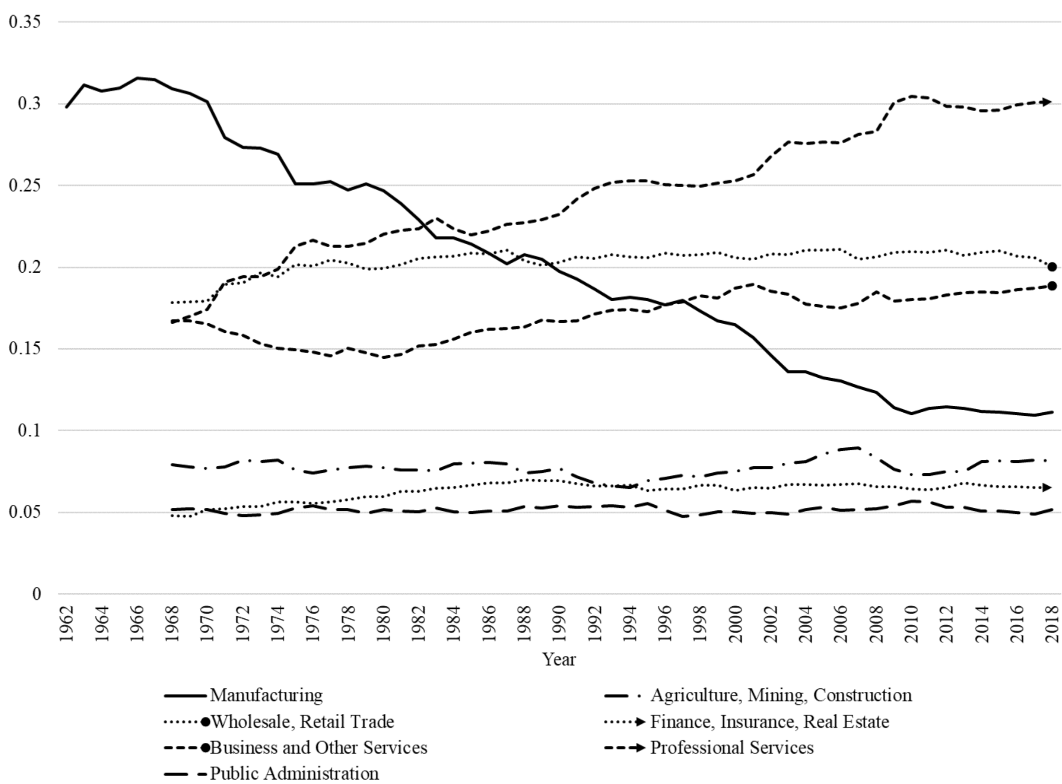

Figure B.1.: Share of Employment in Manufacturing and All Other Sectors

Notes - Author's calculations from the CPS ASEC annual surveys 1962-2018. Non-manufacturing sectors begin in 1968 when time-consistent industry codes are introduced. Sample is employed wage and salary workers ages 16-64, exclusive of self-employed, unpaid family workers, and military workers.

Source - Flood et al. (2018)

Figure B.2.: Real Value Added in Manufacturing and Non-Manufacturing

Notes - Annual manufacturing real value added is from the NBER-CES Manufacturing Industry Database for the years 1962-2011. I deflate value added by the variable PISHIP and then convert it to 2018 dollars following the PCE index (a divisor of .746225). Annual real GDP data are from the U.S. Bureau of Economic Analysis in 2012 dollars which I convert to 2018 dollars also following the PCE index (a divisor of .9257773). Non-manufacturing value added is real GDP less manufacturing real value added. Scale is log USD.
Source - U.S. Bureau of Economic Analysis; NBER CES Manufacturing Database (Becker et al., 2013)

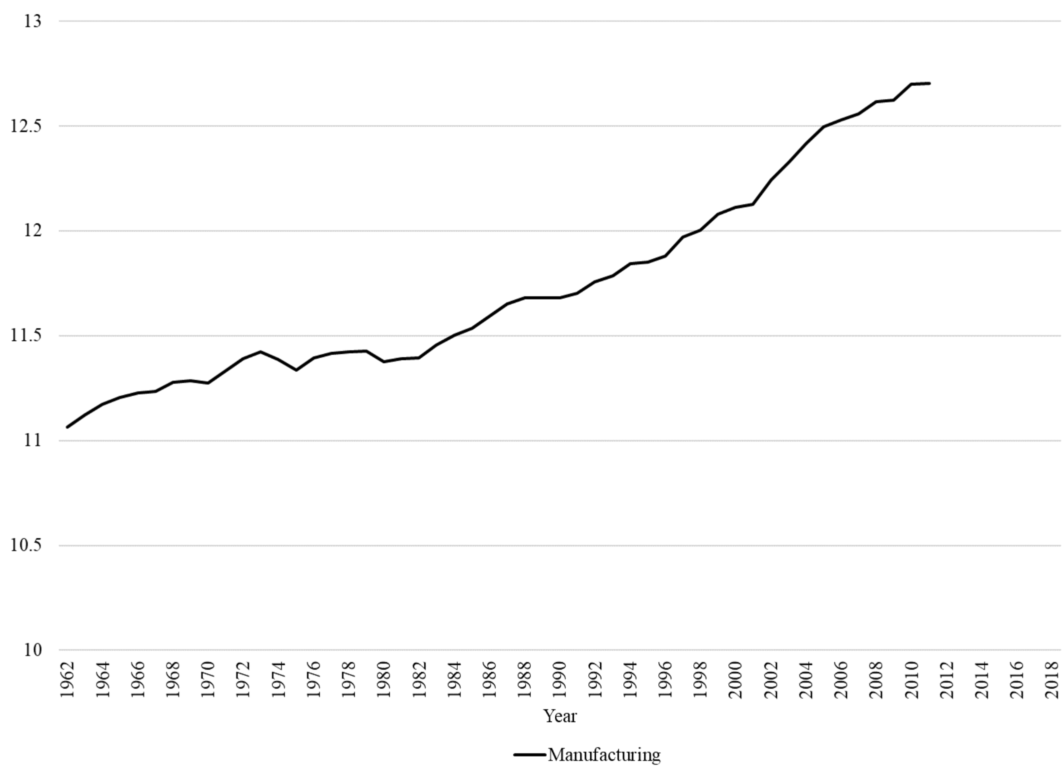

Figure B.3.: Log Real Value Added Per Worker in Manufacturing

Notes - Annual manufacturing real value added and manufacturing employment are from the NBER-CES Manufacturing Industry Database for the years 1962-2011. I deflate value added by the variable PISHIP and then convert it to 2018 dollars following the PCE index (a divisor of .746225). Annual real GDP data are from the U.S. Bureau of Economic Analysis in 2012 dollars which I convert to 2018 dollars also following the PCE index (a divisor of .9257773). Non-manufacturing value added is real GDP less manufacturing real value added. Scale is log USD per worker.
Source - U.S. Bureau of Economic Analysis; NBER CES Manufacturing Database (Becker et al., 2013)

Figure B.4.: Share of Employment with College Degree

Notes - Author's calculations from the CPS ASEC annual surveys 1962-2018. Sample is employed wage and salary workers ages 16-64, exclusive of self-employed, unpaid family workers, and military workers. Prior to 1992, individuals reporting at least 4 years of college are considered bachelor's degree holders. While this figure shows share of workers with college degrees, similar trends are seen for workers with any college education or with post-graduate education. In all cases the manufacturing workforce is becoming more highly skilled, and these trends closely track the remainder of the workforce.

Source - Flood et al. (2018)

A.2.2 Assigning Production Materials to Low- and High-Skill Processes

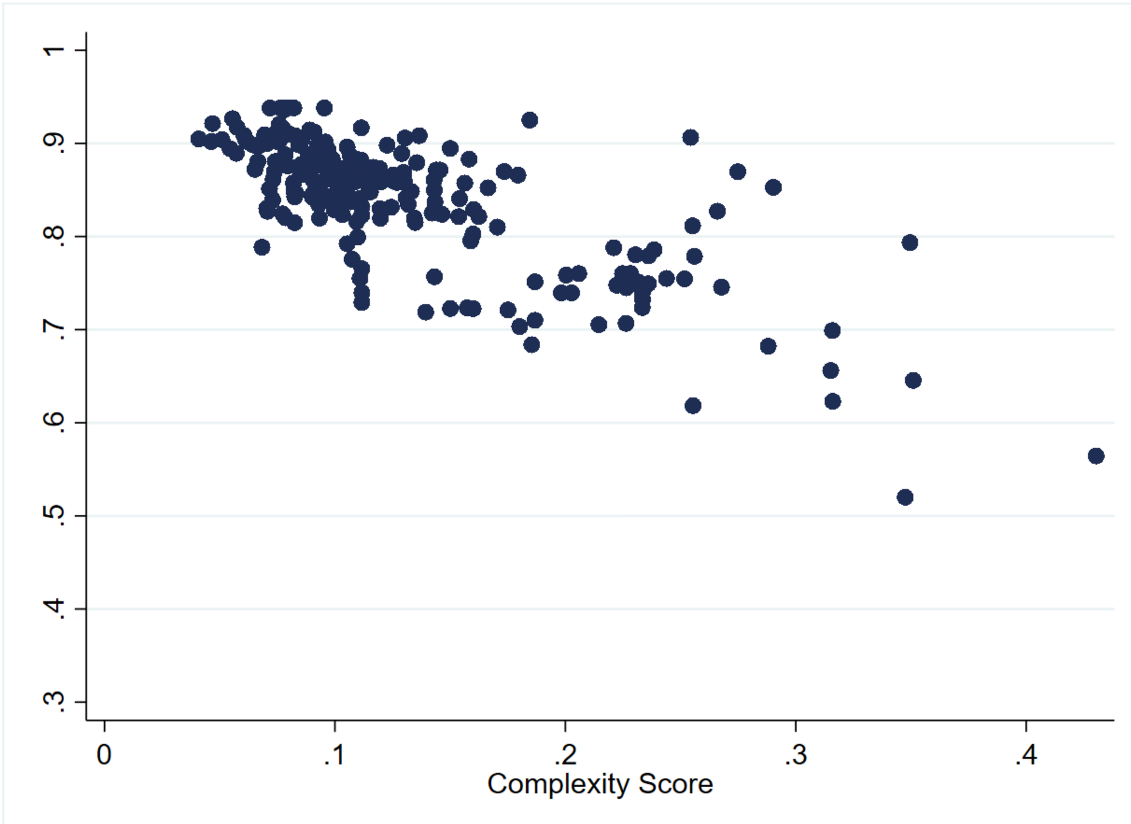

Figure B.5.: Relationship Between Resource Complexity Score and Material Use Prediction of Industry Low-Skill Share in 1990 ($\widehat{\lambda}_{j,1990}$)

Notes - I use spending on production resources to estimate the task share parameter $\alpha_{i,t}$ according to each material's prediction of skilled labor share. This figure shows a negative correlation between a resource's complexity and the share of low-skill workers in the industries that use it. See Section 1.4.1 for details.

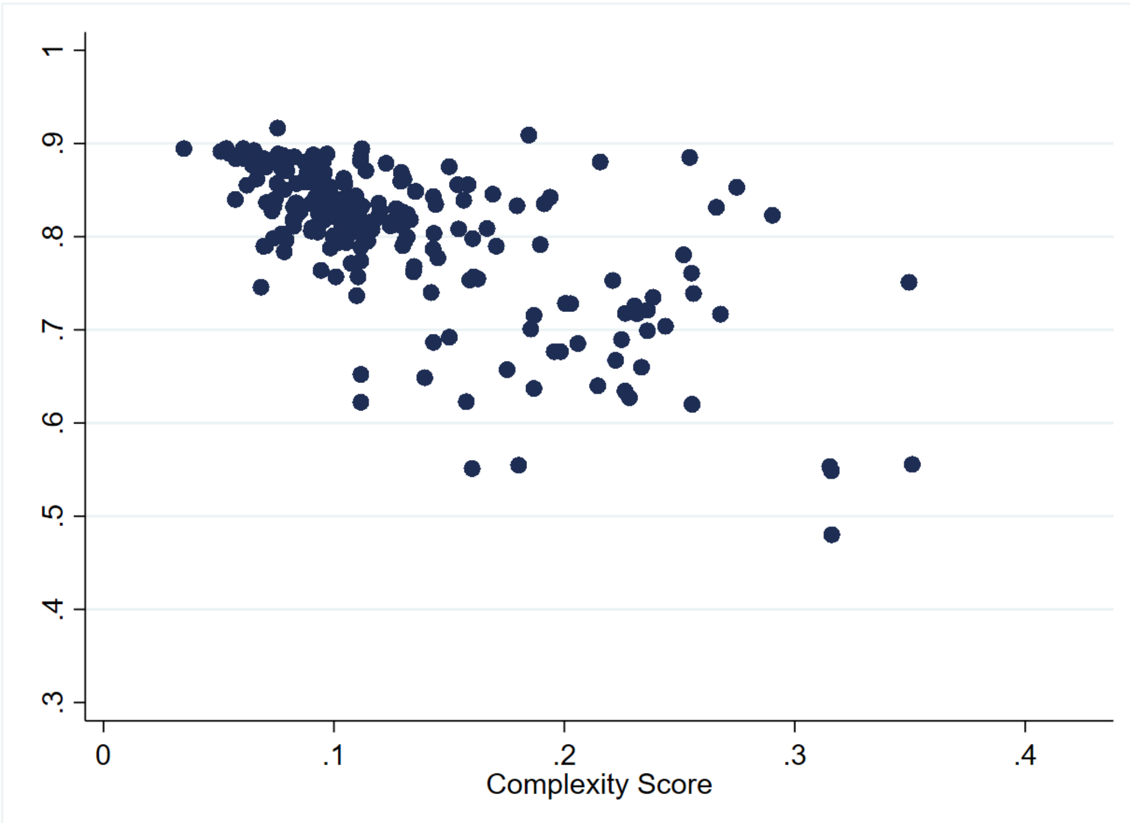

Figure B.6.: Relationship Between Material Complexity Score and Material Use Prediction of Industry Low-Skill Share in 2000 ($\widehat{\lambda}_{j,2000}$)

Notes - I use spending on production materials to estimate task share parameter $\alpha_{i,t}$ according to each material's prediction of skilled labor share. See Section 1.4.1 for details.

