How Occupations Become Gendered: Social Learning and the Long-Run Evolution of Women's Work*

Myera Rashid†

[Click Here For Most Recent Version]

November 1, 2026

Abstract

Occupational segregation remains a persistent source of gender differences in labor market outcomes, with 20 percent of occupations today predominantly female and 45 percent male. These patterns of occupational gender segregation reflect historical processes that unfolded over generations. I study the dynamics underlying the process of occupational gender segregation using a large temporary shock to women's employment, the mass mobilization of men during World War II, which drew women into a broad set of new occupations. I find that, after the war, women were more likely to remain in occupations with lower physical demands and greater schedule flexibility. Women too young to respond directly to wartime labor demand were then more likely to enter these occupations, particularly if their mothers, older sisters, or female neighbors had wartime experience. I rationalize these empirical patterns in an equilibrium model of occupational choice in which women have imperfect information about occupations' gendered non-wage attributes and learn them through work experience and social learning, allowing information to transmit across cohorts of women.

*I thank my advisors Abi Adams, Leah Boustan, Walker Hanlon, and Joel Mokyr for their constant guidance and support. I thank Ellora Derenoncourt, Joseph Ferrie, Claudia Goldin, Rowena Gray, and Joan Martinez for generously sharing data with me. This project also benefited from the helpful feedback of Ran Abramitzky, Lukas Althoff, Eric Auerbach, Andrea Di Giovan Paolo, Katherine Eriksson, Grant Goehring, Elisa Jácome, Franco Malpassi, Giacomo Marcolin, Sarah Quincy, Martin Rotemberg, Molly Schnell, as well as participants at the NBER Gender in the Economy Mentoring Workshop, NBER Development of the American Economy Summer Institute, Economic History Association Annual Meeting, and Northwestern Economic History and Applied Microeconomics graduate workshops. I acknowledge generous financial support from the Russell Sage Foundation, Economic History Association's Sokoloff Fellowship and Early Stage Dissertation Grant, Northwestern Center for Economic History, Northwestern Morton Schapiro Center for Applied Economics, and the Robert Eisner Memorial Fellowship.

†Northwestern University, Department of Economics. E-mail: mrashid@u.northwestern.edu

1 Introduction

One of the most pronounced differences between men and women's labor market outcomes is their sorting into different occupations. Today, women account for roughly 90 percent of nurses but only 2 percent of carpenters. This pattern extends to over 65 percent of all occupations, which are either predominantly male or predominantly female (Online Appendix Figure A.1).[1] This persistent sorting can shape economy-wide outcomes including the gender wage gap, and has been the target of recent policy efforts to recruit and retain workers in occupations in which their gender is underrepresented, including efforts to draw women into the trades and men into nursing (Blau and Kahn, 2017; Blau et al., 2013).[2] Yet designing effective policy requires understanding how these patterns arise in the first place. Today's occupational segregation reflects changes in occupations' gender composition that unfolded over long periods of time: nursing, teaching, and secretarial work, for example, all shifted from men's work to women's work over the past century and a half. These historical shifts in the gender composition of occupations raise the question: how did occupations become gendered and why did occupations follow different gender trajectories?

This paper studies the dynamics underlying the process of occupational gender segregation using the mass mobilization of men during World War II. The draft of men during the war drew women into a broad set of new occupations that were previously predominantly male. Using quasi-exogenous variation in the number of men drafted from each occupation and county, I show that this temporary shock increased women's persistent employment by 1950, and that this increase was concentrated in occupations with lower physical demands and greater schedule flexibility, non-wage attributes measured prior to the war. Consistent with a process of social learning, I find that the effects spill over to the next cohort of women not directly affected by the war, who followed their older sisters, older female neighbors and mothers into treated occupations. Finally, I show that wages fell in occupation-county cells where female employment rose, suggesting that the added supply of new female workers surpassed any heightened labor demand for their work. I then construct a simple model of occupational choice that brings together two forces separately studied in the literature—imperfect information about occupations' non-wage gendered attributes and social learning across cohorts—and show that it can match my empirical findings and provides a general mechanism underlying the persistence and evolution of occupational gender segregation.

The empirical analysis and framework of this paper are anchored in a set of descriptive patterns about

[1] This includes roughly 20 percent of occupations which are over 70% female and 45 percent of occupations which are under 30% female.

[2] Such policies aimed at improving representation of men and women in non-traditional occupations have cost over $500 million federally since 1990. Some of these policies include the Nontraditional Employment for Women Act of 1991, the Women in Apprenticeship and Nontraditional Occupations program, Perkins vocational education acts, the Workforce Innovation and Opportunity Act, and state specific Department of Labor funded policies.

occupational gender segregation. Using 150 years of US Census data, I first document that the gender composition of occupations is not fixed, with occupations' gender compositions changing dynamically over time. Second, newer occupations gender faster than established ones, consistent with pre-existing occupations having entrenched beliefs which take longer to update. Third, an occupation's gender share converges nationally across labor markets, suggesting the role of information diffusion. Lastly, occupations that feminize experience wage declines relative to occupations that did not feminize. The rest of the paper investigates the forces underlying these patterns using a quasi-experimental setting and develops an equilibrium model that rationalizes the empirical findings and accounts for the broader descriptive patterns.

I examine how occupational gender segregation evolves using the mass mobilization of men during World War II as a quasi-experimental setting. This provides a useful setting because wartime mobilization sharply increased demand for women's labor across new occupations, drawing more than 6.3 million women, equivalent to over 15 percent of working-age women in 1940, into the labor force. Before the war, women's participation in the labor force was concentrated in a handful of occupations. Using data from the Palmer Survey (Palmer, 1954; Goldin, 1991), I observe that during the war women were drawn into a broad set of new occupations that were previously predominantly male, including welding, construction, bank-telling, accounting, and work as station agents.

I use variation in the number of men who were drafted from each local occupation during World War II to construct shocks to the demand for women's labor across local occupation-by-county cells. Combining WWII Army enlistment records with the 1940 Census, I observe enlisted men's pre-war civilian occupations and counties of residence and use this information to construct the number of men removed from each local occupation. A key concern is that this measure may not be exogenous, as local draft boards exercised discretion over deferment decisions, including deferring men whose occupations were important to local production. This implies that variation in the number of men who were drafted and enlisted across occupation-by-county cells may reflect local characteristics that also shaped women's employment.[3]

To isolate exogenous variation in draft exposure, I construct a shift-share instrumental variable (SSIV) which predicts the number of men drafted from each occupation-county. The SSIV interacts national draft probabilities for different demographic groups of workers, determined by their age, marital status, race, and place of birth, with each group's predetermined employment across occupation-by-county cells. My instrument thus isolates variation determined by aggregate wartime mobilization rather than by the choices of local draft boards. While similar shift-share designs conventionally generate variation across places, my instrument generates variation at the more granular occupation-by-county level, where the hypothesized

[3]WWII Army enlistment records allow me to observe only the subset of drafted men who were ultimately enlisted.

mechanism of labor replacement operates. This granularity also allows me to account for difficult-to-observe differences across occupations and counties, such as union presence that may have facilitated the return of veterans, by including occupation and county fixed effects.

I first establish that the wartime labor demand shock generated persistent increases in women's employment. Using the 1950 Census, I find that each additional man drafted from an occupation-county cell led to the employment of 0.53 additional women. This persistent response was driven primarily by women who were themselves old enough to enter the labor market during the war.

Next, I show that the effect on women's employment varied systematically with occupations' underlying non-wage attributes. Because historical data on occupational non-wage attributes are limited, I first use a continuous measure of occupations' physical intensity constructed from task content data in the 1939 Dictionary of Occupational Titles, compiled by Gray (2013). I find that the effect of the shock was mainly concentrated in occupations with less physically demanding tasks, with physical demands proxying for attributes that may be less valued by women. Specifically, the effect on women's persistent employment was five times as large in less physically demanding occupations as compared to more physically intensive occupations. I provide additional evidence consistent with sorting on non-wage amenities using measures of occupational cleanliness and flexibility, constructed using alternative data sources with more limited variation. These include data on hours worked collected for a sample of women in the 1940 Census and the use of a large language model trained on text written before World War II to describe occupational amenities (Levine et al., 2026).

The heterogeneous effects on women's persistent employment across occupations were driven by women selectively remaining in occupations with female-valued attributes rather than never having entered occupations without those attributes. Using data from the Palmer Survey to follow the career trajectories of women who entered new occupations during the war, I show that over 90 percent of the women who entered occupations with female-valued attributes during the war, such as foremen and bank-telling, remained in those occupations. In comparison, only 39 percent of the women who entered occupations without female-valued attributes, including welding and construction, remained in those occupations.

I show that the effects of the wartime shock extend to the next cohort of women entering the labor force, who follow older, war-exposed women—including their mothers, older sisters, and neighbors—into occupations affected by the wartime shock. I define the next cohort as women aged 18 to 22 in 1950, who were too young to have participated in the labor force during the war. I first show that each additional man drafted from an occupation-county cell led to the employment of an additional 0.18 young women in that cell. Next, I show that the effect on the next cohort's employment also varied with occupations' non-wage

attributes, with younger women sorting into occupations with more female-valued attributes. Specifically, the employment effect for younger women was 2.5 times as large in occupations that were less physically demanding as compared to occupations that were more physically intensive. Lastly, I connect younger women to older war-exposed women, including their older sisters, older female neighbors, and mothers. I find that the shock increased the number of young women who followed older connected women into their occupations by 10-13 percent. One potential concern with evidence from older connected women is that these patterns could in part reflect referral networks rather than social learning about occupations. To provide evidence consistent with the role of information specifically about occupations and their non-wage attributes, I use data on circulation rates for historical women's career magazines, which contained information about occupations and the experiences of women employed in them. I show that young women's employment response to the wartime shock was larger in places with greater circulation of such magazines. Together, these results suggest that information transmission about occupations and their non-wage attributes that wartime entrants were exposed to helped shape the occupational sorting of the next cohort of women.

Lastly, I show that in occupation-county cells with female-valued attributes, women's wages decreased after the shock. Specifically, per 10 drafted men, relative to their wages in 1940, women's wages in 1950 fell by 27% in occupations that were less physically demanding. In comparison, their wages remained unchanged in occupations that were more physically demanding. Using changes in women's mean years of schooling, I show that the shock did not alter the skill composition of women entering these occupation-county cells, suggesting that the outward shift in women's labor supply to these occupations is not accompanied by the employment of lower-skilled women. The negative effect on women's wages suggests the presence of a supply-side mechanism underlying the process of women's sorting across occupations and helps distinguish it from alternative demand-side channels including reductions in discrimination, which would predict a null or positive effect for women's wages.[4]

Guided by the empirical findings, I develop an equilibrium model of occupational choice which combines two forces studied separately in the literature: imperfect information about occupations' gendered non-wage attributes and social learning across generations of women. Specifically, in this model, women have imperfect information about occupations' non-wage attributes, and these attributes may be valued differently by men and women (Wiswall and Zafar, 2018; Mas and Pallais, 2017; Goldin, 2014b).[5] Women learn an occupation's non-wage attributes by working in it and remain in occupations whose attributes they value. Building on models of intergenerational learning about women's work, I then allow this information to pass across cohorts

[4]See, for example, Hsieh et al. (2019) on the role of declining barriers to occupational entry.

[5]Existing survey evidence shows that workers are misinformed about occupations and their non-wage attributes, including women holding inaccurate beliefs specifically about physical demands of different trades (see for example Gyarmati et al., 2017).

(Fogli and Veldkamp, 2011; Fernández, 2013). Specifically, I model this information as a public signal about the occupation's non-wage attributes, whose precision increases with the number of women who remain employed in the occupation.[6] Together, these forces can account for aggregate patterns of occupational sorting and provide a simple mechanism underlying the persistence and evolution of occupational gender segregation. When information about occupations' non-wage attributes is scarce, women may be uncertain about which occupations they would value, allowing existing patterns of segregation to persist. But when women enter new occupations, discover attributes they value, and transmit this information to subsequent cohorts, occupational gender compositions can shift. As women increasingly sort into occupations with female-valued non-wage attributes, the resulting increase in labor supply lowers women's wages in those occupations.

To match the empirical results, I introduce temporary shocks to the demand for women's labor into the model. This allows me to indirectly trace the mechanism of social learning about occupations' non-wage attributes through the observable effects of the shock on women's subsequent employment. In the model, a temporary increase in labor demand raises the equilibrium wage, drawing women into occupations about which they have limited information. Once the shock subsides, entrants decide whether to remain based on the attributes they have learned, so identical shocks can generate different employment dynamics across occupations depending on their underlying non-wage attributes. I show that the resulting implications of this model for the effects on women's employment across occupations are consistent with my empirical results using the draft of men during World War II as a temporary shock.

My findings shed light on a central question in the literature on occupational gender segregation: how do occupations become gendered? A large body of work in this literature documents the extent and persistence of occupational gender segregation and its contribution to the gender wage gap (Bergmann, 1974; Blau et al., 2013; Blau and Kahn, 2017), and studies the dynamics that reinforce segregation once it is underway, including tipping (Pan, 2015) and men's responses to female entry (Goldin, 2014a; Greenberg et al., 2024).[7] In this paper, I provide empirical evidence and develop a theoretical framework that together provide an explanation for how occupations can become gendered over time. Specifically, my findings highlight that an occupation's gender is not a fixed attribute of the work itself and that social learning across generations of women about imperfectly observed non-wage attributes of occupations can shape occupations' gender compositions.

[6] I model information transmission through women who remain employed because their occupational experiences are readily observable to subsequent cohorts. The experiences of women who leave are less likely to be visible to women considering entry into that occupation.

[7] This literature is also summarized in Blau and Winkler (2018). In related work, Alon et al. (2025) study occupational gender segregation across industries. Cross (2025) studies the spillover effects of desegregation of Army occupations.

The theoretical framework in this paper brings together two bodies of work that have been studied largely separately: gender differences in the valuation of non-wage job attributes and intergenerational learning about women's work. The first literature documents that men and women value workplace amenities differently and studies how these differences shape job choice and earnings (Wasserman, 2023; Sharma, 2023; Le Barbanchon et al., 2021; Cortés and Pan, 2019; Wiswall and Zafar, 2018; Mas and Pallais, 2017; Goldin, 2014b; Goldin and Katz, 2011). The second literature studies how beliefs about the consequences of women's work evolve across generations and shape aggregate women's labor force participation (Olivetti et al., 2020; Fernández, 2013; Fogli and Veldkamp, 2011). I show that a simple model that combines insights from these two literatures can generate the dynamic patterns of occupational gender segregation that I document.

My findings also contribute to a large body of work in economic history examining factors that led to changes in women's labor force outcomes in the US (Goldin, 1990; Greenwood et al., 2005; Goldin and Katz, 2002; Bailey, 2006; Myers, 2017; Albanesi and Olivetti, 2016; Attanasio et al., 2008; Vidart, 2024). Substantial research within this literature has studied the persistent effects of World War II on women's employment (Acemoglu et al., 2004; Goldin and Olivetti, 2013; Goldin, 1991; Rose, 2018). By exploiting plausibly exogenous variation in male mobilization rates at the more granular occupation-by-county level, I am able to provide a more nuanced understanding of the persistent effects of World War II on women's employment. My findings show that while male mobilization drew women into a broad set of occupations, its persistent effects on women's employment were concentrated in occupations with female-valued attributes. More broadly, the framework I develop also suggests that large temporary shocks, such as war mobilization, can accelerate the process of occupational gender segregation by drawing women into new occupations.

Finally, this paper also relates to work on women's entry and employment in individual occupations, including secretarial and clerical work (Rashid, 2026; Rotella, 1981), nursing (Bald, 2025), teaching (Perlmann and Margo, 2001), pharmacy (Goldin and Katz, 2016), public administration (Aneja et al., 2024), and bank telling (Strober and Arnold, 1987). These studies document how particular occupations opened to women and, over time, became feminized. I bring these occupation-specific episodes together within a common framework by showing that shocks to the demand for women's labor draw them into these occupations, while mechanisms such as social learning about non-wage amenities contribute to their feminization over time.[8]

[8] Rashid (2026) also applies the framework of this paper to the case of a specific occupation that became feminized after a technological shock drew women into it. That paper focuses on studying the consequences of such gendering of an occupation.

2 Patterns of Occupational Gendering in the US

To understand the dynamics underlying occupational gendering in the long run, I first document a set of descriptive patterns using 150 years of US Census data (Ruggles et al., 2025, 2024). For the decades 1870 to 1950, I use complete count censuses and for the decades 1960 to 2019, I use samples of Census and ACS data (for details of this data construction see Online Appendix A.1). These data allow me to follow a set of 203 detailed occupations in each decade between 1870 and 2019. This set of occupations includes 137 occupations that can be followed back to 1870 and 66 occupations that emerged in the decades after. This set of occupations does not include any general occupation categories such as "clerical workers" or "technical workers" (for a full list of occupations included see Online Appendix Table A.1). Of the 203 occupations, 188 began as predominantly male occupations.[9] Over the following century and a half, 55 of those occupations (29%) switched permanently from predominantly male to predominantly female, that is, their national workforce became and remained more than 50% female.

Using these data, I document four key descriptive patterns of occupational gendering in the data. First, the gender composition of occupations is not fixed, as the same occupations de-segregate and re-segregate over time. Second, occupations gender at different speeds. Newly emerging occupations, and occupations whose skill requirements shift, become gendered faster than established ones. Third, gendering occurs nationally, with an occupation's female share co-moving across local labor markets. Lastly, wages fall in occupations that feminize. I also detail how these descriptive patterns help rule out several alternative explanations for occupational gendering and help guide the empirical analysis and framework that I develop.

2.1 Occupational Gendering is Dynamic

Figure 1, Panel A summarizes trends in occupational gendering for six occupations: secretaries and typists, practical nurses, bookkeepers, accountants, pharmacists, and carpenters. While practical nursing and carpentry experienced little change in their gender composition over this period, the remaining occupations began predominantly male and de-segregated and re-segregated at different points in time, eventually becoming predominantly female. Similar patterns of occupational gendering appear across a total of 55 occupations which switched from predominantly male to predominantly female over this period, constituting more than a quarter of all occupations in the data (see Online Appendix Table A.2).

One notable exception is computing, which switched from a predominantly female to a predominantly

[9]Here for consistency across decades, predominantly male occupations are defined as those occupations which are less than 50% female. This differs from the definition used in the literature which defines predominantly male or female occupations relative to that gender's labor force participation within a decade (Blau and Winkler, 2018).

male occupation over the mid-twentieth century (Abbate, 2012).[10] However, US Census data do not allow me to track this occupation consistently, as women who worked in computing historically were recorded as general technical workers or operatives and cannot be distinguished from other workers in those broad categories. The model I develop in Section 7 is designed to capture gendering in both directions, from predominantly male to predominantly female, the direction that is more prevalent in the data, and the reverse, as in the case of computing.

2.2 Variation in Speed of Gendering

Figure 1, Panel A also demonstrates that there is significant variation in the number of years occupations take to gender. Secretarial and typing work switched to predominantly female within a few decades, while accounting took more than a century to make the same transition. Figure 1, Panel B summarizes this variation for the 55 occupations that switched, plotting the number of years between each occupation's emergence and the census in which it became predominantly female. New occupations, defined as those first employing more than 500 workers after 1870, include occupations that emerged in the decades after 1870, such as librarians and medical and dental technicians, as well as occupations whose required skills shifted, such as secretarial work after the arrival of the typewriter. These occupations became predominantly female within a median of 60 years of emerging. Old occupations, which already employed more than 500 workers in 1870, took a median of at least 110 years to become more than 50 percent female. See Online Appendix Table A.3 for the full list of new and old occupations and the years each took to become predominantly female.

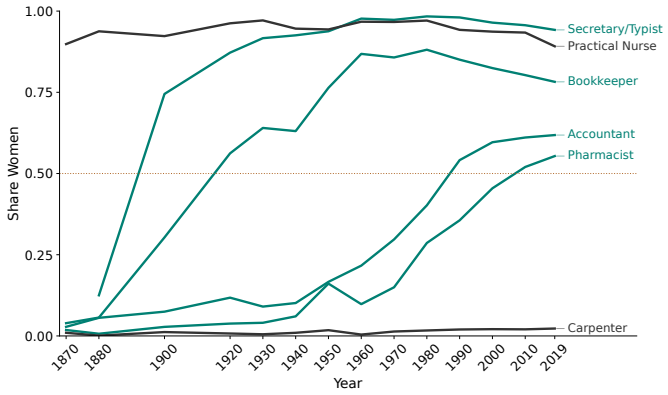
2.3 Occupations Become Nationally Gendered

Figure 1, Panel C illustrates, using the example of the bookkeeping occupation, that occupations become gendered nationally. Specifically, it plots the share of women in bookkeeping over time for each US state and shows that the occupation feminized in every state over the same decades rather than at different times across labor markets. This pattern is not specific to bookkeeping. Online Appendix Figure A.2 summarizes this process for all 55 switching occupations and shows that geographic dispersion in an occupation's female share falls as the occupation feminizes. This suggests that gendering is a national process and does not take place at different times across labor markets, separating it from local processes such as homophily (Pan,

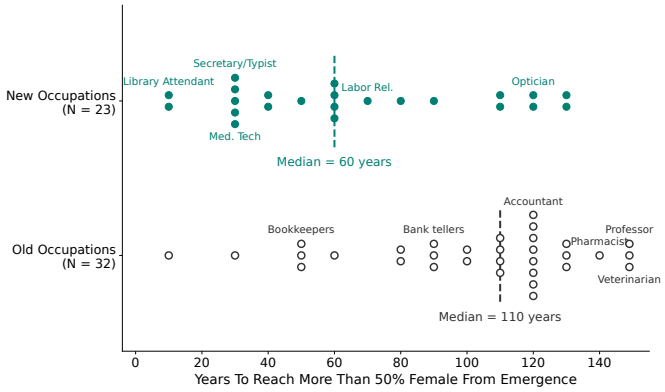
[10] Zhao and Martínez-Marquina (2026) document one channel contributing to the gendering of computing to a predominantly male occupation. They find that women's share of computer science degrees fell as computer science departments moved from colleges of liberal arts and sciences into schools of engineering, and provide experimental evidence that the engineering affiliation deterred women through perceived parental approval.

Figure 1: Patterns of Occupational Gendering in the US: 1870–2019

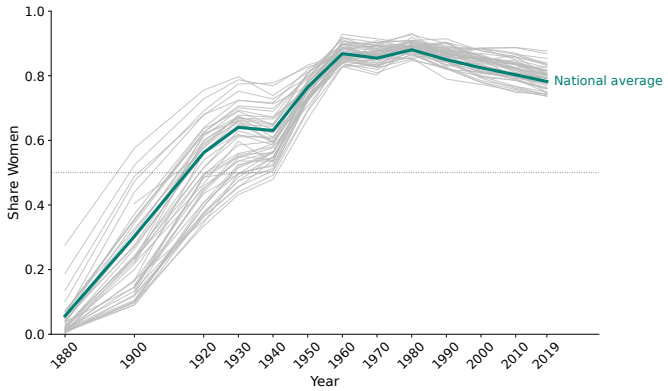
Panel A: Dynamics of Occupational Gendering



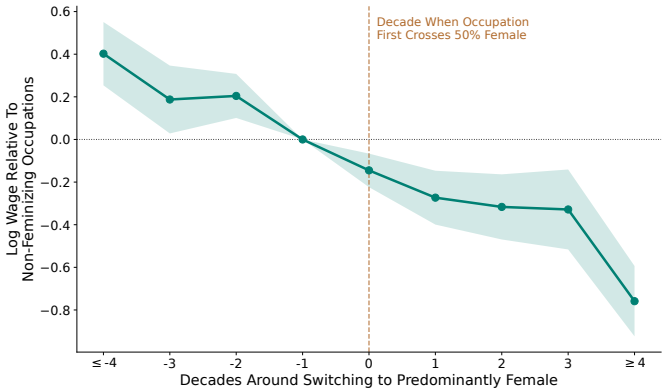
Panel B: Variation in Speed of Gendering



Panel C: National Gendering Example: Bookkeeping



Panel D: Falling Wages for Feminized Occupations



NOTE: This figure documents descriptive patterns of occupational gendering using US Census and ACS data from 1870 to 2019, with occupations classified using the *occ1950* coding. An occupation is predominantly male (female) when women make up less than (more than) 50 percent of its national workforce, and an occupation switches permanently when it crosses 50 percent female without reverting in any subsequent census. In all panels, occupation cells with less than 100 workers are omitted. **Panel A** plots the national female share of employment for six occupations. Occupations in green switched permanently from predominantly male to predominantly female, while occupations in black retained their initial gender composition. **Panel B** plots, for each permanently switching occupation, the number of years between the occupation's first census and the census in which it permanently crossed 50 percent female. New occupations are defined as those which first employ more than 500 workers only after 1870. Dashed lines mark group medians. **Panel C** plots the female share of bookkeepers in each of the 48 continental US states in gray, excluding the District of Columbia, and the employment-weighted national average in green. **Panel D** plots mean income in permanently switching occupations relative to occupations that began predominantly male and never crossed 50 percent female, in decades around each occupation's switch. The outcome is the log of an occupation's mean income, with incomes adjusted to 2019 dollars and measured from 1940 onward. Data from the 1910 Census is not used because women's participation in occupations is recorded differently in that Census (see Goldin, 1990 for additional details).

2015; Schuh, 2024).

2.4 Wages Fall Once Occupations Feminize

Lastly, Figure 1, Panel D shows that occupations that switch from being predominantly male to predominantly female experience falling mean wages after feminizing relative to occupations which do not switch. Specifically, mean wages fall by roughly 30 log points within two decades of an occupation's switch and by more than 70 log points over longer time periods. Online Appendix Figure A.3 shows that wages for men in feminized occupations also fall, suggesting that this pattern does not simply reflect the changing gender composition of workers within occupations. This pattern is also consistent with a body of work which documents that predominantly female occupations pay less (Macpherson and Hirsch, 1995; England et al., 2007; Levanon et al., 2009). Falling wages upon feminization are difficult to reconcile with alternative accounts of the gendering process, including employer learning about women's productivity, which would imply rising wages as women enter. On the other hand, Online Appendix Figure A.4 shows that occupations which do not feminize but still see some women enter them remain stable in their relative mean wages, a pattern less consistent with accounts in which any female entry devalues an occupation's pay (England et al., 2007; Levanon et al., 2009).

2.5 Alternative Mechanisms Inconsistent with Patterns

Together, these descriptive patterns help guide both the empirical analysis and the model that I develop to study the dynamics underlying occupational gendering. These patterns also help assess the plausibility of alternative explanations that might underlie the process of gendering. Below, I first touch upon such alternative explanations and detail how predictions from those explanations would be inconsistent with the patterns documented above. I then outline how the mechanism of the framework that I develop in Section 7 can account for these patterns.

The first descriptive pattern showing that the same occupation can de-segregate and re-segregate over time helps rule out static explanations of occupational gendering, including skill requirements, fixed comparative advantage, and stable gendered preferences, which would predict stable gender compositions of occupations in the long run (Polachek, 1981; Lordan and Pischke, 2022). Dynamic versions of these explanations would imply broader movements of women across occupations rather than particular occupations gendering at specific points in time while other comparable occupations do not.

The documented descriptive patterns are also inconsistent with explanations which attribute changes in

gender segregation to changes in employer discrimination (Hsieh et al., 2019; Alon et al., 2025). Declines in discrimination driven by local attitudes would present as local dynamics rather than nationally synchronized changes. Additionally, changes driven by national legislation would move women into many occupations at once rather than into particular occupations at specific points in time.

Another explanation attributes occupational gendering to homophily, that is, women's preference for working alongside other women, which can generate tipping once an occupation's share of women passes a threshold (Pan, 2015; Schuh, 2024). Such a mechanism, however, would operate at the local level and would suggest that occupations gender across local labor market at different points in time, with occupational gendering starting in markets where women's presence first passes the threshold. Figure 1, Panel C instead shows that gendering takes place nationally, with labor markets moving together at the same time. Local homophily also does not explain how occupations begin to gender in the first place.

Lastly, a channel whereby employers learn about women's productivity, holding female labor supply fixed in the short term, would predict that occupations which feminize should experience rising wages as employers observe women's true productivity, update positively, and continue hiring women (Altonji and Pierret, 2001). However, Figure 1, Panel D shows that wages for all workers, including those for men and women, fall when an occupation feminizes, making this pattern difficult to reconcile with an explanation of occupational gendering where employers learn (see Online Appendix Figure A.3 for patterns for men and women's wages separately).

The empirical analysis and the model of this paper instead highlight a mechanism of social learning about imperfectly observed non-wage attributes of occupations. This mechanism is consistent with each of the four key descriptive patterns. Such a mechanism predicts occupational gendering as dynamic because beliefs about an occupation's attribute update as women enter and learn, allowing an occupation to become feminized when women learn favorable information about its attributes and to remain predominantly male when they do not. In such a framework, newer occupations gender faster because beliefs about them are less entrenched. The social learning mechanism importantly allows for national gendering because an occupation's attributes are relevant across labor markets and information can travel through social networks and media. Lastly, this mechanism is also consistent with falling wages as the occupation feminizes, as women accept lower wages to work in occupations that they learn offer female-valued attributes. I formally develop this framework in Section 3 below.

3 Mass Mobilization of Men During World War II

I study the dynamics underlying the process of occupations becoming gendered in the setting of mass mobilization during World War II (WWII). In this section, I describe the historical process by which men were drafted and inducted into the US Army during WWII. I also discuss how the mobilization of men during the war translated into increased demand for women's labor across occupations.

3.1 Historical Background on WWII Draft

The Selective Training and Service Act of 1940, built on the Selective Service System first introduced during World War I, established the first peacetime draft in U.S. history. At first, all men aged 21–35 were required to register for Selective Service. However, shortly after the US formally entered WWII in December 1941, amendments to the Selective Training and Service Act expanded the registration requirement to all men aged 18–64. However, only men aged 20–44 were ever called to serve through the Selective Service draft.

The mobilization of 9.8 million men through the Selective Service draft represented a large shock to male labor supply at the time. This was equivalent to the removal of over 35% of all men aged 18-44 from the labor force at the time. Men mobilized through the draft represented 65% of all men who served during the war, where the rest included volunteers or those who were already in the armed forces. In my analysis, I focus on men inducted into the Army through the Selective Service draft, which made up approximately 82% of all drafted men. I do not use enlistment records for other branches of the armed forces (e.g., the Navy and Marine Corps) because they make up only a small share of all drafted men and the available records do not allow me to separately identify Selective Service draftees from voluntary enlistees.

The removal of men during the war also translated into a shock to the demand for women's labor resulting in the employment of over 6.3 million women during the war, equivalent to over 15 percent of all working age women in 1940 (U.S. Department of Labor, Women's Bureau, 1952). Demand for women's labor began increasing in particular after the US joined the war and accelerated as "new registrations [took] great numbers of men for the Army away" (U.S. Department of Labor, Women's Bureau, 1942). As a result of the widespread demand shock, women entered new occupation including working as taxi drivers, welders, bank tellers, and medical technicians (U.S. Department of Labor, Women's Bureau, 1942; McMillin, 1943).

3.2 Role of Local Draft Boards

Over 6,400 local draft boards were responsible for processing deferment decisions for men who were called to serve through the Selective Service draft. Despite there being a public national lottery to select order

numbers for registered men, counts of men who were drafted and enlisted from each local occupation were not random.[11] This is because local draft boards were instructed to defer men whose occupations were important to local production. In fact, Selective Service regulations specified that "on the local board is placed the responsibility of deciding which men should be deferred because of their civilian activities" (Executive Order 8560, 1940). This implies that counts of men who were drafted and enlisted may be correlated with local importance of occupations.

4 Data

To construct shocks to the demand for women's labor arising from the mass mobilization of men in each occupation-by-county cell during WWII, I combine WWII Army enlistment records (NARA 2002) with data from the 1940 U.S. Census (Ruggles et al., 2021). To study the mechanisms underlying the process of occupational gendering, I use a variety of additional data sources, including survey data on women's employment during the war, data to construct occupational non-wage attributes, and data about historical women's career magazines. Below, I describe each of these data sources and the relevant information I draw from them. The constructed dataset that I use in my analysis contains 369,185 occupation-by-county cells.

4.1 WWII Army Enlistment Records

The WWII Army Enlistment Records provide information on 9 million men who entered the Army between 1940 and 1945. The information contained in these records was digitized from Army Serial Number punch cards. The cards contain information collected when individuals entered Army service and include both voluntary enlistees and Selective Service inductees. Within these data, 6.1 million men are recorded as Selective Service inductees, representing 68% of all enlistees. From these data, I use information on men's civilian occupations and counties of residence before enlisting in the Army. The data also contain demographic characteristics of enlistees, including their year and place of birth, gender, race, and marital status. I use these characteristics to construct the shift-share instrumental variable described in Section 6. For each variable used in the analysis, I hand-map the categories in the enlistment records to the corresponding categories in the Census data.

[11] The first three waves of the draft used this lottery. Later waves assigned order numbers for men primarily based on their age.

4.2 Census and ACS Data, 1900-2019

My analysis also relies on historical population censuses from 1900 to 1950, as well as samples from the decennial censuses and the American Community Survey (ACS) covering the decades between 1960 and 2019 (Ruggles et al., 2021, 2025). I aggregate these data to the occupation-by-county level using three-digit detailed occupation codes and geographic areas mapped to counties. I use these data to compute the number of men and women employed in each occupation in each county. For the analysis of wages, I restrict the sample to 1940 onward, as 1940 was the first Census to collect information on individual wages. I also use information on workers' demographic characteristics and industries to construct the shift-share instrument and control variables used throughout the analysis.