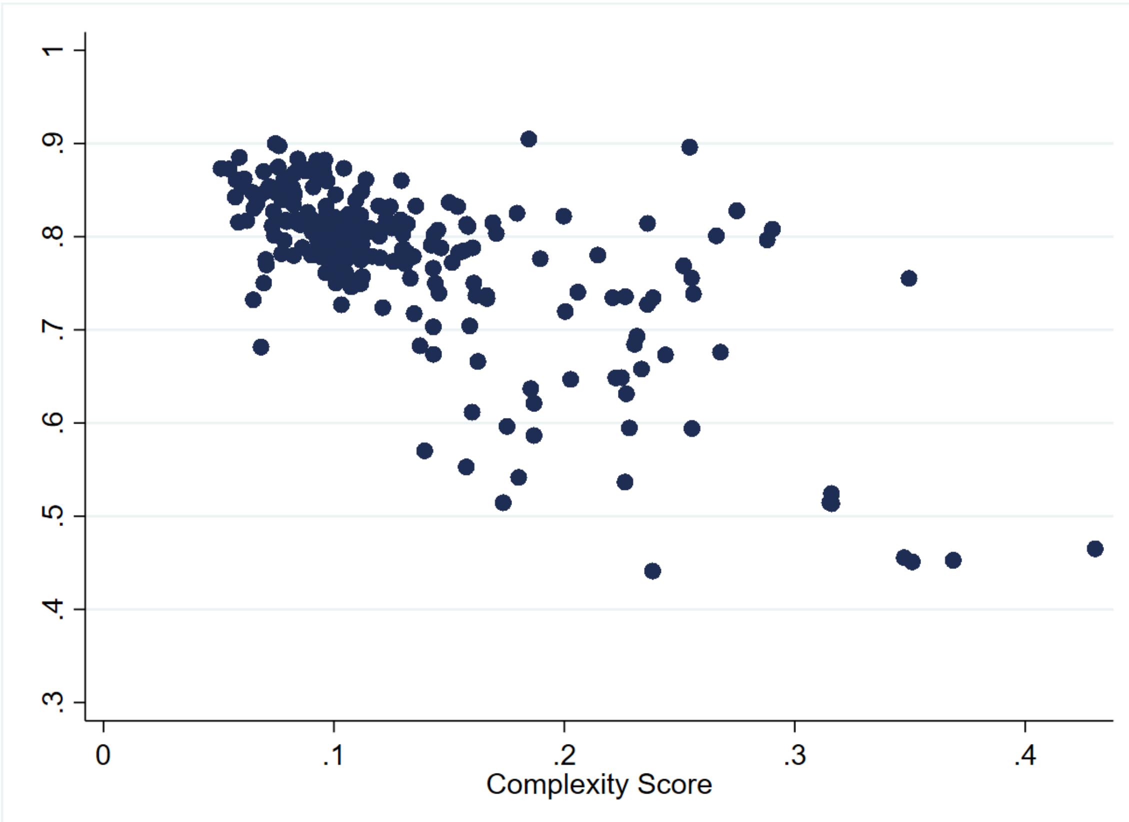

Figure B.7.: Relationship Between Material Complexity Score and Material Use Prediction of Industry Low-Skill Share in 2007 ($\widehat{\lambda}_{j,2007}$)

Notes - I use spending on production materials to estimate task share parameter $\alpha_{i,t}$ according to each material's prediction of skilled labor share. See Section 1.4.1 for details.

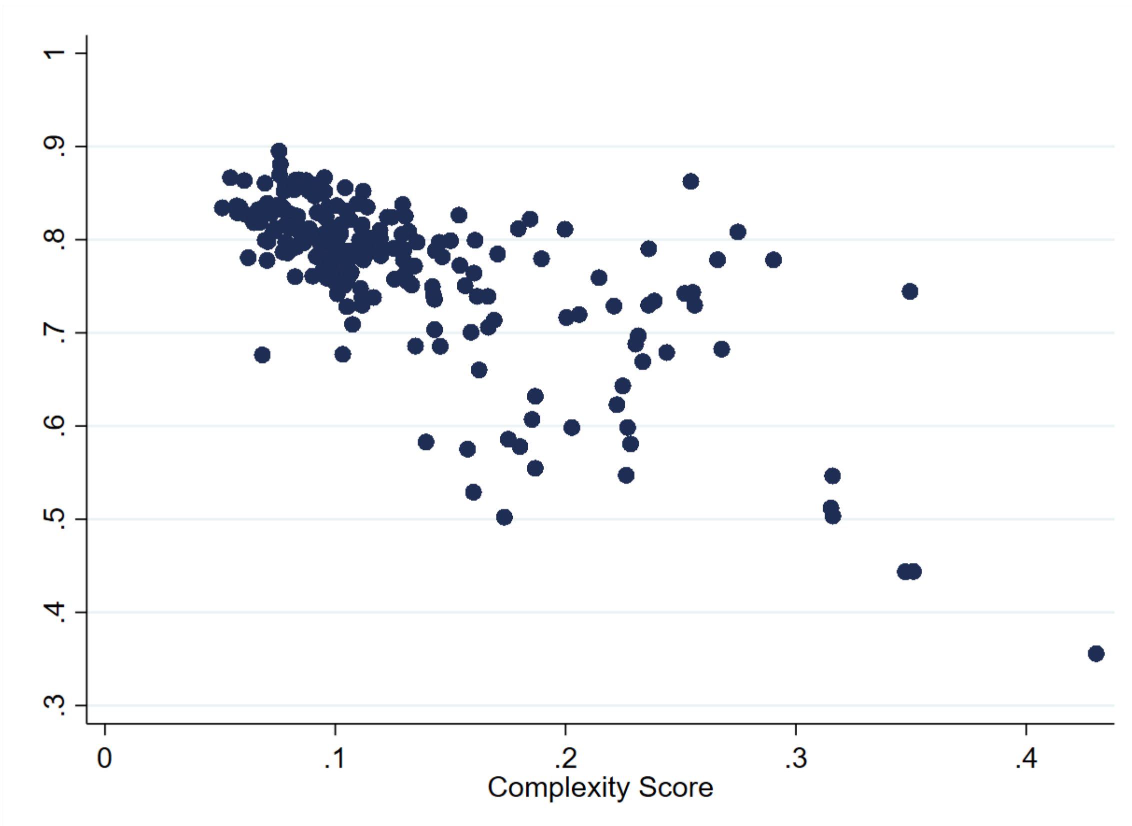

Figure B.8.: Relationship Between Material Complexity Score and Material Use Prediction of Industry Low-Skill Share in 2010($\widehat{\lambda}_{j,2010}$)

Notes - I use spending on production materials to estimate task share parameter $\alpha_{i,t}$ according to each material's prediction of skilled labor share. See Section 1.4.1 for details.

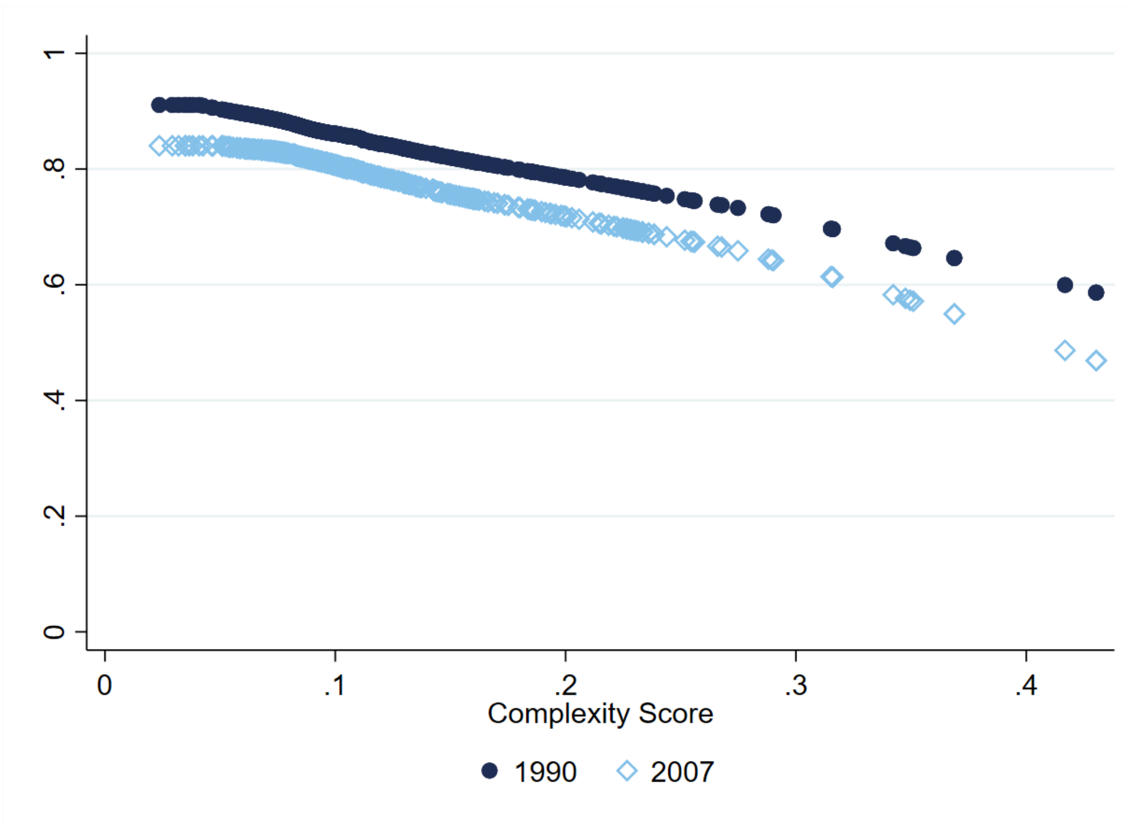

Figure B.9.: Material Complexity Score and Share Allocated to Low-Skill Process ($\widetilde{\lambda}_{j,t}$), 1990 and 2007

Notes - I use spending on production materials to estimate task share parameter $\alpha_{i,t}$ according to each material's prediction of skilled labor share. See Section 1.4.1 for details.

A.2.3 Distributions of Labor-Augmenting Technology Parameters $\widehat{a}_{i,t}$ and $\widehat{b}_{i,t}$ by Year

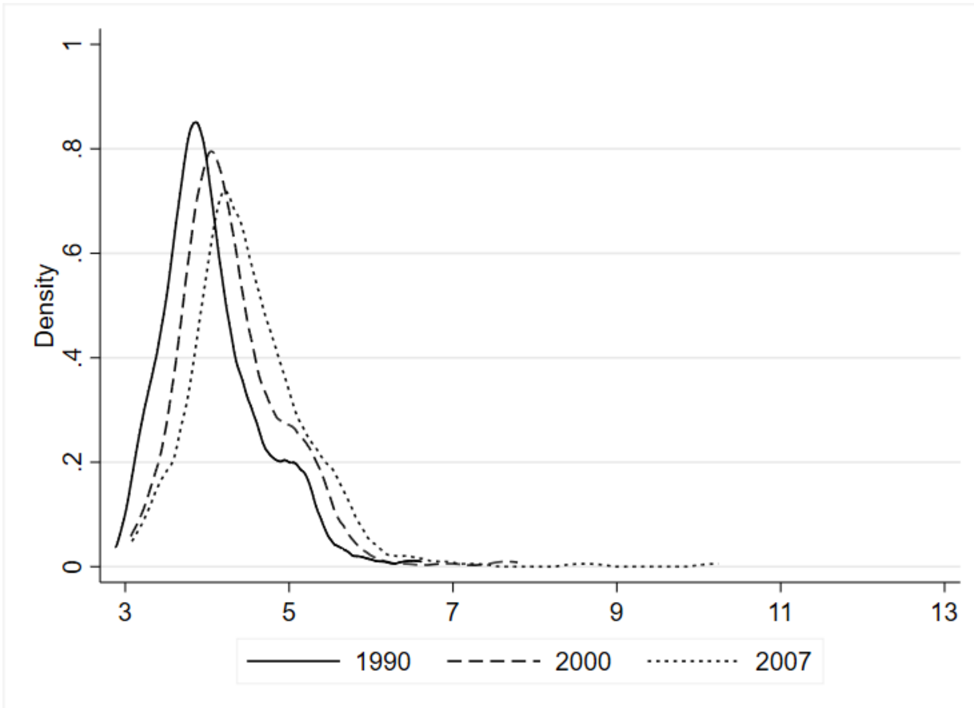

Figure B.10.: Kernel Density Distribution of Log($\widehat{a}_{i,t}$)

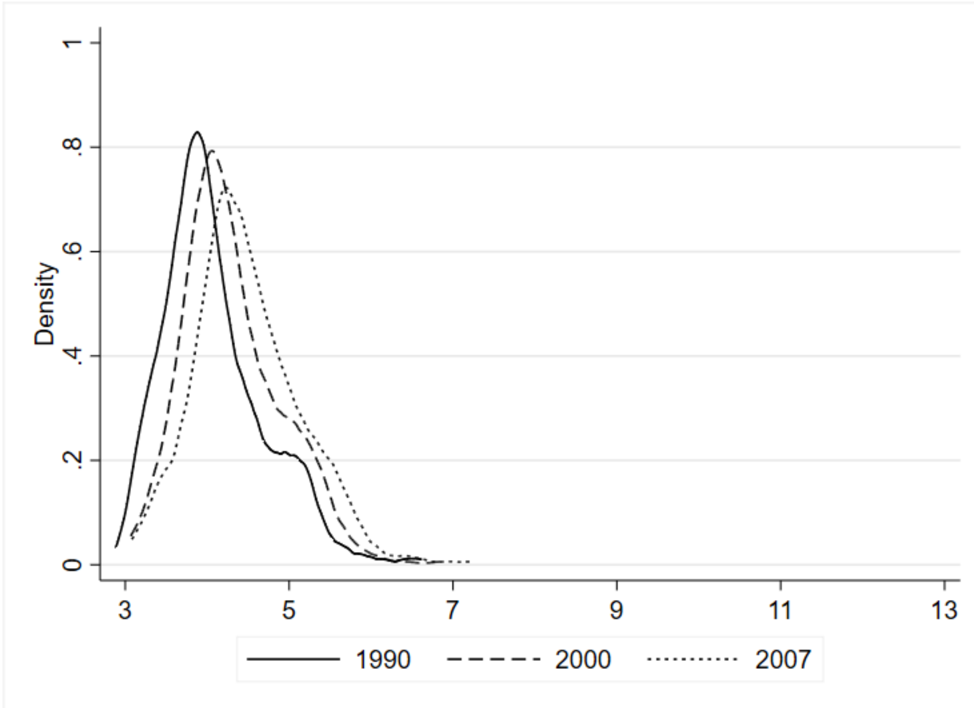

Figure B.11.: Kernel Density Distribution of Log($\widehat{a}_{i,t}$), Omitting Computer Industries

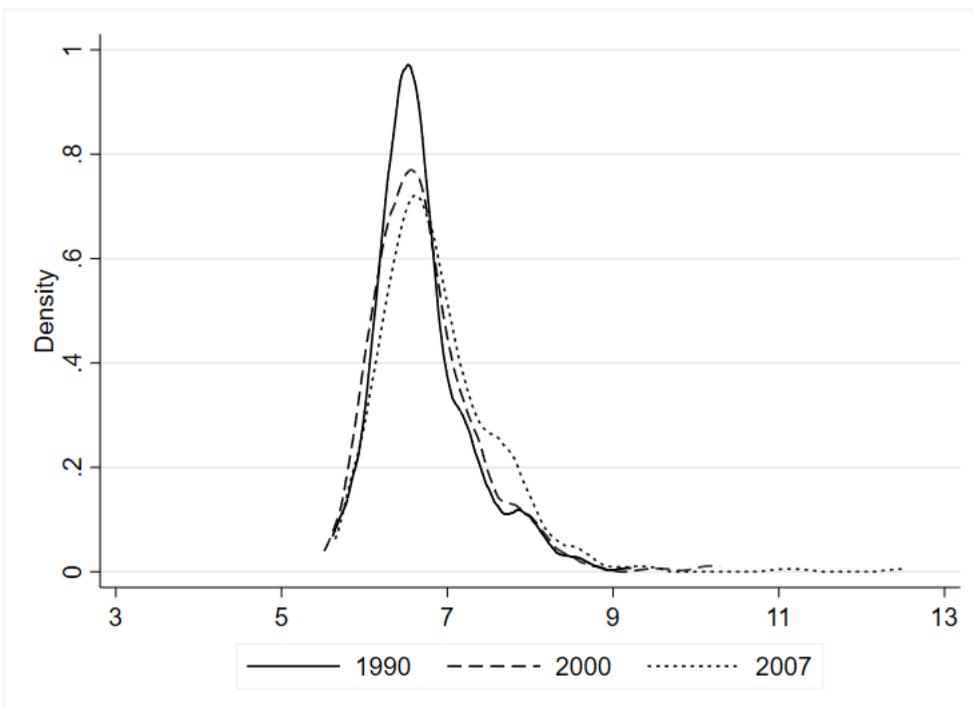

Figure B.12.: Kernel Density Distribution of Log($\widehat{b}_{i,t}$)

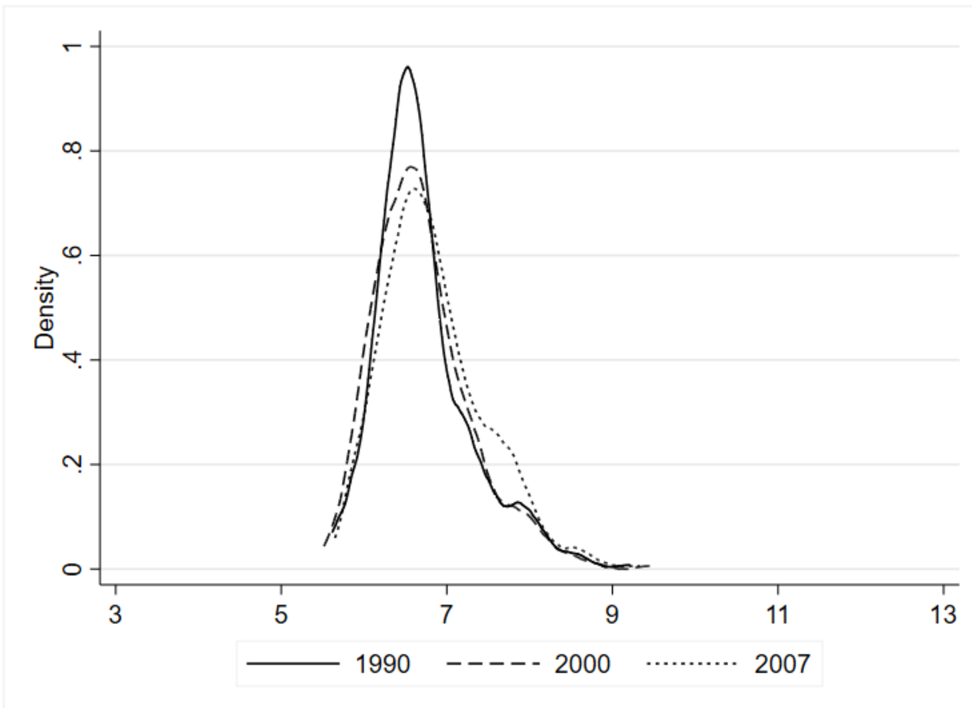

Figure B.13.: Kernel Density Distribution of Log($\widehat{b}_{i,t}$), Omitting Computer Industries

A.2.4 Decomposition Using Alternative Elasticity Estimates

Table B.1.: Decomposition of Manufacturing Employment Changes 1990-2007 Into Four Channels, Using Elasticity of 2.0

	Jobs (Millions)			
	Low-Skill (1)	High-Skill (2)	Total (3)	High-Skill Share (4)
	Panel A. Overall Change			
1990	12.83	2.17	15.00	14.5%
2007	10.00	2.41	12.41	19.4%
Δ	−2.84	+0.24	−2.60	+4.9pp
%Δ	−22.1%	+10.9%	−17.3%	
	Panel B. Decomposition			
Scale	+2.74	+0.62	+3.36	+0.7pp
	+21.3%	+28.7%	+22.4%	
Task	−3.33	+0.93	−2.40	+10.1pp
	−25.9%	+42.7%	−16.0%	
Productivity	−2.90	−1.01	−3.91	−4.0pp
	−22.6%	−46.4%	−26.1%	
Supply	+0.66	−0.31	+0.35	−2.3pp
	+5.1%	−14.1%	+2.3%	

Notes - Decomposition is calculated for 351 manufacturing industries and then summed to national totals. Decomposed changes in Columns (1) through (3) may not exactly sum to total due to rounding. Employment results are converted from annual hours to millions of estimated jobs based on mean annual hours of employed manufacturing workers of the same skill type in the 1980 Census.

Table B.2.: Decomposition of Manufacturing Employment Changes 1990-2007 Into Four Channels, Using Elasticity of 1.5

	Jobs (Millions)			
	Low-Skill (1)	High-Skill (2)	Total (3)	High-Skill Share (4)
	Panel A. Overall Change			
1990	12.83	2.17	15.00	14.5%
2007	10.00	2.41	12.41	19.4%
Δ	−2.84	+0.24	−2.60	+4.9pp
$\%\Delta$	−22.1%	+10.9%	−17.3%	
	Panel B. Decomposition			
Scale	+2.76	+0.60	+3.36	+0.6pp
	+21.5%	+27.6%	+22.4%	
Task	−3.15	+0.49	−2.66	+7.1pp
	−24.5%	+22.3%	−17.7%	
Productivity	−2.95	−0.62	−3.57	−0.9pp
	−23.0%	−28.7%	−23.8%	
Supply	+0.50	−0.22	+0.27	−1.7pp
	+3.9%	−10.3%	+1.8%	

Notes - Decomposition is calculated for 351 manufacturing industries and then summed to national totals. Decomposed changes in Columns (1) through (3) may not exactly sum to total due to rounding. Employment results are converted from annual hours to millions of estimated jobs based on mean annual hours of employed manufacturing workers of the same skill type in the 1980 Census.

A.2.5 Decomposition Figures

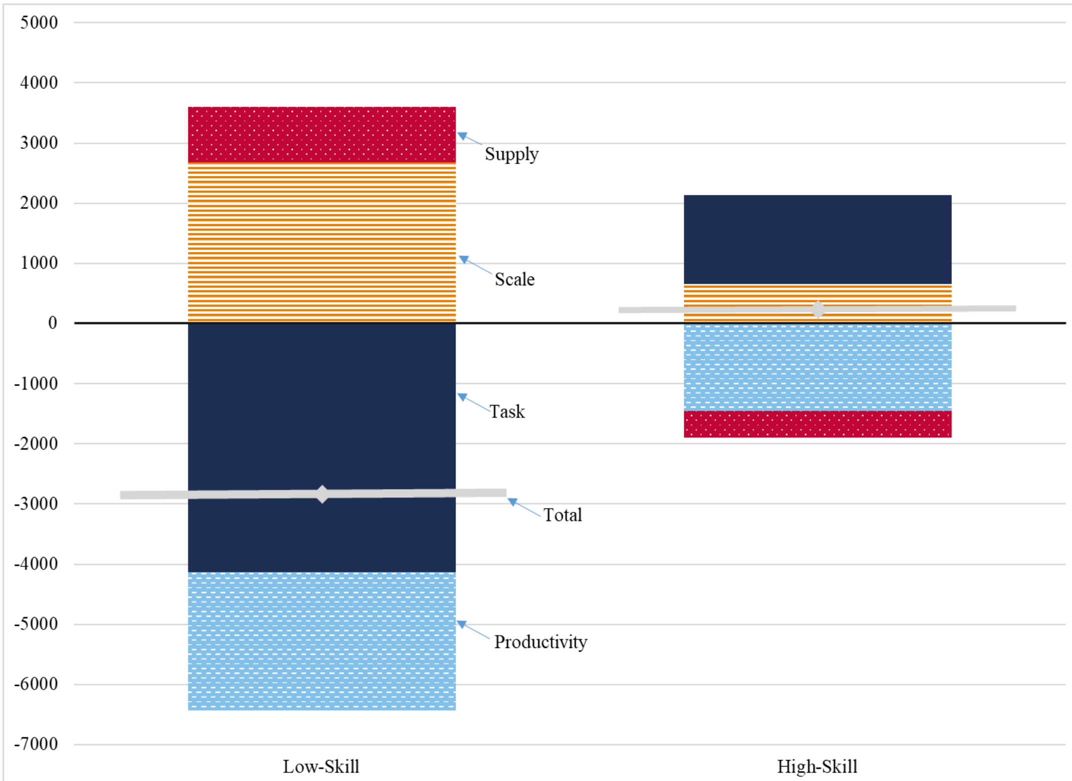

Figure B.14.: Decomposition of Manufacturing Employment Changes, 1990-2007, in Thousands of Job Equivalents

Notes - This figure shows the main decomposition results aggregated over all industries, and corresponds to Table 1.2. See Section 1.5.2 for details.

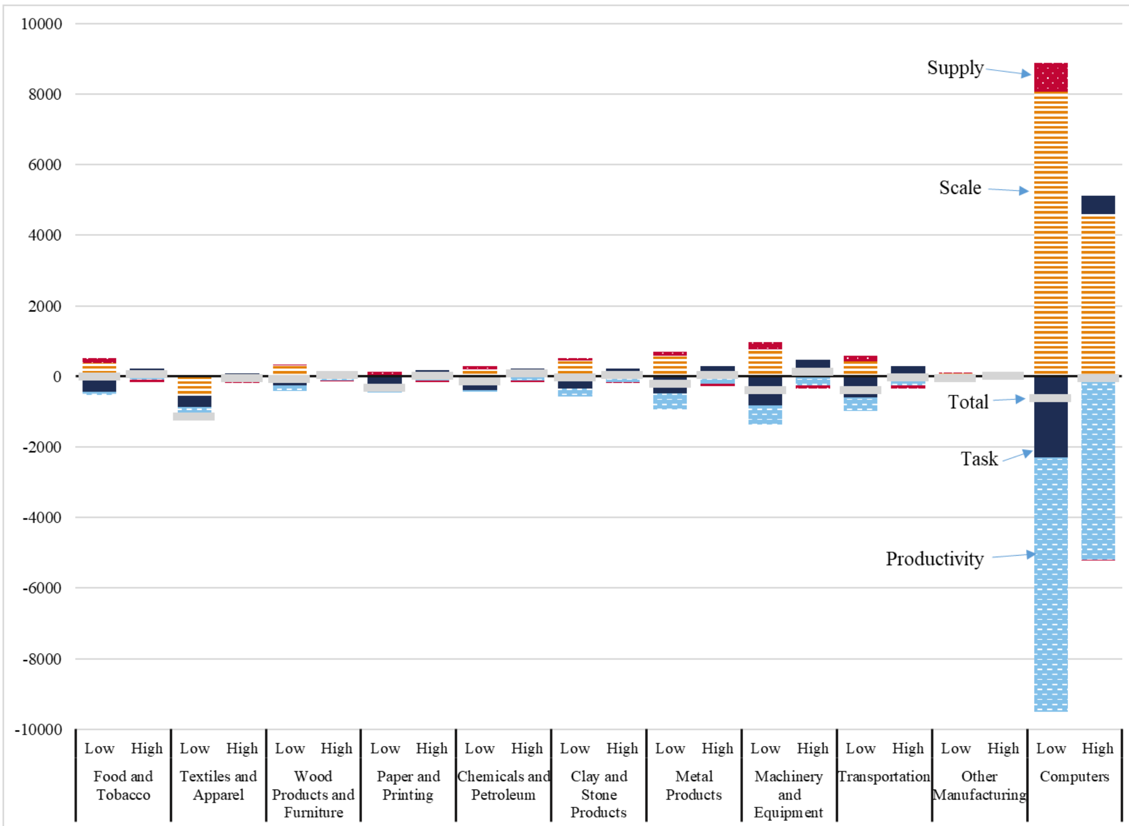

Figure B.15.: Decomposition of Industry Manufacturing Employment Changes in Thousands of Jobs, 1990-2007, by Industry Group, with Computer Industries as Separate Group

Notes - This figure shows the main decomposition results aggregated to 10 industry groups, with the computer industries included as a separate category. Because of the dramatic changes in the price and quality of computers during the 1990s and 2000s, computer industries are omitted from most of the analysis. I define computer industries following Acemoglu et al. (2016). See Section 1.5.2 for an analysis of the main results.

A.2.6 Tobacco Industries

Substantial heterogeneity among the tobacco industries makes them an interesting case study. Although these industries do not represent a large share of employment, a closer look at individual industries sheds light on the mechanics of the decomposition.