4.3 Women's Employment in Occupations During the War

To study the persistent effect of men's mobilization on women's employment within occupations and counties, I use Census data after 1940. However, because the next Census was conducted in 1950, these data do not allow me to observe women's employment during the war. To study women's employment across occupations during the war, I use data from the Palmer Survey (Palmer, 1954; Goldin, 1991). This survey collected retrospective work histories for 4,350 women who worked full-time for at least one month in 1950. I use these data to study the career paths of nearly 1,600 women of working age (16-64) in 1940 who entered new occupations during WWII.[12] I define women as having a new occupation during the war if they report an occupation in 1944 that differs from their reported occupation in 1940.

4.4 Women's Historical Career Magazine

I use data on circulation rates across states of the first women's career magazine, Mademoiselle, which was introduced in 1935 (Hefner et al., 2023). Issues of this magazine covered information about amenities associated with new occupations for women, often through interviews with women working in them. I use these data to compare women's employment responses across states with higher and lower circulation rates, which proxy for women's exposure to information about occupations and their attributes.

[12] Existing studies of the effect of WWII mobilization on women's employment have not examined employment across occupations. For example, the data used by Rose (2018) identify workers' industries of employment but not their occupations.

4.5 Occupational Task Content and Amenities

I proxy for occupational non-wage attributes using historical measures of occupational task content from the 1939 Dictionary of Occupational Titles compiled by Gray (2013); Gray et al. (2025). Following Autor et al. (2003), these measures rank occupations along three dimensions: physical, routine-manual, and nonroutine analytical task intensity. In my analysis, I mainly focus on the physical task content, where a more physical task content proxies for non-wage attributes that are less valued by women. Highly physical occupations in my data include athletes, housekeepers, and laborers. I validate the use of this measure as a proxy for occupational non-wage amenities by showing that attributes such as low physical demands are strongly correlated with more direct measures of female-valued amenities in contemporary O*NET data, including schedule flexibility (see Online Appendix Table C.5).

I complement the analysis based on occupations' physical task content with constructed measures of historical occupational amenities. Using data on hours worked for a sample of women in the 1940 Census, I construct a measure of the average number of hours worked in each occupation. I also use Talkie-1930, a historical language model trained exclusively on text published before 1931 (Levine et al., 2026; Merchant and Levy, 2026), to construct additional measures of historical occupational amenities, including schedule flexibility and workplace cleanliness. In ongoing work, I am digitizing information on occupational amenities contained in historical career magazines.

5 Empirical Strategy

To empirically study the dynamics underlying occupational gender segregation in the setting of mass mobilization in WWII, I use variation in the number of men who were drafted into the Army as part of the Selective Service process from each occupation-county cell. However, ordinary least squares (OLS) estimates of the relationship between the number of men drafted-and-enlisted into the Army and women's employment across occupations and counties can be biased. Local draft boards had considerable discretion over which drafted men were ultimately enlisted based on the local importance of their occupations. The number of men who were drafted-and-enlisted may therefore be correlated with unobserved occupation-county characteristics. Consistent with the historical narrative, I expect OLS estimates to be biased downward if local boards deferred more men in occupations important to local production, as these same cells would also have had greater demand to replace or supplement men with women. Conversely, OLS estimates could be biased upward if boards inducted more men from cells with more progressive gender norms, where women were both more likely to work and could more readily replace men.

To address these biases, I construct a shift-share instrumental variable (SSIV) at the occupation-by-county level which combines variation in national draft risk across demographic groups with the predetermined number of workers from each group in each cell. Importantly, this instrument allows me to preserve the granularity of the analysis at the occupation-by-county level. The instrument isolates variation in the number of men drafted arising from differences in national draft risk across demographic groups, determined by aggregate wartime mobilization rather than by the choice of local draft boards. Thus, the instrument is constructed to capture the component of the draft that is orthogonal to omitted variables at the occupation-by-county level such as unobserved local draft-board considerations that could otherwise bias OLS estimates. The intuition behind this approach is well captured by the machinist occupation in two Pennsylvania counties: Blair County and Delaware County. In 1940, both counties employed a similar number of married male machinists aged 31–35. However, Delaware County employed 10 times as many single male machinists aged 21–25. During the war, single men aged 21–25 were 5.4 times more likely to be drafted and inducted nationally as married men aged 31–35. As a result of the differences in national draft probabilities across the two groups of workers, Delaware County lost a greater number of male machinists to the draft than Blair County (see Online Appendix Figure B.1 for a visual representation of this example from the data).

The shift-share instrumental variable generalizes this intuition and draws on two sources of variation. The first is the number of workers belonging to different demographic groups g in each occupation-county cell, oc, in 1940, which I denote by $\text{Emp}_{ocg,1940}$. I use the following demographic characteristics to split workers into 168 mutually exclusive groups:

$$g = \{\text{Age Group} \times \text{Race} \times \text{Gender} \times \text{Marital Status} \times \text{Foreign Born Status}\} \tag{1}$$

Specifically, I construct 7 different age groups: 16–20, 21–25, 26–30, 31–35, 36–40, 41–45, and 46–64. I use 3 categories for race: White, Black, and Other. Gender, marital status, and foreign-born status are all binary.[13]

The second source of variation that I construct is each demographic group's national probability of being drafted and enlisted as a Selective Service inductee during the war. I define these national shifts across demographic groups, $\Pr(\text{Draft}_{g,1940\text{-}45})$, as:

$$\Pr(\text{Draft}_{g,1940\text{-}45}) = \frac{\text{NumberDrafted}_{g,1940\text{-}45}}{\text{Population}_{g,1940}} \tag{2}$$

[13]Even though the national draft probabilities for many of these demographic groups are zero (e.g., for women), these groups are included so that all workers aged 16–64—the underlying working-age population—are represented. This construction ensures that the shift-share design does not fall into what Borusyak et al. (2022) refer to as the "incomplete shares" case.

The demographic groups g, as defined in equation (1), result in 168 shifts (see Online Appendix Figure B.2 for the distribution of shifts across demographic groups). I show that these shifts are balanced across a range of observable characteristics of the demographic groups in the 1940 Census (see Online Appendix Figure B.3).

Then, interacting predetermined employment of workers of different demographic groups across occupation-county cells in 1940 with their national draft risk during WWII allows me to identify exogenous variation in the number of men drafted from each local labor market. I formally define the SSIV, the predicted number drafted from each occupation-county cell, $\widehat{\text{NumberDrafted}}_{oc}$, as:

$$\widehat{\text{NumberDrafted}}_{oc} = \sum_{g} \Pr(\text{Draft}_{g,\text{1940-45}}) \cdot \text{Emp}_{ocg,1940} \tag{3}$$

Online Appendix Figure B.5 plots the variation in the SSIV across occupation-county cells. The variation in the SSIV closely mirrors the variation in the the number of men drafted from each occupation-county cell (see Online Appendix Figure B.4). Each occupation-county cell had an average number of 8.14 men drafted and 7.60 predicted to be drafted based on the SSIV. Given that both the instrument and the endogenous variable have left-skewed distributions, I show that my main estimates are robust to winsorizing both variables.

To account for differences across counties and across occupations, all specifications include occupation fixed effects and county fixed effects. To focus on variation arising from demographic group-level shifts, rather than baseline differences in a cell's characteristics, I control for predetermined occupation-by-county demographic and industry characteristics measured in 1940. These controls include the number of women, the cell population, the number of married men, and the number of workers in agriculture, mining, manufacturing, and public administration. In robustness tests, I show that my results are robust to including additional controls selected using double LASSO control selection (Belloni et al., 2014).

Estimating Equation. I estimate the relationship between the number of men drafted from each occupation-county cell and women's employment in 1950 using the following empirical framework:

$$y_{oc,1950} = \alpha + \beta \text{MenDrafted}_{oc} + \mathbb{X}_{oc,1940}\gamma + \omega_o + \phi_c + \varepsilon_{oc}. \tag{4}$$

$$\text{First Stage:} \quad \text{MenDrafted}_{oc} = \mu + \lambda \widehat{\text{NumberDrafted}}_{oc} + \mathbb{X}_{oc,1940}\Gamma + \Omega_o + \Phi_c + \epsilon_{oc}. \tag{5}$$

In equation (4), β represents the OLS estimate of the effect of the number of men drafted from each

occupation-county cell, MenDrafted$_{oc}$, on the number of women employed in those cells in 1950, $y_{oc,1950}$, conditional on occupation fixed effects ω_o, county fixed effects ϕ_c, and baseline occupation-by-county characteristics represented by the control vector $\mathbb{X}_{c,1940}$. I measure both the outcome and the endogenous variable as numbers rather than shares because the hypothesized mechanism operates through labor replacement. The relevant shock is therefore the number of men removed from an occupation-county labor market. If employers or local labor markets respond by drawing women into the same occupation-county cell, this response should be proportional to the number of male vacancies created rather than to the percentage change in the number of men drafted. My baseline sample consists of 369,185 cells. I cluster standard errors at the county level and show robustness of standard errors to accounting for different forms of spatial correlation (Conley, 1999; Adao et al., 2019). The reduced form equation is as follows:

$$y_{oc,1950} = \tilde{\alpha} + \tilde{\beta}\widehat{\text{NumberDrafted}}_{oc} + \mathbb{X}_{oc,1940}\tilde{\gamma} + \tilde{\omega}_o + \tilde{\phi}_c + \tilde{\varepsilon}_{oc}, \tag{6}$$

where the two-stage least squares or 2SLS ($\tilde{\beta}/\lambda$) estimate captures the effect of an additional man drafted on the number of women employed in an occupation-county cell in 1950 $y_{oc,1950}$. In the text, I mainly focus on the IV estimates. However, for all main results, I also report the estimated OLS (β), first stage (λ), and reduced form ($\tilde{\beta}$) coefficients.

Identifying Assumption and Validity Checks. In order for this empirical specification to identify the causal impact of male mobilization, conditional on the specified baseline 1940 characteristics, occupation fixed effects, and county fixed effects, my instrument for the number of men drafted during WWII must be orthogonal to omitted characteristics that are correlated with women's employment in 1950. This identifying assumption can be stated formally as:

$$\mathbb{E}[\widehat{\text{NumberDrafted}}_{oc} \times \tilde{\varepsilon}_{oc} | \mathbb{X}_{oc,1940}, \tilde{\omega}_o, \tilde{\phi}_c] = 0 \tag{7}$$

Although this assumption cannot be directly tested, relying on varying national draft probabilities across worker groups based on their predetermined demographics assuages concerns that my instrument for the number of drafted men is correlated with unobserved determinants of women's outcomes across occupation-county cells. Furthermore, a key advantage of my SSIV design is that identification does not require the exogeneity of the number of workers of each group in each occupation-county which, though predetermined, may be endogenous to a cell's employment of women in 1950 through persistent unobserved cell characteristics that jointly shape women's work and 1940 worker composition. In contrast, following

Borusyak et al. (2022), my SSIV can satisfy the exclusion restriction even if the initial numbers of workers are endogenous insofar as the national shifts are exogenous to an occupation-county cell's conditions and unobserved determinants of women's employment.

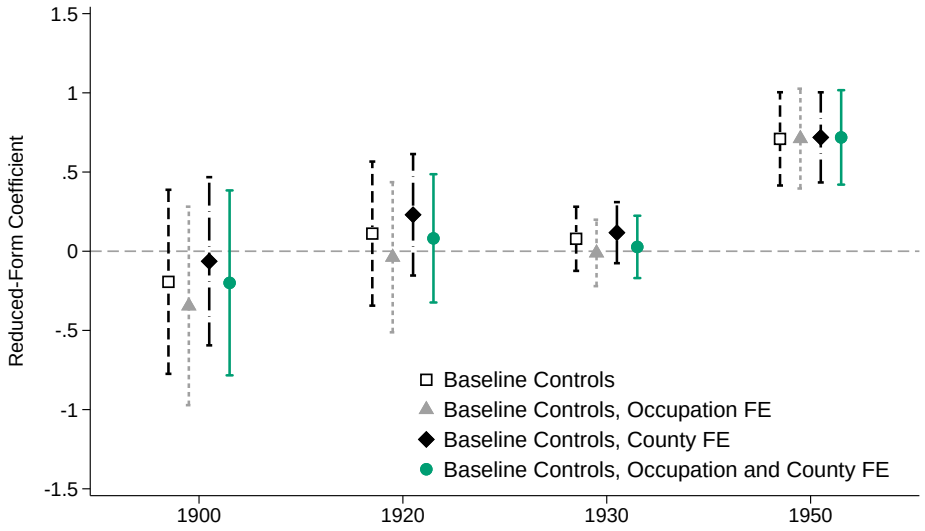
I provide two key pieces of supporting evidence for this shifts-based identification. First, using a procedure developed by Adao et al. (2019), I show that the predetermined number of workers of different demographic groups are not driving SSIV identification. I replace the shifts $\Pr(\text{Draft}_{g,1940\text{-}45})$ with randomly generated shifts $\Pr(\text{Draft}_{g,1940\text{-}45})^{rand}$ and find a negligible share of statistically significant results across 1,000 trials.[14] This provides evidence in favor of the shifts, rather than the potentially endogenous baseline employment numbers alone, driving the identifying variation underlying the SSIV.

Next, I provide a test of pretrends and show that, conditional on the specified baseline 1940 occupation-county characteristics, occupation fixed effects and county fixed effects, the instrument is uncorrelated with women's employment across cells prior to 1940. Figure 2 reports the coefficients on $\widehat{\text{NumberDrafted}}_{oc}$ from the reduced form model specified in equation (6) using the number of women employed in an occupation-county in the 1900, 1920 and 1930 Censuses as outcome. The coefficients are close to 0 in magnitude and statistically insignificant across three decades prior to 1940. I also plot the reduced form effect of the instrument on the number of women employed in 1950 for comparison.[15]

[14] Specifically, I find 2.7% of the 2SLS coefficients are statistically significant (positive or negative) at the 1% level and 11.5% at the 5% level (see Online Appendix Figure C.5). In comparison, Bazzi et al. (2023), who also use a shifts-based SSIV, found that 5.4% of their IV coefficients were significant at the 1% level and 11.5% at the 5% level.

[15] I do not use the 1910 Census as it recorded women's employment differently than across other Censuses. See Goldin (1990) for additional information.

Figure 2: Reduced Form Effect of SSIV on Women's Employment, 1900-1950



NOTE: This figure plots reduced-form coefficients for the effect of the draft instrument on the number of women employed in an occupation-county cell in each census decade from 1900-1950. The coefficients correspond to regressing the number of women employed in a cell in each decade separately on the predicted number of drafted men. The unit of observation is an occupation-county cell. Coefficients for 1900, 1920, and 1930 serve as placebo tests. 1950 coefficient captures the effect of the draft after the war. All regressions include baseline 1940 controls. Baseline 1940 controls include the following demographic and industry controls: number of women, cell population, number of married men, number of workers in the agricultural, mining, manufacturing, and public administration industries. Hollow squares report estimates with baseline controls only. Gray triangles add occupation fixed effects. Black diamonds add county fixed effects. Turquoise circles include both occupation and county fixed effects. 1940 is excluded because baseline 1940 controls include the number of women in the cell in 1940. 1910 is excluded because women's employment in this census is measured differently from other census decades (see Goldin, 1990 for additional details). Standard errors are clustered at the county level and bars show 95% confidence intervals.

6 Empirical Results

In this section I use WWII mobilization as a shock to local labor markets' male labor supply and study its impact on women's long-term employment. I do not interpret the war as the cause of long-run occupational gendering, and the model developed in Section 7 does not require such a wartime shock in particular. Rather, the WWII draft offers a rare setting to study the mechanisms underlying occupational gendering, where such a shock drew women into a broad set of new occupations which were previously predominantly male.

6.1 Effect on Women's Employment in 1950

Table 1 first reports estimates of λ from equation (5) for the first-stage relationship between the number of men drafted from each cell and the SSIV, the predicted number drafted. Conditional on baseline county characteristics, occupation fixed effects, and county fixed effects, an additional predicted man drafted is associated with 1.35 men drafted from each cell (column 4). The F-statistic on the first stage is 64.86. The first stage relationship remains stable and strong across specifications where only baseline controls are used (column 1), only baseline controls and county fixed effects are included (column 2), and when only baseline controls and occupation fixed effects are included (column 3). Also see Online Appendix Figure B.6 for a binned scatterplot depicting the relationship between the actual and predicted numbers drafted from each cell.

Table 1 Panel D reports the persistent effect of a higher shock to the demand for women's labor on their employment in 1950. The 2SLS coefficient shows that conditional on baseline controls, occupation and county fixed effects, an additional man drafted from an occupation-county cell led to the persistent employment of 0.53 additional women by 1950 (column 4). Additionally, the reduced form estimate of $\tilde{\beta}$ from equation (6) imply that an additional predicted man drafted from each cell led to the employment of 0.72 additional women in 1950. Table 1 Panel B also reports the OLS estimates which are nearly 1.7 times smaller than the IV estimates and consistent with the downward bias discussed in Section 6 arising from local draft boards deferring more men in occupations essential to local production. This gap between the OLS and IV estimates is also comparable to the gaps reported in Rashid (2026), Bazzi et al. (2023), and Derenoncourt (2022), who rely on similar shifts-based SSIVs. For brevity and interpretation, in the main text, I focus on interpreting 2SLS estimates but report first stage, OLS, and reduced form coefficients for all main results.

Table 1: Effect of Number Men Drafted on Number of Women Employed in 1950

Panel A: First Stage on MenDrafted$_{oc,1940-1945}$				
$\widehat{\text{NumberDrafted}}_{oc}$	1.296	1.363	1.287	1.353
	(0.272)	(0.259)	(0.290)	(0.278)
F-statistic	160.5	173.9	78.63	64.86
	Number of Women Employed in 1950			
	(1)	(2)	(3)	(4)
Panel B: OLS				
MenDrafted$_{oc,1940-1945}$	0.331	0.321	0.312	0.302
	(0.052)	(0.046)	(0.055)	(0.048)
Panel C: Reduced Form				
$\widehat{\text{NumberDrafted}}_{oc}$	0.709	0.719	0.711	0.719
	(0.150)	(0.145)	(0.161)	(0.152)
Panel D: 2SLS				
MenDrafted$_{oc,1940-1945}$	0.548	0.527	0.553	0.531
	(0.110)	(0.108)	(0.120)	(0.116)
Outcome Mean	25.13	25.13	25.13	25.13
Observations	369,185	369,185	369,185	369,185
Baseline Controls	Y	Y	Y	Y
County FE	N	Y	N	Y
Occ FE	N	N	Y	Y

NOTE: Regressions in Panel A document the first stage relationship between the SSIV and the number of men drafted from each occupation-county cell. Panels B, C and D use the number of women employed in 1950 as outcome. The unit of observation is an occupation-county cell. Column 1 only includes baseline occupation-county level controls, column 2 includes these baseline controls and county fixed effects, column 3 includes baseline controls and occupation fixed effects, and column 4 uses baseline controls, county fixed effects and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, number of workers in the agricultural, mining, manufacturing, and public administration industries. Robust standard errors clustered at the county level are reported in parentheses.

6.1.1 Robustness

Before providing further results on the mechanisms underlying occupational gendering, I conduct additional checks to solidify a causal interpretation of my results. I report these checks in the Online Appendix.

1. *Alternative Standard Errors.* I show that main estimates from Table 1 (column 4) remain precise when accounting for spatial autocorrelation based on standard errors obtained using the Conley (1999) approach with cutoffs at 50km and 100km, and based on standard errors using the Adao et al. (2019) correction for SSIV estimators.

2. *Varying Control Sets.* I show that my main results are robust to using the Belloni et al. (2014) double LASSO procedure to select controls in addition to the baseline occupation-by-county level characteristics included in the control vector $\mathbb{X}_{oc,1940}$.

3. *Accounting for Left-Skewed Distribution of number of men drafted and SSIV.* In order to show that results are not driven by the left-skewed distribution of both the SSIV and endogenous variables, I show robustness of main results to winsorizing both variables at the bottom 1st, 5th and 10th percentiles.

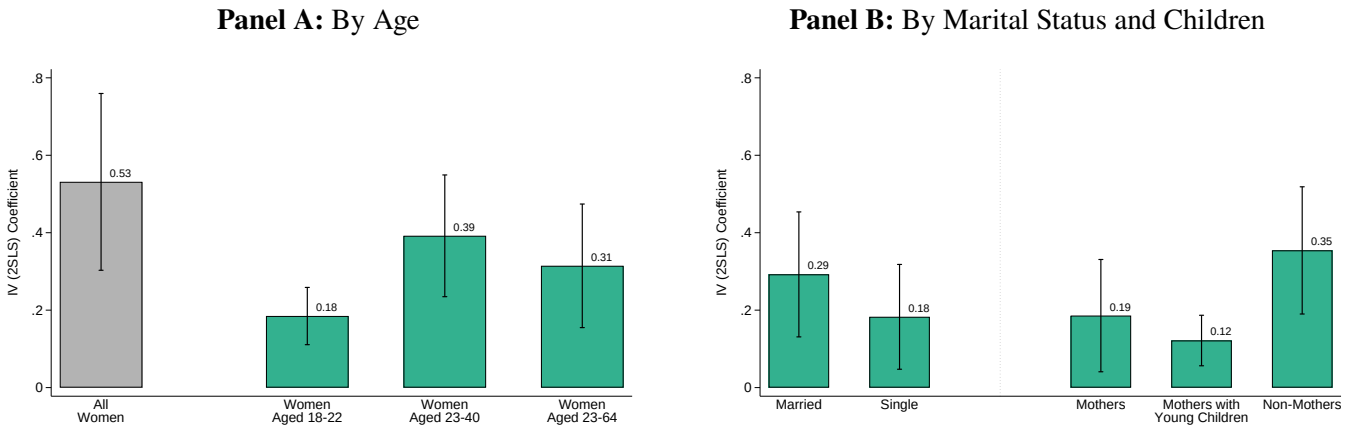
4. *Weak IV Tests.* Though all first stage F-statistics documented in Table 1 are well above conventional levels, I provide a number of tests for weak-IV identification including reporting Anderson and Rubin (1949) confidence intervals.

6.1.2 Heterogeneity

I next ask how the effect on women's 1950 employment varied with a woman's direct exposure to the wartime shock, which I proxy using her birth cohort. I split women observed in 1950 into three age groups: 18-22, 23-40, and 23-64. Women aged 18 to 22 in 1950 were young during the war years and were unlikely to have entered wartime employment themselves. Effect of the draft shock on their employment across occupations and counties likely operated through information transmission. Women aged 23 to 40 in 1950 were old enough to work during 1940 to 1945 and young enough that wartime entry could plausibly reshape their long-term employment patterns, making them the group for whom the shock should matter most directly. Finally, women aged 23 to 64 in 1950 comprise all cohorts old enough to be directly exposed to the wartime shock.

Figure 3 presents the effect of an additional man drafted from an occupation-county cell on women's employment by their age in 1950 and the baseline effect on all women from Table 1 for comparison. Women aged 23-40, who were old enough to work during the war and young enough for wartime entry to reshape their long-term employment, experienced the largest persistent effect on their employment. Importantly for the mechanism outlined by the model, women aged 18-22 in 1950, who themselves were too young to have entered occupations during the war, also respond to the shock. I revisit this result in Section 6.3, where I provide evidence of the social learning mechanism underlying it. Lastly, all women aged 23-64 in 1950, who were old enough to have responded to labor scarcity during the war years, see an effect smaller than that of prime-age women but larger than that of the youngest, unexposed cohort. Online Appendix Table C.2 presents the corresponding first stage, reduced form, and OLS estimates.

Figure 3: Heterogeneity in the Effect of War Mobilization on Number of Women Employed in 1950



NOTE: This figure plots IV (2SLS) estimates of the effect of the number of men drafted from each occupation-county cell on the number of women employed in 1950 for different groups of women. Panel A reports estimates for the number of women employed overall (gray) and by age group, and Panel B reports estimates for the number of employed women by marital status and by presence of children. Each bar corresponds to a separate regression in which the number of men drafted is instrumented with the SSIV. The unit of observation is an occupation-county cell. All regressions include baseline occupation-county level controls, county fixed effects, and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, and number of workers in the agricultural, mining, manufacturing, and public administration industries. Confidence intervals are based on robust standard errors clustered at the county level. First stage results and the corresponding OLS and reduced form estimates are reported in Online Appendix Tables C.1 and C.2.

Next, Figure 3 Panel B shows that women's employment response also varies by marital status and the presence of children. Women who are married in 1950 experience a greater effect in response to war mobilization. This pattern partly reflects cohort exposure, as women who were old enough to be exposed to the war shock are also more likely to be of marriage age by 1950. Women without children in 1950 see the largest employment effect, while women with young children see the smallest, consistent with young children imposing the highest costs of market work on mothers (Kleven et al., 2019). Consistent with Goldin and Olivetti (2013) and the historical narrative (Anderson, 1981; Campbell, 1984), these results also suggest that labor scarcity from war mobilization helped reshape norms surrounding the employment of married women and women with children. Online Appendix Table C.2 presents this result and corresponding coefficients as a table.

Lastly, I show variation in women's employment response by race. Consistent with the historical narrative, I find that both White and Black women experienced a similar overall effect on their employment in 1950 arising from labor scarcity during the war (see Online Appendix Table C.3).[16] Consistent with

[16]The historical account suggests that wartime scarcity allowed Black women to leave domestic service and farm labor for industrial jobs that had previously excluded them (Anderson, 1982). This result is also consistent with Bailey and Collins (2006), who document a narrowing of the wage gap between White and Black women in the 1940s, in part due to changes in the occupations Black women gained access to amid wartime labor scarcity.

results documented in Martinez and Derenoncourt (2025), I also find a large effect on the employment of Black men due to increases in demand for their labor generated by wartime mobilization of White men (see Online Appendix Table C.3). Additionally, the null result on White men serves as a useful placebo check, showing that the labor scarcity due to the draft of mainly White men did not have an effect on White men's employment and the resulting excess demand was instead met by groups previously excluded from these jobs, including women and Black men.

6.2 Effect on Women's Employment by Occupational Characteristics

I next study how the effect on women's persistent employment by 1950 varies by the underlying non-wage attributes of occupations. First, I use information on the baseline task content of each occupation in 1939. I rank occupations as above or below median across the physical dimension of task content (see Section 4 for additional information these data). Examples of occupations which are ranked as above median physical task content include engineers, laborers, and nursing, and occupations that are ranked as below median in their physical intensity include bank-teller, optician, and teacher. Online Appendix Table C.5 shows that high physical intensity of occupations is negatively correlated with occupational amenities valued by women such as schedule flexibility. To estimate how the effect of war mobilization on women's employment in 1950 varies by an occupation's physical demands, I use the following specification:

$$\begin{aligned} y_{oc,1950} = \alpha + \beta_1 \text{MenDrafted}_{oc} + \beta_2 \text{MenDrafted}_{oc} \times \mathbb{1}[\text{Physical}_o \leq \text{median}] \\ + \mathbb{X}_{oc,1940}\gamma + \omega_o + \phi_c + \varepsilon_{oc}, \end{aligned} \tag{8}$$

where Physical_o is the physical task content measure for occupation o and $\mathbb{1}[\text{Physical}_o \leq \text{median}]$ equals one if occupation o is at or below the median of physical task content across occupations. Its main effect is absorbed by the occupation fixed effects, ω_o. Both MenDrafted_{oc} and its interaction with $\mathbb{1}[\text{Physical}_o \leq \text{median}]$ are treated as endogenous and instrumented with $\widehat{\text{NumberDrafted}}_{oc}$ and $\widehat{\text{NumberDrafted}}_{oc} \times \mathbb{1}[\text{Physical}_o \leq \text{median}]$, respectively. β_1 measures the effect of the number of men drafted on the number of women employed in 1950 in occupations with above-median physical content. β_2 measures how this effect differs in occupations with below-median physical content, so that $\beta_1 + \beta_2$ gives the effect of war mobilization in occupations with below-median physical content.

Figure 4 plots the 2SLS estimates of the effect of the number of men drafted on the number of women aged 23 and above employed in 1930 and 1950, separately for occupations below and above the median of physical task content measure. I focus on women who were older than 23 as we expect these women to

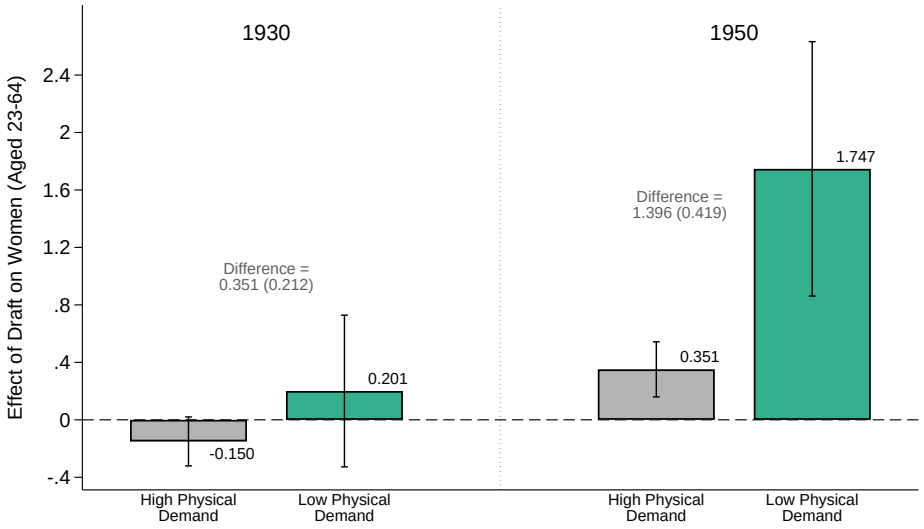
respond directly to the shock. In Section 6.3, I focus separately on younger unexposed women aged 18-22. Consistent with the placebo results on the shock overall, Figure 4 shows that in 1930, women do not sort into less versus highly physical occupations any differently in response to the shock.

Figure 4 shows that in 1950 women's employment per man drafted was roughly 5 times as large in occupations that were less physical (0.35 vs 1.74). The smaller effect on women's employment in more physical occupations is consistent with women not remaining in occupations with low levels of the amenities they value. High physical demands correlate negatively with flexibility, an amenity that is valued by women (see Online Appendix Table C.5). Online Appendix Figure C.1 presents results for other task content measures including non-routine analytic and routine manual content. I find that women's employment in occupations with above-median non-routine analytic content was twice as large as in occupations that ranked below-median in this dimension. I don't find any significant differences in women's sorting response across occupations that had above versus below-median routine manual content. In Section 6.2.1, using data on women's employment during the war, I show that women entered all types of occupations during the war, including more physically intensive occupations. Online Appendix Table C.4 reports the corresponding first stage F–statistics along with Sanderson and Windmeijer (2016) conditional first stage F-statistics for each of the two endogenous variables.

Next, I use reported hours worked within a week in the 1940 census to proxy for occupations with more flexible hours, an amenity valued by women. I classify occupations as offering flexible hours if their mean weekly hours are less than 40. Using this measure in equation (8) instead of $\mathbb{1}[\text{Task}_o \leq \text{median}]$, I find that more women were employed in occupations that offered more flexible hours. Specifically, each drafted man led to 1.22 additional women employed in flexible-hours occupations, compared to 0.76 in occupations with more demanding hours. However, results using underlying hours data from the Census are nosier given limited variation in the underlying variable and a small sample that reports their hours in the 1940 Census. See Online Appendix Table C.6 for this result.

Lastly, I also build explicit measures of historical occupational amenities valued by women, including schedule flexibility and cleanliness of work environment, using Talkie-1930, an LLM trained to rely only on pre-1930 sources (Levine et al., 2026). Consistent with the model's prediction, I find that women's employment response in occupations coded as more flexible and cleaner was higher than in those coded as inflexible or unclean (see Online Appendix Table C.7).

Figure 4: Effect on Number of Exposed Women Employed in 1950, by Occupational Physical Content



NOTE: This figure plots IV (2SLS) estimates of the effect of the number of men drafted from each occupation-county cell on the number of women aged 23-64 employed in 1930 and 1950, separately for occupations below and above the median of physical task content measure, along with 95% confidence intervals. Each pair of bars corresponds to a separate regression in which the number of men drafted and its interaction with an indicator for below-median physical content are instrumented with the SSIV and its interaction with the same indicator. The below-median bar reports the main effect and the above-median bar reports the sum of the main effect and the interaction. The unit of observation is an occupation-county cell. All regressions include baseline occupation-county level controls, county fixed effects, and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, and number of workers in the agricultural, mining, manufacturing, and public administration industries. Confidence intervals are based on robust standard errors clustered at the county level.

6.2.1 Evidence Consistent with Selective Retention

In this section, I provide evidence consistent with selective retention, showing that women entered all types of occupations during the war, those with attributes preferred by women and those without, but stayed selectively in those with better attributes. Using data from the Palmer Survey (Palmer, 1954; Goldin, 1991), I am able to study the occupational choices of a small sample of women in 1940, 1944, and 1949. Online Appendix Figure C.3 plots the occupational choices during the war for women who report having an occupation for the first time in 1944. It shows that in 1944 women report working in both occupations with more female-valued attributes, such as accounting, bank telling, and medical technician work, as well as occupations with fewer female-valued attributes, such as welding, metal work, and truck driving. Online Appendix Figure C.3 also highlights examples of occupations with female-valued attributes which were predominantly male before the war and became feminized after the war.