Figure A.2.6 shows the decomposition for the four industries in this subgroup. First to note are the scale-induced employment losses in both cigarette manufacturing and in tobacco stemming and redrying, an upstream industry that separates the tobacco leaf from its stem in preparation for further processing. These industries experienced little to no imports from China during this time period, making import competition an unlikely cause of the scale losses. The phenomenon may instead be demand driven as output declines coincide with a downward trend in cigarette consumption in the U.S. and other developed countries (Drope and Schluger, 2018; US Department of Health and Human Services, 2014). The 1990s was a time of many law suits against the cigarette industry for its hitherto denial that nicotine was addictive (Scott, 1999). It was also a time of aggressive public health campaigns against smoking, and the onset of a series of federal regulations on the sales, advertising, and manufacturing of cigarettes.[6] Scale losses may also link to supply factors as a federal price support and quota program for tobacco farming ended in 2004.[7]

While both cigarettes and tobacco stemming and redrying exhibit scale declines, only the cigarette industry shows additional displacement from the productivity channel. Technology continued to advance during the 1990s and 2000s, and cigarette rolling and packaging has become highly automated.[8] Cigarette rolling machines doubled in efficiency between 1988 and 2006, from 10,000 to 20,000 cigarettes per minute (Cross et al., 2014). The low-skill loss from the productivity channel is consistent with this automation. Finally, this industry also exhibits task shifts. These shifts are consistent with increased need for high-skill workers for supervision and quality control, possibly in response to heightened scrutiny by federal agencies.

[6] In 1992, Congressional action prompted all states to increase their minimum legal age for smoking to at least 18 years by the following year (Apollonio and Glantz, 2016). In 1996, the Food and Drug Administration established its authority to regulate the industry (Federal Register, 1996), provoking subsequent litigation by the tobacco industry (Meier, 1998). The American Legacy Foundation (later renamed the Truth Initiative) was established in 1999 and began a nationwide campaign targeting teen smoking. See www.truthinitiave.org.

[7] This is the Fair and Equitable Tobacco Reform Act of 2004, also known as the Tobacco Buyout.

[8] Philip Morris International describes its cigarette production process on its website at https://www.pmi.com/our-business/about-us/products/how-cigarettes-are-made.

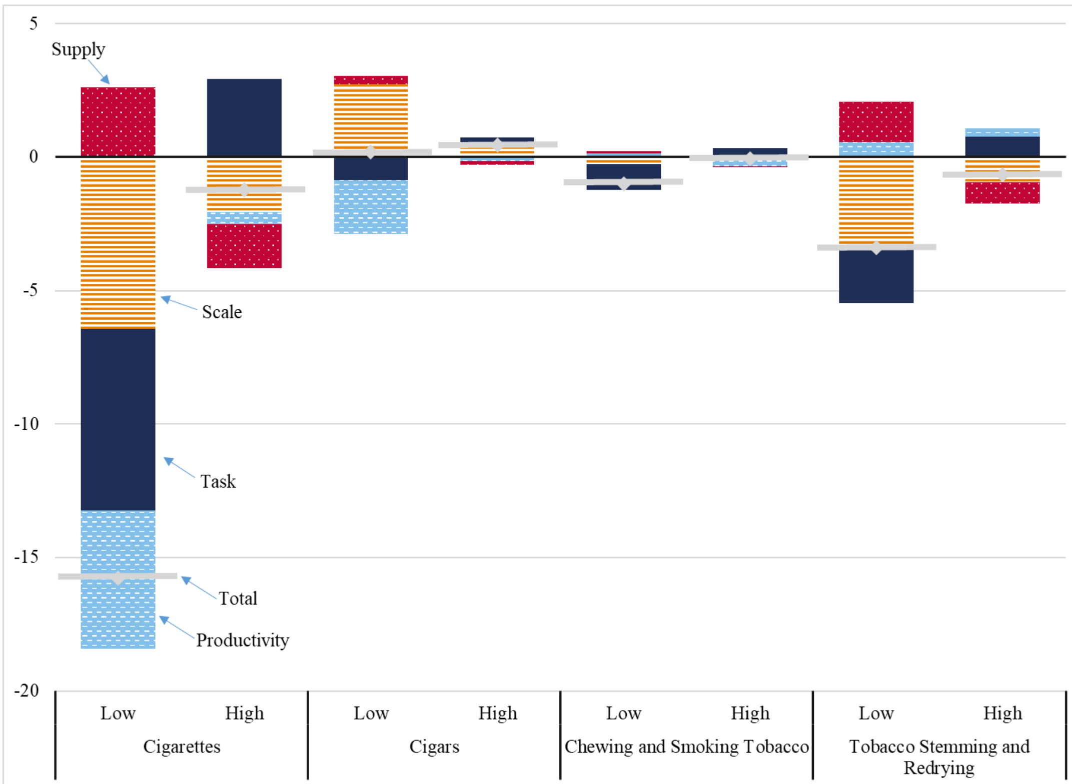

Figure B.16.: Decomposition of Employment Changes in Tobacco Manufacturing Industries, 1990-2007, in Thousands of Job Equivalents

Notes - This figure shows the main decomposition results for the four tobacco manufacturing industries. See Section A.2.6 for details.

In the tobacco stemming and redrying industry there is likewise some task shifting, but a notable lack of advancement in labor-augmenting technology. Tobacco stemming, in which the leaf is separated from the stem before it is aged, flavored, and rolled into cigarettes, is a notoriously labor-intensive process. Despite ongoing efforts to mechanize, it is largely done by hand even today (Wilhoit et al., 2013; Sperry et al., 2013).

An interesting contrast to both these industries is the scale *increases* in the cigar industry. U.S. cigar consumption had been on a downward trend for decades, like cigarettes, until 1993 when it suddenly pivoted upward (US Department of Health and Human Services, 2014).[9] This has been attributed in part to marketing including the use of cigars by celebrities (US Department of Health and Human Services, 1998; Delnevo, 2006).[10] The difference in excise taxes between cigars and cigarettes may also have lead to substitution as cigarette consumption continued to decline (Delnevo, 2006; US Department of Health and Human Services, 2014). Scale gains for low-skill workers are offset by productivity displacement, consistent with automation.

[9] According to this report, cigar consumption tripled over the next two decades.

[10] The popular magazine Cigar Aficionado, often featuring celebrities with cigars, began publishing in September 1992.

A.2.7 Applications

China Shock Application Data

In this application I use bilateral trade data from the UN Comtrade Database for the years 1991, 2000, and 2007 for my import penetration variables described in Section 1.6.1.[11] From this database I use imports from China as reported by the U.S. For the associated instrumental variable I use imports from China as reported by Australia, Denmark, Finland, Germany, Japan, New Zealand, Spain, and Switzerland, following Autor et al. (2013). These data are converted from the Harmonized System to the slightly aggregated SIC industries used in Autor et al. (2013) with the modifications described in Appendix Section A.1.1. For U.S. consumption by industry in 1988 and 1991, which are also components of the import penetration variables, I use U.S. total imports and exports by SIC made available by Peter K. Schott (Schott, 2008), and value of shipments from NBER-CES Manufacturing Database (Becker et al., 2013). These data are brought to 2018 USD using the PCE index.

China Shock Application: Import Penetration and Its Instrument

To determine an industry's exposure to Chinese import competition, I follow Acemoglu et al. (2016) to define import penetration in industry i in time t as

$$\Delta IP_{i,t} = \frac{\Delta M_{i,t}^{UC}}{Y_{i,91} + M_{i,91} - E_{i,91}} \tag{A.1}$$

where $\Delta M_{i,t}^{UC}$ is the change in the value of imports from China to the U.S. The denominator is U.S. consumption of industry j, which is output plus imports less exports, calculated in base year 1991. I bring all values to 2018 USD using the PCE index.

As rising imports are endogenous domestic productivity shocks, instrumentation for import penetration is essential. Again following Acemoglu et al. (2016), I instrument for $\Delta IP_{i,t}$ by

[11] I use 1991 because that is the first year in which the Harmonized System is consistently available for all countries in the sample. In alternative specifications I also include 2010 data.

$$\Delta OIP_{i,t} = \frac{\Delta M_{i,t}^{OC}}{Y_{i,88} + M_{i,88} - E_{i,88}} \tag{A.2}$$

where $\Delta M_{i,t}^{OC}$ is the change in imports from China to a group of eight other high-income countries. The denominator is U.S. consumption in industry i, lagged to 1988 to avoid any prediction by firms of China's impending export boom. This instrumental variable approach isolates imports into the U.S. that can be predicted by China's domestic productivity increase.

China Shock Application: Summary Statistics

Table B.3.: Descriptive Statistics for Chinese Import Penetration Application, Low Skill

	1990-2007	1990-2000	2000-2007
	(1)	(2)	(3)
Δ IP	0.502	0.248	0.767
	(0.95)	(0.68)	(1.46)
Δ OIP	0.279	0.194	0.360
	(0.52)	(0.50)	(0.62)
Annual Δ Log Employment,			
Total	-2.258	-1.003	-3.263
	(3.51)	(3.04)	(4.29)
Scale	0.747	1.660	0.147
	(3.16)	(2.57)	(4.33)
Task	-2.279	-2.072	-2.199
	(1.29)	(1.01)	(1.66)
Productivity	-1.229	-0.723	-2.269
	(1.58)	(1.73)	(2.41)
Supply	0.503	0.132	1.058
	(0.50)	(0.56)	(1.06)
Observations	351	351	351

Notes - Δ IP is 100× the annual change in Chinese import penetration in the U.S. as defined by Equation A.1. Δ OIP is 100× the annual change in Chinese import penetration in IV countries as defined by Equation A.2. Annual Δ Log Employment is multiplied by 100 for interpretation as log points.

Table B.4.: Descriptive Statistics for Chinese Import Penetration Application, High Skill

	1990-2007	1990-2000	2000-2007
	(1)	(2)	(3)
Δ IP	0.396	0.214	0.642
	(0.86)	(0.62)	(1.37)
Δ OIP	0.217	0.134	0.328
	(0.45)	(0.38)	(0.63)
Annual Δ Log Employment,			
Total	-0.009	0.693	-0.759
	(2.89)	(3.52)	(3.85)
Scale	1.173	1.568	0.858
	(2.50)	(2.63)	(3.64)
Task	3.987	4.213	3.739
	(1.33)	(1.60)	(1.63)
Productivity	-4.012	-4.750	-3.051
	(2.13)	(2.80)	(3.63)
Supply	-1.157	-0.338	-2.304
	(0.94)	(1.43)	(2.02)
Observations	351	351	351

Notes - Δ IP is 100× the annual change in Chinese import penetration in the U.S. as defined by Equation A.1. Δ OIP is 100× the annual change in Chinese import penetration in IV countries as defined by Equation A.2. Annual Δ Log Employment is multiplied by 100 for interpretation as log points.

Decomposition Approach

In Section 1.6, I apply my framework to assess the channels through which Chinese import competition, automation, and offshoring led to employment loss. To operationalize my decomposition in terms of logs rather than levels, I use a modified identity in place of the standard identity expressed in Equation 1.11. This modified identity is

$$\frac{L_{i,t+1}}{L_{i,t}} = \frac{L_{i,t+1}}{L_{i,scale,1}} \times \frac{L_{i,scale,1}}{L_{i,task,1}} \times \frac{L_{i,task,1}}{L_{i,productivity,1}} \frac{L_{i,productivity,1}}{L_{i,t}} \tag{A.3}$$

I then take logs so that the log total change in employment is the sum of log changes due to each channel:

$$\ln\Big(\frac{L_{i,t+1}}{L_{i,t}}\Big) = \ln\Big(\frac{L_{i,t+1}}{L_{i,scale,1}}\Big) + \ln\Big(\frac{L_{i,scale,1}}{L_{i,task,1}}\Big) + \ln\Big(\frac{L_{i,task,1}}{L_{i,productivity,1}}\Big) + \ln\Big(\frac{L_{i,productivity,1}}{L_{i,t}}\Big) \tag{A.4}$$

where the right hand side is the sum of effects due to scale, task, productivity, and supply, respectively. Just as described in Section 1.4.2, I calculate Equation A.4 for all 24 possible combinations. I then take the mean of these 24 combinations as my estimate. As these estimates are in terms of log changes, I do not make any adjustments to convert the interpretation from annual hours to effective job counts. Therefore the interpretation of the outcome variables in the applications in Section 1.6 is log thousands of annual hours.

Alternative Time Horizons

Tables B.5 and B.6 report the results for the impact of Chinese import penetration on employment for different time horizons. These are supplemental to the main results reported in Section 1.6.1.

Table B.5.: Effects of Direct Exposure to Chinese Imports on Low-Skill Employment: 2SLS Estimates, Alternative Time Horizons

	Total	Scale	Task	Productivity	Supply
	(1)	(2)	(3)	(4)	(5)
	Panel A: 1990-2007 Long Different (N=351)				
100 × annual Δ Chinese import penetration	-2.126***	-1.644**	0.042	-0.405*	-0.120**
	(0.761)	(0.670)	(0.124)	(0.207)	(0.051)
	Panel B: 1990-2000 Difference (N=351)				
100 × annual Δ Chinese import penetration	-4.736**	-2.766**	-0.003	-1.610*	-0.357*
	(2.247)	(1.392)	(0.226)	(0.849)	(0.195)
	Panel C: 2000-2010 Difference (N=351)				
100 × annual Δ Chinese import penetration	-1.389***	-1.254***	-0.014	-0.082	-0.039
	(0.466)	(0.453)	(0.103)	(0.173)	(0.068)
	Panel D: 1990-2010 Stacked First Differences (N=702)				
100 × annual Δ Chinese import penetration	-2.160***	-1.602***	-0.012	-0.434*	-0.113
	(0.769)	(0.592)	(0.117)	(0.234)	(0.072)

Notes - Regressions are weighted by start-of-period labor hours of the relevant skill group. Also included is an indicator for time period. Robust standard errors are clustered at the 3-digit SIC. * p< .1, ** p< .05, *** p< .01

Table B.6.: Effects of Direct Exposure to Chinese Imports on High-Skill Employment: 2SLS Estimates, Alternative Time Horizons

	Total	Scale	Task	Productivity	Supply
	(1)	(2)	(3)	(4)	(5)
	Panel A: 1990-2007 Long Different (N=351)				
100 × annual Δ Chinese import penetration	-0.826*	-0.749*	0.268	-0.499*	0.154**
	(0.456)	(0.384)	(0.165)	(0.257)	(0.073)
	Panel B: 1990-2000 Difference (N=351)				
100 × annual Δ Chinese import penetration	-0.579	-0.729	0.795	-1.226*	0.581*
	(0.914)	(0.759)	(0.490)	(0.703)	(0.320)
	Panel C: 2000-2010 Difference (N=351)				
100 × annual Δ Chinese import penetration	-0.652**	-1.099***	0.081	0.279	0.087
	(0.299)	(0.395)	(0.107)	(0.187)	(0.119)
	Panel D: 1990-2010 Stacked First Differences (N=702)				
100 × annual Δ Chinese import penetration	-0.638**	-1.027***	0.220	-0.013	0.183
	(0.302)	(0.376)	(0.148)	(0.157)	(0.118)

Notes - Regressions are weighted by start-of-period labor hours of the relevant skill group. Also included is an indicator for time period. Robust standard errors are clustered at the 3-digit SIC. * p< .1, ** p< .05, *** p< .01

A.3 Theory Appendix

The canonical SBTC model can be nested in the task allocation framework developed by Acemoglu and Autor (2011) in various ways. I present one way here. For more background on the theoretical links between the two models, see Acemoglu and Autor (2011) and Autor (2013).

Suppose output for industry i and time t is produced by

$$Y_{i,t} = \left(\frac{1}{N} \sum_{n=1}^{N} \tau_{n,i,t} \right)^{\frac{1}{\rho}}$$

where $\tau_{n,i,t}$ is the output of task n and ρ governs the constant elasticity of substitution between tasks. Task n is carried out by the skill type that has comparative advantage in that task, determined by

$$\max \left\{ \left[(a_{n,i,t} L_{i,t})^{\rho} - w_{L,i,t} L_{i,t} \right], \left[(b_{n,i,t} H_{i,t})^{\rho} - w_{H,i,t} H_{i,t} \right] \right\}$$

where $L_{i,t}$ is low-skill labor, $H_{i,t}$ is high-skill labor, and $a_{n,i,t}$ and $b_{n,i,t}$ are their respective relative productivities in task n. Wages are represented by $w_{L,i,t}$ and $w_{H,i,t}$. Assume high-skill comparative advantage is increasing in n, so that there is some threshold task $\tilde{n}$ above which it is optimal to employ high-skill labor, and below which it is optimal to employ low-skill labor. Then output can be written

$$\begin{aligned} Y_{i,t} &= \left(\frac{1}{N} \sum_{n=1}^{\tilde{n}_{i,t}} (a_{n,i,t} L_{i,t})^{\rho} + \frac{1}{N} \sum_{n=\tilde{n}_{i,t}}^{N} (b_{n,i,t} H_{i,t})^{\rho} \right)^{\frac{1}{\rho}} \\ &= \left[\alpha_{i,t} \left(a_{i,t} L_{i,t} \right)^{\rho} + (1 - \alpha_{i,t}) \left(b_{i,t} H_{i,t} \right)^{\rho} \right]^{\frac{1}{\rho}} \end{aligned}$$

where $\alpha_{i,t} \equiv \tilde{n}_{i,t}/N$, and $a_{i,t}$ and $b_{i,t}$ are the mean productivity of low- and high-skill labor within their set of tasks.

B. APPENDIX FOR: STALLED RACIAL PROGRESS AND JAPANESE TRADE IN THE 1970S AND 1980S

B.1 Data Appendix

B.1.1 Current Population Survey

We use the 1962-1999 Annual Social and Economic Supplement (ASEC) of the Current Population Survey (CPS). To focus on black-white differences, we exclude non-white, non-black individuals. To be consistent across years, we include Hispanic whites, as the CPS did not ask for Hispanic ethnicity until the 1971 survey. We further restrict the sample to non-military men ages 16-64 who were not living in group quarters and were not full time students. We define full time, full year employment as having worked more than 48 weeks in the previous year, and at least 35 hours in the previous week.

Topcoding of income varies through the duration of the CPS. In real terms, the lowest top code is in 1981, at $101,100 1999 Dollars. Across all years, we replace every individual with a reported real income above $101,100 1999 Dollars with 1.5 times that amount ($151,650).

From 1968-1999, we classify workers as being in manufacturing based on 1990 Census occupation codes provided by the CPS. Prior to 1968, industry codes are only available in a small number of general categories. For 1963-1967 we classify codes 5-21 as manufacturing, and for 1962 we use codes 4-20.

B.1.2 Labor Market Variables

We omit individuals living in institutions and unpaid family workers throughout. Following Autor et al. (2013), we impute weeks worked last year for those who report wage income but not weeks. The imputed value is set equal to the mean value for those we observe with the same years of education and 1990 Census occupation code; if that value is not available, the imputed value is set equal to the mean value for those we observe with the same years of education. As the 1970 sample only provides intervalled weeks worked last year, we replace those intervals with the averages of weeks worked within those intervals in the 1980 sample.

To compute weekly wages, we first account for topcoding by replacing values of annual wage income above the 98th percentile to 1.5 times the 98th percentile value. We then divided by the number of weeks worked in the previous year. We replace any values that exceed 150 percent of the topcoded value of annual wage income divided by 50 to this value, and convert to 1999 Dollars using the CPI deflator.

We define annual earned income as the sum of wage income, business income, and farm income. Here we face the challenge that topcoding is both inconsistent across years and income categories. Prior to summing these three sources, we replace values above the 95th percentile of each by year with 1.5 times the 95th percentile value. For business and farm income, which can take on negative values, we replace values below the 3rd percentile of each year with 1.5 times the 3rd percentile value. We then adjust these values to reflect 1999 USD using the CPI deflator.

We define a household as in the 1970 Census, and the race of the household by the race of the household head. Earned household income is calculated analogously to individual level earned income. Total household income includes all income sources, for which we address topcoding by replacing (for each component) values above the 95th percentile of each by year with 1.5 times the 95th percentile value. All household-level income variables are averaged over the number of adults ages 16-64 in the household.

B.1.3 SITC to HS Crosswalk

We constructed a new, country-specific crosswalk from SITC to HS product codes in order to utilize the crosswalk provided in Autor et al. (2013) which maps HS product codes to SIC industries. We describe that crosswalk here.

We utilize Comtrade imports data for years in which both HS codes and SITC codes are available and connected them using the correspondence tables available from the UN (https://unstats.un.org/unsd/trade/classifications/correspondence-tables.asp). Since SITC Rev. 2 codes were not available for 1970, we use SITC Rev. 1 codes. We calculated the shares of Japanese import values of each detailed SITC code (5-digit when consistently available) that mapped to its corresponding 6-digit HS codes for each importer for each year from 1991 through 1994.

For each importer we then averaged the SITC-to-HS value of imports shares across these years. If an SITC code had positive Japanese imports values between 1970 and

1990, but did not have positive values for at least one of the years between 1991 and 1994, we instead calculated the shares from the importer's global imports from 1991 to 1994. The resulting shares are used as origin-destination specific weights to convert imports data between 1970 and 1990 from SITC to HS codes. An analagous crosswalk was made for exports for the net imports robustness exercise in Table 2.6.

B.1.4 CBP Imputation Procedure

In each county-industry cell, the CBP reports the total number of employees as well as a count of establishments by brackets of employment totals. As a precaution to avoid disclosing the operations of any individual employer, the CBP suppresses some total employment counts in some county-industry cells, while it always reports the number of establishments by firm size bracket. For these cases, we impute employment following a procedure analogous to that described in Autor et al. (2013), which multiplies the number of establishments in each bracket by the average firm size in that bracket that can be observed in the CBP. A minor difference between the 1970 CBP and the later series used by Autor et al. (2013) is that suppressed employment totals are also bracketed in the later series, while there is no additional information provided in 1970. In the 1970 CBP, 12 Kansas counties were omitted from the data. These counties are in three of our regression sample CZs. Robustness checks omitting these three CZs are available upon request.

B.1.5 1960 Census Industry Disaggregation

To construct our 1960 instrument, we rely on CZ industry composition calculated from the 1960 5% Census sample, since CBP data for 1960 does not exist. In order to utilize the Census data, we needed to disaggregate the fairly coarse 1960 Census industry codes into the SIC industries we use for the rest of our variables. To do this, we developed the following procedure.

We took the employed population in the 1960 5% Census sample and calculated each CZ industry composition based on 1990 Census industry codes. We connected this in a one-to-many mapping to CZ-level SIC-based industry employment data from the 1970 CBP, using the crosswalk available in Census Bureau Technical Paper #65. Because we use a slight aggregation of SIC industries, we had to combine two Census codes, 140 and 142, giving us 61 Census codes. 58 of these are trade exposed since the remaining 3 are "not specified" industries that do not connect to SIC codes in the Census crosswalk (122, 332, 392). This resulting mapping of 58 Census industry codes to 380 SIC codes allows us to calculate CZ-specific employment shares of each SIC code within each Census code.

This disaggregation relies on the assumption that within each Census code, a CZ's SIC-based employment composition is constant between 1960 and 1970. If the associated SIC industries were not present in the 1970 data for a particular CZ-Census industry cell, we use state-level (or, as a last resort, national-level) employment shares.

B.1.6 Routine- and Offshorability- Indices

We include as controls the share of employment in routine-intensive occupations and the average offshorability index of occupations in a CZ in 1970. Both of these measures use task data from Autor and Dorn (2013). To identify routine-intensive occupations, we follow Autor and Dorn (2013) to create a routine-intensive index for each occupation: log(routine score) - log(manual score) - log(abstract score). As they do, we recode the bottom 5 percent of the population in the base year for manual and abstract to be the 5th percentile. After ranking occupations by the routine-intensive index, those which take the top third of employment in the base year are classified as routine-intensive.

The offshorability index is derived from the variables *face-to-face contact* and *on-site job* by occupation in O*NET data; more details on its construction can be found in Autor and Dorn (2013). Our control is the mean index in the CZ according to its composition of occupations in 1970. In the regression sample, the CZ mean offshorability index is standardized to have a mean of 0 and standard deviation of 10 in 1970.

B.1.7 Import Competition in Final Goods

In column (6) of Table 2.6, we report the result for an alternative measure of import exposure that seeks to isolate the effect of final goods imports, rather than intermediates to be used as inputs by firms. In order to remove the share of imports that are used as intermediates, we follow the approach of Autor et al. (2013) by exploiting the 1972 input-output data provided by the BEA. We convert the codes used in the BEA's IO Data files to SIC 1972 via their crosswalk and then to the slightly aggregated set of SIC 1987 codes that are used in the rest of the analysis. Since these are not one-to-one mappings, we incorporate weights based on the NBER-CES Manufacturing Industry Database concordance.

Assuming that make and use values in domestic production also reflect the nature of imports, we use these tables to construct shares of each industry's make value that is then used as inputs in other manufacturing industries, and deduct this share from the change in imports value when constructing IPW. We likewise construct a modified instrument.

B.2 Empirical Appendix

In this appendix we provide additional descriptive statistics and supplemental results to those presented in the main text.

B.2.1 Product Composition of Japanese Imports

Table B.1 shows the growth in imports for the SITC product categories that received the ten largest increases from 1970-1990. Values are listed in millions of 1999 U.S. Dollars.

B.2.2 Import Penetration Ratio

In our main analysis we measure exposure to Japanese imports by change imports per worker, following Autor et al. (2013) and related papers, so that we can readily compare our results with theirs. Here we use an alternative measure of exposure to import competition: import penetration ratio (IPR), similar to Acemoglu et al. (2016). IPR is the change in imports from Japan by industry as a share of initial domestic consumption of that industry. For each CZ, we construct a variable that is a weighted average of IPR based on its 1960 industry composition. Specifically, for each CZ i we calculate

$$\Delta IPR_{uit} = \sum_j \frac{L_{ijt}}{L_{ujt}} \frac{\Delta M_{ujt}}{Y_{it} + M_{it} - X_{it}} \tag{B.1}$$

where L_{ijt} is the number of workers in commuting zone i in industry j at the beginning of period t; L_{ujt} is that same value for the United States; Y_{it}, M_{it}, and $X_{i,t}$ are output (shipments), imports, and exports, respectively, of industry i in time period t; and ΔM_{ujt} is the change in imports from Japan in the industry's product space (in \$1000s) during the time period. Analogous to our main measure, we construct an instrument based on 1960 industry composition and Japanese imports by other countries. Our output, imports, and exports data used to calculate consumption by industry are from Feenstra (1996) and Feenstra (1997). These data are in 1972-basis SIC codes which we convert to our slightly aggregated set of 4-digit SIC industries.

We report our employment results using IPR in Table B.2. Generally our findings are similar to our main findings, although our result finding an increase in white manufacturing employment loses statistical significance.

B.2.3 Agricultural Employment Trends

In Table 3.4 we saw that the changes in the racial gap in non-manufacturing employment actually grew at a faster pace than the gap in manufacturing employment during our period of interest. One concern then is that our focus on manufacturing is misplaced, and the large effects we find on earnings through this channel must be spurious. These trends however, appear to be driven by a continuation of the long-term secular decline of black agricultural employment, as well as the growth in black non-employment. In Figure B.2.10, we show that especially in the 1960s and 1970s there is a sharp decline in black employment in agriculture relative to whites. In contrast, in Figure B.2.10, we see that, as a share of employment, black growth in non-manufacturing kept pace with whites outside of agriculture.

B.2.4 Geographic Dispersion of Japanese Trade

In Figures B.2.10 and B.3 we show heat maps for the geographic dispersion of trade for all CZs in the continental United States and our regression sample, respectively. We see that the largest trade increases took place across the northern "Rust Belt" region, as well as into New England and southern California. In contrast, the heavily black regions of the Deep South were less exposed. Most of the regions we exclude due to sample size of black working age males are sparsely populated Western commuting zones. In general, these CZs were less exposed than those in our regression sample.

In Table B.3 we list the ten most and least affected commuting zones among the 40 largest CZs in our sample. The hardest hit areas were large Midwestern manufacturing cities like Detroit and Buffalo, though San Jose, California also makes the list. The smallest growth areas were primarily in the Sun Belt and West Coast; Pittsburgh, Pennsylvania is one notable exception.

B.2.5 Alternative Time Horizons

Due to concerns about the timing of the 1980 Census and the receding of Japanese trade in the late-1980s, we used the 1990-1970 long difference approach throughout the main text. In Table B.4 we explore different time horizons for our main results on manufacturing share. Column (1) repeats column (1) of Table 2.5. Column (2) and (3) look just at 1980-1970 and 1990-1980, respectively. From 1980-1970 we see a strong negative effect on black manufacturing employment without any evidence for the positive effect on white outcomes that we saw in our main results. In contrast from 1990-1980 we see strong positive effects on manufacturing employment overall, with little evidence for a differential effect. One interpretation of these results is that the initial influx of import competition in the late 1970s led to large layoffs of black manufacturing workers, and that it was not until the 1980s that firms adjusted and began re-employing (higher skill, white) workers. But, as we state in the main text, the 1990-1980 difference may be unreliable due to the 1980 recession and the decline in Japanese imports at the end of the 1980s. Column (4) presents the preferred specification from Autor et al. (2013), which stacks the 1980-1970 and 1990-1980 differences. These results are consistent with those from our preferred specification, but with less precision.