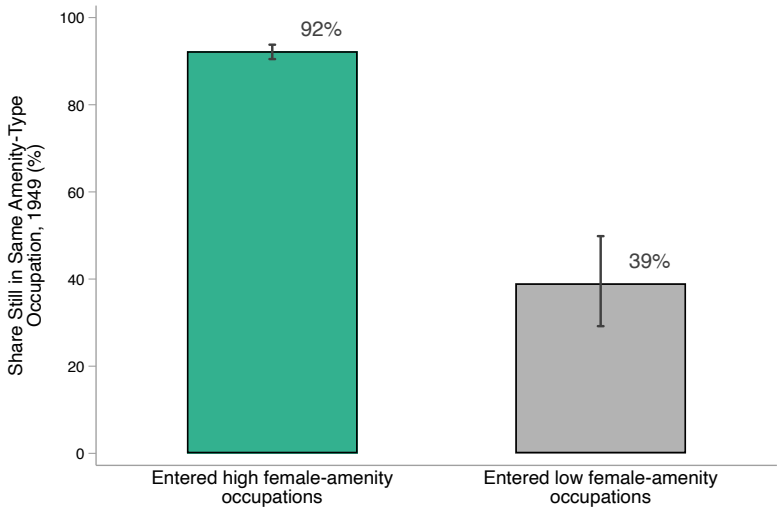
Figure 5 plots the share of women who remained in occupations with more versus less female-valued attributes that they entered during the war. Online Appendix Table C.8 presents the classification of occupations in the survey into those with high and low female-valued attributes. Among women who entered high-attribute occupations, 92 percent remained employed in a high-type occupation by 1949. In comparison, only 39 percent of those who entered low female-attribute occupations, remained in them by 1949. This suggests that, consistent with the selective retention, women entered all types of occupations but disproportionately remained in those occupations which provided better non-wage attributes. Online Appendix Figure C.4 plots trajectories of women who entered different example occupations with female-valued attributes and Online Appendix Figure C.5 plots trajectories for example occupations without such female-valued attributes.

6.3 Effects on Entry of Next Cohort

I next study effects on the employment of younger cohorts of women, who themselves would not have been exposed to the war shock, but were connected to older cohorts of women who were exposed directly. This builds on the earlier result that women aged 18 to 22 in 1950, who were too young to have worked during the war, also entered and remained in more draft-exposed occupations (see Figure 3).

Figure 6 plots the effect of an additional man drafted on the number of women aged 18-22 in occupations with higher versus lower levels of female-valued attributes, as measured by physical task content. First, consistent with there being no effect of the instrument on young women's employment before the shock, young women are not sorting differently in response to the shock across occupations in 1930.

Figure 5: Selective Retention of Women Entering New Occupations During War

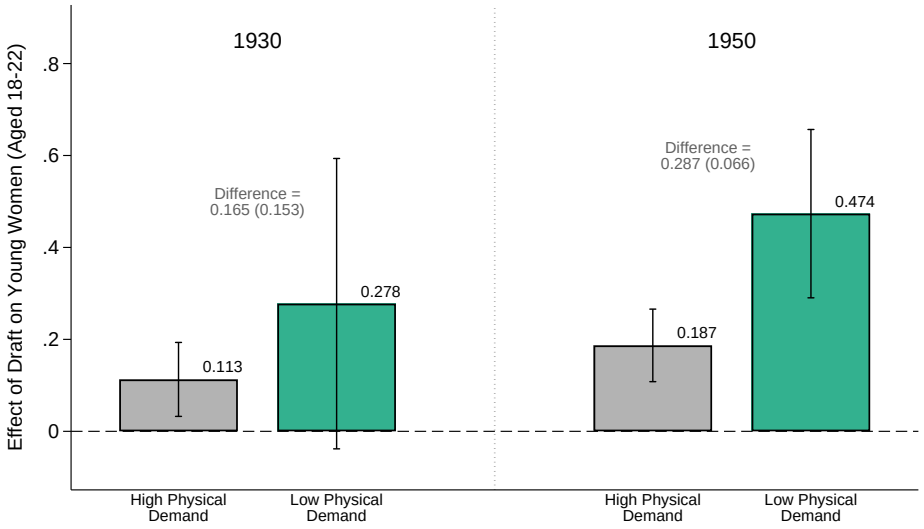


NOTE: This figure uses work history data from the Palmer survey (Palmer, 1954) with occupation codes from Goldin (1991). The sample consists of women who entered a new occupation during the war, defined as reporting an occupation in 1944 that differs from their reported occupation in 1940. Each occupation is classified as high or low female amenity. See Online Appendix Table C.8 for the full classification. Bars show the share of wartime entrants into each occupation type still employed in an occupation of the same amenity type in 1949. Women not employed in 1949 are counted as not retained. Confidence intervals are obtained using the Wilson (1927) score method for binomial proportions.

Next, Figure 6 shows that there were 2.5 times as many younger unexposed women employed in less physical occupations as in more physically demanding occupations. Online Appendix Table C.10 also shows this effect for occupations that had more flexible hours as measured by an occupation having less than 40 mean hours per week. I find a similar difference in the employment response of younger women as when comparing less vs more physically intensive occupations. This pattern is also similar when using LLM coded amenities, including schedule flexibility and cleanliness. However, due to limited variation in the data underlying additional measures of flexibility and cleanliness, these estimates are less precise than estimates comparing more versus less physical occupations.

I argue that the main mechanism underlying the effects on the next generation of women is social learning, whereby younger women learn about occupations and their non-wage attributes from older cohorts of women. I provide evidence consistent with this mechanism by linking women in the 1950 Census to sets of older women they could be directly connected to and learn from, namely their sisters, female neighbors and mothers. Table 2 first uses the number of young women with at least one older sister in the same occupation as them in 1950 as an outcome in equation (4). It shows that each additional drafted man increased the number of young women working in the same occupation as an older sister by 9.5 percent of the outcome's mean. Next, column (2) shows that each drafted man increased young women's employment in an older female neighbor's occupation by 0.0300 (11.7% of the outcome mean). Lastly, column (3) uses

Figure 6: Effect on Young Women's Employment Per Man Drafted, by Occupational Physical Content



NOTE: This figure plots IV (2SLS) estimates of the effect of the number of men drafted from each occupation-county cell on the number of women aged 18-22 employed in 1930 and 1950, separately for occupations below and above the median of physical task content measure, along with 95% confidence intervals. Each pair of bars corresponds to a separate regression in which the number of men drafted and its interaction with an indicator for below-median physical content are instrumented with the SSIV and its interaction with the same indicator. The below-median bar reports the main effect and the above-median bar reports the sum of the main effect and the interaction. The unit of observation is an occupation-county cell. All regressions include baseline occupation-county level controls, county fixed effects, and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, and number of workers in the agricultural, mining, manufacturing, and public administration industries. Confidence intervals are based on robust standard errors clustered at the county level.

a small sample of women aged 18-22 co-residing with their mothers who are of working age in 1950, to construct the number of young women in the same occupations as their mothers as an outcome, and shows that younger women were also more likely to follow their mothers into more draft exposed occupations.

Table 2: Effect of Number Men Drafted on Young Women Following Connected Women into Occupations, 2SLS Estimates

	Dependent Variable: Number of Women Aged 18–22 Employed in the Occupation of a Connected Woman		
	Older Sister (1)	Older Female Neighbor (2)	Mother (3)
MenDrafted_{oc}	0.0055	0.0300	0.0129
	(0.0012)	(0.0080)	(0.0040)
Effect Relative to Outcome Mean (%)	9.5	11.7	13.4
Outcome Mean	0.058	0.256	0.097
First Stage F-statistic	64.86	64.86	64.86
Anderson and Rubin (1949) test (p-value)	0.000	0.000	0.000
Observations	369,185	369,185	369,185
Baseline Controls	Y	Y	Y
County FE	Y	Y	Y
Occ FE	Y	Y	Y

NOTE: This table reports 2SLS estimates of the effect of the number of men drafted from an occupation-county cell on the number of women aged 18 to 22 employed in the cell's occupation who are connected to an older woman employed in the same occupation. In column 1 the connected woman is an older sister aged 23 to 64 co-residing with a parent, in column 2 an older woman aged 23 to 64 in an immediately neighboring household, and in column 3 the young woman's co-residing mother. The number of men drafted is instrumented with the SSIV. The row "Effect Relative to Outcome Mean" reports the coefficient divided by the outcome mean, in percent, since the outcomes differ in their base levels across columns. The first stage F-statistic is from the first-stage regression, which is identical across columns, and I also report the weak-instrument-robust Anderson and Rubin (1949) test of the joint significance of the endogenous regressor. The unit of observation is an occupation-county cell. All specifications include baseline occupation-county level controls, county fixed effects and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, number of workers in the agricultural, mining, manufacturing, and public administration industries. Robust standard errors clustered at the county level are reported in parentheses.

Lastly, I use a sample of women aged 18-22 in 1950 linked to their younger selves in 1940. The linked data allow me to descriptively compare the outcomes in 1950 of women who grew up with older sisters to those without as observed using their 1940 household structure. Table 3, column (1) first shows that younger women who grew up with older sisters were more likely to be employed in 1950 when compared to other women in the same age cohort and from the same county. Next, column (2) shows that conditional on being employed, women with older sisters were more likely to be employed in occupation-county cells that had more men drafted from them. Lastly, column (3) uses 1950 county fixed effects in addition to 1940 county fixed effects to show that women who grew up with older sisters were employed in occupations that were more draft exposed when compared to similarly aged women who grew up without an older sister, were from the same county in 1940, and were employed in the same county in 1950. Together with results from Table 2, I interpret these descriptive results as additional evidence that being directly connected to older women who were exposed to the drafted provided younger women with greater information about occupations. Because these results can in part reflect the role of referral networks from older connected women, in the next section

I provide evidence consistent with the role of information diffusion about occupational non-wage attributes.

Table 3: Growing Up with an Older Sister and Occupational Choices of Linked Young Women

	Employed, 1950	Draft Exposure of 1950 Cell	
	(1)	(2)	(3)
Older Sister in Household, 1940	0.032	14.625	12.877
	(0.002)	(4.153)	(4.582)
Outcome Mean	0.542	139.450	139.589
Observations	428,689	140,280	140,134
1940 County FE	Y	Y	Y
1950 County FE	N	N	Y

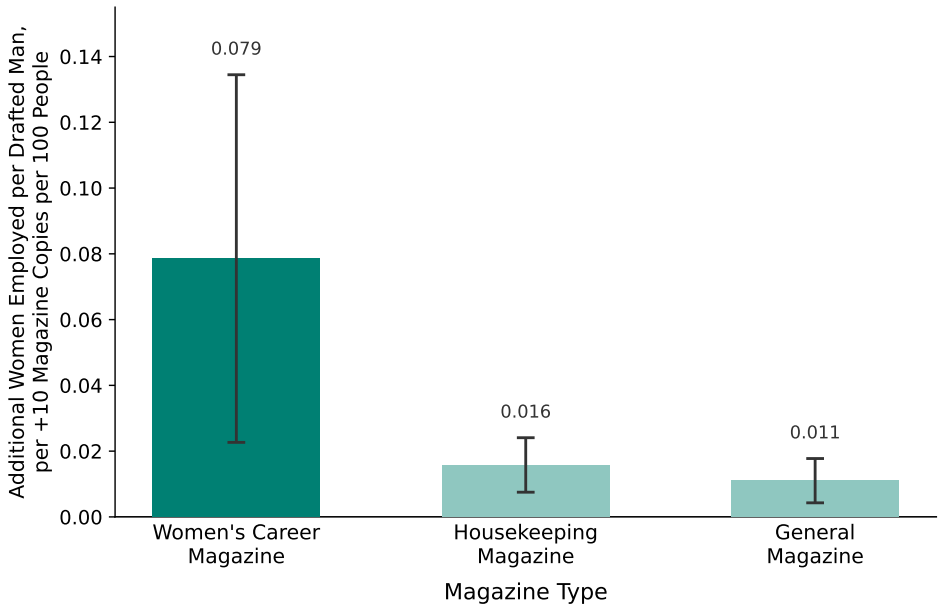
NOTE: This table uses the sample of women linked between the 1940 and 1950 complete count censuses, aged 8 to 12 in 1940 and 18 to 22 in 1950. The regressor is an indicator for co-residing with a sister older than age 13 in 1940. Column 1 uses an indicator for being employed in 1950 as outcome. Columns 2 and 3 restrict to women employed in 1950 and use the predicted number of men drafted from the woman's 1950 occupation-county cell as outcome, so that a positive coefficient indicates sorting into more draft-exposed occupations. All regressions include 1940 county fixed effects, so comparisons are between women who grew up in the same county. Column 3 adds 1950 county fixed effects, comparing women employed in the same county in 1950, so that the estimate reflects sorting across occupations within a destination labor market rather than sorting across places. Standard errors clustered at the 1950 county level are reported in parentheses.

6.3.1 Evidence Consistent with Information Diffusion about Occupational Non-Wage Attributes

So far, I present evidence consistent with the role of older cohorts of women in transmitting information to younger cohorts of women. However because this can in part reflect the role of referral networks, in this section, I provide evidence consistent with the role of public information sources about occupations and their non-wage attributes and the diffusion of such information over time. Public sources of information can then contribute to the dissemination of information and help generate national patterns of occupations becoming gendered as observed in Figure 1 Panel D.

Figure 7 plots the number of additional young women employed per drafted man in places with higher circulation of women's career magazines as compared to places where such magazines had less circulation. Career magazines for women in the late 1930s featured experiences of women in new occupations frequently through interviews. I find that places that had a higher circulation of such career magazines had more younger women employed per drafted man. Despite accounting for differences across places through county fixed effects, because circulation of such magazines can still be endogenous, I benchmark this result to that obtained when using the circulation rates of housekeeping or general magazines. I show that the difference in employment response between places with high versus low circulation is 5-7x larger when using women's career magazines as compared to using housekeeping or general magazine titles.

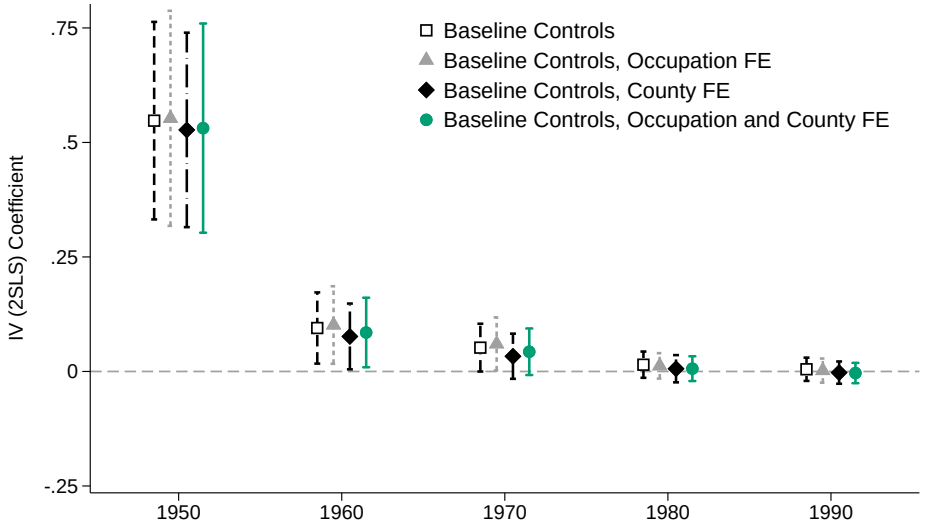
Figure 7: Effect on Women's Persistent Employment, 1950-1990



NOTE: This figure shows how the effect of the WWII draft on women's employment varies with the local circulation of three types of magazines. The women's career magazine is *Mademoiselle*, the housekeeping magazine is *Ladies' Home Journal*, and the general magazine is *Life*. Each bar reports the coefficient on the interaction between the predicted number of men drafted from an occupation-county cell and the state's circulation of the magazine, from a reduced-form regression of the number of women employed in the cell in 1950. Coefficients are scaled to an increase of 10 magazine copies per 100 workers in the state in 1940. Circulation is the total circulation of each magazine in the state from 1935-1940, from Hefner et al. (2023). Each magazine is estimated in a separate regression that also includes the predicted number of men drafted, baseline occupation-by-county controls, and occupation and county fixed effects. The sample includes 369,185 occupation-county cells. 95% confidence intervals are plotted, based on standard errors clustered at the county level.

Figure 8 uses samples of the 1960-1990 censuses to plot IV estimates of the effect of an additional man drafted from an occupation-county cell on women's employment in decades after 1950. It shows that the effect of the draft as a temporary shock to women's employment across occupation-by-county cells falls to zero within 4-5 decades. This evidence is consistent with information about attributes of occupations that women had the opportunity to try during the war having diffused. This is also matched by the patterns of occupations such as bank-telling and accounting becoming feminized within 4-5 decades of WWII. For example, accountants became persistently close to 60% women in its share of workers by 1990 and its speed of becoming feminized slowed down, consistent with beliefs converging to the true value of the amenities through the learning process (see Figure 1 Panel A). Online Appendix Figure C.7 takes random samples of the 1950 estimate and shows that while the size of the estimate decreases, it remains significantly different from 0 across 100 simulations. Thus, the effect size going to 0 for later decades is not simply driven by attrition arising from using smaller samples of the data.

Figure 8: Effect on Women's Persistent Employment, 1950-1990



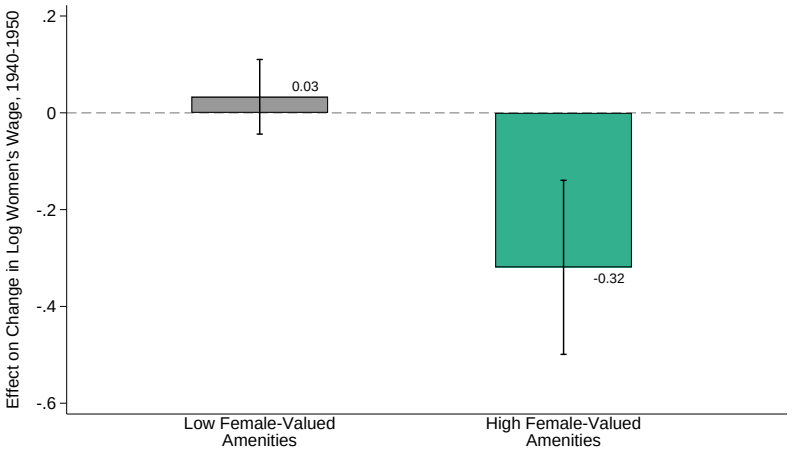
Note: This figure plots IV (2SLS) estimates of the effect of an additional man drafted from an occupation-county cell during WWII on the number of women employed in the cell in each census year from 1950 to 1990. The number of men drafted is instrumented with the predicted number of men drafted from the cell. Each marker reports a separate regression. The four specifications include baseline controls only, baseline controls with occupation fixed effects, baseline controls with county fixed effects, and baseline controls with both occupation and county fixed effects. Baseline controls are measured in 1940 and include the cell's population, number of employed women, number of married men, and number of workers in agriculture, mining, manufacturing, and public administration. The 1950 estimates use the full-count 1950 census and correspond to Panel D of Table 1. Estimates for 1960 to 1990 use the 1960 5% sample, the 1970 Form 1 and Form 2 metro 1% samples, and the 1980 and 1990 1% samples. 95% confidence intervals are shown, based on standard errors clustered at the county level.

6.4 Effect on Women's Wages

Lastly, I show that the shock from mass mobilization of men had a negative effect on women's wages in occupation-county cells that offer female-valued attributes. I use self-reported income data from the 1940 and 1950 Censuses by gender to study the effect of 10 additional men drafted on the change in women's log wages across cells by occupations which offered female-valued attributes and those occupations which did not. Specifically, I estimate equation (8) using the 1940 to 1950 change in log women's mean wage across cells as the outcome.

In Figure 9, I show that 10 additional men drafted led to a reduction in women's wages in occupations that provided better female-valued attributes. Specifically, women's wages in such occupations decreased by 32 log points relative to their wages in 1940, equivalent to a 27 percent reduction in wages. However, their wages remained unchanged in occupations without female-valued attributes, as measured by the physical task content of occupations. This result also helps rule-out alternative demand-side explanations including changing discrimination patterns or employer learning about women's productivity which would predict a null or positive effect on wages.

Figure 9: Effect on Change in Log Mean Women's Wage, by Occupational Characteristics



NOTE: This figure plots IV (2SLS) estimates of the effect of 10 additional men drafted from an occupation-county cell on the change in women's log mean wage in the cell between 1940 and 1950, separately for occupations with low and high female-valued amenities. Occupations with high female-valued amenities are those with below-median physical task content in the 1939 Dictionary of Occupational Titles (Gray 2013), and occupations with low female-valued amenities are those above the median. Both estimates come from a single regression in which the number of men drafted and its interaction with the high-amenity indicator are instrumented with the predicted number of men drafted and its interaction with the same indicator. The bar for low-amenity occupations reports the main effect, and the bar for high-amenity occupations reports the sum of the main effect and the interaction. Wages are annual wage and salary income from the 1940 census and the 1950 census sample line, with 1940 wages converted to 1950 dollars using the CPI. The regression includes baseline controls and occupation and county fixed effects. 95% confidence intervals are shown, based on standard errors clustered at the county level.

To rule out changes in the skill composition of women drawn in by the shock, I use women's reported education in the 1940 and 1950 Censuses. I find that the shock has no effect on the change in women's mean years of schooling across occupation-by-county cells between 1940 and 1950. Online Appendix Table C.9 presents these results.

Lastly, it's possible that learning can operate in the other direction as well, whereby women switch out of occupations for which the amenities proved worse than their prior beliefs. Such occupations can have excess demand for labor as women switch out of them and draw men back in, resulting in increases in men's wages in such occupations. Consistent with this, Online Appendix Figure C.6 shows that men's wages increased in occupations without female-valued amenities.

7 A Model of Occupational Gendering

To formalize the role of social learning about non-wage attributes in the process of occupational gendering, I develop a simple equilibrium model of occupational choice. The goal of the model is to match the empirical results and the descriptive patterns of aggregate occupational gender segregation. Workers in my model learn about amenities from work experience and social learning. Shocks to the demand for women's labor induce female entry into previously untested occupations. Entrants into an occupation decide whether to stay based on the amenities. Women who stay in an occupation transmit information to the next cohort, allowing an occupation to continue feminizing after the shock reverts. I outline the basic framework of the model and the predictions that it generates below. For additional details see the Online Appendix.

7.1 Preliminaries

7.1.1 Environment

The economy comprises a unit measure of workers, indexed by i, in a continuum of local labor markets c. Time is discrete with $t \in \{0, 1, 2\}$. Workers in each time period choose an occupation $o \in \{1, ..., O\}$. In each time period, a new cohort of women enters. Women in this model live for two time periods, as young entrants E_{oct} who choose an occupation and as older stayers R_{oct} who decide whether to stay in an occupation in which they entered in the last period. I assume that exogenous male labor M_{oct} is supplied inelastically. In the model that follows, I focus on a single labor market because the framework's mechanism and predictions do not depend on distinguishing across labor markets. I therefore suppress the label c. Note that, however, distinguishing between local labor markets is important for the empirical analysis.

I maintain two assumptions about the information structure. First, I assume that labor markets are competitive spot markets without search frictions, so that workers are fully informed about wages. Second, I assume that workers do not observe the non-wage amenities of an occupation and instead hold beliefs over the amenities.

7.1.2 Worker Preferences

Workers have preferences over wages and non-wage amenities of an occupation. Let w_{ot} be the offered wage for an occupation at time t. For a worker i, the utility from working in occupation o at time t is:

$$u_{iot} = \beta_w \ln w_{ot} + \beta_A a_{io} + \epsilon_{iot}. \tag{9}$$

In this utility function, a_{io} denotes each worker's amenity value from an occupation and is comprised of A_o, the true non-wage amenity of an occupation and v_{io}, an idiosyncratic match component that workers only learn after having worked in an occupation. Specifically, $a_{io} = A_o + \nu_{io}$. I assume that non-wage amenities A_o are exogenous and occupation specific. In the context of the empirical application and descriptive patterns, I view amenities as "female-valued" amenities, those which are valued more by women, such as schedule flexibility or lower physical demands. I assume v_{io} is distributed normally with variance σ_v^2, that is, $\nu_{io} \sim \mathcal{N}(0, \sigma_v^2)$. This idiosyncratic match component allows for there to be heterogeneity in the number of workers who stay or leave an occupation after having entered it.

In this utility function, ϵ_{iot} is a worker's idiosyncratic taste for an occupation, representing personal preferences over an occupation's characteristics. I assume ϵ_{iot} is i.i.d with a type-I extreme value distribution. Lastly, β_A and β_w govern the marginal rate of substitution between wages and amenities. Specifically, a worker with higher β_A than β_w will be willing to trade off wages for better non-wage occupational characteristics.

7.1.3 Production

Let Y_{ot} be the value added from hiring a worker in occupation o at time t:

$$Y_{ot} = \theta_{ot} L_{ot}^{\alpha}. \tag{10}$$

In this production function, θ_{ot} represents occupation specific productivity of labor and L_{ot} denotes the total labor in an occupation in each time period, comprised of male labor M_{ot} and female labor F_{ot}. Because

the focus of my model is a labor supply mechanism, for simplicity I abstract away from labor demand channels and assume men and women are perfect substitutes. Specifically, $L_{ot} = M_{ot} + F_{ot}$. Note that, in each period male labor is supplied inelastically and female labor is comprised of new entrants E_{ot} and stayers R_{ot} who entered an occupation in the last period and decided to stay. I also assume decreasing returns to labor, that is, $\alpha \in (0, 1)$.

7.1.4 Information

I assume that workers do not observe non-wage characteristics A_o. This implies that A_o can only be learned by women after entering an occupation. Instead, each cohort of women share a common belief about amenities:

$$A_o \mid \mathcal{I}_t \sim \mathcal{N}\Big(\mu_{ot}, \tau_{ot}^{-1}\Big), \tag{11}$$

where $\mathcal{I}_t$ represents the information available to a cohort at time t. I assume this information is distributed normally with mean μ_{ot} and precision τ_{ot}.

In each period, every woman who stays in an occupation, after having entered it and learning its non-wage characteristics, transmits an independent report conveying information about the non-wage characteristics to the next cohort of entrants. Averaging over independent reports gives the next cohort the public signal s_{ot}:

$$s_{ot} = A_o + \xi_{ot}, \qquad \xi_{ot} \sim \mathcal{N}\Big(0, \frac{\sigma_\xi^2}{R_{ot}}\Big).$$

I assume that the signal is centered around the true amenity value A_o and has precision $\pi_{ot} = \frac{R_{ot}}{\sigma_\xi^2}$ which is increasing in the stock of women who stay in the occupation.

7.2 Excess Labor Demand for Women

From the production function in (2), I derive the labor demand for each occupation as:

$$L_{ot}^d(w; \theta) = (\alpha \theta_{ot} / w_{ot})^{\frac{1}{1-\alpha}}. \tag{12}$$

In the model, shocks to the demand for women's labor are shocks to the primitives: productivity θ_{ot} or male labor supply M_{ot}. The empirical analysis leverages shocks to male labor supply. Such shocks generate excess demand for women's labor X_o, defined as the gap between labor demanded and labor supplied when both are evaluated at the pre-shock wage of an occupation $w_{o,0}$ with the shocked primitives M_{o1} and θ_{o1}:

$$X_o \equiv L_o^d(w_{o0};\theta_{o1}) - \big[\, M_{o1} + F_{o0} \,\big]. \tag{13}$$

In Online Appendix D I show that this excess demand to a first order approximation can be used to summarize the effect of a shock to the supply of male labor or the productivity of labor on the equilibrium wage and entry response of women.

7.3 Entry Decision

In each period, each woman evaluates the expected systematic utility using her cohort's common belief mean μ_{ot} for each occupation. A woman chooses occupation o if the expected utility from that occupation V_{ot} and her occupation o-specific idiosyncratic taste ϵ_{ot} are greater than or equal to that from all other occupations $k \neq o$. That is, $V_{ot} + \epsilon_{ot} \geq V_{kt} + \epsilon_{kt}$. I derive the probability that a woman enters occupation o in any given period p_{ot} as:

$$p_{ot} = \frac{\exp(V_{ot}/\sigma_\epsilon)}{\sum_{k=0}^{O} \exp(V_{kt}/\sigma_\epsilon)}. \tag{14}$$

This derivation follows from the type I extreme value distribution of ϵ_{ot}. See Online Appendix D for a complete derivation of this expression. Without loss of generality, I assume a unit mass of women, so that the number of women who enter an occupation in a given period E_{ot} equals this probability, that is $E_{ot} = p_{ot}$

7.4 Retention Decision

A woman who entered an occupation o in the last period must decide in period t whether to remain in that occupation. She makes this decision after having learned her amenity value from this occupation, $a_{io} = A_o + v_{io}$. Then, each woman who entered occupation o, will remain in it if her amenity value exceeds a reservation threshold $\bar{a}$, that is, $A_o + v_{io} \geq \bar{a}$. Combining this with the normality of v_{io}, I derive the rate at which each female entrant remains in occupation o, or the retention rate $r(A_o)$, as a function of the true occupation specific amenity A_o:

$$r(A_o) = \Phi\Big(\frac{A_o - \bar{a}}{\sigma_\nu}\Big), \tag{15}$$

where, this retention rate is strictly increasing in the underlying occupation specific amenity. The number of stayers in an occupation from one period to the next evolves as $R_{o,t+1} = E_{ot} r(A_o)$.

7.5 Learning and Belief Updating

In each period, the new cohort of women who are deciding to enter an occupation o, observe the public signal s_{ot}, and use it to update their beliefs. Beliefs about non-wage characteristics are updated according to Bayes rule, such that the weight on information about amenities from the public signal, $\lambda_{ot} = \pi_{ot}/(\tau_{ot} + \pi_{ot})$, is increasing in stock of stayers:

$$\mu_{o,t+1} = (1 - \lambda_{ot})\,\mu_{ot} + \lambda_{ot}\,s_{ot}. \tag{16}$$

Online Appendix D outlines the derivation of this belief updating rule.

7.6 Model Mechanics and Predictions

The model proceeds in three periods, where period $t = 0$ is the period before a temporary shock has occurred, period $t = 1$ is the shock period, and period $t = 2$ is the post-shock period after a given shock has reversed. This is intended to map to my empirical analysis, in which the temporary shock from mass mobilization during WWII largely reversed after the war had ended.

Conditional on the beliefs and returning stocks at the entry stage in a period, the model has a unique competitive equilibrium in wages and employment, with all occupational wages determined jointly. Young women choose occupations according to the entry probabilities p_{ot}, and labor markets clear. Retention and the Gaussian report-processing rule determine the state inherited by the next cohort. Online Appendix D.8 gives the formal equilibrium definition and proves existence and uniqueness.

The model allows me to derive a set of predictions for the after-shock period $t = 2$, which allow me to match the model to the results of my empirical analysis. To map to the empirical analysis, each prediction is derived for occupations that were previously untested by women and experience a temporary shock that reverts by $t = 2$. Thus, each prediction is written comparing the shocked path S to a no-shock or less-shocked counterfactual C. In my empirical analysis, I compared previously untested local occupations which had greater versus less exposure to the shock from mass mobilization.[17]

Prediction 1. Consider a temporary shock to $t = 1$ primitives $M_{o,1}$ or $\theta_{o,1}$ for occupation o. Excess demand X_o will increase total female employment F_{o2} to this occupation in period 2. That is, $\frac{\partial \mathbb{E}[F_{o2}]}{\partial X_o} > 0$.

[17] The predictions apply under the sufficient conditions stated in Appendix E, which require sufficiently pessimistic initial beliefs and, for the amenity comparisons, a sufficiently large information advantage over the counterfactual. Positive entry further requires learning to outweigh wage crowding.

Prediction 2. For the same shock at $t = 1$, the effect on total female employment at $t = 2$ is increasing in the true amenity value A_o. That is, $\frac{\partial\left(\mathbb{E}[F_{o2}^S]-\mathbb{E}[F_{o2}^C]\right)}{\partial A_o} > 0$.

Prediction 3. For the same shock at $t = 1$, the effect on young women's entry at $t = 2$ is increasing in the true amenity value A_o. That is, $\frac{\partial\left(\mathbb{E}[E_{o2}^S]-\mathbb{E}[E_{o2}^C]\right)}{\partial A_o} > 0$.

Prediction 4. For the same shock at $t = 1$, the effect on women's wages at $t = 2$ is decreasing in the true amenity value A_o. That is, $\frac{\partial\left(\mathbb{E}[\ln w_{o2}^S]-\mathbb{E}[\ln w_{o2}^C]\right)}{\partial A_o} < 0$.

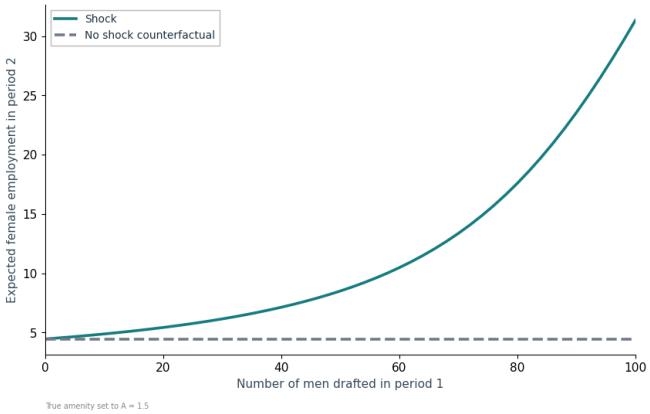
Mechanism. Intuitively, these predictions follow from wages rising in equilibrium in period $t = 1$ in response to excess demand for labor. This draws in new female entrants $E_{o,1}$. In period $t = 2$, for an occupation for which underlying characteristics A_o are high and prior beliefs over them μ_o are sufficiently pessimistic, a shock will result in higher retention $R_{o,2}$ as compared to a no-shock counterfactual. This increases the precision of the signal s_{o2}, leading to learning by the next cohort of women, who then update their beliefs. This leads to higher entry $E_{o,2}$ in period $t = 2$, even after the shock has reversed. Higher $R_{o,2}$ and $E_{o,2}$ result in higher total female employment F_{o2} to the occupation. See Online Appendix E for detailed proofs for each of the predictions.

I depict illustrations for each of the predictions using simulations of the model. For an occupation with sufficiently high underlying amenities, Figure 2 Panel A illustrates Prediction 1 and plots expected values of female employment at $t = 2$ against increasing values of a shock corresponding to the removal of men from the labor force. It depicts higher expected female employment increasing in the value of the shock. Panel B plots the result of simulations consistent with Prediction 2 and depicts increasing values of female employment at $t = 2$ against values of underlying occupational characteristic A_o for an occupation that experienced a shock relative to a no-shock path. Consistent with prediction 3, Panel C shows that for the same shock, the effect on entry of the next cohort is increasing in the underlying female-valued characteristics. Lastly Panel D depicts prediction 4 and shows that the effect on women's wages is decreasing with respect to female-valued occupational characteristics.

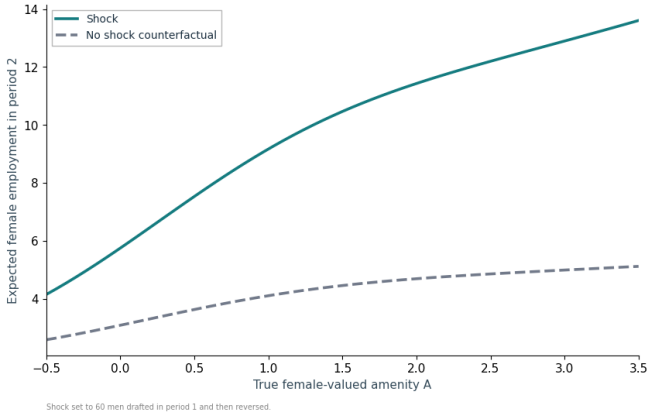
In the Online Appendix, I also derive predictions that map to the descriptive patterns of national gendering of occupations and variation in speed of gendering across occupations. Specifically, to map to these patterns, I reintroduce local labor markets c and introduce a parameter κ which captures speed of information spread. Historically, and in my empirical analysis, κ represents spread of information about occupational characteristics across local markets through word of mouth, migration, as well as women's career magazines which I rely on in my analysis.

Figure 10: Simulations of Model Depicting Predictions

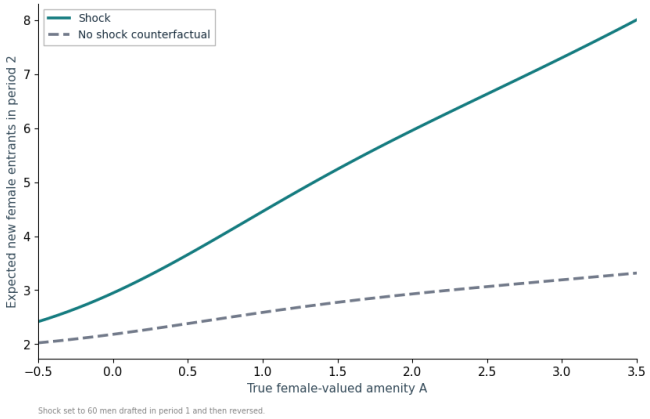
Panel A: Effect on Employment



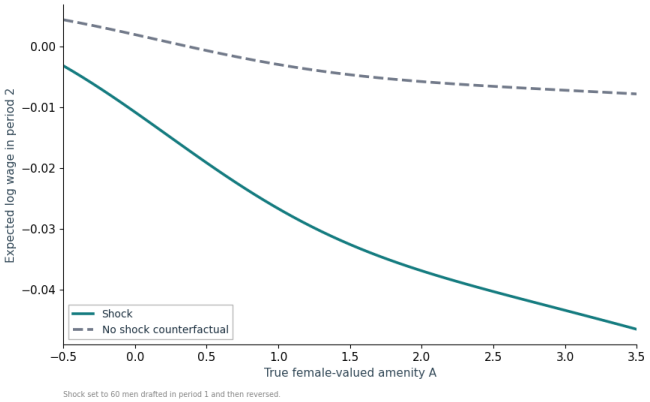
Panel B: Effect on Employment by True Amenity



Panel C: Effect on Entry by True Amenity



Panel D: Effect on Wages by True Amenity



NOTES: Each panel plots simulations of the model, averaged across 5000 draws of the public signal and starting from a pessimistic belief about occupational amenities. Panel A plots expected female employment at $t = 2$ in an occupation with high female-valued amenities against the size of the shock, measured as the number of men removed from the labor supply at $t = 1$. Panel B plots expected female employment at $t = 2$ against the occupation's true female-valued amenity A_o, with and without the shock. Panel C plots the entry of the next cohort of women at $t = 2$ against A_o for the shock and no-shock simulations. Panel D plots women's log wages at $t = 2$ against A_o. Panel A sets initially amenity level to 1.5. Panels B-D use a shock that removes 60 of 100 men from the labor supply at $t = 1$

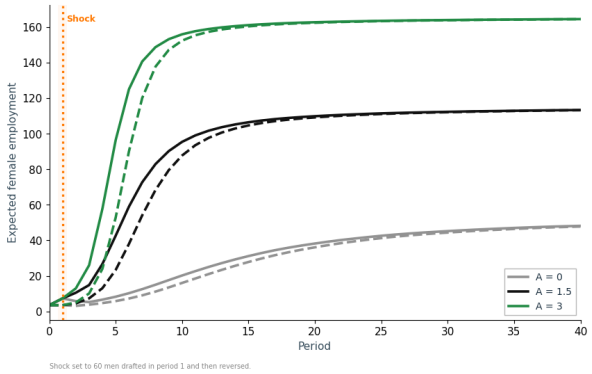
7.7 Occupational Gendering in the Absence of the Shock

According to the framework, we expect that temporary shocks to the demand for women's labor accelerate learning about occupations and their hidden non-wage characteristics. In this section, I average across simulations of the model to depict what women's expected employment trajectories across occupations with varying true underlying characteristics look like with and without a shock to male labor supply in period $t = 1$. Specifically, I take 5000 random draws of the public signal that women perceive before choosing occupations in each period, simulate female employment for each draw, and average across the different simulated employment paths for each period. I repeat this procedure with and without the shock for each occupation type. I use a shock that removes 60 out of 100 men from the labor supply temporarily only in period $t = 1$. All simulations use the same pessimistic starting belief over occupational characteristics.

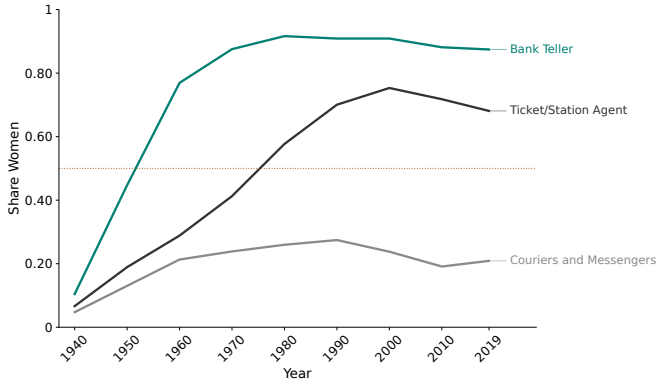
Figure 11 Panel A plots the simulated female employment trajectories across occupations with three different levels of underlying true female-valued amenities: high (green), medium (black), and low (gray). It plots trajectories with the shock in solid lines and trajectories without the shock in dashed lines. It shows that without a shock in period $t = 1$, over longer horizons occupations still converge to the same expected female employment level, however, this convergence is slower. This is due to beliefs about occupational characteristics taking longer to converge to true values in the absence of shocks. Online Appendix Figure C.8 plots the corresponding simulated trajectories of beliefs.

Figure 11: Long-Run Female Employment Trajectories Across Occupation Types

Panel A: Simulated Female Employment Trajectories



Panel B: Actual Example Trajectories



NOTE: Panel A plots simulated expected female employment in occupations with high ($A = 3$, green), medium ($A = 1.5$, black), and low ($A = 0$, gray) true female-valued amenities. Solid lines show employment with a temporary shock that removes 60 of 100 men from the labor supply in period $t = 1$ only, marked by the vertical line, and dashed lines show employment without the shock. Each line averages employment across 5,000 simulated draws of the public signal, and all simulations start from the same pessimistic belief about occupational amenities. Panel B plots the share of women among bank tellers, ticket and station agents, and couriers and messengers from 1940 to 2019 using US Census and American Community Survey data. Couriers and messengers correspond to the 1950 census category of messengers and office boys.

Figure 11 Panel B plots actual trajectories for three different types of occupations which all employed

few women before the war in 1940 and faced a shock to the demand for women's labor during the war. These three occupations, bank-tellers, ticket and station agents, and couriers and messengers, followed different trajectories of women's employment after the war. The underlying historical female-valued characteristics and unobserved amenities of these three occupations are also different. For example, the courier and messenger occupation ranks highest out of the three on the physical task content measure. It also required women to travel long distances and far away from home. The occupation of ticket and station agent ranks between the other two on it's physical content. It also required women to work shifts at different hours but offered the benefit of working behind a screen. The bank-teller occupation ranks lowest in terms of its physical content and offered women a clean environment to work in with fixed time schedules. The long term female employment and gender composition trajectories of these three distinct occupations follow closely the simulated paths of hypothetical occupations in Panel A and their underlying amenities.

8 Conclusion

Occupational gender segregation remains a feature of labor markets today, well beyond the historical US setting studied in this paper, with policy efforts increasingly seeking to draw workers into occupations in which their gender has been underrepresented, including recent efforts to increase men's representation in nursing and women's participation in the trades. The findings of this paper point to imperfect information as one reason why such segregation can persist even when occupational opportunities are available. When few women have worked in an occupation, they may have limited information about characteristics such as its physical demands, schedule flexibility, cleanliness, or working environment. Temporary shocks can give women an opportunity to experience previously untested occupations, learn their characteristics, and correct initially imprecise or pessimistic beliefs. Women then remain where they discover favorable amenities and transmit this information to subsequent cohorts, allowing temporary entry to generate persistent changes in an occupation's gender composition.

The results of this paper also carry different policy implications than explanations centered on discrimination or other barriers to occupational entry. Such explanations point toward policies that reduce those barriers, including anti-discrimination enforcement or changes in hiring practices (Hsieh et al., 2019). By contrast, if segregation partly reflects imperfect information about occupational characteristics, policy may need to improve workers' information or create opportunities to experience occupations they would otherwise not consider. If this information then diffuses through social networks, temporary interventions may generate persistent changes in occupational sorting.

The framework and setting of my paper allow me to study social learning about occupations and their characteristics indirectly, through how information generated by exogenous shocks can shape occupational choices across cohorts of women. I show that in response to a demand shock that pulled a cohort of women into new occupations, younger women disproportionately followed their older sisters and female neighbors into more shocked occupations in which those women remained. Future work could test this mechanism more directly by experimentally providing men and women with information about difficult-to-observe occupational characteristics and studying how it changes their sorting across occupations.

Finally, my framework and analysis focus on a labor-supply mechanism and treat occupational amenities as exogenous. The empirical patterns could also reflect a demand-side channel in which employers adjust amenities as women enter. Such adjustment however may be limited by technological constraints on job design in the short-term (Scott, 2026). Future work could study how learning on both sides of the labor market—workers learning about occupations and employers learning about worker preferences—jointly shapes occupational gendering in the long-run.

References

Abbate, J. (2012). *Recoding Gender: Women's Changing Participation in Computing*. Cambridge, MA: MIT Press.

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

Adao, R., M. Kolesár, and E. Morales (2019). Shift-share designs: Theory and Inference. *The Quarterly Journal of Economics 134*(4), 1949–2010.

Albanesi, S. and C. Olivetti (2016). Gender Roles and Medical Progress. *Journal of Political Economy 124*(3), 650–695.

Alon, T., S. Coskun, and J. Olmstead-Rumsey (2025). Gender Divergence in Sectors of Work. CEPR Discussion Paper No. 20076.

Altonji, J. G. and C. R. Pierret (2001). Employer Learning and Statistical Discrimination. *Quarterly Journal of Economics 116*(1), 313–350.

Anderson, K. (1981). *Wartime Women: Sex Roles, Family Relations, and the Status of Women during World War II*. Westport, CT: Greenwood Press.

Anderson, K. T. (1982). Last hired, first fired: Black women workers during World War II. *The Journal of American History 69*(1), 82–97.

Anderson, T. W. and H. Rubin (1949). Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of mathematical statistics 20*(1), 46–63.

Aneja, A., S. Farina, and G. Xu (2024). Beyond The War: Public Service And The Transmission Of Gender Norms. NBER Working Paper.

Attanasio, O., H. Low, and V. Sánchez-Marcos (2008). Explaining Changes in Female Labor Supply in a Life-Cycle Model. *American Economic Review 98*(4), 1517–1552.

Autor, D., F. Levy, and R. J. Murnane (2003). The Skill Content Of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics 118*(4), 1279–1333.

Bailey, M. J. (2006). More Power to the Pill: The Impact of Contraceptive Freedom on Women's Life Cycle Labor Supply. *The Quarterly Journal of Economics 121*(1), 289–320.

Bailey, M. J. and W. J. Collins (2006). The wage gains of African-American women in the 1940s. *The Journal of Economic History 66*(3), 737–777.

Bald, A. (2025). The Birth of an Occupation: Professional Nursing in the Era of Public Health. *Working Paper*.

Bazzi, S., A. Ferrara, M. Fiszbein, T. Pearson, and P. A. Testa (2023). The other great migration: Southern whites and the new right. *The Quarterly Journal of Economics 138*(3), 1577–1647.

Belloni, A., V. Chernozhukov, and C. Hansen (2014). Inference on Treatment Effects After Selection Among High-Dimensional Controls. *Review of Economic Studies 81*(2), 608–650.

Bergmann, B. R. (1974). Occupational Segregation, Wages and Profits When Employers Discriminate by Race or Sex. *Eastern Economic Journal 1*(2), 103–110.

Blau, F. D., P. Brummund, and A. Y.-H. Liu (2013). Trends in Occupational Segregation by Gender 1970–2009: Adjusting for the Impact of Changes in the Occupational Coding System. *Demography 50*(2), 471–492.

Blau, F. D. and L. M. Kahn (2017). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature 55*(3), 789–865.

Blau, F. D. and A. E. Winkler (2018). *The Economics of Men, Eomen and Work* (Eighth ed.). Oxford University Press.

Borusyak, K., P. Hull, and X. Jaravel (2022). Quasi-experimental Shift-Share Research Designs. *The Review of Economic Studies 89*(1), 181–213.

Campbell, D. (1984). *Women at War with America: Private Lives in a Patriotic Era*. Cambridge, MA:

Harvard University Press.
Conley, T. G. (1999). GMM Estimation with Cross Sectional Dependence. *Journal of econometrics 92*(1), 1–45.
Cortés, P. and J. Pan (2019). When time binds: Substitutes for household production, returns to working long hours, and the skilled gender wage gap. *Journal of Labor Economics 37*(2), 351–398.
Cross, A. B. (2025). Women's Entry into Nontraditional Occupations: Impacts of the Gender Desegregation of the U.S. Army. Working Paper, University of Alaska Anchorage.
Derenoncourt, E. (2022). Can You Move To Opportunity? Evidence From The Great Migration. *American Economic Review 112*(2), 369–408.
England, P., P. Allison, and Y. Wu (2007). Does bad pay cause occupations to feminize, does feminization reduce pay, and how can we tell with longitudinal data? *Social Science Research 36*(3), 1237–1256.
Executive Order 8560 (1940, October). Selective service regulations. Issued by President Franklin D. Roosevelt, October 4, 1940. Volume III: Classification and Selection, Section XXII: Occupational Deferments, paragraph 350(a).
Fernández, R. (2013). Cultural change as learning: The evolution of female labor force participation over a century. *American Economic Review 103*(1), 472–500.
Fogli, A. and L. Veldkamp (2011). Nature or nurture? learning and the geography of female labor force participation. *Econometrica 79*(4), 1103–1138.
Goldin, C. (1990). *Understanding the Gender Gap: An Economic History of American Women*. NBER.
Goldin, C. (1991). The Role of World War II in the Rise of Women's Employment. *American Economic Review 81*(4), 741–756.
Goldin, C. (2014a). A Pollution Theory of Discrimination: Male and Female Differences in Occupations and Earnings. In L. P. Boustan, C. Frydman, and R. A. Margo (Eds.), *Human Capital in History: The American Record*, pp. 313–348. Chicago: University of Chicago Press.
Goldin, C. (2014b). A grand gender convergence: Its last chapter. *American Economic Review 104*(4), 1091–1119.
Goldin, C. and L. F. Katz (2002). The Power of the Pill: Oral Contraceptives and Women's Career and Marriage Decisions. *Journal of Political Economy 110*(4), 730–770.
Goldin, C. and L. F. Katz (2011). The Cost of Workplace Flexibility for High-Powered Professionals. *The Annals of the American Academy of Political and Social Science 638*(1), 45–67.
Goldin, C. and L. F. Katz (2016). A most egalitarian profession: Pharmacy and the evolution of a family-friendly occupation. *Journal of Labor Economics 34*(3), 705–746.
Goldin, C. and C. Olivetti (2013). Shocking Labor Supply: A Reassessment of the Role of World War II on Women's Labor Supply. *American Economic Review: Papers & Proceedings 103*(3), 257–262.
Gray, R. (2013). Taking Technology To Task: The Skill Content Of Technological Change In Early Twentieth Century United States. *Explorations In Economic History 50*(3), 351–367.
Gray, R., S. M. O'Keefe, S. Quincy, and Z. Ward (2025). Tasks and black-white inequality over the long twentieth century. *Journal of Human Resources*. Forthcoming.
Greenberg, K., M. Wasserman, and E. A. Weber (2024). The Effects of Gender Integration on Men: Evidence from the US Military. *NBER Working Paper*.
Greenwood, J., A. Seshadri, and M. Yorukoglu (2005). Engines of Liberation. *The Review of Economic Studies 72*(1), 109–133.
Gyarmati, D., B. Pakula, C. Nguyen, and D. Leonard (2017). Enhancing the retention and advancement of women in trades in British Columbia. Technical report, Social Research and Demonstration Corporation.
Hefner, B. E., E. Timke, and K. J. Hegg (2023, October). Circulating american magazines.
Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019). The Allocation of Talent and U.S. Economic Growth. *Econometrica 87*(5), 1439–1474.
Kleven, H., C. Landais, and J. E. Søgaard (2019). Children and gender inequality: Evidence from denmark.

American Economic Journal: Applied Economics 11(4), 181–209.

Le Barbanchon, T., R. Rathelot, and A. Roulet (2021). Gender differences in job search: Trading off commute against wage. *The Quarterly Journal of Economics 136*(1), 381–426.

Levanon, A., P. England, and P. Allison (2009). Occupational feminization and pay: Assessing causal dynamics using 1950–2000 u.s. census data. *Social Forces 88*(2), 865–891.

Levine, N., D. Duvenaud, and A. Radford (2026, April). Introducing talkie: A 13b vintage language model from 1930.

Lordan, G. and J.-S. Pischke (2022). Does Rosie Like Riveting? Male and Female Occupational Choices. *Economica 89*(353), 110–130.

Macpherson, D. A. and B. T. Hirsch (1995). Wages and gender composition: Why do women's jobs pay less? *Journal of Labor Economics 13*(3), 426–471.

Martinez, J. J. and E. Derenoncourt (2025). Market Forces and Employer Racial Preferences: Evidence from Wartime Shortages. Working Paper, University of California, Berkeley.

Mas, A. and A. Pallais (2017). Valuing Alternative Work Arrangements. *American Economic Review 107*(12), 3722–3759.

McMillin, L. F. (1943). The second year: A study of women's participation in war activities of the federal government. Technical report, United States Civil Service Commission, Washington, DC.

Merchant, H. and B. Levy (2026). From frontier models to point-in-time models. https://frontiertopit.com/. University of Chicago Booth School of Business.

Myers, C. K. (2017). The Power Of Abortion Policy: Reexamining The Effects Of Young Women's Access To Reproductive Control. *Journal of Political Economy 125*(6), 2178–2224.

Olivetti, C., E. Patacchini, and Y. Zenou (2020). Mothers, Peers, and Gender-Role Identity. *Journal of the European Economic Association 18*(1), 266–301.

Palmer, G. L. (1954). *Labor Mobility in Six Cities: A Report on the Survey of Patterns and Factors in Labor Mobility, 1940–1950*. New York: Social Science Research Council.

Pan, J. (2015). Gender Segregation in Occupations: The Role of Tipping and Social Interactions. *Journal of Labor Economics 33*(2), 365–408.

Perlmann, J. and R. A. Margo (2001). *Women's Work? American Schoolteachers, 1650–1920*. Chicago: University of Chicago Press.

Polachek, S. W. (1981). Occupational Self-Selection: A Human Capital Approach to Sex Differences in Occupational Structure. *Review of Economics and Statistics 63*(1), 60–69.

Rashid, M. (2026). Keys to upward mobility: Typewriter adoption and women's economic outcomes. *Working Paper*.

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.

Rotella, E. J. (1981). The Transformation Of The American Office: Changes In Employment And Technology. *The Journal of Economic History 41*(1), 51–57.

Ruggles, S., C. A. Fitch, R. Goeken, D. Hacker, M. A. Nelson, M. Schouweiler, and M. Sobek (2021). IPUMS Ancestry Full Count Data: Version 3.0 [dataset]. Minneapolis, MN: IPUMS, 2021.

Ruggles, S., S. Flood, M. Sobek, D. Backman, G. Cooper, J. A. Rivera Drew, S. Richards, R. Rodgers, J. Schroeder, and K. C. W. Williams (2025). IPUMS USA: Version 16.0. Dataset.

Ruggles, S., M. A. Nelson, M. Sobek, C. A. Fitch, R. Goeken, J. D. Hacker, E. Roberts, and J. R. Warren (2024). IPUMS Ancestry Full Count Data: Version 4.0. Dataset.

Sanderson, E. and F. Windmeijer (2016). A weak instrument F-test in linear IV models with multiple endogenous variables. *Journal of Econometrics 190*(2), 212–221.

Schuh, R. (2024). Miss Allocation: The Value of Workplace Gender Composition and Occupational Segregation. Working Paper, Federal Reserve Bank of New York.

Scott, D. (2026). Workplace amenities and labor market inequality. Working Paper.

Sharma, G. (2023, April). Monopsony and gender. Working paper.

Strober, M. H. and C. L. Arnold (1987). The Dynamics of Occupational Segregation among Bank Tellers. In C. Brown and J. A. Pechman (Eds.), *Gender in the Workplace*, pp. 107–148. Washington, DC: Brookings Institution.

U.S. Department of Labor, Women's Bureau (1942). Your questions as to women in war industries. Women's Bureau Bulletin 194, U.S. Government Printing Office, Washington, DC. Prepared by Mary Elizabeth Pidgeon.

U.S. Department of Labor, Women's Bureau (1952). 1952 handbook of facts on women workers. Women's Bureau Bulletin 242, U.S. Department of Labor, Women's Bureau, Washington, DC.

Vidart, D. (2024). Human Capital, Female Employment, and Electricity: Evidence from the early 20th-century United States. *Review of Economic Studies 91*(1), 560–594.

Wasserman, M. (2023). Hours Constraints, Occupational Choice, and Gender: Evidence from Medical Residents. *Review of Economic Studies 90*(3), 1535–1568.

Wilson, E. B. (1927). Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association 22*(158), 209–212.

Wiswall, M. and B. Zafar (2018). Preference for the Workplace, Investment in Human Capital, and Gender. *Quarterly Journal of Economics 133*(1), 457–507.

Zhao, Y. and A. Martínez-Marquina (2026). Engineering the Gender Gap: Fall of Women's Share in Computer Science. Working Paper.

Online Appendix

How Occupations Become Gendered: Social Learning and the Long-Run Evolution of Women's Work

Myera Rashid

Northwestern University

Online Appendix Table of Contents

A Descriptive Patterns Data and Additional Figures and Tables **1**

A.1 Data Construction for Descriptive Patterns 1

A.2 Additional Descriptive Patterns 3

B SSIV Construction and Robustness **9**

B.1 Shift Share IV Variation 9

B.2 Placebo Test 12

B.3 Random Shifts Placebo Test 13

C Additional Empirical Results **14**

C.1 Heterogeneity 14

C.2 Analysis Using Amenities 17

C.3 Palmer Survey Analysis Additional Figures and Tables 22

C.4 Wages 26

C.5 Social Learning 28

C.6 Information Diffusion 29

C.7 Simulated Belief Trajectories 30

D Model Derivations **30**

D.1 Labor demand 30

D.2 Occupational entry 31

D.3 Retention 33

D.4 Reports and signal precision 34

D.5 Learning and belief updating 35

D.6 Market clearing and the equilibrium system 37

D.7 Excess labor demand and the local wage and entry response 38

D.8 Within-period equilibrium: definition, existence, and uniqueness 41

E Model Predictions and Proofs **43**

E.1 Comparison, timing, and scope 44

E.2 The period-1 entry response 45

E.3 Equilibrium responses and noisy learning 46
E.4 Prediction 1: post-shock total female employment 48
E.5 Prediction 2: amenities and the total-employment effect 49
E.6 Prediction 3: entry by the next cohort 50
E.7 Prediction 4: amenities and the log-wage effect 52
E.8 Interpreting and assessing the sufficient conditions 54

A Descriptive Patterns Data and Additional Figures and Tables

A.1 Data Construction for Descriptive Patterns

To construct data that would allow me to follow occupations from 1870 to 2019 in the US, I use both complete count Censuses and samples of Census and ACS data. Specifically, I combine complete count US censuses for 1870 through 1950, accessed through the NBER server (Ruggles et al., 2024), with IPUMS samples of Census and ACS data for later years. For samples, I use the 1960 and 2000 5 percent samples, the 1970 form 1 and form 2 metro samples, the 1980 and 1990 1 percent samples, and the 2006–2010 and 2015–2019 ACS five-year samples, which I assign to 2010 and 2019 respectively (Ruggles et al., 2025). Occupations are classified using the *occ1950* variable. I hand-correct the categorization of occupations in the 1940 complete count Census, and I use corrected occupation codes for 1880 following Rashid (2026). In constructing a set of occupations that can be followed consistently over time, I hand-combine several detailed occupation categories into single occupations. I also exclude military occupations, occupations not elsewhere classified, and general occupation categories, resulting in a set of 203 occupations that I follow across decades. See Online Appendix Table A.1 below for the list of occupations included in the descriptive analysis. I compute the share of women in each occupation-decade cell as the number of female workers divided by the total number of workers in the cell, excluding cells with fewer than 100 workers. For the geographic analysis, I compute the same measure at the occupation-state level. The 1890 census records do not exist, and I omit 1910 because women's participation in occupations is recorded differently in that census (see Goldin, 1990 for additional details). For analysis using wages, I only rely on data from 1940 onward as individual income was not recorded prior to the 1940 Census.

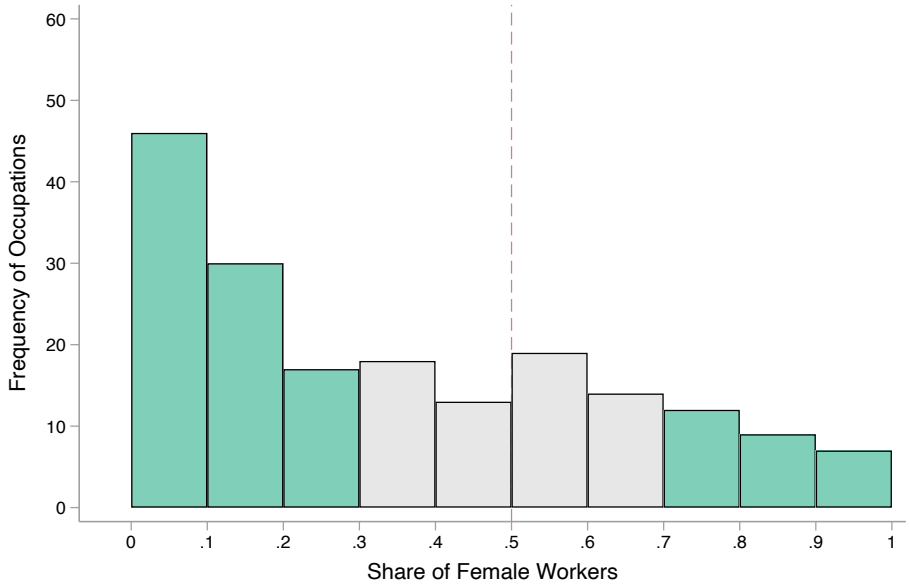
Table A.1: Occupations Included in the Descriptive Analysis

Accountant	Construction Painter	Furrier	Medical Attendant	Physician	Social scientist
Actor	Cooks	Garage Wkr	Medical/Dental Tech	Piano repair	Spinner
Advertising agent	Counter Wkr	Gardener	Metal Filer	Pilot	Stationary engineer
Amusement Attendant	Cranemen	Glazier	Metal Heater Op	Plasterer	Stationary firemen
Architect	Credit men	Guard	Metal Job Setter	Plumber	Steward
Artist	Crossing Watchmen	Heat treater	Metal Molder	Policemen	Stock Salesmen
Sportsperson	Dancer	Home Advisor	Metal Roller	Porter	Stone cutter
Auctioneer	Decorator	Hospital attendant	Midwive	Postmaster	Store Buyer
Author	Deliverymen	House Laundress	Millers	Power station Ops	Structural metal worker
Auto Attendant	Demonstrator	House Wkr	Milliner	Practical nurse	Superintendent
Baggagemen	Dentist	Housekeeper	Millwrights	Printing Pressmen	Survey Chainmen
Baker	Designer	Huckster	Mine Motormen	Professor/Instructor	Surveyor
Bank tellers	Dietitian	Insulation Wkr	Mine Op	Public Admin Inspector	Tailor
Bartender	Dispatcher	Insurance agent	Model maker	Public Admin Official	Teacher
Beautician/Barber	Draftsmen	Janitor	Movie Projectionist	Purchasing agent	Teamsters
Blacksmith	Dyer	Jeweler	Musician	Radio operator	Telegraph messenger
Blaster	Electrician	Labor Relations	Natural Scientist	Railroad Brakemen	Telegraph Op
Boatmen	Electrotyper	Laundry/Cleaning Op	Newsboy	Railroad Switchmen	Telephone Op
Boilermaker	Elevator Op	Lawyer	Nurse	Railway Conductor	Testing Tech
Bookbinder	Engineer	Librarian	Office boys	Railway messenger	Therapist
Bookkeepers	Engraver	Library Attendant	Office machine Op	Real estate agent	Ticket agent
Bootblack	Entertainer	Linemen	Oiler	Recreation	Tool Maker
Brickmason	Farm Buyer	Locomotive engineer	Optician	Religious Wkr	Trade Apprentice
Bus Conductor	Farm manager	Locomotive firemen	Optometrist	Road Machine Op	Train Motormen
Cabinetmaker	Farm service Wkr	Lodge/Union Official	Osteopath	Roofer	Truck Driver
Carpenter	Farmer	Lodging Housekeeper	Other Agent	Sailor	Typist/Secretary
Cement finisher	Firemen	Longshoremen	Other Inspector	Sawyer	Upholsterer
Chemist	Fishermen	Loom fixer	Other Tech	Seamstress	Usher
Chiropractor	Floormen	Lumbar Inspector	Painter	Service Attendant	Veterinarian
Cleaner	Food Grader	Lumbermen	Paperhanger	Sheet Metal Wkr	Waiter
Clergymen	Foremen	Machinist	Pharmacist	Sheriff	Weaver
Collectors	Forester	Mail carrier	Photoengraver	Ship Officer	Welder
Compositor/Typesetter	Funeral Director	Meat cutter	Photographer	Shipping clerk	Welfare
Constable	Furnacemen	Mechanics	Photographic Wkr	Shoemaker	

NOTE: This table lists the 203 occupations included in the descriptive analysis, based on the 1950 Census occupational classification harmonized across census years. The set excludes military occupations, occupations not elsewhere classified (n.e.c.) or not reported, and general occupation categories. Occupations are included if they employ at least 100 workers in at least one census year. For harmonization across census decades multiple specific detailed classification categories are combined into single occupations.

A.2 Additional Descriptive Patterns

Figure A.1: Occupational Gender Segregation, 2000 Census



NOTE: This figure plots the distribution of occupations by the share of women among their workers in 2000. Each bar shows the number of occupations whose share of women falls within a 10 percentage point interval. Green bars mark occupations that are predominantly male, with less than 30% women, or predominantly female, with more than 70% women. Predominantly male and predominantly female are defined using men and women's labor force participation at the time. For instance, occupations which have women over-represented in them, relative to their labor force participation at the time which was close to 70 percent, are labeled as predominantly female. Gray bars mark occupations with 30% to 70% women, considered as integrated occupations. The dashed line marks 50% women. The sample includes 203 occupations from the 2000 Census sample data (Ruggles et al., 2025).