In column (5) we estimate a placebo regression of the increase in Japanese imports from 1970-1990 on the change in manufacturing employment from 1960-1970. Here we include as controls the 1960 CZ manufacturing share and census division fixed effects, as well as their interactions with the black indicator. We find little evidence of an effect of future imports on past manufacturing changes, providing support for the validity of our instrumental variable strategy. If anything, the pre-trends were towards black growth in manufacturing in areas that would receive higher Japanese import competition.

In column (6) we present an alternative specification of our preferred long differences strategy where we use as our left-hand side variable the change in the black-white manufacturing employment *gap*. The coefficient on our import exposure variable is thus analogous to the coefficient on the black interaction in the fully-interacted specification. The result is very similar.

B.2.6 Instrumental Variable Robustness Exercises

Denote $\overline{X_j}$ as the industry average of some CZ-level variable X_i weighted based on industry j's employment distribution across CZs. That is,

$$\overline{X_j} = \frac{\sum_i s_{ij} X_i}{\sum_i s_{ij}} \tag{B.2}$$

where s_{ij} is the share of CZ i's employment belonging to industry j. To provide conditions for the consistency of "shift-share" IV estimators, Borusyak et al. (2018) show that the CZ-level regression in the main text is equivalent to estimating the two-stage least squares regression

$$\overline{\Delta Y}^{\perp}_{jk,1990-1970} = \alpha_k + \beta_k \overline{\Delta IPW}^{\perp}_{uj,1990-1970} + \epsilon^{\perp}_{jk} \tag{B.3}$$

instrumented by $\frac{\Delta M_{oj,1990-1970}}{L_{i1960}}$, where k is race, and the superscript $\perp$ represents a variable that has been residualized over the set of controls. In other words, the CZ-level regressions we report in the main text are equivalent to a set of industry-level regressions where the variables are exposure-weighted averages of CZ characteristics and outcomes.

The industry-level approach provides several benefits, including an alternative set of standard errors (see Appendix B.2.8) and more transparency of how industries influence identification. We first plot the relationship between industry-level import exposure (as measured by our instrument) and changes in CZ-manufacturing share by race in Figure B.2.10. Industries are binned based on their percentile of import exposure, excluding automobiles and computers, which we highlight in red and green, respectively.[1] As we discussed in Section 2.3.2 these industries present unique concerns for identification.

The "as-if" random assignment framework of Borusyak et al. (2018) requires that, from the hypothetical distribution of shocks that led to the Japanese export boom, each industry was expected to receive the same shock. Consistent with the concerns raised in the main text, Figure B.2.10 shows that both automobiles and computers are far to the right of the shock distribution, suggesting the possibility their realized trade values were the result of a different underlying process. We also see an additional set of outlying industries that raise concerns.

In Table B.5 we use the industry-level approach to relax the assumption of mean independence of the shock distribution for these industries. First, column (1) repeats column (1) of Table 2.5 and verifies the approaches yield identical point estimates. Columns (2) and (3) exclude the computer and automobile industries, respectively. We also include industry-level controls for the exposure to these industries. While with

[1] We follow Acemoglu et al. (2016) and classify computer industries as SIC87dd 3571, 3572, 3577, 3578, 3651, 3652, 3661, 3663, 3669, 3671, 3672, 3674, 3675, 3676, 3677, 3678, 3679, 3695, 3812, 3822, 3823, 3824, 3825, 3826, 3829, 3844, 3845, and 3873. For automobiles, we use SIC87dd 3-digit grouping 371.

automobiles in particular, we lose substantial precision, the estimates are consistent with our main results. In column (4) we exclude any outlying industries whose import exposure (as measured by the instrument) was more than \$30,000 per worker (roughly three standard deviations above the mean), and include controls for exposure to each of these outlying industries.[2] Our results are essentially unchanged.

In columns (5) and (6), we include controls for exposure to 1- and 2-digit manufacturing clusters, respectively, which allows the mean of the shock generating process to differ at these levels.[3] The 2-digit cluster controls are especially demanding given that 1960 manufacturing sector data is imputed based on broader industry classifications (see Appendix B.1.5). Thus it is not surprising that their inclusion leads to a noticeably smaller coefficient on the black interaction term. Nonetheless, the result remains negative and statistically significant across both specifications.

In Table B.6 we implement an additional set of robustness exercises recommended by Borusyak et al. (2018) within this industry-level framework. We use each individual developed country's imports as a separate instrument, which allows us to perform a test of overidentifying restrictions. We use three different estimation methods: two-stage least squares, limited information maximum likelihood, and generalized method of moments. The method chosen has little impact on our results and we fail to reject the overidentifying restrictions ($p = .62$).

B.2.7 Intraclass Correlations

As discussed in the Section 2.3.2, Borusyak et al. (2018) show that shift-share instruments are consistent provided the industry-level shocks are orthogonal to CZ-level unobservables, and are sufficiently dispersed across industries. They further show that the latter condition can be relaxed to allow for correlation within industry clusters, and to be conditional on observables. We test this assumption here.

Following the approach in Borusyak et al. (2018), we estimate the hierarchical random effects model

$$\hat{g}_n = a_{1,n} + b_{2,n} + c_{3,n} + e_n \tag{B.4}$$

where $\hat{g}_n$ is, for industry n, the residual of a regression of Japanese exports to other

[2] This affects four SIC87dd industries: 3751 (motorcycles, bicycles, and parts); 3827 (optical instruments and lenses); 3844 (X-ray apparatus and tubes); and 3845 (Electromedical equipment)

[3] Note that column (5) is equivalent to column (2) of Table 2.6.

countries on a set of industry-level controls; and a, b, and c are random effects specific to industry n's 1-digit, 2-digit, and 3-digit classification, respectively.[4] Note that the residual e_n represents variation at the 4-digit industry level, the level at which the instrument is computed. For 1-digit classifications, we follow the system used by Autor et al. (2014); for 2- and 3-digit classifications, we follow the SIC system. To avoid distorting our estimates with variation caused by large outliers, we winsorize all values of $\hat{g}$ above \$30,000 per worker to be \$30,000 (roughly three standard deviations above the mean).[5] Following convention, we impose a normal distribution for the random effects, and estimate the model using maximum likelihood.

Table B.7 reports intraclass correlation coefficients from this exercise. We find a moderate amount of clustering at the 3-digit level, but given the large number of 3-digit industries in our data (135), this presents less of a concern for consistency. When estimating industry-level regressions in Appendices B.2.6 and B.2.8, we cluster our standard errors at the 3-digit level to account for this correlation. At the higher 2-digit and 1-digit levels, though larger than what Borusyak et al. (2018) find for China, the correlation is much more mild. This is consistent with similar industries receiving different levels of shock exposure, and the dispersion assumption necessary for the consistency of the IV.

B.2.8 Borusyak-Hull-Jaravel Standard Errors

In the main text, we report standard errors that are clustered at the state level to account for correlations within proximate geographies. As the identification from "shift-share" instruments is driven by shocks at the industry-level, Adão et al. (2018) note that correlated errors within industries across different geographies may be a larger concern, and derive an alternative set of standard errors to account for this. Borusyak et al. (2018) show that the standard errors produced by the industry-level regressions discussed in Appendix B.2.6 are asymptotically equivalent to those constructed by Adão et al. (2018). In Table B.8 we reproduce Table 2.5 using the industry-level approach, clustering at the 3-digit SIC-level, which is the level of clus-

[4]For controls, we follow our main specification and use industry-level exposure to CZ-level percentage of employment in manufacturing, college percentage of population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and census divisions. See Appendix B.2.6 for more details of the industry-level approach.

[5]Because large outliers increase the variation within clusters, the winsorization produces larger and more conservative estimates of the amount of within cluster correlation.

tering suggested by our intraclass correlation exercises in Appendix B.2.7. We find this approach produces universally smaller standard errors than state level clustering.

B.2.9 Mean Earnings

As we discuss in the main text, we prefer working with median income and earnings rather than means due to concerns about topcoding and the susceptibility to outliers in small samples. In Table B.9 we estimate the effects of import competition on disparities in mean log income and earnings for males and households. Note that we are unable to compute an analogue of the median log earnings of all working age males, since we cannot take the log of 0.

Just as in medians, we find little evidence for change in the wage or earnings gap among those with positive earnings. However, we do find negative and statistically significant effects on the household earnings gap, albeit smaller than that estimated in Table 2.12. We also find a smaller but non-trivial impact on the household income gap, although it is not statistically significant at conventional levels. Note that unlike for our mean male earnings regressions, mean household earnings is sensitive to changes in non-labor force participation for working age males.

B.2.10 Nationally Representative Descriptive Statistics

In section 2.4.4 we used nationally representative statistics for performing back of the envelope calculations. Table B.10 provides a full set of descriptive statistics for this sample.

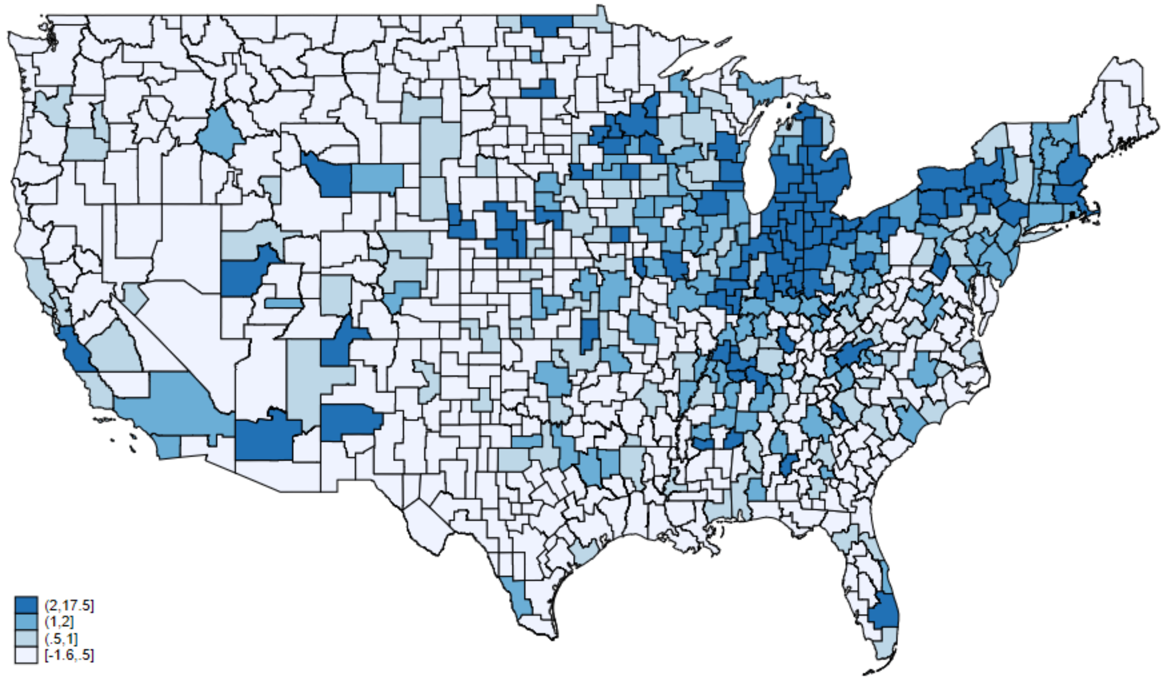

Figure B.1.: Change in Import Exposure Intensity, 1990-1970: All CZs

Notes - Change in IPW from 1970 to 1990 for each commuting zone in the continental United States.

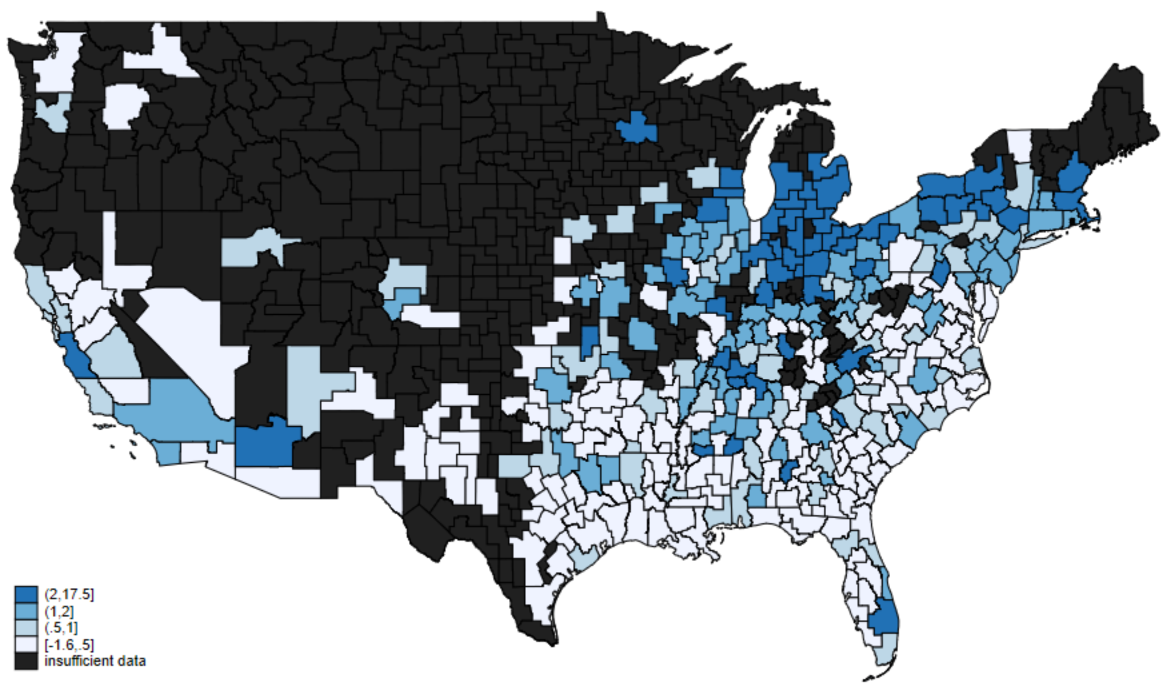

Figure B.2.: Change in Import Exposure Intensity, 1990-1970: Regression Sample

Notes - Change in IPW from 1970 to 1990 for commuting zones in the regression sample.