Table A.2: Occupations Switching to More than 50% Female, 1870–2019

Occupation	Initial Share Women	Decade Crossed 50%	Share Women 2019	Occupation	Initial Share Women	Decade Crossed 50%	Share Women 2019
Library Attendant	0.31	1880	0.81	Author	0.05	1980	0.60
Nurse	0.44	1880	0.86	Collectors	0.02	1980	0.72
Cleaner	0.49	1880	0.88	Weaver	0.47	1970	0.69
Librarian	0.29	1900	0.83	Designer	0.05	1980	0.58
Medical Attendant	0.16	1900	0.87	Furrier	0.04	1980	0.69
Spinner	0.33	1900	0.84	Ticket agent	0.03	1980	0.68
Telephone Op	0.14	1900	0.79	Accountant	0.03	1990	0.62
Typist/Secretary	0.21	1900	0.94	Advertising agent	0.02	1990	0.53
Bookkeepers	0.04	1920	0.78	Baker	0.02	1990	0.65
Dancer	0.31	1920	0.78	Bartender	0.03	1990	0.60
Laundry/Cleaning Op	0.32	1920	0.65	Entertainer	0.21	1990	0.60
Steward	0.35	1920	0.76	Labor Relations	0.17	1990	0.68
Welfare	0.38	1920	0.78	Optician	0.01	1990	0.72
Waiter	0.29	1930	0.72	Postmaster	0.12	1990	0.63
Hospital attendant	0.49	1950	0.87	Public Admin Official	0.05	1990	0.53
Medical/Dental Tech	0.27	1950	0.71	Real estate agent	0.01	1990	0.55
Musician	0.40	1950	0.65	Social scientist	0.03	1990	0.59
Bank tellers	0.01	1960	0.87	Tailor	0.32	1990	0.52
Beautician/Barber	0.06	1960	0.85	Dispatcher	0.13	2000	0.57
Counter Wkr	0.09	1960	0.67	Insurance agent	0.03	2000	0.60
Decorator	0.07	1960	0.66	Newsboy	0.03	2000	0.62
Floormen	0.10	1960	0.60	Purchasing agent	0.22	2000	0.54
Food Grader	0.23	1960	0.82	Superintendent	0.10	2000	0.51
Home Advisor	0.27	1960	0.59	Pharmacist	0.02	2010	0.55
Huckster	0.07	1960	0.72	Artist	0.17	2019	0.51
Recreation	0.05	1960	0.64	Professor/Instructor	0.13	2019	0.50
Therapist	0.12	1960	0.79	Veterinarian	0.04	2019	0.59
Crossing Watchmen	0.15	1970	0.63				

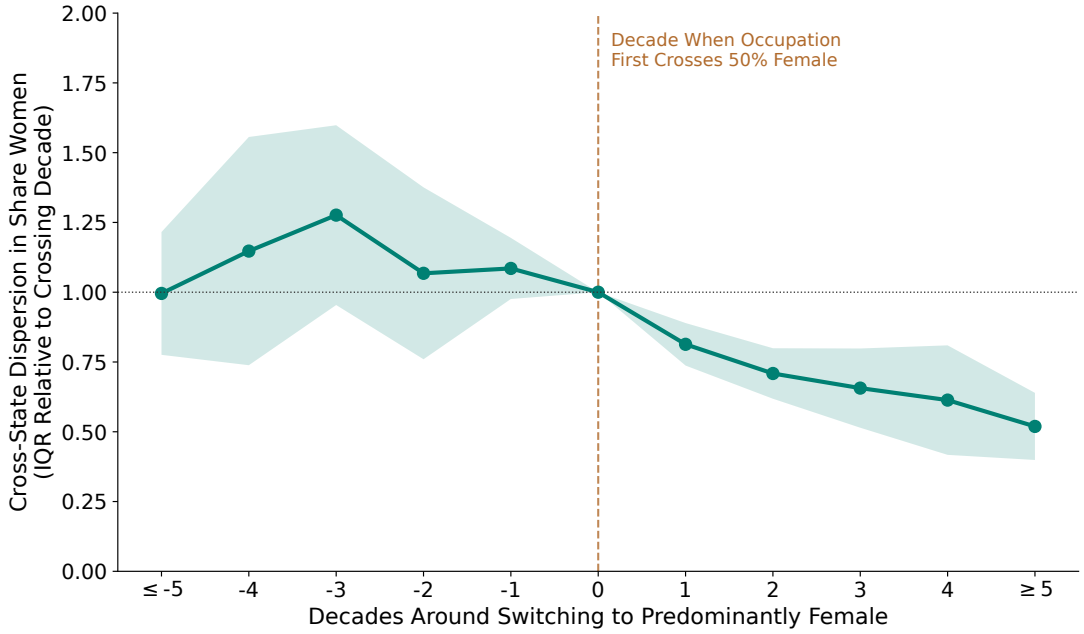
NOTE: This table lists the 55 occupations that switched permanently from predominantly male to predominantly female between 1870 and 2019, ordered by the decade in which each occupation first crossed 50 percent female without reverting in any subsequent census. Occupation-year cells with fewer than 100 workers nationally are excluded, and initial shares refer to each occupation's first included census year. Occupations crossing 50 percent in 2019 are classified using their final observation. Values from the 1910 Census are omitted because women's participation in occupations is recorded differently in that Census.

Table A.3: Years To Reach More Than 50% Female From Emergence, New versus Old Occupations

New Occupations (N = 23)		**Old Occupations (N = 32)**	
Occupation	Years	Occupation	Years
Library Attendant	10	Cleaner	10
Nurse	10	Spinner	30
Librarian	30	Bookkeepers	50
Medical Attendant	30	Laundry/Cleaning Op	50
Medical/Dental Tech	30	Steward	50
Telephone Op	30	Waiter	60
Typist/Secretary	30	Hospital attendant	80
Home Advisor	40	Musician	80
Welfare	40	Bank tellers	90
Dancer	50	Beautician/Barber	90
Counter Wkr	60	Huckster	90
Decorator	60	Crossing Watchmen	100
Floormen	60	Weaver	100
Recreation	60	Author	110
Labor Relations	70	Collectors	110
Food Grader	80	Furrier	110
Therapist	90	Ticket agent	110
Designer	110	Accountant	120
Social scientist	110	Baker	120
Advertising agent	120	Bartender	120
Optician	120	Entertainer	120
Dispatcher	130	Postmaster	120
Purchasing agent	130	Public Admin Official	120
		Real estate agent	120
		Tailor	120
		Insurance agent	130
		Newsboy	130
		Superintendent	130
		Pharmacist	140
		Artist	149
		Professor/Instructor	149
		Veterinarian	149
Median	60	Median	110

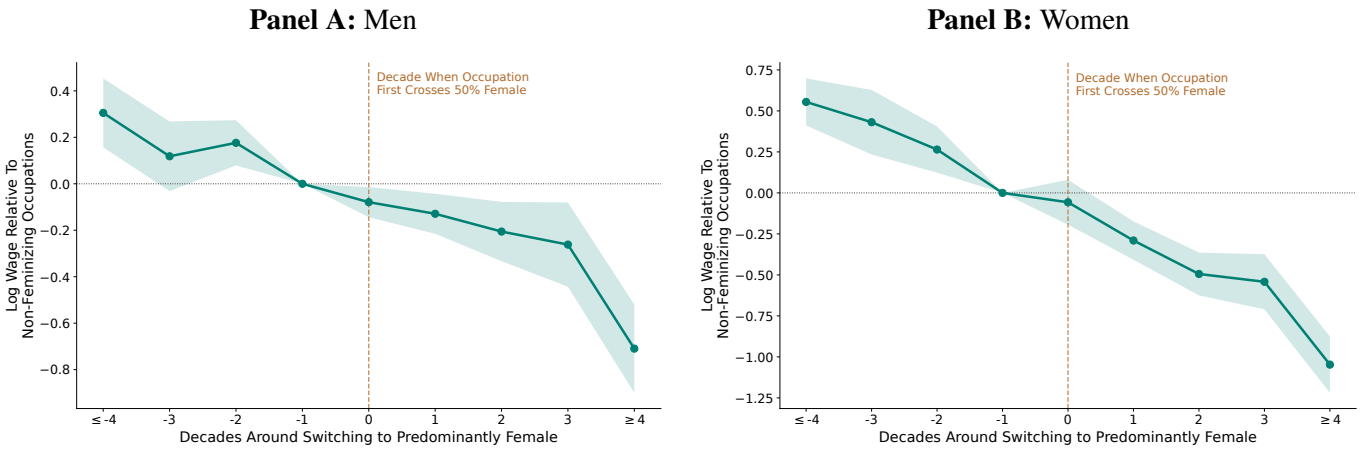
NOTE: This table lists, for each occupation that switched permanently from predominantly male to predominantly female, the number of years between the occupation's first census with at least 100 workers and the census in which it permanently crossed 50 percent female, corresponding to Figure 1, Panel B. New occupations are those first employing more than 500 workers after 1870, and old occupations employed more than 500 workers already in 1870, so old occupations' durations are measured from 1870 and are lower bounds. Occupations crossing 50 percent in 2019 are classified using their final observation.

Figure A.2: Geographic Dispersion in Share Women of Occupations Around Feminization



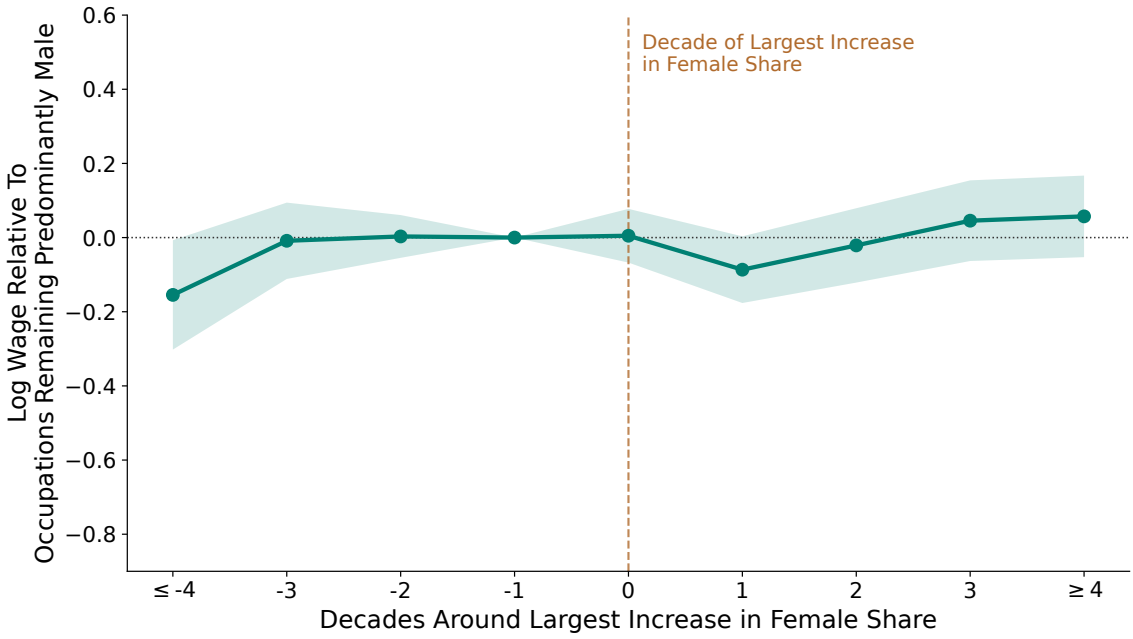
NOTE: This figure plots the evolution of the geographic dispersion of occupations' female shares around the decade in which each occupation switched to predominantly female. The sample comprises the occupations that switched permanently from predominantly male to predominantly female between 1870 and 2019, with event time zero defined as the census in which an occupation first crossed 50 percent female without reverting in any subsequent census. For each occupation and census year, dispersion is measured as the interquartile range of the occupation's female share across the 48 continental US states, excluding the District of Columbia. State-year cells with fewer than 100 workers in the occupation are excluded, and occupation-year observations are retained only when at least ten states contribute. Each occupation's interquartile range is expressed relative to its own value in the crossing decade, so that all occupations equal one at event time zero, and a value of 0.5 indicates that dispersion has fallen to half its crossing-decade level. The solid line shows the average across occupations at each event time and the shaded region the corresponding 95 percent confidence interval. Event times of five or more decades before and after crossing are binned at -5 and 5. Values from the 1910 Census are omitted because women's participation in occupations is recorded differently in that Census.

Figure A.3: Relative Wages of Men and Women in Occupations That Feminize



NOTE: This figure plots mean incomes of men (Panel A) and women (Panel B) in occupations that switched permanently from predominantly male to predominantly female, relative to occupations that began predominantly male and never crossed 50 percent female, in decades around each occupation's switch. In each panel, the outcome is the log of the occupation's mean income among the corresponding group minus the log of the employment-weighted mean income for that group among never-switching occupations in the same year, normalized to zero one decade before the occupation crosses 50 percent female. Incomes are adjusted to 2019 dollars and measured from census years 1940 through 2019. Occupations predominantly female in 1940 and occupations that crossed 50 percent female only temporarily are excluded. Lines show averages across occupations, shaded regions 95 percent confidence intervals, and event times of four or more decades are binned at -4 and 4. Occupation-year cells in which the group has no recorded income are excluded from that panel.

Figure A.4: Relative Wages in Occupations With Female Entry That Did Not Feminize



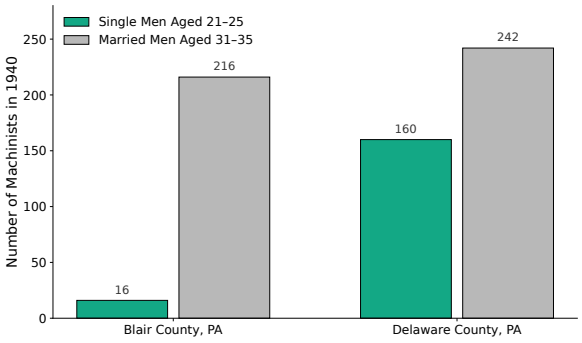
NOTE: This figure plots mean incomes in occupations that began predominantly male, experienced an increase of at least 10 percentage points in their female share, but never crossed 50 percent female, relative to occupations that began predominantly male and whose female share rose by less than 5 percentage points. The event is each occupation's decade of largest one-decade increase in female share. The outcome is the log of the occupation's mean income minus the log of the employment-weighted mean income among benchmark occupations in the same year, normalized to zero one decade before the event, with incomes adjusted to 2019 dollars and measured from census years 1940 through 2019. The line shows the average across occupations, the shaded region the 95 percent confidence interval, and event times of four or more decades are binned at -4 and 4.

B SSIV Construction and Robustness

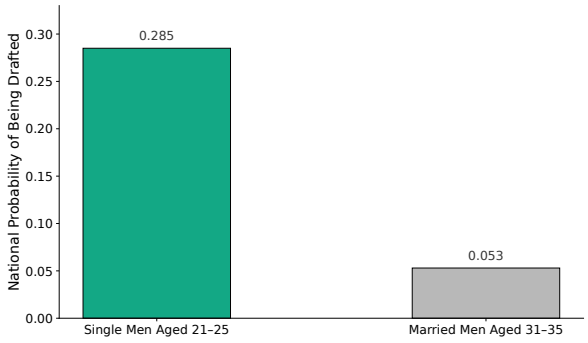
B.1 Shift Share IV Variation

Figure B.1: Intuition Underlying Shift Share Approach

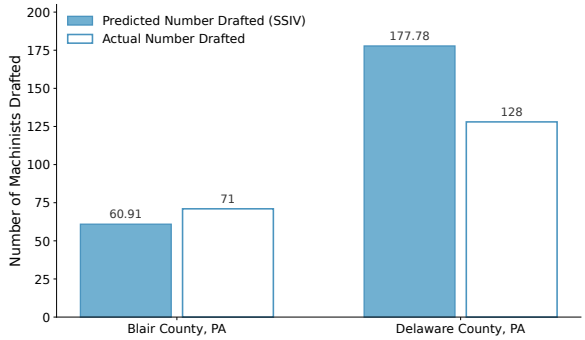
Panel A: Local Demographic Composition (Shares)



Panel B: National Draft Probabilities (Shifts)

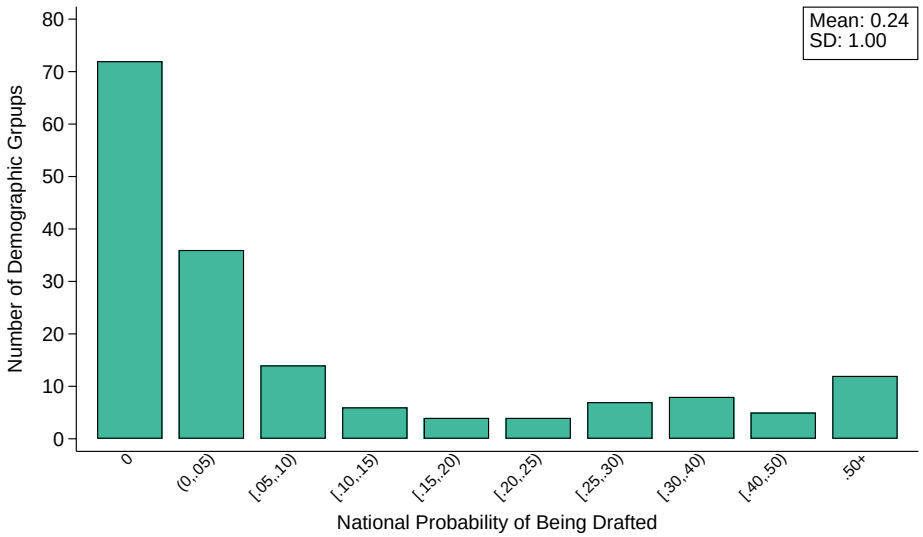


Panel C: Predicted and Actual Number Drafted



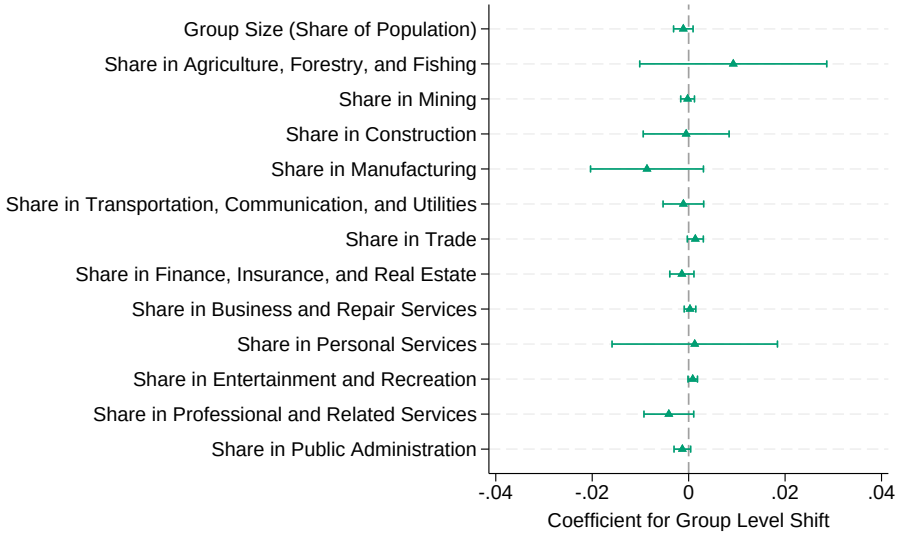
NOTE: This figure illustrates the intuition behind the shift share IV using the occupation machinists in two counties: Blair county, PA and Delaware county, PA, and two demographic groups: White US-born single men aged 21–25 and White US-born married men aged 31-35. The two demographic groups differ in the marital status and age. **Panel A** plots the number of machinists in each county in 1940 belonging to the two demographic groups. **Panel B** plots the national probability of being drafted for each demographic group, computed as the number in each group drafted and enlisted nationally divided by the group's population in the 1940 census. **Panel C** plots the resulting SSIV, the predicted number of drafted machinists in each county, after aggregating over all demographic groups for each machinist-county cell. It also plots the endogenous variable, the actual number of machinists drafted from each county.

Figure B.2: Variation in National Draft Probabilities Across Demographic Groups



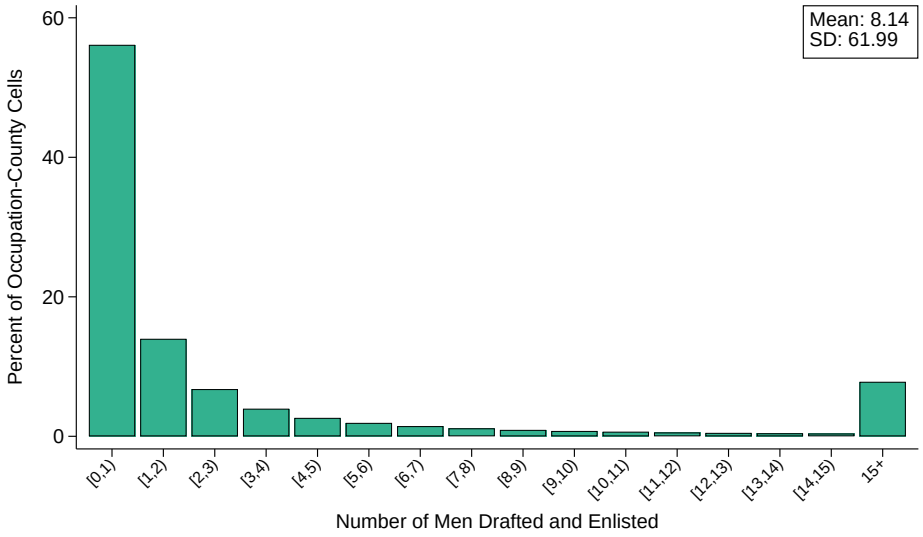
Note: This figure plots the distribution of the shifts used in the construction of the shift share IV, the national probability of being drafted for different demographic groups of worker.

Figure B.3: Balance of Shifts on 1940 Demographic Group Characteristics



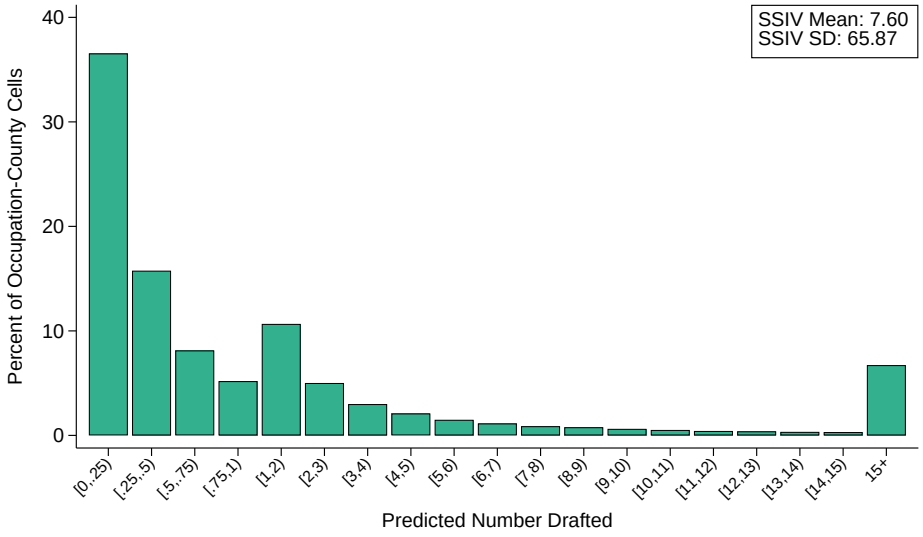
Note: This figure plots the coefficients obtained from regressing each of the listed 1940 demographic group characteristics on the national draft probabilities (shifts). Group size is the group's share of the 1940 population aged 16–64. The industry characteristics are the share of the group's population aged 16–64 employed in each major industry category in 1940. Industry categories are constructed using the *ind1950* variable from Ruggles et al. (2021).

Figure B.4: Variation in the Number of Men who are Drafted and Enlisted



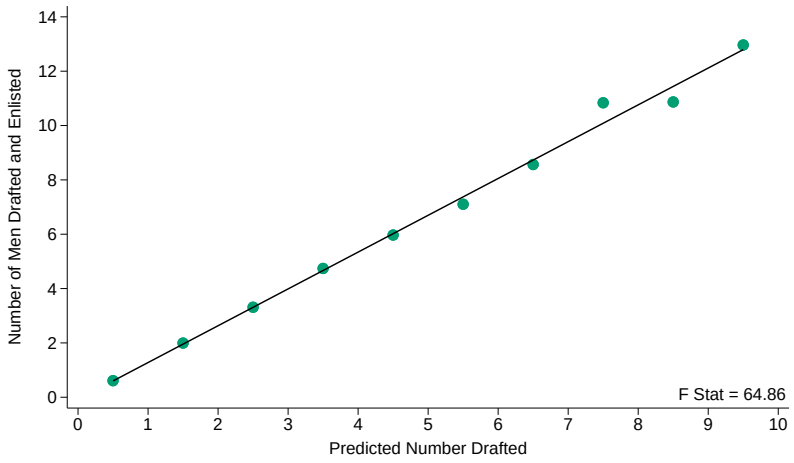
NOTE: This figure plots the distribution of the endogenous variable, the number of men who are drafted and enlisted from each occupation-county cell, and reports the mean and standard deviation. Cells that appear in 1940 Census but do not have any enlistment from them in the Army enlistment records are assigned 0 number of men drafted and enlisted.

Figure B.5: Variation in the Predicted Number Drafted



NOTE: This figure plots the distribution of the shift share IV, the predicted number drafted from each occupation-county cell, and reports the mean and standard deviation.

Figure B.6: Variation Underlying the First-Stage Relationship Between the SSIV and Number of Drafted and Enlisted Men



Note: This figure plots the mean number of men drafted and enlisted from occupation county cells for each bin of the instrument, the predicted number of drafted, with a line of best fit. The F-statistic is obtained from the first stage regression which includes baseline occupation-county level controls, occupation fixed effects, and county fixed effects. Cells in the top decile of the instrument, which have a predicted number drafted of greater than 10, are omitted from the display. These cells have a mean number of men drafted and enlisted of 76.19 and a mean predicted number drafted of 76.06.

B.2 Placebo Test

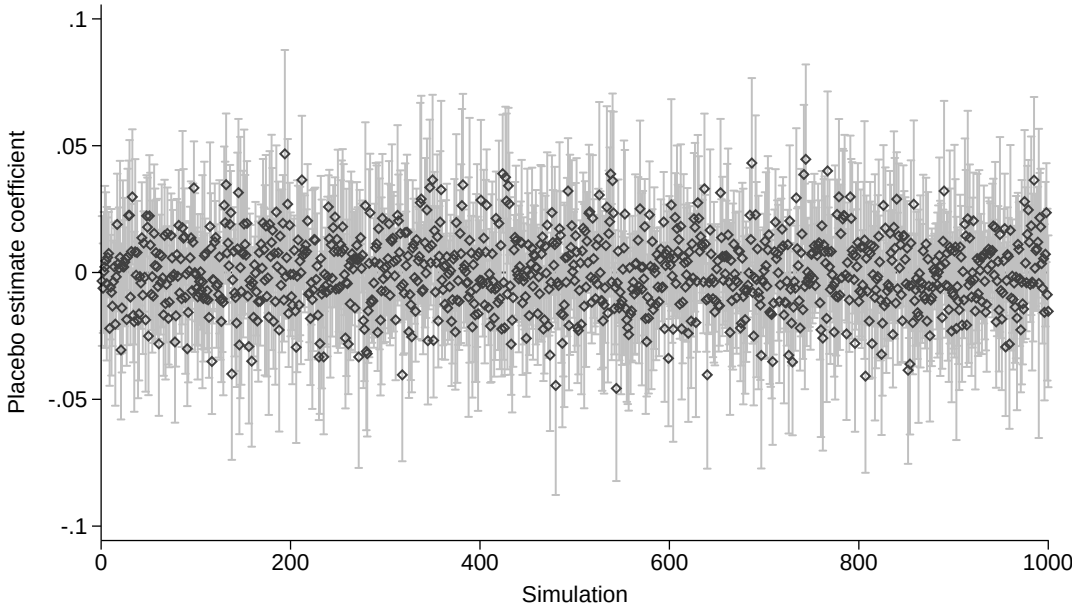
Table B.1: Placebo Test of the Identification Strategy:
Reduced-Form Effect of SSIV on Pre-Period Women's Employment in Each Occupation-County

	Number of Women$_{oc}$, 1900	Number of Women$_{oc}$, 1920	Number of Women$_{oc}$, 1930
$\widehat{\text{Number Drafted}}_{oc}$	-0.200	0.081	0.027
	(0.298)	(0.207)	(0.100)
Outcome Mean	18.5	21.3	24.5
Observations	208,031	270,883	297,755
Baseline Controls	Y	Y	Y
County FE	Y	Y	Y
Occ FE	Y	Y	Y

Note: This table reports the reduced-form effect of the SSIV on the number of women employed in each occupation-county prior to 1940. The unit of observation is an occupation-county cell. The dependent variable is the number of women employed in each occ-county in 1900 (column 1), 1920 (column 2) and 1930 (column 3). 1940 is excluded because baseline 1940 controls include the number of women in the cell in 1940. 1910 is excluded because women's employment in this census is measured differently from other census decades (see Goldin, 1990 for additional details). Baseline 1940 controls include the following demographic and industry controls: number of women, cell population, number of married men, number of workers in the agricultural, mining, manufacturing, and public administration industries. All regressions include county and occupation fixed effects. Robust standard errors clustered at the county level are reported in parentheses.

B.3 Random Shifts Placebo Test

Figure B.7: Coefficients Obtained from SSIV Using Random Placebo Shifts



NOTE: Instrumental variables regressions of the number of men drafted on women's employment in 1950. The number of men drafted is instrumented using a shift-share instrument based on the predetermined shares and a randomly generated shift. The random shift was generated based on a normal distribution with mean zero and variance five as in Adao et al. (2019). The figure shows the coefficients and 95% confidence intervals from instrumental variables regressions where the instrument was generated with 1,000 random shifts. Regressions include the baseline set of controls with census region fixed effects and assume robust standard errors.

C Additional Empirical Results

C.1 Heterogeneity

Table C.1: Heterogeneity in the Effect of Number Men Drafted on Number of Women Employed in 1950 by Age

Panel A: First Stage on MenDrafted$_{oc,1940-1945}$				
$\widehat{\text{NumberDrafted}}_{oc}$	1.353	1.353	1.353	1.353
	(0.278)	(0.278)	(0.278)	(0.278)
F-statistic	64.84	64.84	64.84	64.84
	Number of Women Employed in 1950			
	All Women	Aged 18-22	Aged 23-40	Aged 23-64
	(1)	(2)	(3)	(4)
Panel B: OLS				
MenDrafted$_{oc,1940-1945}$	0.302	0.128	0.187	0.170
	(0.048)	(0.016)	(0.034)	(0.037)
Panel C: Reduced Form				
$\widehat{\text{NumberDrafted}}_{oc}$	0.719	0.250	0.530	0.425
	(0.152)	(0.058)	(0.107)	(0.106)
Panel D: 2SLS				
MenDrafted$_{oc,1940-1945}$	0.531	0.185	0.392	0.314
	(0.116)	(0.038)	(0.080)	(0.081)
Outcome Mean	25.19	3.940	11.33	20.61
Observations	368,291	368,291	368,291	368,291
Baseline Controls	Y	Y	Y	Y
County FE	Y	Y	Y	Y
Occ FE	Y	Y	Y	Y

NOTE: Regressions in Panel A document the first stage relationship between the SSIV and the number of men drafted from each occupation-county cell. Panels B, C and D use the number of women employed in 1950 in each age group as outcome: all women in column 1, women aged 18-22 in column 2, women aged 23-40 in column 3, and women aged 23-64 in column 4. The unit of observation is an occupation-county cell. All columns include baseline occupation-county level controls, county fixed effects and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, number of workers in the agricultural, mining, manufacturing, and public administration industries. Robust standard errors clustered at the county level are reported in parentheses.

Table C.2: Heterogeneity in the Effect of Number Men Drafted on Number of Women Employed in 1950 by Marital Status and Children

Panel A: First Stage on MenDrafted$_{oc,1940-1945}$					
$\widehat{\text{NumberDrafted}}_{oc}$	1.353	1.353	1.353	1.353	1.353
	(0.278)	(0.278)	(0.278)	(0.278)	(0.278)
F-statistic	64.84	64.84	64.84	64.84	64.84
	Number of Women Employed in 1950				
	Married	Single	Mothers	Mothers with Young Children	Non-Mothers
	(1)	(2)	(3)	(4)	(5)
Panel B: OLS					
MenDrafted$_{oc,1940-1945}$	0.124	0.218	0.040	0.028	0.266
	(0.027)	(0.030)	(0.014)	(0.006)	(0.043)
Panel C: Reduced Form					
$\widehat{\text{NumberDrafted}}_{oc}$	0.395	0.247	0.251	0.164	0.479
	(0.087)	(0.113)	(0.072)	(0.022)	(0.145)
Panel D: 2SLS					
MenDrafted$_{oc,1940-1945}$	0.292	0.183	0.186	0.122	0.354
	(0.082)	(0.069)	(0.074)	(0.033)	(0.084)
Outcome Mean	12.21	8.095	7.501	2.180	17.58
Observations	368,291	368,291	368,291	368,291	368,291
Baseline Controls	Y	Y	Y	Y	Y
County FE	Y	Y	Y	Y	Y
Occ FE	Y	Y	Y	Y	Y

NOTE: Regressions in Panel A document the first stage relationship between the SSIV and the number of men drafted from each occupation-county cell. Panels B, C and D use the number of women employed in 1950 in each group as outcome: married women in column 1, single women in column 2, mothers in column 3, mothers with young children in column 4, and women without children in column 5. The unit of observation is an occupation-county cell. All columns include baseline occupation-county level controls, county fixed effects and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, number of workers in the agricultural, mining, manufacturing, and public administration industries. Robust standard errors clustered at the county level are reported in parentheses.