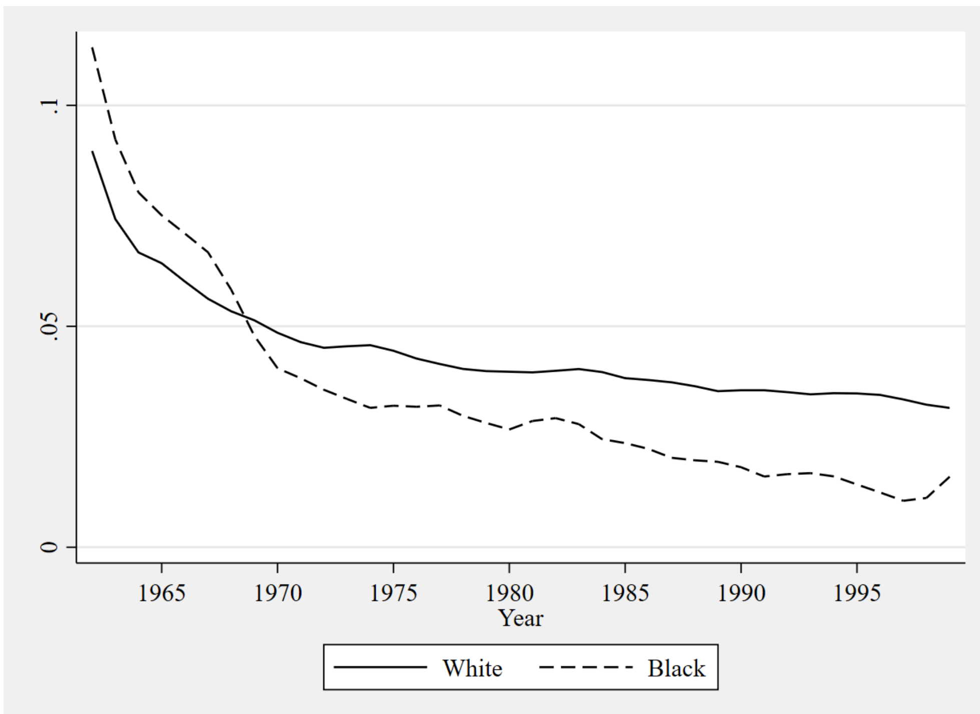

Figure B.3.: Fraction of Employment in Agriculture: Working Age Men, 1962-1999

Notes - Yearly scatterplot data smoothed using LOWESS with bandwidth=0.15

Source - Current Population Survey (1962-1999).

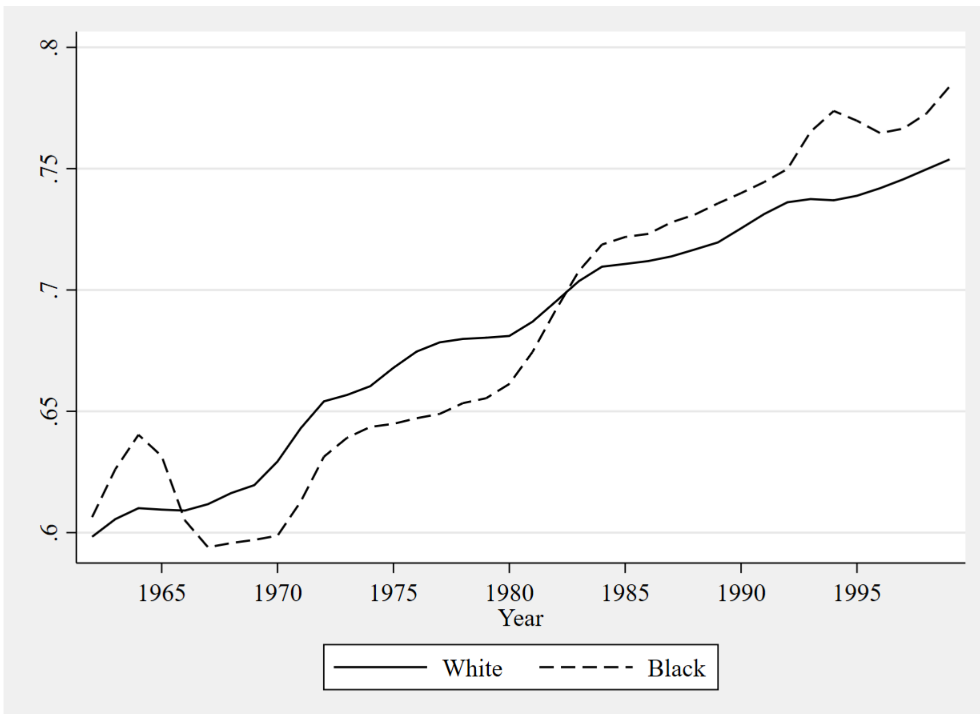

Figure B.4.: Fraction of Employment in Non-Agricultural Non-Manufacturing: Working Age Men, 1962-1999

Notes - Yearly scatterplot data smoothed using LOWESS with bandwidth=0.15

Source - Current Population Survey (1962-1999).

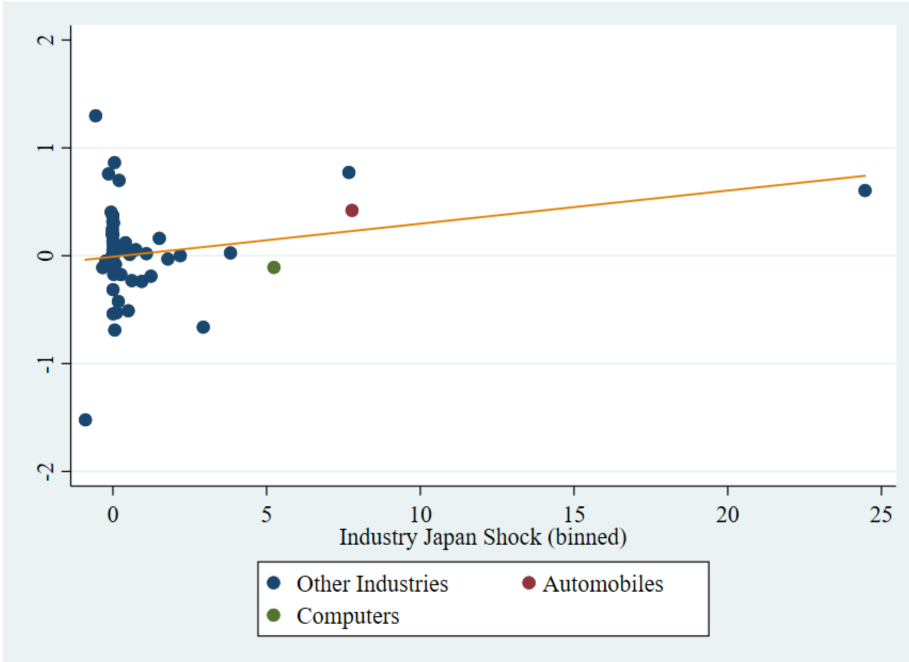

(a) White

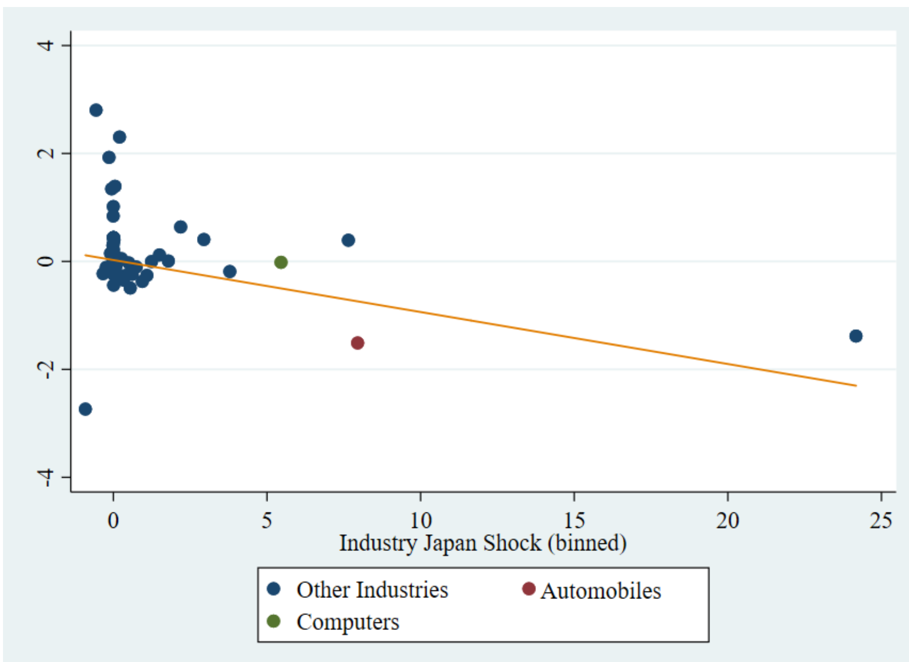

(b) Black

Figure B.5.: Industry-level Japanese Import Exposure and Average Residualized Change in Manufacturing Employment by Race, 1990-1970

Notes - Each non-automobile, non-computer industry bin represents 2% of industries. Y-axis is (for each bin) the average CZ-level change in manufacturing employment. Yellow line is weighted least squares best fit.

Table B.1.: Growth of Japanese Imports to U.S. by Product, 1970-1990: Ten Largest

	SITC	1970	1990	Growth
	(1)	(2)	(3)	(4)
Passenger motor cars (excluding buses)	7321	2,123.13	26,644.29	24,521.17
Statistical machines cards or tapes	7143	5.98	7,718.19	7,712.21
Bodies & parts of motor vehicles (excluding motorcycles)	7328	142.37	7,431.99	7,289.62
Thermionic valves and tubes, transistors, etc.	7293	132.73	4,868.01	4,735.28
Other telecommunications equipment	72499	265.63	4,074.20	3,808.57
Parts of office machinery, n.e.s.	71492	56.96	3,647.93	3,590.96
Internal combustion engines, not for aircraft	7115	168.42	3,350.47	3,182.05
Equipment for indoor games	89424	29.93	2,977.28	2,947.36
Phonographs, tape & other sound recorders etc.	8911	1,437.22	4,295.48	2,858.26
Lorries and trucks, including ambulances, etc.	7222	134.92	1,960.21	1,825.29

Notes - In millions of 1999 U.S. Dollars

Source - UN Comtrade

Table B.2.: Japanese Import Penetration Ratio and Change in Racial Employment Status Gap, 1990-1970 Long Difference: 2SLS Estimates

	Mfg emp	Non-mfg emp	Unemp	NILF
	(1)	(2)	(3)	(4)
(ΔJapanese import penetration)	0.408	-0.178	-0.036	-0.187
	(0.327)	(0.321)	(0.089)	(0.156)
(Δ Japanese import penetration) $\times$ Black	-2.564***	0.904	-0.197	1.519***
	(0.500)	(0.853)	(0.369)	(0.550)
Observations	716	716	716	716

Notes - Robust standard errors clustered at the state-level in parentheses. Models are weighted by race-specific 1970 population. Each regression includes census division fixed effects; commuting zone-level controls for percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960; a black indicator; and interactions of the black indicator with all of these variables. $p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table B.3.: Growth of Imports Exposure Per Worker Across CZs, 1990-1970: 40 Largest CZs (Regression Sample)

Ten Largest Increases		Ten Smallest Increases	
Detroit, MI	9.292	New Orleans, LA	0.097
San Jose, CA	5.952	Sacramento, CA	0.124
Buffalo, NY	5.692	San Antonio, TX	0.311
Minneapolis, MN	3.362	Tampa, FL	0.362
Cleveland, OH	3.203	Arlington, VA	0.397
Cincinnati, OH	2.829	Seattle, WA	0.412
Dayton, OH	2.516	Houston, TX	0.458
Syracuse, NY	2.308	Pittsburgh, PA	0.466
Indianapolis, IN	2.221	New York, NY	0.647
Boston, MA	2.197	Denver, CO	0.889

Table B.4.: Japanese Imports on Change in Manufacturing Employment/ Working Age Population in CZs, Alternative Time Horizons: 2SLS Estimates

	Change in Manufacturing Employment					Chng in Man Gap
	1990-1970 LD (1)	1980-1970 FD (2)	1990-1980 FD (3)	1990-1970 Stacked (4)	1970-1960 LD (5)	1990-1970 LD (6)
(Δ Imports from Japan to US)/worker	0.193* (0.117)	-0.152 (0.173)	0.598*** (0.167)	0.310 (0.211)	0.110 (0.155)	-0.813*** (0.223)
(Δ Imports from Japan to US)/worker × Black	-0.785*** (0.173)	-1.605*** (0.226)	0.151 (0.218)	-0.554** (0.237)	0.331 (0.216)	
Observations	716	716	716	1432	716	358

Notes - Robust standard errors clustered at the state-level in parentheses. Columns (1)-(5) are weighted by race-specific 1970 population, while column (6) is weighted by 1970 population. Each regression includes census division fixed effects and commuting zone-level controls for percentage of employment in manufacturing. Columns (1)-(4) and (6) include additional controls for college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960. Columns (1)-(5) include a black indicator and interactions of the black indicator with all of control variables. $p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table B.5.: Japanese Imports and Change in Manufacturing Employment / Working Population in CZs, Industry-Level Regressions, 1990-1970 Long Difference: 2SLS Estimates

	All (1)	No Comp (2)	No Autos (3)	No Out (4)	1-dig Shares (5)	2-dig Shares (6)
(Δ Imports from Japan to US)/worker	0.193** (0.081)	0.225** (0.108)	0.283 (0.461)	0.189** (0.083)	0.268** (0.128)	0.234* (0.127)
(Δ Imports from Japan to US)/worker × Black	-0.785*** (0.104)	-0.841*** (0.135)	-1.311* (0.757)	-0.754*** (0.073)	-0.509*** (0.140)	-0.227** (0.106)
Observations	762	706	756	754	762	762

Notes - Robust standard errors clustered at the 3-digit SIC-level in parentheses. Models are weighted by race-specific CZ industry exposure. Each regression includes controls for census division exposure; exposure to commuting zone-level percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960; a black indicator; and interactions of the black indicator with all of these variables. Column (2) excludes computer industries and includes a control for CZ-level exposure to computer industries and its interaction with a black indicator. Column (3) excludes automobile industries and includes a control for CZ-level exposure to automobile industries and its interaction with a black indicator. Column (4) excludes industries with outlying trade IV and includes controls for CZ-level exposure to each of these industries and their interactions with a black indicator. Column (5) includes controls for CZ-level exposure to 1-digit manufacturing industries. Column (6) includes controls for CZ-level exposure to 2-digit SIC manufacturing industries. $p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table B.6.: Japanese Imports and Change in Manufacturing Employment / Working Population in CZs, Industry-Level Regressions, 1990-1970 Long Difference: Overidentification Tests

	2SLS	LIML	GMM
	(1)	(2)	(3)
(Δ Imports from Japan to US)/worker	0.115*	0.113	0.128**
	(0.068)	(0.071)	(0.058)
(Δ Imports from Japan to US)/worker $\times$ Black	-0.759***	-0.758***	-0.718***
	(0.064)	(0.064)	(0.048)
Observations	762	762	762
J-statistic	8.057	8.051	8.057
p-value on J-test	0.623	0.624	0.623

Notes - Robust standard errors clustered at the 3-digit SIC-level in parentheses. Models are weighted by race-specific CZ industry exposure. Each regression includes controls for census division exposure; exposure to commuting zone-level percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960; a black indicator; and interactions of the black indicator with all of these variables. J-statistics are from Hansen test of instrument overidentifying restrictions.