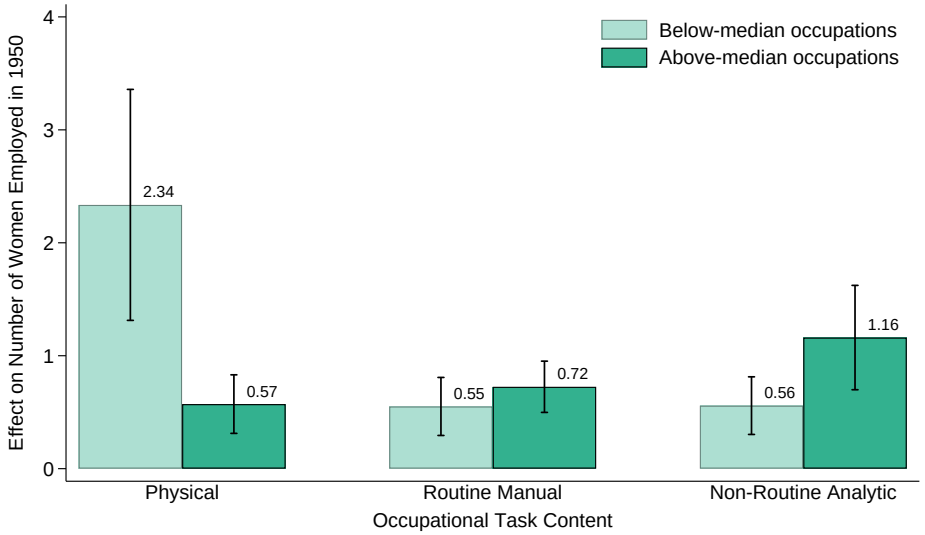
Table C.3: Heterogeneity in the Effect of Number Men Drafted on Employment in 1950 by Race and Gender

Panel A: First Stage on MenDrafted$_{oc,1940-1945}$				
$\widehat{\text{NumberDrafted}}_{oc}$	1.353	1.353	1.353	1.353
	(0.278)	(0.278)	(0.278)	(0.278)
F-statistic	64.84	64.84	64.84	64.84
	Number Employed in 1950			
	Women		Men	
	White	Black	White	Black
	(1)	(2)	(3)	(4)
Panel B: OLS				
MenDrafted$_{oc,1940-1945}$	0.419	-0.117	1.986	-0.194
	(0.068)	(0.023)	(0.140)	(0.064)
Panel C: Reduced Form				
$\widehat{\text{NumberDrafted}}_{oc}$	0.392	0.312	0.514	3.038
	(0.224)	(0.096)	(1.343)	(0.609)
Panel D: 2SLS				
MenDrafted$_{oc,1940-1945}$	0.290	0.231	0.380	2.246
	(0.140)	(0.107)	(0.925)	(0.842)
Outcome Mean	20.95	4.059	62.51	7.597
Observations	368,291	368,291	368,291	368,291
Baseline Controls	Y	Y	Y	Y
County FE	Y	Y	Y	Y
Occ FE	Y	Y	Y	Y

NOTE: Regressions in Panel A document the first stage relationship between the SSIV and the number of men drafted from each occupation-county cell. Panels B, C and D use the number of individuals employed in 1950 in each group as outcome: white women in column 1, Black women in column 2, white men in column 3, and Black men in column 4. The unit of observation is an occupation-county cell. All columns include baseline occupation-county level controls, county fixed effects and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, number of workers in the agricultural, mining, manufacturing, and public administration industries. Robust standard errors clustered at the county level are reported in parentheses.

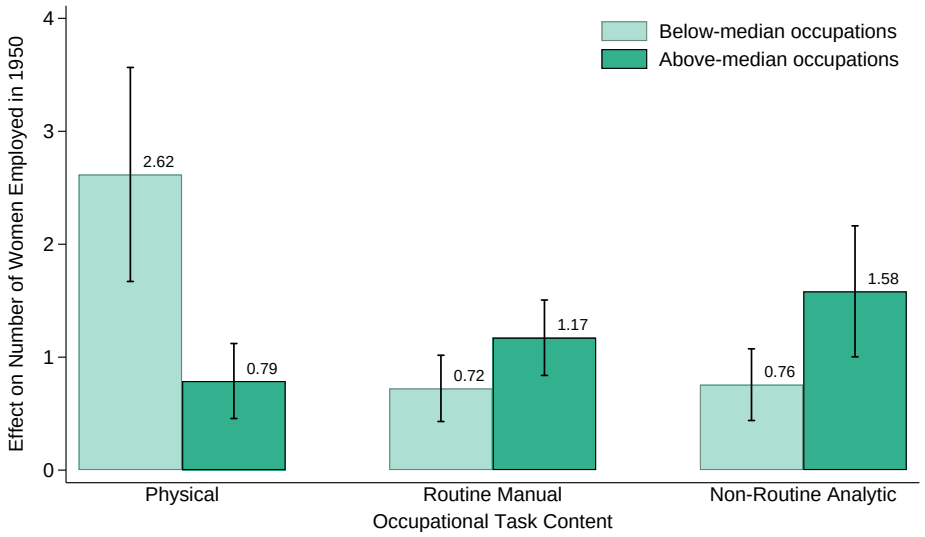
C.2 Analysis Using Amenities

Figure C.1: Variation in the Effect of War Mobilization on Number of Women Employed in 1950 by Occupational Task Content, 2SLS Estimates



NOTE: This figure plots 2SLS estimates of the effect of the SSIV on the number of women employed in 1950, separately for occupations below and above the median of each task content measure, along with 95% confidence intervals. Each pair of bars corresponds to a separate regression of the outcome on the SSIV and its interaction with an indicator for above-median task content. The below-median bar reports the main effect and the above-median bar reports the sum of the main effect and the interaction. The unit of observation is an occupation-county cell. All regressions include baseline occupation-county level controls, county fixed effects, and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, and number of workers in the agricultural, mining, manufacturing, and public administration industries. Confidence intervals are based on robust standard errors clustered at the county level.

Figure C.2: Variation in the Effect of War Mobilization on Number of Women Employed in 1950 by Occupational Task Content, Reduced Form Estimates



NOTE: This figure plots reduced form estimates of the effect of the SSIV on the number of women employed in 1950, separately for occupations below and above the median of each task content measure, along with 95% confidence intervals. Each pair of bars corresponds to a separate regression of the outcome on the SSIV and its interaction with an indicator for above-median task content. The below-median bar reports the main effect and the above-median bar reports the sum of the main effect and the interaction. The unit of observation is an occupation-county cell. All regressions include baseline occupation-county level controls, county fixed effects, and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, and number of workers in the agricultural, mining, manufacturing, and public administration industries. Confidence intervals are based on robust standard errors clustered at the county level.

Table C.4: Effect of Number Men Drafted on Number of Women Employed in 1950 by Occupational Task Content, 2SLS Estimates

	Number of Women Employed in 1950		
	Physical (1)	Routine Manual (2)	Non-Routine Analytic (3)
MenDrafted_{oc}	2.336	0.551	0.558
	(0.522)	(0.131)	(0.130)
$\text{MenDrafted}_{oc} \times \mathbb{1}\{AboveMedian_o\}$	-1.764	0.174	0.603
	(0.464)	(0.083)	(0.210)
Total Effect (Above Median)	0.572	0.725	1.161
	(0.132)	(0.116)	(0.236)
First Stage F-statistic (Main)	101.6	70.20	63.45
First Stage F-statistic (Interaction)	74.09	102.9	67.48
Sanderson and Windmeijer (2016) F-statistic (Main)	87.87	29.69	49.50
Sanderson and Windmeijer (2016) F-statistic (Interaction)	59.56	73.66	86.61
Outcome Mean	26.45	26.45	26.45
Observations	345,572	345,572	345,572
Baseline Controls	Y	Y	Y
County FE	Y	Y	Y
Occ FE	Y	Y	Y

NOTE: This table reports 2SLS estimates of Equation (**??**), using the number of women employed in 1950 as outcome. Each column interacts the number of men drafted with an indicator for whether the occupation is above the median of the task content measure indicated in the column header: physical intensity in column 1, routine manual content in column 2, and non-routine analytic content in column 3. The number of men drafted and its interaction are instrumented with the SSIV and its interaction with the same indicator. The main effect of the above-median indicator is absorbed by the occupation fixed effects. The row "Total Effect (Above Median)" reports the sum of the main effect and the interaction, with its standard error in parentheses. First stage F-statistics are reported separately for the first stage of each endogenous variable, and Sanderson and Windmeijer (2016) F-statistics are the conditional first stage F-statistics for each endogenous variable, which account for the presence of multiple endogenous regressors. The unit of observation is an occupation-county cell. The sample is restricted to cells in occupations with available task content measures. All columns include baseline occupation-county level controls, county fixed effects and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, number of workers in the agricultural, mining, manufacturing, and public administration industries. Robust standard errors clustered at the county level are reported in parentheses.

Table C.5: Correlation Between Occupational Flexibility and Task Content

	Physical (1)	Routine Manual (2)	Non-Routine Analytic (3)
Flexibility Index (Goldin, 2014b)	-0.119	-0.034	0.426
	(0.109)	(0.169)	(0.129)

NOTE: This table reports coefficients from regressions of each standardized task content measure on the standardized flexibility index. The flexibility index is drawn from Goldin (2014b) and measures the degree of temporal flexibility in each occupation, with higher values indicating greater flexibility. Task content measures are drawn from Gray (2013); Gray et al. (2025) and capture the physical intensity (column 1), routine manual content (column 2), and non-routine analytic content (column 3) of each occupation. All variables are standardized to have mean zero and standard deviation one, so each coefficient gives the standard deviation change in the task measure associated with a one standard deviation increase in the flexibility index. The unit of observation is an occupation. Standard errors are reported in parentheses.

Table C.6: Effect of Number Men Drafted on Number of Women Employed in 1950 by Occupational Hours Flexibility, 2SLS Estimates

	Dependent Variable: Number of Women Employed
	(1)
MenDrafted_{oc}	0.763
	(0.243)
$\text{MenDrafted}_{oc} \times \mathbb{1}\{\textit{FlexibleHours}_o\}$	0.456
	(0.471)
Effect in Flexible-Hours Occupations	1.220
	(0.592)
First Stage F-statistic (Main)	15.57
First Stage F-statistic (Interaction)	13.91
Sanderson and Windmeijer (2016) F-statistic (Main)	21.28
Sanderson and Windmeijer (2016) F-statistic (Interaction)	7.41
Anderson and Rubin (1949) test (p-value)	0.000
Outcome Mean	25.13
Observations	369,185
Baseline Controls	Y
County FE	Y
Occ FE	Y

NOTE: This table reports 2SLS estimates interacting the number of men drafted with an indicator for occupations with flexible hours, $\mathbb{1}\{\textit{FlexibleHours}_o\}$. Occupations are classified using mean weekly hours worked in the 1940 complete count census, winsorized at the 99th percentile, with flexible-hours occupations defined as those with mean weekly hours of 40 or fewer. The number of men drafted and its interaction with the indicator are instrumented with the SSIV and its interaction with the same indicator. The main effect of the indicator is absorbed by the occupation fixed effects. The main effect gives the effect of the number of men drafted in occupations with long hours, and the row "Effect in Flexible-Hours Occupations" reports the sum of the main effect and the interaction, with its standard error in parentheses. First stage F-statistics are the cluster-robust joint tests of the excluded instruments in the first stage of each endogenous variable, and Sanderson and Windmeijer (2016) F-statistics are the conditional first stage F-statistics for each endogenous variable, which account for the presence of multiple endogenous regressors. I also report the weak-instrument-robust Anderson and Rubin (1949) test of the joint significance of the endogenous regressors. The unit of observation is an occupation-county cell. The outcome mean is unweighted. All specifications include baseline occupation-county level controls, county fixed effects and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, number of workers in the agricultural, mining, manufacturing, and public administration industries. Robust standard errors clustered at the county level are reported in parentheses.

Table C.7: Effect of Number Men Drafted on Number of Women Employed in 1950 by Occupational Amenities, 2SLS Estimates

	Dependent Variable: Number of Women Employed		
	Flexibility Index (1)	Cleanliness Index (2)	Composite Index (3)
MenDrafted_{oc}	0.537	0.590	0.578
	(0.127)	(0.164)	(0.180)
$\text{MenDrafted}_{oc} \times \textit{Amenity}_{o}$	0.064	0.100	0.097
	(0.179)	(0.128)	(0.181)
Effect at $+1$ SD	0.601	0.691	0.675
	(0.287)	(0.278)	(0.350)
First Stage F-statistic (Main)	58.22	96.48	57.97
First Stage F-statistic (Interaction)	49.57	44.67	72.44
Sanderson and Windmeijer (2016) F-statistic (Main)	10.14	49.39	12.54
Sanderson and Windmeijer (2016) F-statistic (Interaction)	15.55	30.10	16.17
Anderson and Rubin (1949) test (p-value)	0.000	0.000	0.000
Outcome Mean	25.19	25.31	25.37
Observations	367,932	366,158	364,905
Baseline Controls	Y	Y	Y
County FE	Y	Y	Y
Occ FE	Y	Y	Y

NOTE: This table reports 2SLS estimates interacting the number of men drafted with a continuous occupation-level amenity score, $\textit{Amenity}_{o}$, standardized to mean zero and standard deviation one: a flexibility index in column 1, a cleanliness index in column 2, and a composite index averaging the two in column 3. These indices are constructed using Talkie-1930 LM Levine et al. (2026). The number of men drafted and its interaction with the amenity score are instrumented with the SSIV and its interaction with the same score. The main effect of the amenity score is absorbed by the occupation fixed effects. The main effect gives the effect of the number of men drafted at the mean of the amenity score, the interaction gives the change in this effect per standard deviation of the score, and the row "Effect at $+1$ SD" reports the sum of the two, with its standard error in parentheses. First stage F-statistics are reported separately for the first stage of each endogenous variable, and Sanderson and Windmeijer (2016) F-statistics are the conditional first stage F-statistics for each endogenous variable, which account for the presence of multiple endogenous regressors. Because the specifications are weakly identified, I also report the weak-instrument-robust Anderson-Rubin test of the joint significance of the endogenous regressors. The unit of observation is an occupation-county cell. Samples differ slightly across columns because the amenity scores are available for slightly different sets of occupations. All specifications include baseline occupation-county level controls, county fixed effects and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, number of workers in the agricultural, mining, manufacturing, and public administration industries. Robust standard errors clustered at the county level are reported in parentheses.

C.3 Palmer Survey Analysis Additional Figures and Tables

Table C.8: Classification of Palmer Survey Occupations by Female Amenity Type

High female amenity	Low female amenity
Accountant, Acct Collector, Actor, Advertising agent, Agent, Artist, Author, Baker, Bank teller, Beautician, Bookbinder, Bookkeeper, Building Mngr, Chemist, Cleaner, Clerical, Cook, Counter Wkr, Credit Wkr, Dancer, Decorator, Demonstrator, Designer, Draftsmen, Dressmaker, Entertainer, Food grader, Foremen, Funeral Related, Home Advisor, Hospital Attendant, Household Wkr, Housekeeper, Instructor, Insurance agent, Jeweler, Laundress, Laundry Op, Lawyer, Librarian, Med/Dental Attendant, Med/Dental Tech, Messenger, Milliner, Musician, Nurse, Nutritionist, Office Machine Op, Official/Managerial Wkr, Personnel, Photographer, Physician, Practical nurse, Professional Wkr, Public Admin Inspector, Public Admin Official, Radio Op, Real estate, Recreation Wkr, Religious Wkr, Salesperson, Secretary, Service Attendant, Service Wkr, Shipping Clerk, Social Wkr, Society Official, Spinner, Sports Wkr, Store Buyer/Head, Store Floor Wkr, Tailor, Teacher, Technician, Telegraph Op, Telephone Op, Testing tech, Therapist, Ticket/Station Agent, Waiter	Apprentice, Auto Attendant, Bartender, Bus Conductor, Compositor, Construction, Construction Painter, Delivery Wkr, Electrician, Elevator Op, Farm laborer, Farmer, Fisherman, Forester, Furnaceman, Garage Wkr, Guard, Inspector, Janitor, Laborer, Lithographer, Lumber Inspector, Machinist, Mail carrier, Meat cutter, Mechanic, Metal Wkr, Model maker, Operative, Painter, Peddler, Photographic Process Wkr, Police Wkr, Pressmen, Sawyer, Telegraph messenger, Truck Driver, Vehicle Dispatcher, Weaver, Welder

NOTE: This table lists the classification of the occupations recorded in the Palmer survey into high and low female-amenity types. High female-amenity occupations comprise clerical, professional, sales, communications, care, and personal service work, characterized by clean, safe, indoor work. Low female-amenity occupations comprise factory operative work, manual trades, transport, and outdoor work. The classification is applied uniformly to occupations reported in all survey years. Women not in the labor force are excluded from classification. General non-classified operatives are also excluded from the classification as they can comprise of both high and low female amenity occupations.

Figure C.3: New Occupations Entered During War

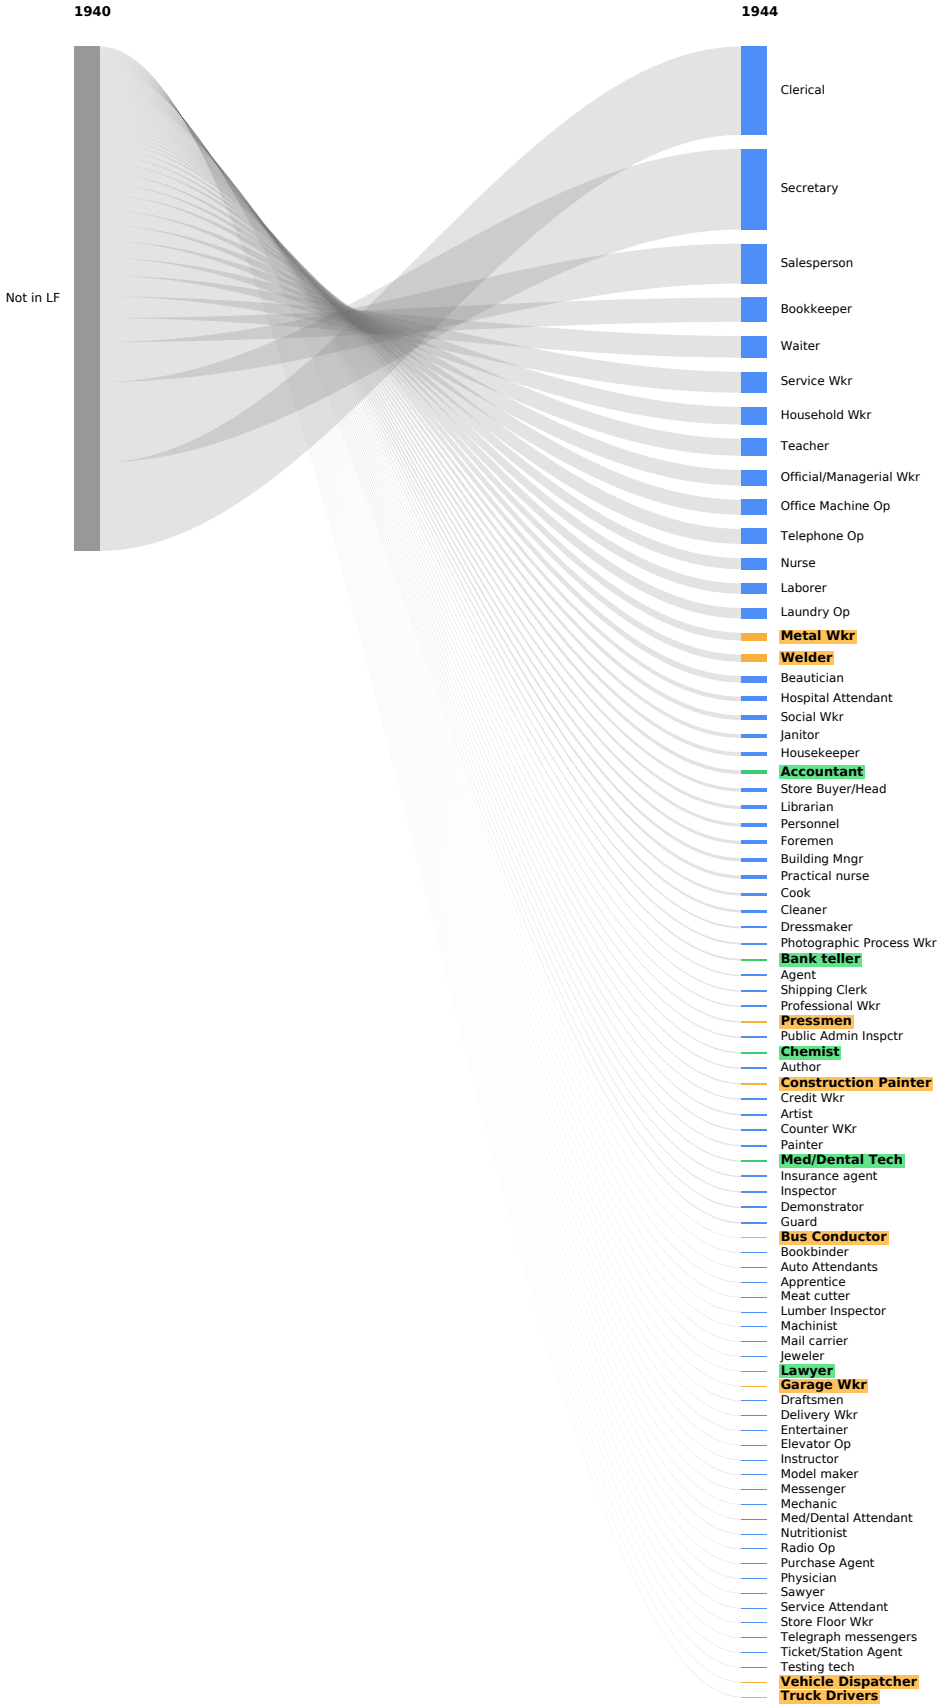
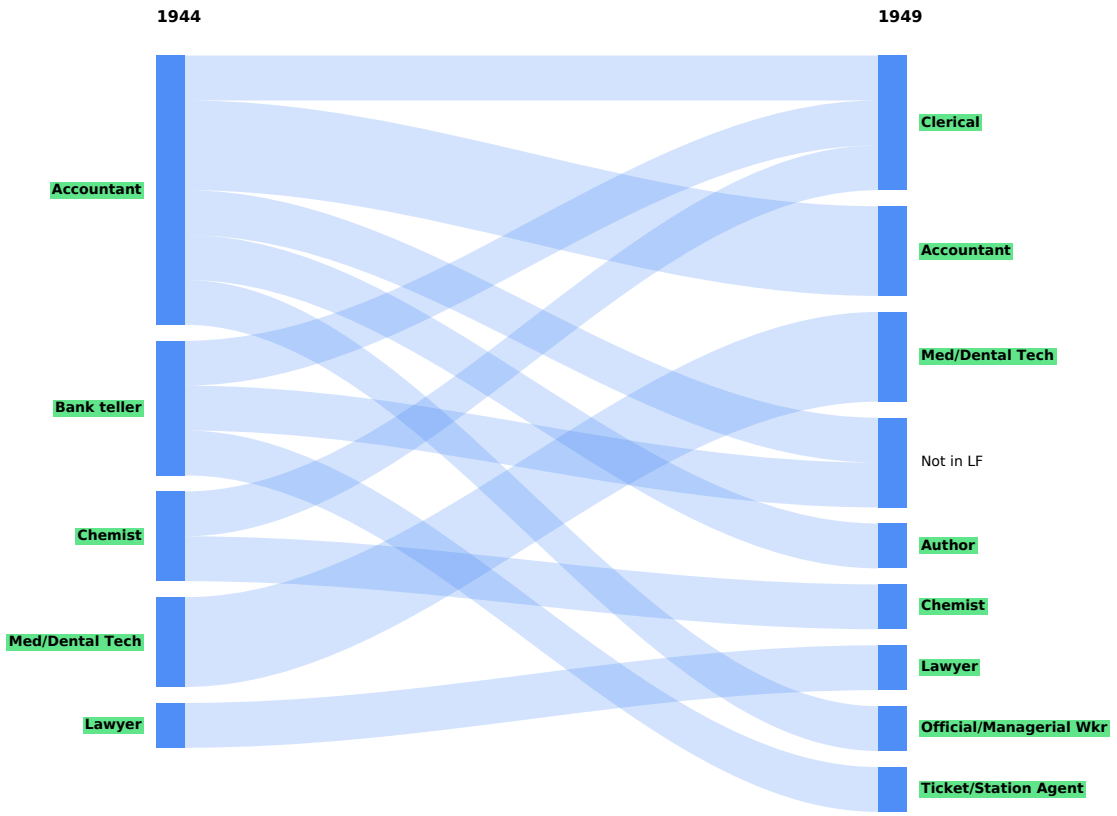
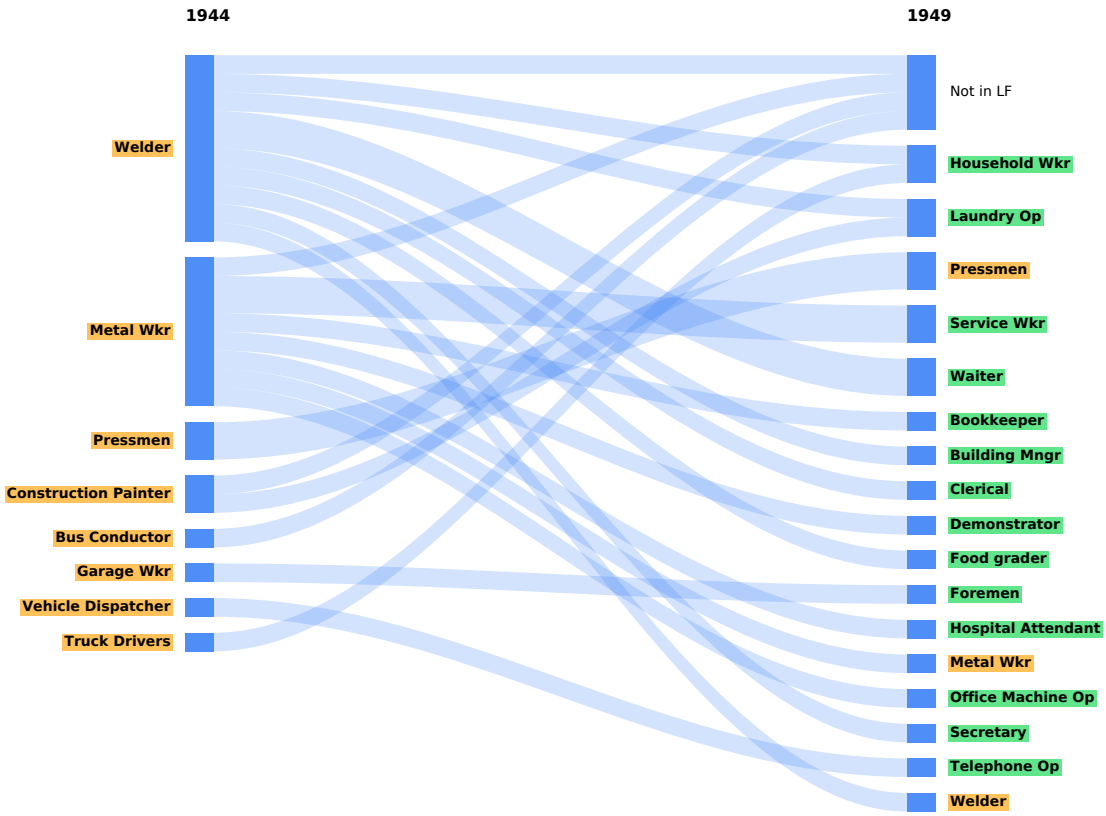


Figure C.4: Occupational Trajectories from 1944 to 1949 for Occupations with Female-Valued Attributes



NOTE: This figure traces the 1944–1949 occupational trajectories of women in the Palmer Survey who were not in the labor force in 1940 and were employed in the five selected occupations with female-valued attributes shown on the left in 1944. Flow widths are proportional to the unweighted number of women following each trajectory. Green label highlights identify occupations with female-valued attributes.

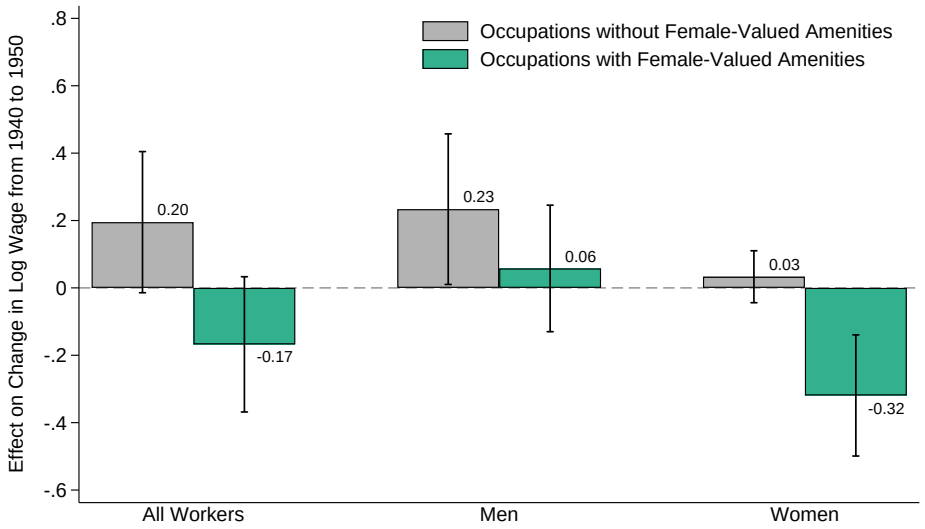
Figure C.5: Occupational Trajectories from 1944 to 1949 for Occupations without Female-Valued Attributes



NOTE: This figure traces the 1944–1949 occupational trajectories of women in the Palmer Survey who were not in the labor force in 1940 and were employed in the five selected occupations without female-valued attributes shown on the left in 1944. Flow widths are proportional to the unweighted number of women following each trajectory. Green label highlights identify occupations with female-valued attributes and orange without.

C.4 Wages

Figure C.6: Effect on Wages, by Gender and Occupational Characteristics



NOTE: This figure plots IV (2SLS) estimates of the effect of 100 additional men drafted from an occupation-county cell on the change in log mean wage in the cell between 1940 and 1950, for all workers, men, and women. Within each group, the gray bar reports the effect in occupations without female-valued amenities and the green bar reports the effect in occupations with female-valued amenities. Occupations with female-valued amenities are those with below-median physical task content in the 1939 Dictionary of Occupational Titles (Gray 2013). Each group is estimated in a separate regression in which the number of men drafted and its interaction with the amenity indicator are instrumented with the predicted number of men drafted and its interaction with the same indicator. The gray bar reports the main effect, and the green bar reports the sum of the main effect and the interaction. Each regression uses occupation-county cells in which the group's mean wage is observed in both 1940 and 1950. Wages are annual wage and salary income from the 1940 census and the 1950 census sample line, with 1940 wages converted to 1950 dollars using the CPI. All regressions include baseline controls and occupation and county fixed effects. The estimates for women are the same as in Figure 9. 95% confidence intervals are shown, based on standard errors clustered at the county level.

Table C.9: Effect of Number of Men Drafted on Change in Log Women's Years of Schooling, 1940–1950, by Occupational Physical Requirements: 2SLS Estimates

	Dependent Variable: Change in Log Women's Years of Schooling
	(1)
MenDrafted_{oc}	−0.0000579
	(0.0000351)
$\text{MenDrafted}_{oc} \times \mathbf{1}\{BelowMedianPhsyical_o\}$	−0.0000255
	(0.0000291)
Effect in Below-Median-Physical Occupations	−0.0000835
	(0.0000517)
First Stage F-statistic (Main)	15.89
First Stage F-statistic (Interaction)	37.75
Sanderson and Windmeijer (2016) F-statistic (Main)	21.00
Sanderson and Windmeijer (2016) F-statistic (Interaction)	55.26
Anderson and Rubin (1949) test (p-value)	0.114
Kleibergen–Paap rk Wald F-statistic	10.194
Outcome Mean	0.07770
Observations	99,015
Baseline Controls	Y
County FE	Y
Occ FE	Y

NOTE: This table reports 2SLS estimates interacting the number of men drafted with an indicator for occupations with below-median physical requirements, $\mathbf{1}\{BelowMedianPhsyical_o\}$. The number of men drafted and its interaction with the indicator are instrumented with the SSIV and its interaction with the same indicator. The main effect of the indicator is absorbed by occupation fixed effects. The main coefficient gives the effect of the number of men drafted when the indicator equals zero. The row "Effect in Below-Median-Physical Occupations" reports the sum of the main effect and the interaction, with its standard error in parentheses accounting for the covariance between the two coefficients. First stage F-statistics are cluster-robust joint tests of the excluded instruments in the first stage of each endogenous variable. The Sanderson and Windmeijer (2016) F-statistics are conditional first stage F-statistics that account for the presence of multiple endogenous regressors. The weak-instrument-robust Anderson and Rubin (1949) test evaluates the joint null that both endogenous coefficients equal zero. The Kleibergen–Paap rk Wald F-statistic is also reported. The unit of observation is an occupation-county cell, and the outcome mean is unweighted. All specifications include baseline controls, county fixed effects, and occupation fixed effects. Baseline 1940 controls include the number of women, cell population, number of married men, and employment counts in agriculture, mining, the two manufacturing categories, and public administration. Robust standard errors clustered at the county level are reported in parentheses.

C.5 Social Learning

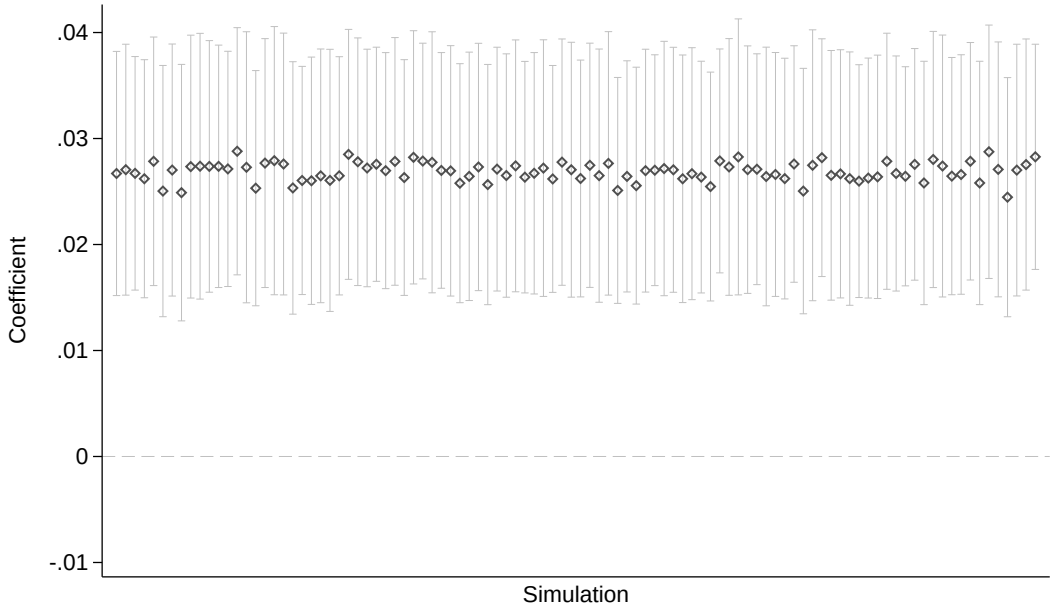
Table C.10: Effect of Number Men Drafted on Number of Women Aged 18–22 Employed in 1950 by Occupational Amenities, 2SLS Estimates

	Dependent Variable: Number of Women Aged 18–22 Employed				
	Less Physical	Flexible Hours	Flexible Schedule	Clean	Clean and Flexible
	(1)	(2)	(3)	(4)	(5)
MenDrafted_{oc}	0.187	0.204	0.190	0.217	0.224
	(0.040)	(0.045)	(0.044)	(0.047)	(0.053)
$\text{MenDrafted}_{oc} \times \textit{Amenity}_o$	0.287	0.248	0.090	0.054	0.086
	(0.066)	(0.115)	(0.046)	(0.032)	(0.046)
Total Effect (High Amenity)	0.474	0.452	0.280	0.271	0.310
	(0.093)	(0.153)	(0.085)	(0.074)	(0.096)
First Stage F-statistic (Main)	15.32	49.74	39.06	53.89	45.84
First Stage F-statistic (Interaction)	36.74	33.86	32.06	32.07	32.16
Sanderson and Windmeijer (2016) F-statistic (Main)	23.54	11.85	10.14	49.39	12.54
Sanderson and Windmeijer (2016) F-statistic (Interaction)	59.66	13.61	15.54	30.10	16.16
Anderson and Rubin (1949) test (p-value)	0.000	0.000	0.000	0.000	0.000
Outcome Mean	4.178	3.940	3.951	3.970	3.982
Observations	344,678	368,291	367,038	365,264	364,011
Baseline Controls	Y	Y	Y	Y	Y
County FE	Y	Y	Y	Y	Y
Occ FE	Y	Y	Y	Y	Y

NOTE: This table reports 2SLS estimates interacting the number of men drafted with occupation-level amenity measures, with the number of women aged 18 to 22 employed in the occupation-county cell in 1950 as outcome. In columns 1 and 2, $\textit{Amenity}_o$ is an indicator: below-median physical task content in column 1 and flexible hours in column 2, defined as mean weekly hours of 40 or fewer in the 1940 complete count census. In columns 3 through 5, $\textit{Amenity}_o$ is a continuous score standardized to mean zero and standard deviation one: a schedule flexibility index in column 3, a cleanliness index in column 4, and a composite index averaging the two in column 5, constructed using Talkie-1930 LM (Levine et al., 2026). The number of men drafted and its interaction with the amenity measure are instrumented with the SSIV and its interaction with the same measure. The main effect of the amenity measure is absorbed by the occupation fixed effects. In columns 1 and 2 the main effect gives the effect of the number of men drafted in occupations without the amenity and the row "Total Effect" reports the effect in occupations with the amenity. In columns 3 through 5 the main effect gives the effect at the mean of the amenity score and the row "Total Effect" reports the effect at one standard deviation above the mean. Standard errors of the total effects are in parentheses beneath them. First stage F-statistics are the cluster-robust joint tests of the excluded instruments in the first stage of each endogenous variable, and Sanderson and Windmeijer (2016) F-statistics are the conditional first stage F-statistics for each endogenous variable, which account for the presence of multiple endogenous regressors. I also report the weak-instrument-robust Anderson and Rubin (1949) test of the joint significance of the endogenous regressors. The unit of observation is an occupation-county cell. Samples differ across columns because the amenity measures are available for different sets of occupations. All specifications include baseline occupation-county level controls, county fixed effects and occupation fixed effects. Baseline 1940 controls include the following demographic and industry controls at the occupation-county level: number of women, cell population, number of married men, number of workers in the agricultural, mining, manufacturing, and public administration industries. Robust standard errors clustered at the county level are reported in parentheses.

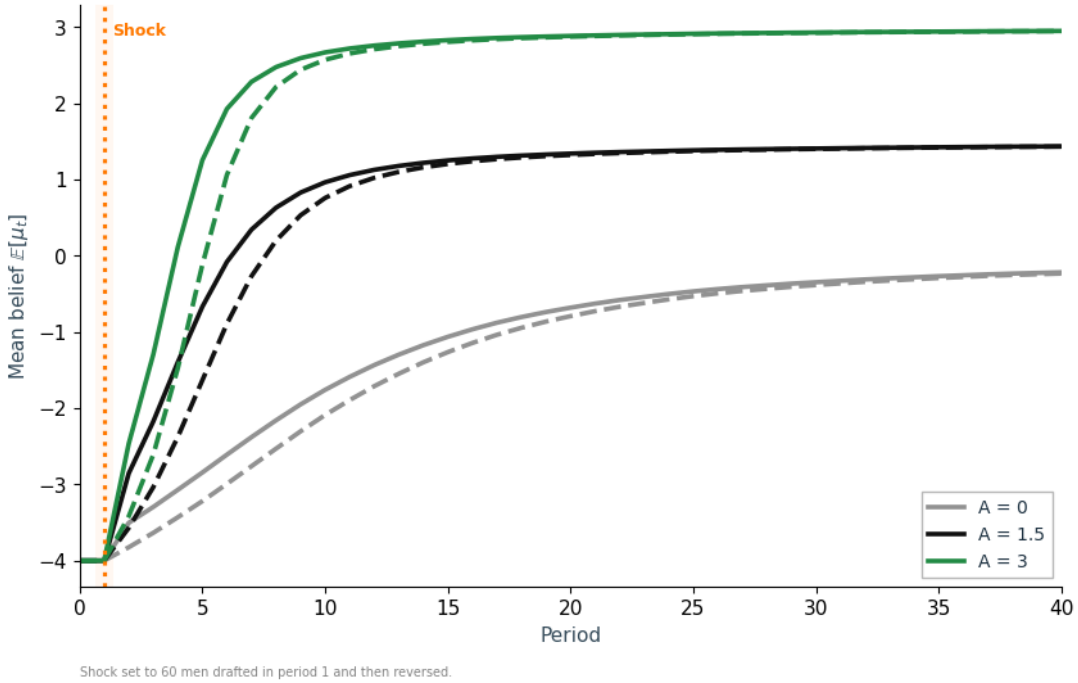
C.6 Information Diffusion

Figure C.7: Main Result with 5% Random Samples of 1950 Census



C.7 Simulated Belief Trajectories

Figure C.8: Simulated Belief Trajectories Corresponding to 11 Panel A



D Model Derivations

This appendix derives the labor-demand schedule, occupational entry probabilities, retention rate, and belief-updating equations used in the model. It then states the equilibrium conditions and derives the local relationship between excess labor demand, wages, and female entry. Throughout, the local labor-market index is suppressed.

D.1 Labor demand

Output is measured in units of the numeraire, and the occupation's production function is $Y_{ot} = \theta_{ot} L_{ot}^{\alpha}$, with $\theta_{ot} > 0$ and $\alpha \in (0,1)$. Here Y_{ot} is total occupational output, rather than output per worker. Men and women are perfect substitutes. A competitive producer takes the wage as given and solves

$$\max_{L \geq 0} \left\{ \theta_{ot} L^{\alpha} - w_{ot} L \right\}. \tag{1}$$

For $L > 0$, the first and second derivatives of profit are

$$\alpha\theta_{ot}L^{\alpha-1} - w_{ot}, \qquad \alpha(\alpha-1)\theta_{ot}L^{\alpha-2} < 0.$$

The marginal product tends to infinity as L approaches zero and to zero as L approaches infinity. Hence the unique optimum is interior and satisfies

$$w_{ot} = \alpha\theta_{ot}L_{ot}^{\alpha-1}. \tag{2}$$

Solving for employment yields the labor-demand schedule

$$L_{ot}^{d}(w_{ot};\theta_{ot}) = \left(\frac{\alpha\theta_{ot}}{w_{ot}}\right)^{1/(1-\alpha)}. \tag{3}$$

Taking logarithms and differentiating gives

$$\ln L_{ot}^{d} = \frac{\ln\alpha + \ln\theta_{ot} - \ln w_{ot}}{1-\alpha}, \tag{4}$$

$$d\ln L_{ot}^{d} = \eta_D\left(d\ln\theta_{ot} - d\ln w_{ot}\right), \qquad \eta_D \equiv \frac{1}{1-\alpha} > 0. \tag{5}$$

Thus η_D is the absolute wage elasticity of labor demand and also its elasticity with respect to productivity. Equivalently, the elasticity of inverse demand with respect to total employment is $-(1-\alpha)$.

D.2 Occupational entry

Expected utility. Using the maintained subjective beliefs and the individual's observed taste shock, the expected utility of entering occupation o is

$$\mathbb{E}[u_{iot} \mid \mathcal{I}_t, w_{ot}, \epsilon_{iot}] = \beta_w \ln w_{ot} + \beta_A\mu_{ot} + \epsilon_{iot} = V_{ot} + \epsilon_{iot}, \tag{6}$$

$$V_{ot} \equiv \beta_w \ln w_{ot} + \beta_A\mu_{ot}, \qquad V_{0t} \equiv b_0. \tag{7}$$

This expression uses the zero conditional mean of the unobserved match component. Because utility is linear in amenities and entry is based on current utility, amenity uncertainty enters the choice through its mean. Belief precision does not enter current utility separately.

Choice probabilities. Suppress the time subscript. The distribution and density of a taste shock are

$$G(e) = \exp\left\{-\exp\left(-\frac{e}{\sigma_\epsilon}\right)\right\}, \qquad g(e) = \frac{1}{\sigma_\epsilon}\exp\left(-\frac{e}{\sigma_\epsilon}\right) G(e). \tag{8}$$

Woman i chooses o if $V_o + \epsilon_{io} \geq V_k + \epsilon_{ik}$ for every $k \neq o$, including home production. Ties have probability zero. Conditional on $\epsilon_{io} = e$, independence implies

$$\Pr(o \text{ chosen} \mid \epsilon_{io} = e) = \prod_{\substack{k=0 \\ k\neq o}}^{O} G(V_o - V_k + e).$$

Integrating over the woman's own taste shock therefore gives

$$\begin{aligned} p_o &= \int_{-\infty}^{\infty} g(e) \prod_{\substack{k=0 \\ k\neq o}}^{O} G(V_o - V_k + e)\, de \\ &= \int_{-\infty}^{\infty} \frac{e^{-e/\sigma_\epsilon}}{\sigma_\epsilon} \exp\left\{-e^{-e/\sigma_\epsilon} \sum_{k=0}^{O} e^{(V_k - V_o)/\sigma_\epsilon}\right\} de. \end{aligned} \tag{9}$$

In the second line, the $k = o$ term incorporates the exponential term in $g(e)$. Use the change of variable $z = e^{-e/\sigma_\epsilon}$, for which $dz = -\sigma_\epsilon^{-1} e^{-e/\sigma_\epsilon} de$, and reverse the integration limits. It follows that

$$\begin{aligned} p_o &= \int_0^{\infty} \exp\left\{-z \sum_{k=0}^{O} e^{(V_k - V_o)/\sigma_\epsilon}\right\} dz \\ &= \left[\sum_{k=0}^{O} e^{(V_k - V_o)/\sigma_\epsilon}\right]^{-1} = \frac{e^{V_o/\sigma_\epsilon}}{\sum_{k=0}^{O} e^{V_k/\sigma_\epsilon}}. \end{aligned} \tag{10}$$

For a continuum of women, the fraction choosing each occupation equals its choice probability. Restoring time subscripts and using the unit cohort mass,

$$E_{ot} = p_{ot} = \frac{\exp[(\beta_w \ln w_{ot} + \beta_A \mu_{ot})/\sigma_\epsilon]}{\exp(b_0/\sigma_\epsilon) + \sum_{k=1}^{O} \exp[(\beta_w \ln w_{kt} + \beta_A \mu_{kt})/\sigma_\epsilon]}. \tag{11}$$

For a cohort of mass N, the same derivation gives $E_{ot} = N p_{ot}$. All choice probabilities are unchanged by adding the same constant to every alternative's systematic utility. Only utility differences relative to the scale σ_ϵ determine choices.

Entry elasticities and wage–amenity trade-offs. Let $D_t = \sum_{k=0}^{O} \exp(V_{kt}/\sigma_\epsilon)$. Differentiating $\ln p_{ot} = V_{ot}/\sigma_\epsilon - \ln D_t$ gives

$$\frac{\partial \ln p_{ot}}{\partial V_{kt}} = \frac{\mathbf{1}\{k = o\} - p_{kt}}{\sigma_\epsilon}. \tag{12}$$

Consequently, holding the other occupations' wages and beliefs fixed,

$$\varepsilon^S_{ot} \equiv \frac{\partial \ln E_{ot}}{\partial \ln w_{ot}} = \frac{\beta_w}{\sigma_\epsilon}(1 - p_{ot}), \qquad \frac{\partial \ln E_{ot}}{\partial \mu_{ot}} = \frac{\beta_A}{\sigma_\epsilon}(1 - p_{ot}). \tag{13}$$

The second expression is a semi-elasticity. For a different work occupation $k \neq o$, the cross-wage elasticity is $-\beta_w p_{kt}/\sigma_\epsilon$, and the cross-belief semi-elasticity is $-\beta_A p_{kt}/\sigma_\epsilon$.

The wage change that compensates for a change in the perceived amenity follows directly from $dV_{ot} = \beta_w \, d\ln w_{ot} + \beta_A \, d\mu_{ot} = 0$:

$$\left.\frac{d \ln w_{ot}}{d\mu_{ot}}\right|_{V_{ot}} = -\frac{\beta_A}{\beta_w}. \tag{14}$$

The ratio β_A/β_w therefore measures the compensating change in log wages per unit of perceived amenity; its interpretation depends on the units in which amenities are measured.

Finally, with male employment, returning stocks, and beliefs fixed, total labor supply is $L^s_{ot} = M_{ot} + R_{ot} + E_{ot}$. Its own-wage elasticity is

$$\eta^S_{ot} \equiv \frac{\partial \ln L^s_{ot}}{\partial \ln w_{ot}} = \frac{E_{ot}}{L^s_{ot}} \varepsilon^S_{ot}. \tag{15}$$

Only new entrants adjust at this stage, so the relevant employment share is E_{ot}/L^s_{ot}, rather than F_{ot}/L^s_{ot}.

D.3 Retention

The model imposes a reservation amenity $\bar{a}$ that is common across women, occupations, and dates. A woman who entered occupation o in period t remains for period $t+1$ if and only if

$$A_o + \nu_{io} \geq \bar{a}. \tag{16}$$

The cutoff is held fixed when wages change. This is a maintained retention rule, distinct from the young woman's entry problem.[1]

[1] If an older woman instead compared the utility with home production using contemporaneous wages and taste shocks, the cutoff would be $(b_0 - \beta_w \ln w_{o,t+1} + \epsilon_{i0,t+1} - \epsilon_{io,t+1})/\beta_A$. That cutoff generally depends on wages and individual tastes. The fixed rule in (16) abstracts from these margins.

Because the match component is unobserved at entry and independent of entry tastes, its distribution among entrants remains $\nu_{io} \sim \mathcal{N}(0, \sigma_\nu^2)$. Standardizing this normal random variable gives

$$\begin{aligned} r(A_o) &\equiv \Pr(A_o + \nu_{io} \geq \bar{a} \mid i \text{ entered } o) \\ &= \Pr\left(\frac{\nu_{io}}{\sigma_\nu} \geq \frac{\bar{a} - A_o}{\sigma_\nu}\right) \\ &= 1 - \Phi\left(\frac{\bar{a} - A_o}{\sigma_\nu}\right) = \Phi\left(\frac{A_o - \bar{a}}{\sigma_\nu}\right), \end{aligned} \tag{17}$$

where Φ is the standard normal distribution function. Its derivative is

$$r'(A_o) = \frac{1}{\sigma_\nu}\phi\left(\frac{A_o - \bar{a}}{\sigma_\nu}\right) > 0, \tag{18}$$

with ϕ the standard normal density. Aggregating across the entrant cohort yields

$$R_{o,t+1} = E_{ot}\, r(A_o). \tag{19}$$

Workers who are already old in period t exit at its end. Thus R_{ot} is not carried forward a second time in (19).

D.4 Reports and signal precision

To derive the signal technology, first consider an integer number $n \geq 1$ of reports. Each reporter conveys

$$x_{jot} = A_o + e_{jot}, \qquad e_{jot} \sim \mathcal{N}(0, \sigma_\xi^2), \qquad j = 1, \ldots, n, \tag{20}$$

where $\sigma_\xi^2 > 0$. Report errors are independent across reporters and dates, and independent of match components, entry tastes, and previous report errors. In particular, report errors remain mean zero conditional on the reporter satisfying the retention criterion. Reports concern A_o, not $A_o + \nu_{io}$. [2]

The average report is

$$s_{ot} = \frac{1}{n}\sum_{j=1}^{n} x_{jot} = A_o + \frac{1}{n}\sum_{j=1}^{n} e_{jot}. \tag{21}$$

[2] A report of the stayer's own match amenity would generally be selected. With $z_o = (A_o - \bar{a})/\sigma_\nu$, $\mathbb{E}[A_o + \nu_{io} \mid A_o + \nu_{io} \geq \bar{a}] = A_o + \sigma_\nu \phi(z_o)/\Phi(z_o) > A_o$. The reporting assumption removes this selection in the content of the report; selection still determines how many women report.

Conditional on A_o and n, independence and normality imply

$$\mathbb{E}[s_{ot} \mid A_o, n] = A_o, \qquad \mathrm{Var}(s_{ot} \mid A_o, n) = \frac{1}{n^2}\sum_{j=1}^{n}\sigma_\xi^2 = \frac{\sigma_\xi^2}{n}. \tag{22}$$

Thus $s_{ot} \mid A_o, n \sim \mathcal{N}(A_o, \sigma_\xi^2/n)$. Conditional on a given report count, the average also contains all the information about A_o in the individual reports: the identity

$$\sum_{j=1}^{n}(x_{jot} - A_o)^2 = \sum_{j=1}^{n}(x_{jot} - s_{ot})^2 + n(s_{ot} - A_o)^2$$

shows that the remaining term in the Gaussian likelihood is independent of A_o.

With continuous employment masses, the model adopts the corresponding *effective-report technology*, with reporting mass equal to $R_{o,t}$:

$$s_{ot} = A_o + \xi_{ot}, \qquad \xi_{ot} \mid A_o, R_{o,t} \sim \mathcal{N}\left(0, \frac{\sigma_\xi^2}{R_{o,t}}\right), \qquad R_{o,t} > 0. \tag{23}$$

The continuous-mass specification is a maintained aggregate noise technology; it does not literally average independent reports over every member of an atomless population. Rescaling employment units requires rescaling σ_ξ^2 proportionally to preserve signal variance. Signal precision is consequently

$$\pi_{ot} \equiv [\mathrm{Var}(\xi_{ot} \mid A_o, R_{o,t})]^{-1} = \frac{R_{o,t}}{\sigma_\xi^2} = \frac{E_{ot}r(A_o)}{\sigma_\xi^2}. \tag{24}$$

When $R_{o,t+1} = 0$, no report arrives and $\pi_{ot} = 0$ by convention.

D.5 Learning and belief updating

Information-processing convention. Each cohort applies the Gaussian updating rule to the report, using its supplied precision as a weight. The model does not incorporate a separate inference from how report volume, precision, retention, or equilibrium wages were generated. The derivation below is Bayesian updating for the maintained Gaussian prior and report likelihood, conditional on the supplied precision. It does not characterize full Bayesian inference from all endogenous observables.[3]

[3]This restriction matters because $\pi_{ot} = R_{o,t}/\sigma_\xi^2$ itself reveals the reporting stock when σ_ξ^2 is known. If entrants also observe $E_{ot} > 0$ and know the retention technology, they could infer $A_o = \bar{a} + \sigma_\nu\Phi^{-1}(R_{o,t}/E_{ot})$ in the continuum model. Consequently, merely withholding direct observations of $R_{o,t+1}$ would not justify ignoring this information while supplying π_{ot}. The baseline explicitly restricts inference to report content under the stated processing rule.

Posterior density. Suppress subscripts. The prior is $A \sim \mathcal{N}(\mu, \tau^{-1})$. With supplied signal precision $\pi > 0$, the likelihood used by the cohort is

$$\ell(s \mid A; \pi) \propto \exp\Big[-\frac{\pi}{2}(s - A)^2\Big].$$

Multiplying this likelihood by the prior density, and omitting factors that do not depend on A, gives

$$f^+(A; s, \pi) \propto \exp\left\{-\frac{1}{2}\left[\tau(A - \mu)^2 + \pi(s - A)^2\right]\right\}. \tag{25}$$

Expanding the quadratic terms,

$$\tau(A - \mu)^2 + \pi(s - A)^2 = (\tau + \pi)A^2 - 2(\tau\mu + \pi s)A + \tau\mu^2 + \pi s^2.$$

Completing the square gives

$$\tau(A - \mu)^2 + \pi(s - A)^2 = (\tau + \pi)\left(A - \frac{\tau\mu + \pi s}{\tau + \pi}\right)^2 + C, \tag{26}$$

where $C = \tau\mu^2 + \pi s^2 - (\tau\mu + \pi s)^2/(\tau + \pi)$ is independent of A. The posterior kernel is therefore Gaussian, with

$$\tau^+ = \tau + \pi, \qquad \mu^+ = \frac{\tau\mu + \pi s}{\tau + \pi}, \qquad \mathrm{Var}^+(A) = \frac{1}{\tau + \pi}. \tag{27}$$

Updating across cohorts. Restoring subscripts and defining the signal weight gives

$$\lambda_{ot} \equiv \frac{\pi_{ot}}{\tau_{ot} + \pi_{ot}} = \frac{R_{o,t}}{\sigma_\xi^2 \tau_{ot} + R_{o,t}}, \tag{28}$$

$$\mu_{o,t+1} = (1 - \lambda_{ot})\mu_{ot} + \lambda_{ot} s_{ot}, \tag{29}$$

$$\tau_{o,t+1} = \tau_{ot} + \frac{R_{o,t}}{\sigma_\xi^2}. \tag{30}$$

At fixed prior precision and report-noise variance,

$$\frac{\partial \lambda_{ot}}{\partial R_{o,t}} = \frac{\sigma_\xi^2 \tau_{ot}}{(\sigma_\xi^2 \tau_{ot} + R_{o,t})^2} > 0. \tag{31}$$

When no signal arrives, $\lambda_{ot} = 0$, and both belief parameters remain unchanged. Repeated application of the rule adds the precision of each new report; previously incorporated reports are not counted again.

Expected updates and posterior uncertainty. For a fixed true amenity and a given pre-report state, substitute $s_{ot} = A_o + \xi_{ot}$ into (29):

$$\mu_{o,t+1} - \mu_{ot} = \lambda_{ot}(A_o - \mu_{ot}) + \lambda_{ot}\xi_{ot}. \tag{32}$$

For $\pi_{ot} > 0$, taking moments over report noise yields

$$\mathbb{E}[\mu_{o,t+1} - \mu_{ot} \mid A_o, \mu_{ot}, \tau_{ot}, \pi_{ot}] = \lambda_{ot}(A_o - \mu_{ot}), \tag{33}$$

$$\mathrm{Var}(\mu_{o,t+1} \mid A_o, \mu_{ot}, \tau_{ot}, \pi_{ot}) = \frac{\lambda_{ot}^2}{\pi_{ot}}. \tag{34}$$

The variance in (34) describes variation in updated belief means across report realizations. It differs from the subjective posterior variance $(\tau_{ot} + \pi_{ot})^{-1}$ within a given realization. Neither formula replaces the realized posterior mean by its expectation in the nonlinear entry equation.

D.6 Market clearing and the equilibrium system

At the entry stage in period t, take the vectors of beliefs, male employment, and returning female employment as given. Wages clear every occupational market:

$$\left(\frac{\alpha\theta_{ot}}{w_{ot}}\right)^{1/(1-\alpha)} = M_{ot} + R_{ot} + p_{ot}(\mathbf{w}_t, \boldsymbol{\mu}_t), \qquad o = 1, \ldots, O, \tag{35}$$

where p_{ot} is given by (11). Equivalently,

$$w_{ot} = \alpha\theta_{ot}(M_{ot} + R_{ot} + E_{ot})^{\alpha-1}. \tag{36}$$

Entry and wages must be determined jointly: wages affect occupational choice, and entry affects the marginal product of labor. With multiple occupations, their wages are also linked through the common denominator of the choice probabilities.

Given initial belief means, precisions, and returning stocks, an equilibrium under the maintained information-processing rule satisfies (11) and (35) in each period, retention according to (19), signals according to (23), and belief transitions according to (29)–(30). The endogenous state linking successive entry stages is

$$\{\mu_{ot}, \tau_{ot}, R_{ot}\}_{o=1}^{O}.$$

For the three-date version of the model, these equations apply to the transitions that determine outcomes through $t = 2$. They do not require an assumption that beliefs have reached a steady state.

D.7 Excess labor demand and the local wage and entry response

This section derives the first-order calculation used to summarize the initial effect of a shock. The calculation holds the beliefs and returning stocks at the entry stage fixed. To use the period-0 allocation as the reference for period 1, take period 0 to be an initialized snapshot and let no additional report arrive before period-1 entry, so $\mu_{o1} = \mu_{o0}$ and $\tau_{o1} = \tau_{o0}$. Initialize the returning stock at $R_{o0} = E_{o0} r(A_o)$, which also gives $R_{o1} = R_{o0}$. Employment is stationary conditional on the initial beliefs; beliefs need not be stationary under ongoing learning. The report recursion starts with cohort 1. If the no-shock economy would already change between dates 0 and 1, the appropriate reference is instead its period-1 equilibrium, with both economies sharing the same inherited state.

The gap at the reference wage. At a period-0 equilibrium, $L_{o0} = M_{o0} + F_{o0} = L_o^d(w_{o0}; \theta_{o0})$. Evaluate labor demand under the shocked primitives at the reference wage and subtract reference female employment plus shocked male employment:

$$X_o \equiv L_o^d(w_{o0}; \theta_{o1}) - [M_{o1} + F_{o0}]. \tag{37}$$

Let $d_o = M_{o0} - M_{o1}$ and $g_o = \theta_{o1}/\theta_{o0} - 1$, with $1 + g_o > 0$. The demand schedule and reference market clearing imply the exact identity

$$\begin{aligned} X_o &= L_{o0}(1 + g_o)^{\eta_D} - M_{o1} - F_{o0} \\ &= d_o + L_{o0}\left[(1 + g_o)^{\eta_D} - 1\right]. \end{aligned} \tag{38}$$

For small shocks, $(1 + g_o)^{\eta_D} = 1 + \eta_D g_o + o(|g_o|)$, so

$$X_o \simeq d_o + \eta_D L_{o0} g_o. \tag{39}$$

A pure male-removal shock has $X_o = d_o$ exactly. For a productivity shock, X_o measures the additional labor demanded at the reference wage.

One occupation with other alternatives fixed. Hold other occupations' wages and beliefs fixed, and evaluate all elasticities at the reference allocation. To first order, a log-wage adjustment $\Delta \ln w_o = \ln w_{o1} - \ln w_{o0}$ reduces the initial labor-demand gap by $\eta_D L_{o0} \Delta \ln w_o$ and raises female entry by $E_{o0}\varepsilon^S_{o0}\Delta \ln w_o$. Market clearing therefore requires

$$X_o \simeq \left(\eta_D L_{o0} + E_{o0}\varepsilon^S_{o0}\right) \Delta \ln w_o. \tag{40}$$

Using (15), the resulting approximations are

$$\Delta \ln w_o \simeq \frac{X_o}{L_{o0}(\eta_D + \eta^S_{o0})}, \qquad \Delta E_o \simeq \frac{\eta^S_{o0}}{\eta_D + \eta^S_{o0}} X_o. \tag{41}$$

Since the returning stock is fixed at this stage, $\Delta F_o = \Delta E_o$. These are local approximations with coefficients evaluated at the reference equilibrium; they do not make X_o a global sufficient statistic for arbitrarily large combinations of shocks.

Joint adjustment across occupations. When all occupational wages adjust, the cross-wage derivatives must also enter the calculation. Let $\mathbf{p} = (p_{10}, \ldots, p_{O0})^\mathsf{T}$ contain the reference probabilities for work occupations, and let $\boldsymbol{\ell} = (\ln w_1, \ldots, \ln w_O)^\mathsf{T}$. Define

$$B \equiv \frac{\beta_w}{\sigma_\epsilon}\left[\operatorname{diag}(\mathbf{p}) - \mathbf{p}\mathbf{p}^\mathsf{T}\right], \qquad D \equiv \operatorname{diag}(\eta_D L_{10}, \ldots, \eta_D L_{O0}). \tag{42}$$

Here B is the Jacobian of entrant quantities with respect to log wages for a unit cohort. For cohort mass N, multiply B by N. At fixed beliefs, $d\mathbf{E} = B\, d\boldsymbol{\ell}$, while labor demand satisfies $d\mathbf{L}^d = D(d\ln \boldsymbol{\theta} - d\boldsymbol{\ell})$. Differentiating market clearing with returning stocks fixed gives

$$(D + B)d\boldsymbol{\ell} = D\, d\ln \boldsymbol{\theta} - d\mathbf{M} \equiv d\mathbf{X}. \tag{43}$$

The matrix D is positive definite, and $\operatorname{diag}(\mathbf{p}) - \mathbf{p}\mathbf{p}^\mathsf{T}$ is positive semidefinite: its quadratic form is the variance of a random variable taking value v_o with probability p_{o0} and value zero with the home probability. Thus $D + B$ is invertible. The first-order responses are

$$d\boldsymbol{\ell} = (D + B)^{-1} d\mathbf{X}, \qquad d\mathbf{E} = B(D + B)^{-1} d\mathbf{X}. \tag{44}$$

For small finite shocks, $\mathbf{X}$ from (37) can replace $d\mathbf{X}$ to first order. The scalar expressions in (41) apply when other alternatives are held fixed; (44) accounts for their equilibrium wage adjustment.

D.8 Within-period equilibrium: definition, existence, and uniqueness

Fix a period after its public reports have been processed and older workers have made their retention decisions. The belief means and returning stocks at the entry stage are therefore given. Suppress the time subscript, and let $O \geq 1$ be finite. For this subsection, allow the entering female cohort to have any finite mass $N > 0$; the normalization elsewhere in the appendix is $N = 1$. Take as given finite $M_o, R_o \geq 0$, $\mu_o \in \mathbb{R}$, $\theta_o \in (0, \infty)$ for every occupation, and $b_0 \in \mathbb{R}$. Maintain $\alpha \in (0, 1)$ and finite $\beta_w, \beta_A, \sigma_\epsilon > 0$. The belief precisions can be any finite positive values: conditional on the updated means, they do not enter the current entry problem separately.

Definition (Within-period equilibrium). A within-period equilibrium consists of a wage vector $\mathbf{w} \in (0, \infty)^O$, entrant quantities $(E_0, E_1, \ldots, E_O) \in [0, N]^{O+1}$, and employment quantities $\{F_o, L_o\}_{o=1}^O$ satisfying the following conditions.