Table B.7.: Intraclass Correlations of Residualized Japanese Trade Shock

	ICC	SE
	(1)	(2)
1-digit	0.044	(0.026)
2-digit	0.065	(0.041)
3-digit	0.160	(0.049)
4-digit Industries	380	380

Notes - Robust standard errors in parentheses. Intraclass correlation coefficients from hierarchical random effects model. Japanese trade shock residual computed from regression of industry-level exports to six other highly developed countries on industry-level exposure to CZ-level percentage of employment in manufacturing, college percentage of population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and census divisions. 1-digit industry classifications follow system from **?**. 2- and 3-digit industry classifications are SIC87.

Table B.8.: Japanese Imports and Change in Racial Employment Status Gap, Industry-Level Regressions, 1990-1970 Long Difference: 2SLS Estimates

	Mfg emp	Non-mfg emp	Unemp	NILF
	(1)	(2)	(3)	(4)
(Δ Imports from Japan to US)/worker	0.193**	-0.097	-0.010	-0.086**
	(0.081)	(0.080)	(0.028)	(0.036)
(Δ Imports from Japan to US)/worker $\times$ Black	-0.785***	0.228**	-0.071	0.542***
	(0.104)	(0.108)	(0.081)	(0.078)
Observations	762	762	762	762

Notes - Robust standard errors clustered at the 3-digit SIC-level in parentheses. Models are weighted by race-specific CZ industry exposure. Each regression includes controls for census division exposure; exposure to commuting zone-level percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960; a black indicator; and interactions of the black indicator with all of these variables

Table B.9.: Japanese Imports and Changes in Mean Log Earnings, 1990-1970 Long Difference: 2SLS Estimates

	Working Age Males		Households	
	Weekly Wage (1)	Annual Earnings (2)	Annual Earnings (3)	Annual Income (4)
(Δ Imports from Japan to US)/worker	-0.039	-0.104	-0.308	-0.074
	(0.405)	(0.397)	(0.333)	(0.307)
(Δ Imports from Japan to US)/worker $\times$ Black	0.204	0.024	-1.128**	-0.956
	(0.310)	(0.510)	(0.556)	(0.597)
Observations	716	716	716	716

Notes - Robust standard errors clustered at the state-level in parentheses. Models are weighted by race-specific 1970 population. Each regression includes census division fixed effects; commuting zone-level controls for percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960; a black indicator; and interactions of the black indicator with all of these variables. $p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Table B.10.: Descriptive Statistics: Nationally Representative Sample

	Black		White		ΔGap
	1970 (1)	1990 (2)	1970 (3)	1990 (4)	(5)
Percentage of population employed in manufacturing	19.412 (39.552)	12.857 (33.472)	22.970 (42.064)	17.756 (38.214)	-1.341
Percentage of population employed in non-manufacturing	51.607 (49.974)	49.588 (49.998)	59.396 (49.109)	63.242 (48.215)	-5.866
Unemployed share of population	4.379 (20.464)	9.575 (29.425)	2.858 (16.663)	4.351 (20.400)	3.703
Labor force non-participation rate	24.602 (43.069)	27.981 (44.890)	14.776 (35.486)	14.651 (35.361)	3.504
Median log weekly wage, male earners	612.317 (82.554)	614.253 (83.535)	654.631 (79.295)	647.104 (84.943)	9.464
Median log annual earned income, male earners	999.966 (104.302)	995.961 (120.389)	1044.223 (105.975)	1038.146 (113.675)	2.072
Median log annual earned income, all working-age males	979.440 (381.748)	955.478 (424.266)	1037.371 (303.740)	1024.793 (324.687)	-11.384
Median log HH earned income	935.678 (330.072)	950.599 (362.998)	989.468 (268.972)	1000.677 (284.588)	3.713
Median log HH total income	941.393 (205.765)	960.638 (198.622)	990.371 (181.279)	1008.681 (153.246)	0.934
HH welfare rate	14.351 (35.059)	17.682 (38.152)	2.892 (16.758)	4.491 (20.711)	1.732

Notes - Standard deviations in parentheses.

Table B.11.: Japanese Imports and Change in Employment/Working Age Population in Non-Manufacturing Sectors by Race and Skill Group, 1990-1970 Long Difference: 2SLS Estimates

	Ag/ mining/ const	Trans- port	Whlsle & retl trade	Serv- ice	Pub- lic/ othr
	(1)	(2)	(3)	(4)	(6)
			Panel A: HS Dropouts		
All Workers					
(Δ Imports from Japan	-0.091	0.081	0.197***	-0.053	0.076
to US)/worker	(0.092)	(0.051)	(0.058)	(0.034)	(0.082)
Black Workers					
(Δ Imports from Japan	-0.142	0.190***	0.220***	-0.125	0.175*
to US)/worker	(0.126)	(0.061)	(0.092)	(0.082)	(0.098)
White Workers					
(Δ Imports from Japan	-0.166	0.093**	0.220***	-0.009	0.052
to US)/worker	(0.112)	(0.044)	(0.070)	(0.040)	(0.073)
			Panel B: HS Grads		
All Workers					
(Δ Imports from Japan	0.154**	-0.078	0.216***	-0.109**	-0.101
to US)/worker	(0.073)	(0.064)	(0.101)	(0.051)	(0.109)
Black Workers					
(Δ Imports from Japan	-0.109	0.298***	0.288**	-0.179	-0.276
to US)/worker	(0.071)	(0.076)	(0.142)	(0.117)	(0.253)
White Workers					
(Δ Imports from Japan	0.073	-0.125	0.188*	-0.105**	-0.085
to US)/worker	(0.085)	(0.077)	(0.104)	(0.052)	(0.110)
			Panel C: College Educated		
All Workers					
(Δ Imports from Japan	0.010	-0.107**	0.027	-0.115	0.034
to US)/worker	(0.040)	(0.043)	(0.069)	(0.130)	(0.083)
White Workers					
(Δ Imports from Japan	0.015	-0.083**	0.052	-0.095	-0.005
to US)/worker	(0.044)	(0.039)	(0.074)	(0.131)	(0.087)
Observations	358	358	358	358	358

Notes - Robust standard errors clustered at the state-level in parentheses. Models are weighted by race-specific 1970 population. Each entry represents a separate regression for that race and/or skill group. Each regression includes census division fixed effects; commuting zone-level controls for percentage of employment in manufacturing, college percentage of the population, average offshorability index of occupations, percentage of employment in routine occupations, black percentage of population, and foreign-born percentage of population in 1960; a black indicator; and interactions of the black indicator with all of these variables. $*p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

C. APPENDIX FOR: END-OF-LIFE MEDICAL SPENDING: EVIDENCE FROM PET INSURANCE

C.1 End-of-Life Spending Patterns by Age and Year

In Tables B.1 and B.2 we report some statistics on end-of-life spending by age and year for dogs in our sample.

Table B.1.: Spending Patterns by Age at Diagnosis

	Last 12 Months Spending		Last 12 Months Spending, Share		First 12 Spending		First 12 Months Spending, Full Sample	
	mean (1)	st. dev. (2)	mean (3)	st. dev. (4)	mean (5)	st. dev. (6)	mean (7)	st. dev. (8)
2	4486.63	3871.09	0.63	0.28	4103.92	3728.30	3970.47	3693.59
3	4788.84	4047.29	0.60	0.27	4554.25	4129.64	4455.73	4006.89
4	4649.80	4072.49	0.55	0.24	4211.44	3625.88	4195.07	3678.86
5	4446.72	3728.47	0.52	0.24	4095.77	3683.28	4073.91	3615.21
6	4241.56	3696.45	0.47	0.21	3708.26	3349.09	3751.74	3403.71
7	4123.94	3619.04	0.44	0.21	3755.84	3500.51	3813.82	3501.81
8	4129.50	3647.24	0.41	0.19	3673.98	3529.48	3691.58	3501.87
9	3918.54	3422.73	0.39	0.18	3426.75	3280.63	3463.77	3290.27
10	3825.47	3434.50	0.36	0.17	3273.93	3266.11	3325.98	3277.99
11	3671.68	3091.46	0.33	0.16	3149.56	3002.78	3205.65	3035.13
12	3727.19	3513.39	0.31	0.15	3048.74	3111.44	3072.47	3108.58
Total	3983.01	3520.05	0.40	0.20	3484.40	3350.51	3525.44	3357.05
Obs.	31,484		31,484		31,484		33,899	

Notes - Columns show mean and standard deviation of each variable. Last 12 Months Spending is spending during the last 12 months of life. Last 12 Months Spending, Share is this end-of-life spending as a share of total lifetime spending on the pet's health care. First 12 Months Spending is spending during the first 12 months upon diagnosis. Share of lifetime spending in Columns (3) and (4) uses some imputed annual spending for cases where the dog does not have insurance for every year of life. To impute, we use median annual claims for pre-diagnosis dogs within the same breed size/breed category/gender/age cell as the annual claim for any missing dog/year observation.

Table B.2.: Spending Patterns by Year of Diagnosis

	Last 12 Months Spending		Last 12 Months Spending, Share		First 12 Months Spending		First 12 Months Spending, Full Sample	
	mean (1)	st. dev. (2)	mean (3)	st. dev. (4)	mean (5)	st. dev. (6)	mean (7)	st. dev. (8)
2009	3298.74	3048.75	0.36	0.19	3124.00	2782.78	3118.08	2776.60
2010	3594.35	3204.45	0.39	0.19	3229.24	2981.29	3213.86	2958.61
2011	3709.64	3185.78	0.39	0.19	3301.91	3066.80	3303.17	3058.99
2012	3891.07	3283.37	0.39	0.19	3405.94	3168.81	3400.92	3159.73
2013	3933.16	3412.79	0.39	0.20	3449.91	3252.07	3467.48	3278.68
2014	4107.77	3482.87	0.40	0.20	3593.28	3389.95	3592.60	3371.41
2015	4197.53	3780.65	0.41	0.21	3577.89	3606.52	3609.81	3547.78
2016	4489.28	4045.94	0.42	0.22	3844.71	3861.90	3937.31	3842.97
2017	4524.03	3916.74	0.42	0.22	3775.23	3759.60	3992.00	3805.47
Total	3983.01	3520.05	0.40	0.20	3484.40	3350.51	3525.44	3357.05
Obs.	31,484		31,484		31,484		33,899	

Notes - Columns show mean and standard deviation of each variable. Last 12 Months Spending is spending during the last 12 months of life. Last 12 Months Spending, Share is this end-of-life spending as a share of total lifetime spending on the pet's health care. First 12 Months Spending is spending during the first 12 months upon diagnosis. Share of lifetime spending in Columns (3) and (4) uses some imputed annual spending for cases where the dog does not have insurance for every year of life. To impute, we use median annual claims for pre-diagnosis dogs within the same breed size/breed category/gender/age cell as the annual claim for any missing dog/year observation.

VITA

EDUCATION

Ph.D. Economics, Department of Economics, Purdue University August 2020

Committee: Timothy N. Bond (chair), Kevin Mumford, Jack Barron, David Hummels

M.S. Economics, Purdue University May 2016

B.A. Economics, University of Notre Dame May 2012

Minor in Hesburgh Program in Public Service

FIELDS OF RESEARCH

Labor Economics, International Economics

WORKING PAPERS

"Stalled Racial Progress and Japanese Trade in the 1970s and 1980s" with Timothy N. Bond

"Skill-Biased Technical Change and Employment in U.S. Manufacturing"

"End-of-Life Medical Spending: Evidence from Pet Insurance" with Kevin J. Mumford

RESEARCH VISITS

Federal Reserve Bank of Minneapolis December 2019

University of Notre Dame June-July 2019

CONFERENCES AND PRESENTATIONS

2020: Compass Lexecon, Chicago, IL; University of Notre Dame; University of Nebraska – Lincoln; Mississippi State University; Oregon State University; Federal Trade Commission; College of William Mary; Bureau of Labor Statistics; U.S. Department of the Treasury; Joint Committee on Taxation; Bates White Economic Consulting, Washington, DC

2019: Federal Reserve Bank of Minneapolis; Southern Economic Association 89th Annual Meetings, Fort Lauderdale, FL; 31st EALE Conference, Uppsala, Sweden; Purdue University; University of Notre Dame; Midwest Economics Association's Annual Meetings, St. Louis, MO

2018: 30th EALE Conference, Lyon, France (poster session); NBER Summer Institute Labor Studies Workshop, Cambridge, MA (co-author presented); Econometric Society North American Summer Meeting, Davis, CA; Society of Labor Economics Annual Meetings, Toronto, ON; Royal Economic Society Junior Symposium, University of Sussex; Midwest Economics Association's Annual Meetings, Evanston, IL

2017: Purdue University

2016: 19th Annual Conference on Global Economic Analysis, World Bank, Washington DC

TEACHING EXPERIENCE

Economics of Race and Gender Discrimination (Undergraduate), *Teaching Assistant* Spring 2019, 2020
Principles of Microeconomics (Undergraduate), *Instructor* Summer 2016
Econometrics (Undergraduate), *Teaching Assistant* Spring 2016
Econometrics (PhD), *Teaching Assistant* Spring 2016
Econometrics (Masters), *Teaching Assistant* Spring 2015
Principles of Microeconomics (Undergraduate), *Teaching Assistant* Fall 2014, Spring 2015

REFEREEING EXPERIENCE

Journal of Human Capital

HONORS AND AWARDS

PURCE Award for Graduate Study	August 2014-August 2020
Purdue Graduate Student Government Travel Grant	2018
Krannert Distinguished Teacher	Summer 2016
Purdue Graduate Student Government Travel Grant	2016
Lyman Internship in Public Policy Award (University of Notre Dame)	2011

NON-ACADEMIC EMPLOYMENT

Legal Assistant, *Davis Polk & Wardwell LLP*, Washington DC	2012-2014