First, almost every young woman chooses an alternative maximizing her expected current utility at the wage vector and the given beliefs. Aggregating these choices gives

$$E_o = N p_o(\mathbf{w}), \qquad p_o(\mathbf{w}) = \frac{\exp[(\beta_w \ln w_o + \beta_A \mu_o)/\sigma_\epsilon]}{\exp(b_0/\sigma_\epsilon) + \sum_{j=1}^O \exp[(\beta_w \ln w_j + \beta_A \mu_j)/\sigma_\epsilon]}, \quad o \geq 1, \tag{45}$$

and $E_0 = N p_0(\mathbf{w})$, where $p_0(\mathbf{w}) = 1 - \sum_{o=1}^O p_o(\mathbf{w})$ is the home probability. Thus $\sum_{o=0}^O E_o = N$. Second, producers maximize profits, so

$$L_o = L_o^d(w_o; \theta_o) = \left(\frac{\alpha \theta_o}{w_o}\right)^{1/(1-\alpha)}. \tag{46}$$

Third, employment satisfies accounting and every occupational labor market clears:

$$F_o = R_o + E_o, \qquad L_o = M_o + F_o, \qquad o = 1, \ldots, O. \tag{47}$$

The returning stocks are inherited from the retention rule. They are not additional choices in the current entry-stage equilibrium. Likewise, subsequent reports and belief updates determine the next period's state rather than an additional market-clearing condition in this period.

Lemma D.1 (Existence and uniqueness). Under the conditions above, a within-period equilibrium exists for every given vector of belief means and returning stocks. Its wage vector, entrant quantities, and employment quantities are unique. This conclusion holds with all O occupational wages determined simultaneously.

Proof: existence. Write $a_o = M_o + R_o$, $q = \beta_w/\sigma_\epsilon > 0$, $c_o = \exp(\beta_A \mu_o/\sigma_\epsilon) > 0$, and $c_0 = \exp(b_0/\sigma_\epsilon) > 0$. Then

$$p_o(\mathbf{w}) = \frac{c_o w_o^q}{c_0 + \sum_{j=1}^{O} c_j w_j^q}, \qquad Z_o(\mathbf{w}) \equiv L_o^d(w_o;\theta_o) - a_o - N p_o(\mathbf{w}). \tag{48}$$

An equilibrium wage vector is a simultaneous zero of $\mathbf{Z} = (Z_1, \ldots, Z_O)$. Each Z_o is continuously differentiable on $(0,\infty)^O$. Holding the other wages fixed,

$$\frac{\partial Z_o(\mathbf{w})}{\partial w_o} = -\frac{\eta_D}{w_o} L_o^d(w_o;\theta_o) - \frac{Nq}{w_o} p_o(\mathbf{w})[1 - p_o(\mathbf{w})] < 0, \qquad \eta_D = \frac{1}{1-\alpha}. \tag{49}$$

This own-wage monotonicity will provide a unique clearing wage conditional on the other wages. To establish simultaneous clearing, construct bounds that hold uniformly across those other wages.

Since $L_o^d(w;\theta_o) \to \infty$ as $w \downarrow 0$ and there are finitely many occupations, choose $\underline{w} > 0$ sufficiently small that

$$L_o^d(\underline{w};\theta_o) > a_o + N \qquad \text{for every } o. \tag{50}$$

Because $N p_o(\mathbf{w}) \le N$, excess demand is strictly positive whenever $w_o = \underline{w}$, regardless of the other wages.

For an upper bound, take $\overline{w} \ge 1$ with $\overline{w} > \underline{w}$. On a face where $w_o = \overline{w}$ and $w_j \le \overline{w}$ for all j, dividing the logit numerator and denominator by $\overline{w}^q$ gives

$$p_o(\mathbf{w}) = \frac{c_o}{c_0 \overline{w}^{-q} + \sum_{j=1}^{O} c_j (w_j/\overline{w})^q} \ge \frac{c_o}{c_0 + \sum_{j=1}^{O} c_j} \equiv \underline{p}_o > 0. \tag{51}$$

The lower bound $\underline{p}_o$ does not depend on $\overline{w}$. Since $L_o^d(w;\theta_o) \to 0$ as $w \to \infty$, enlarge $\overline{w}$ until

$$L_o^d(\overline{w};\theta_o) < a_o + N\underline{p}_o \qquad \text{for every } o. \tag{52}$$

The right-hand side is strictly positive even when $M_o = R_o = 0$. Hence excess demand is strictly negative on every upper face of the wage box $K = [\underline{w}, \overline{w}]^O$.

For each $\mathbf{w} \in K$, let $T_o(\mathbf{w})$ be the wage x solving $Z_o(x, \mathbf{w}_{-o}) = 0$, where $\mathbf{w}_{-o}$ denotes all wages except w_o. The face inequalities, continuity, and (49) imply that this solution exists, is unique, and lies in $(\underline{w}, \overline{w})$. The implicit-function theorem and the strict derivative in (49) imply that T_o is continuous in the other wages. Thus $T = (T_1, \ldots, T_O)$ is a continuous map from the nonempty compact convex set K into itself. Brouwer's fixed-point theorem gives a vector $\mathbf{w}^* = T(\mathbf{w}^*)$. By construction, $Z_o(\mathbf{w}^*) = 0$ for every

occupation. Defining entry and employment by (45)–(47) therefore produces a within-period equilibrium.

Proof: uniqueness. Suppose that two distinct positive wage vectors $\mathbf{w}$ and $\widetilde{\mathbf{w}}$ clear all markets at the same state. Interchanging their labels if necessary, select occupation k such that

$$\rho \equiv \frac{\widetilde{w}_k}{w_k} = \max_{1 \le o \le O} \frac{\widetilde{w}_o}{w_o} > 1. \tag{53}$$

Thus $\widetilde{w}_k > w_k$, while $\widetilde{w}_j / \widetilde{w}_k \le w_j / w_k$ for every $j \neq k$. Rewrite occupation k's entry probability as

$$p_k(\mathbf{w}) = \left[1 + \frac{c_0}{c_k w_k^q} + \sum_{j \neq k} \frac{c_j}{c_k} \left(\frac{w_j}{w_k} \right)^q \right]^{-1}. \tag{54}$$

Moving from $\mathbf{w}$ to $\widetilde{\mathbf{w}}$ weakly reduces every term in the sum and strictly reduces the home-option term $c_0/(c_k w_k^q)$. Consequently, $p_k(\widetilde{\mathbf{w}}) > p_k(\mathbf{w})$. Labor demand in k, however, strictly decreases because its wage increases. Market clearing at $\mathbf{w}$ would therefore imply

$$L_k^d(\widetilde{w}_k; \theta_k) < L_k^d(w_k; \theta_k) = a_k + N p_k(\mathbf{w}) < a_k + N p_k(\widetilde{\mathbf{w}}), \tag{55}$$

contradicting market clearing at $\widetilde{\mathbf{w}}$. The wage vector is therefore unique. Given that vector and the fixed state, all choice probabilities, entrant quantities, and employment quantities are uniquely determined. For given individual taste draws, entry choices are also unique except on the measure-zero set of draws that produce ties. ∎

The lemma establishes a unique allocation conditional on the state at a given date, with all occupational wage adjustments accounted for. Retention and learning determine how that state changes across dates. Uniqueness of this within-period allocation does not by itself establish a unique dynamic steady state or convergence to one.

E Model Predictions and Proofs

This appendix establishes four post-shock predictions under explicit sufficient conditions. The comparisons allow both the shocked and counterfactual economies to learn from noisy reports. Wages in the occupation adjust to clear its labor market in both periods. Throughout, the occupation index is suppressed.

E.1 Comparison, timing, and scope

Consider an occupation with true amenity A. The entering female cohort has unit mass, and its retention rate is

$$r(A) = \Phi\left(\frac{A - \bar{a}}{\sigma_\nu}\right), \qquad r'(A) = \frac{1}{\sigma_\nu}\phi\left(\frac{A - \bar{a}}{\sigma_\nu}\right) > 0. \tag{56}$$

All preference, production, and noise parameters satisfy the positivity conditions of the model, with $\alpha \in (0, 1)$. All amenity values and belief means considered below are finite.

Paths S and C share the same state immediately before period-1 entry: belief mean μ, belief precision $\tau > 0$, and returning stock R_1. Path S experiences a greater temporary reduction in male labor, an increase in productivity, or both, relative to C. The shock reverses by period 2, when both paths have the same male employment $M \geq 0$ and productivity $\theta > 0$. No information generated by period-1 entry is available before that entry decision. Let

$$e^j \equiv E_1^j, \qquad R^j \equiv R_2^j = e^j r(A), \qquad j \in \{S, C\}. \tag{57}$$

The signal s_1^j is generated by these stayers and is observed before period-2 entry. Its precision is R^j/σ_ξ^2. Thus s_1^j, rather than a signal generated by period-2 entrants, determines the period-2 posterior.

Conditional occupational comparison. The predictions hold the systematic utilities of other alternatives fixed in each period, across paths, and across the amenity comparisons. Their contribution to the logit denominator is a constant

$$B_t = \exp(b_0/\sigma_\epsilon) + \sum_{\substack{k=1 \\ k \neq o}}^{O} \exp(V_{kt}/\sigma_\epsilon) > 0.$$

Write $B = B_2$. This permits multiple alternative occupations, while abstracting from their wage and learning responses to the shock. The existence and uniqueness of the full multi-occupation equilibrium does not by itself establish these conditional comparative statics in the full joint equilibrium.

When differentiating with respect to A, hold the initial state, period-1 shock, period-2 fundamentals, and outside alternatives fixed. Consequently, e^S and e^C are held fixed: period-1 entrants choose on the basis of the common prior, before experiencing the amenity. In particular, the initial returning stock is not recalibrated as A varies. Returning stocks $R^j = e^j r(A)$ in period 2 do vary with A.

Information and expectations. Maintain the model's Gaussian report-processing rule: cohorts use the public report and its supplied precision, without separately inferring amenities from reporting volume,

retention, or wages. Experience reveals the common amenity, reports concern that common component, and retention follows the fixed cutoff in (56). Expectations below average over report noise conditional on the true A, the common initial state, and the specified shock. They do not average over the prior distribution of A. The two paths need not receive the same report realization or use any particular coupling of their report errors.

Finite logit utilities imply $e^C > 0$. Thus an "untested" occupation here means one with little prior female entry and information, rather than a literal zero entry probability. All restrictions below are sufficient, not necessary.

E.2 The period-1 entry response

At the common prior and returning stock, period-1 excess demand is

$$Z_1(w; M_1, \theta_1) = \left(\frac{\alpha\theta_1}{w}\right)^{1/(1-\alpha)} - M_1 - R_1 - E_1(w; \mu),$$

where $E_1(w; \mu)$ is the logit entry schedule with outside term B_1. Demand is strictly decreasing in w, whereas entry is strictly increasing. Hence Z_1 is strictly decreasing. A reduction in M_1 or increase in θ_1 raises Z_1 at every wage. Its unique zero moves to a higher wage, which increases female entry. Therefore, for a componentwise greater shock, with at least one strict change,

$$e^S > e^C > 0, \qquad R^S > R^C > 0. \tag{58}$$

To interpret derivatives with respect to excess demand, let x index a differentiable admissible shock path, with $M_1(x) \geq 0$, $\theta_1(x) > 0$, $M_1'(x) \leq 0$, and $\theta_1'(x) \geq 0$. Assume at least one derivative inequality is strict at the point considered. Define the gap at the fixed counterfactual wage by

$$X(x) = L^d(w_1^C; \theta_1(x)) - [M_1(x) + F_1^C]. \tag{59}$$

Both $X'(x) > 0$ and $e'(x) > 0$, so $de/dX = e'(x)/X'(x) > 0$ along this path. For example, differentiating period-1 clearing, with $q_1 = E_1(1 - E_1)$ and $\eta_D = (1 - \alpha)^{-1}$, gives

$$\frac{d \ln w_1}{dx} = \frac{\eta_D L_1 \theta_1'(x)/\theta_1 - M_1'(x)}{\eta_D L_1 + (\beta_w/\sigma_\epsilon) q_1} > 0, \qquad \frac{de}{dx} = \frac{\beta_w}{\sigma_\epsilon} q_1 \frac{d \ln w_1}{dx} > 0. \tag{60}$$

For a pure male-removal shock, X equals the number of men removed relative to C. For mixed shocks,

X provides a local index along the specified path; it does not globally order arbitrary combinations of productivity and male-supply changes. If the unshocked period-1 allocation equals the initialized period-0 allocation, (59) coincides with the main text's definition using w_0 and F_0.

E.3 Equilibrium responses and noisy learning

The period-2 equilibrium map. For a posterior mean m and a returning stock R, let $F = G(m, R)$ denote total female employment at the occupation's unique clearing wage. Define $E = G(m, R) - R$ and $L = M + G(m, R)$. The equilibrium equations are

$$w = \alpha\theta L^{\alpha-1}, \qquad E = \frac{\exp[(\beta_w \ln w + \beta_A m)/\sigma_\epsilon]}{B + \exp[(\beta_w \ln w + \beta_A m)/\sigma_\epsilon]}. \tag{61}$$

Introduce

$$b = \frac{\beta_A}{\sigma_\epsilon}, \qquad c = \frac{\beta_w(1-\alpha)}{\sigma_\epsilon}, \qquad p = E, \qquad q = p(1-p), \qquad z = \frac{q}{L}, \qquad Q = 1 + cz. \tag{62}$$

Here $p \in (0, 1)$ is the entrant fraction of the new female cohort. Since $L = M + R + p \geq p$, $0 < z \leq 1 - p \leq 1$.

Differentiating entry and inverse demand gives

$$dE = q\left(\frac{\beta_w}{\sigma_\epsilon} d\ln w + b\, dm\right), \qquad d\ln w = -\frac{1-\alpha}{L}(dE + dR).$$

Substitution and collection of terms imply $Q\, dE = bq\, dm - cz\, dR$. Adding dR yields

$$G_m = \frac{bq}{Q} > 0, \qquad G_R = \frac{1}{Q} > 0, \qquad G_R - 1 = -\frac{c}{b}\frac{G_m}{L}. \tag{63}$$

At a fixed belief, extra stayers increase total female employment but crowd out some entry by young women.

Curvature in beliefs. At fixed R, $dq/dE = 1 - 2p$ and $dL/dE = 1$. Differentiating G_m therefore gives

$$\frac{G_{mm}}{G_m} = b\frac{1 - 2p + cz^2}{(1+cz)^2}. \tag{64}$$

Since

$$-(1+cz)^2 \leq -1 \leq 1 - 2p + cz^2 \leq 1 + cz^2 \leq (1+cz)^2,$$

the following bounds hold at every realized posterior:

$$|G_{mm}| \le bG_m, \qquad 0 < G_m \le \frac{b}{4}, \qquad \frac{1}{1+c} \le G_R \le 1. \tag{65}$$

These inequalities do not impose global convexity or concavity of employment in beliefs.

Distribution of updated beliefs. Let $\gamma = 1/\sigma_\xi^2 > 0$ and $\delta = A - \mu$. At reporting stock $R > 0$, signal precision is γR. Applying the Gaussian updating formula to $s_1 = A + Z/\sqrt{\gamma R}$ gives

$$m(R, A, Z) = \mu + \lambda(R)\delta + \sqrt{v(R)}Z, \qquad Z \sim \mathcal{N}(0,1), \tag{66}$$

$$\lambda(R) = \frac{\gamma R}{\tau + \gamma R}, \qquad v(R) = \frac{\gamma R}{(\tau + \gamma R)^2}. \tag{67}$$

The quantity $v(R)$ is the variance of posterior means across report realizations. The subjective posterior variance of the amenity is instead $(\tau + \gamma R)^{-1}$. Direct differentiation yields

$$\lambda_R = \frac{\gamma\tau}{(\tau + \gamma R)^2} > 0, \qquad v_R = \frac{\gamma(\tau - \gamma R)}{(\tau + \gamma R)^3}, \qquad \frac{|v_R|}{\lambda_R} \le \frac{1}{\tau}. \tag{68}$$

Expected outcomes. Define expected total female employment and expected young entry by

$$H(R, A) = \mathbb{E}_Z[G(m(R, A, Z), R)], \qquad Y(R, A) = H(R, A) - R. \tag{69}$$

Differentiating with respect to R and integrating the Gaussian term by parts,

$$\begin{aligned} H_R &= \mathbb{E}[G_R] + \lambda_R\delta\,\mathbb{E}[G_m] + \frac{v_R}{2\sqrt{v}}\mathbb{E}[ZG_m] \\ &= \mathbb{E}[G_R] + \lambda_R\delta\,\mathbb{E}[G_m] + \frac{v_R}{2}\mathbb{E}[G_{mm}], \end{aligned} \tag{70}$$

$$H_A = \lambda\mathbb{E}[G_m], \qquad Y_A = H_A. \tag{71}$$

The integration-by-parts identity is $\mathbb{E}[ZG_m] = \sqrt{v}\,\mathbb{E}[G_{mm}]$. On any compact positive-R interval, the derivative bounds in (65) and an integrable multiple of $1 + |Z|$ justify differentiation under the expectation. The same bounds justify integration by parts. All actual comparisons below have $R^C > 0$.

Combining (65) and (68) gives the useful lower bound

$$H_R \ge \mathbb{E}[G_R] + \lambda_R\left(\delta - \frac{b}{2\tau}\right)\mathbb{E}[G_m]. \tag{72}$$

The dispersion term in (70) is retained throughout. In particular, expected outcomes are not obtained by solving the equilibrium at the expected posterior mean.

E.4 Prediction 1: post-shock total female employment

Prediction 1 (Total female employment). Maintain the comparison in Section E.1. If the occupation is sufficiently underappreciated that

$$A - \mu \geq \frac{\beta_A}{2\sigma_\epsilon \tau} = \frac{b}{2\tau}, \tag{73}$$

then a greater temporary shock raises expected period-2 total female employment:

$$\mathbb{E}[F_2^S] - \mathbb{E}[F_2^C] > 0. \tag{74}$$

Along a differentiable shock path as in Section E.2,

$$\frac{d\mathbb{E}[F_2]}{dX} > 0. \tag{75}$$

Proof. The period-1 argument gives $R^S > R^C$. Under (73), the bracket in (72) is nonnegative, while $\mathbb{E}[G_R] > 0$. Hence $H_R(R, A) > 0$ for every $R > 0$. It follows that

$$\mathbb{E}[F_2^S] - \mathbb{E}[F_2^C] = H(R^S, A) - H(R^C, A) = \int_{R^C}^{R^S} H_R(R, A)\, dR > 0. \tag{76}$$

At fixed A, a marginal shock changes the returning stock by $dR/dX = r(A)\, de/dX$. Therefore

$$\frac{d\mathbb{E}[F_2]}{dX} = H_R\, r(A) \frac{de}{dX} > 0.$$

This establishes both claims. ∎

Condition (73) bounds any adverse effect from changes in the dispersion of posterior beliefs. It is not a claim that every report realization raises beliefs or employment.

E.5 Prediction 2: amenities and the total-employment effect

For the amenity comparisons, define

$$h(A) = \frac{r'(A)}{r(A)}, \qquad \lambda_j = \lambda(R^j), \qquad j \in \{S, C\}, \tag{77}$$

$$K_S = \mathbb{E}\left[\frac{p_2^S(1-p_2^S)}{1 + cp_2^S(1-p_2^S)/L_2^S}\right] = \frac{\mathbb{E}[G_m^S]}{b} > 0, \tag{78}$$

$$\Omega(A) = 1 + h(A)\left(\delta + \frac{b}{2\tau}\right). \tag{79}$$

The moment K_S measures the expected responsiveness of period-2 entry to beliefs on the shocked path, including equilibrium wage adjustment. It is determined by the model's primitives and the report distribution; it is not an additional preference parameter.

Let $D_F(A) = \mathbb{E}[F_2^S] - \mathbb{E}[F_2^C]$ denote the total-employment effect.

Prediction 2 (Amenity complementarity of total employment). At the amenity value under consideration, suppose that

$$\delta \geq \frac{b}{2\tau}, \tag{80}$$

$$e^S \geq (1+c)e^C, \tag{81}$$

$$\frac{\lambda_C}{\lambda_S} < \frac{4K_S}{\Omega(A)}. \tag{82}$$

Then $D_F(A) > 0$ and $D_F'(A) > 0$. If these conditions hold throughout an amenity interval, the total-employment effect is strictly increasing on that interval.

Proof. Since $R^j = e^j r(A)$ and e^j is fixed in this comparison, the chain rule gives

$$\frac{d\mathbb{E}[F_2^j]}{dA} = e^j r'(A) H_R(R^j, A) + \lambda_j \mathbb{E}[G_m^j]. \tag{83}$$

The first term includes the wage and information effects of changing retention; the second changes the mean report at a fixed reporting stock. Under (80), (72) and $G_R \geq (1+c)^{-1}$ imply

$$\frac{d\mathbb{E}[F_2^S]}{dA} \geq \frac{e^S r'(A)}{1+c} + b\lambda_S K_S. \tag{84}$$

To bound the counterfactual derivative, use (70) and the curvature bound to obtain

$$H_R \leq \mathbb{E}[G_R] + \left(\lambda_R \delta + \frac{b}{2}|v_R|\right)\mathbb{E}[G_m].$$

The identities

$$R\lambda_R = \lambda(1-\lambda) \leq \lambda, \qquad R|v_R| = \frac{\lambda(1-\lambda)|1-2\lambda|}{\tau} \leq \frac{\lambda}{\tau}, \qquad e^j r' = hR^j \tag{85}$$

together with $\delta \geq 0$, $G_R \leq 1$, and $G_m \leq b/4$ therefore give

$$\frac{d\mathbb{E}[F_2^C]}{dA} \leq e^C r'(A) + \frac{b}{4}\lambda_C \Omega(A). \tag{86}$$

Subtracting the two bounds yields

$$D_F'(A) \geq r'(A)\left(\frac{e^S}{1+c} - e^C\right) + b\left(\lambda_S K_S - \frac{\lambda_C \Omega(A)}{4}\right) > 0. \tag{87}$$

The first bracket is nonnegative by (81), and the second is strictly positive by (82). Prediction 1 supplies $D_F(A) > 0$. Integrating the positive derivative establishes the interval statement. ∎

The conditions require both an underappreciated occupation and a substantial information advantage relative to the counterfactual. Small initial female employment alone is insufficient: the counterfactual learning weight must be small relative to the shocked weight, and the new cohort must remain responsive to beliefs. No restriction requires every Gaussian report to produce an interior entry share bounded away from zero and one.

E.6 Prediction 3: entry by the next cohort

Let $D_E(A) = \mathbb{E}[E_2^S] - \mathbb{E}[E_2^C]$. Since returning stocks are nonrandom conditional on the period-1 comparison,

$$D_E(A) = D_F(A) - (R^S - R^C). \tag{88}$$

A positive total-employment effect is therefore not enough to establish positive new entry. Define the bound on wage crowding

$$C_E(A) = \frac{c}{b} \max_{R \in \{R^C, R^S\}} \frac{(\tau + \gamma R)^2}{\gamma\tau(M+R)}. \tag{89}$$

This bound is finite because $M + R^C > 0$.

Prediction 3 (Entry and its amenity complementarity). If

$$\delta > \frac{b}{2\tau} + C_E(A), \tag{90}$$

then $D_E(A) > 0$. If, in addition,

$$\frac{\lambda_C}{\lambda_S} < \frac{4K_S}{\Omega(A)}, \tag{91}$$

then $D'_E(A) > 0$. When both conditions hold throughout an amenity interval, the shock raises young women's entry and its effect on entry is strictly increasing in the true amenity throughout that interval.

Proof of the positive entry effect. Subtracting one from (70) and using (63) gives

$$Y_R = -\frac{c}{b}\mathbb{E}\left[\frac{G_m}{L}\right] + \lambda_R \delta\, \mathbb{E}[G_m] + \frac{v_R}{2}\mathbb{E}[G_{mm}]. \tag{92}$$

The first term is wage crowding; the remaining terms are the effects of the change in the distribution of beliefs. Since $L \geq M + R > 0$,

$$Y_R \geq \lambda_R \left[\delta - \frac{b}{2\tau} - \frac{c}{b} f(R)\right] \mathbb{E}[G_m], \qquad f(R) = \frac{(\tau + \gamma R)^2}{\gamma\tau(M+R)}. \tag{93}$$

Here $f(R) = [(M+R)\lambda_R]^{-1}$. Direct differentiation gives

$$f''(R) = \frac{2(\tau - \gamma M)^2}{\gamma\tau(M+R)^3} \geq 0. \tag{94}$$

Thus $f(R) \leq \max\{f(R^C), f(R^S)\}$ throughout $[R^C, R^S]$. Condition (90) makes $Y_R > 0$ on this whole interval. Consequently,

$$D_E(A) = \int_{R^C}^{R^S} Y_R(R, A)\, dR > 0. \tag{95}$$

Proof of amenity complementarity. The full amenity derivative of expected entry on path j is

$$\frac{d\mathbb{E}[E_2^j]}{dA} = e^j r'(A) Y_R(R^j, A) + \lambda_j \mathbb{E}[G_m^j]. \tag{96}$$

The level condition has established $Y_R(R^S, A) > 0$. Dropping the positive retention contribution on the shocked path therefore gives

$$\frac{d\mathbb{E}[E_2^S]}{dA} \geq b\lambda_S K_S. \tag{97}$$

For the counterfactual, drop the negative crowding term in (92) and bound the curvature term:

$$Y_R \leq \left(\lambda_R \delta + \frac{b}{2}|v_R|\right) \mathbb{E}[G_m].$$

Using (85) in (96), and then $G_m \leq b/4$, gives

$$\frac{d\mathbb{E}[E_2^C]}{dA} \leq \lambda_C \Omega(A) \mathbb{E}[G_m^C] \leq \frac{b}{4} \lambda_C \Omega(A). \tag{98}$$

It follows that

$$D_E'(A) \geq b\left(\lambda_S K_S - \frac{\lambda_C \Omega(A)}{4}\right) > 0 \tag{99}$$

under (91). This proves the stronger entry claim as well as the positive level effect. ■

The restriction on underappreciation is stronger than in Prediction 1 because learning must offset the wage loss caused by additional stayers. The entry-complementarity proof does not require the separate condition $e^S \geq (1+c)e^C$ used in Prediction 2.

E.7 Prediction 4: amenities and the log-wage effect

Define the expected log-wage effect

$$D_w(A) = \mathbb{E}[\ln w_2^S] - \mathbb{E}[\ln w_2^C], \tag{100}$$

and the bounds on the two workforces

$$\ell_C = M + R^C > 0, \qquad \ell_S = M + R^S + 1, \qquad \rho = \frac{\ell_S}{\ell_C} > 1. \tag{101}$$

The added one in ℓ_S is the maximum entry from the unit young cohort.

Prediction 4 (Amenity complementarity of log-wage declines). Suppose that

$$\delta \geq \frac{b}{2\tau}, \tag{102}$$

$$e^S \geq (1+c)\rho e^C, \tag{103}$$

$$\frac{\lambda_C}{\lambda_S} < \frac{4K_S}{\rho \Omega(A)}. \tag{104}$$

Then $D_w(A) < 0$ and $D_w'(A) < 0$. If the conditions hold throughout an amenity interval, the shock-induced

difference in expected log wages is strictly decreasing throughout that interval.

Proof: a direct log-wage comparison. Let $d = 1 - \alpha > 0$ and define

$$U(m, R) = \ln[M + G(m, R)], \qquad J(R, A) = \mathbb{E}_Z[U(m(R, A, Z), R)]. \tag{105}$$

Inverse demand implies

$$D_w(A) = -d \left[J(R^S, A) - J(R^C, A) \right]. \tag{106}$$

This identity compares expected log wages directly; an ordering of mean female employment alone would not suffice for the result.

The derivatives of the log workforce are

$$U_m = \frac{G_m}{L} > 0, \qquad U_R = \frac{G_R}{L} > 0, \qquad \frac{U_{mm}}{U_m} = \frac{G_{mm}}{G_m} - \frac{G_m}{L} = b\frac{1 - 2p - z}{Q^2}. \tag{107}$$

Because $0 \le z \le 1 - p$, the numerator satisfies $-1 \le -p \le 1 - 2p - z \le 1$. Since $Q \ge 1$,

$$|U_{mm}| \le bU_m. \tag{108}$$

The same Gaussian differentiation argument used for H yields

$$J_R = \mathbb{E}[U_R] + \lambda_R \delta \, \mathbb{E}[U_m] + \frac{v_R}{2}\mathbb{E}[U_{mm}], \tag{109}$$

$$J_A = \lambda \mathbb{E}[U_m]. \tag{110}$$

These derivatives are justified on $[R^C, R^S]$ because $L \ge M + R^C > 0$ bounds the log-workforce derivatives needed for dominated convergence and integration by parts. Under (102),

$$J_R \ge \mathbb{E}[U_R] + \lambda_R \left(\delta - \frac{b}{2\tau} \right) \mathbb{E}[U_m] > 0. \tag{111}$$

It follows immediately from $R^S > R^C$ and (106) that $D_w(A) < 0$.

Proof: amenity complementarity. Write $T_j(A) = J(e^j r(A), A)$. The full derivative is

$$T_j'(A) = e^j r'(A) J_R(R^j, A) + \lambda_j \mathbb{E}[U_m^j]. \tag{112}$$

On the shocked path, $L_2^S \le \ell_S$, $G_R^S \ge (1+c)^{-1}$, and $\mathbb{E}[G_m^S] = bK_S$. Using (111),

$$T_S'(A) \ge \frac{1}{\ell_S}\left[\frac{e^S r'(A)}{1+c} + b\lambda_S K_S\right]. \tag{113}$$

For the counterfactual, the upper curvature bound gives

$$J_R \le \mathbb{E}[U_R] + \left(\lambda_R \delta + \frac{b}{2}|v_R|\right)\mathbb{E}[U_m].$$

Apply (85), $L_2^C \ge \ell_C$, $G_R^C \le 1$, and $G_m^C \le b/4$ to obtain

$$T_C'(A) \le e^C r'(A)\mathbb{E}[U_R^C] + \lambda_C \Omega(A)\mathbb{E}[U_m^C] \le \frac{1}{\ell_C}\left[e^C r'(A) + \frac{b}{4}\lambda_C \Omega(A)\right]. \tag{114}$$

Subtracting these bounds and using $\rho = \ell_S/\ell_C$ yields

$$T_S' - T_C' \ge \frac{1}{\ell_S}\left\{r'(A)\left[\frac{e^S}{1+c} - \rho e^C\right] + b\left[\lambda_S K_S - \frac{\rho\lambda_C\Omega(A)}{4}\right]\right\} > 0. \tag{115}$$

The first bracket is nonnegative by (103); the second is strictly positive by (104). Therefore

$$D_w'(A) = -(1-\alpha)(T_S' - T_C') < 0.$$

Integrating establishes monotonicity over any interval on which the conditions hold. ∎

The workforce factor ρ strengthens the total-employment conditions because log wages depend on proportional changes in total labor. Its use is conservative: it compares the largest possible shocked workforce with the smallest possible counterfactual workforce. The result concerns $\mathbb{E}[\ln w]$, not $\ln \mathbb{E}[w]$. Under perfect substitution, women and men receive the same occupational wage, so it also describes the wage received by women in the occupation.

E.8 Interpreting and assessing the sufficient conditions

A joint sufficient region. A convenient set of conditions that establishes all four predictions, including both entry claims in Prediction 3, is

$$\delta > \frac{b}{2\tau} + C_E(A), \qquad e^S \ge (1+c)\rho e^C, \qquad \frac{\lambda_C}{\lambda_S} < \frac{4K_S}{\rho\Omega(A)}. \tag{116}$$

Because $\rho > 1$, these inequalities imply the sufficient conditions for Predictions 1–4. Each quantity is evaluated at the amenity under consideration. To compare an interval of amenity values using monotonicity, the relevant conditions must hold throughout that interval.

These restrictions capture three requirements. First, the common prior must understate the true amenity enough to permit a substantial learning gain. Second, the shock must induce enough exploration relative to the counterfactual. Third, the information advantage must translate into a responsive new cohort; entry that is nearly zero or nearly saturated makes the conservative responsiveness condition harder to satisfy. The additional term $C_E(A)$ requires the learning gain to dominate wage crowding of young entrants.

Computing the responsiveness term. The moment K_S can be evaluated by one-dimensional Gaussian integration after solving (61) at each posterior. Alternatively, if $0 < u < v < 1$ and $\Pr\{u \le p_2^S \le v\} \ge \omega > 0$, then

$$K_S \ge \frac{\omega \min\{u(1-u), v(1-v)\}}{1+c}. \tag{117}$$

Substituting this lower bound gives a stronger sufficient condition that can be checked using a probability of interior entry, rather than an assumption that every report realization yields interior entry bounded away from zero and one.

That probability can itself be calculated from the primitives. The posterior mean supporting an equilibrium entrant share $p \in (0,1)$ at returning stock R is

$$T(p,R) = \frac{1}{b}\ln\left(\frac{Bp}{1-p}\right) - \frac{\beta_w}{\beta_A}\ln\left[\alpha\theta(M+R+p)^{\alpha-1}\right]. \tag{118}$$

This follows by inverting logit odds and substituting the clearing wage. Since entry is strictly increasing in the posterior mean, write $\bar{m}_S = \mu + \lambda_S\delta$ and $v_S = v(R^S)$ to obtain

$$\Pr\{u \le p_2^S \le v\} = \Phi\left(\frac{T(v,R^S) - \bar{m}_S}{\sqrt{v_S}}\right) - \Phi\left(\frac{T(u,R^S) - \bar{m}_S}{\sqrt{v_S}}\right). \tag{119}$$

Thus the equilibrium moment in the conditions is fully determined by the specified economy. The sufficient conditions do not characterize all parameter values for which the predictions hold. Their failure alone does not establish that a prediction is false. The results concern the post-shock period; they make no assertion of permanence or monotonicity over later generations.