Essays on Marriage and Labor Markets

by

Ariel J. Binder

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in the University of Michigan
2020

Doctoral Committee:

Professor John Bound, Co-Chair
Assistant Professor Ana Reynoso, Co-Chair
Professor David Lam
Professor Pamela Smock

Ariel J. Binder

ajbinder@umich.edu

ORCID iD: 0000-0002-0726-4936

ACKNOWLEDGEMENTS

It takes a village to mint a new PhD. I would not have come half as far had it not been for the guidance and support of a rich network of colleagues, friends and family.

This dissertation reflects the dedicated advising of my committee: John Bound, Ana Reynoso, David Lam and Pamela Smock. David sparked my interest in the economics of marriage, and saw to it that that interest was the start of a rich education. I appreciated the opportunity to co-author the third chapter of this dissertation with him. John, a pre-eminent labor market scholar, was intrigued by my idea that the changing U.S. marriage market was affecting men's labor market prospects in an important manner. Without his encouragement and belief in the core of my idea—which he maintained even as he levied strong "suggestions" for improving its execution—the first chapter never would have gone as far as it did. I am thankful that the Economics Department hired Ana in time for my critical fifth year in the program: her work defines the cutting edge of the economics-of-marriage literature, and her expert advising was instrumental in transforming my first chapter from speculation into rigorous econometric evidence. Pam went above and beyond her duties as cognate adviser to expose me to how sociologists and family demographers think about marriage markets: a perspective which greatly informed my work and which I appreciate having as my career as a demographic economist begins. I owe each of them a debt of gratitude, and hope I continue to represent them well in my career ahead.

My professional development began during my second year of graduate school, when I took Mel Stephens' class on labor supply. This class sparked my interest in household bargaining models and female labor supply, which I began honing that following summer through further reading and conversations with Mel. David and Martha Bailey's fantastic class on economic demography, which I took during my third year, further fueled my interests. Over the ensuing years, I received invaluable feedback from an active labor group within the Economics Department. In addition to those already mentioned, I acknowledge Frank Stafford, Justin Wolfers, Betsey Stevenson, Kevin Stange, Mike Mueller-Smith, Melanie Wasserman and, especially, Charlie Brown for their attendance and participation during my seminar talks.

For dedicated administrative support and responses to my various questions I thank Laura Flak, Hiba Baghdadi, Lisa Neidert, Carol Bowen, Lauren Pulay, Julie Heintz and, especially, Miriam Rahl. I quite literally would not have completed my work and various degree requirements were it not for them.

In addition to granting access to these and other highly competent people, the Economics Department and Population Studies Center at the Institute for Social Research provided invaluable research support. I acknowledge the William Haber Graduate Fellowship Fund, the George Katona Economic Behavior Research Award Fund, Summer Research Apprenticeship funding and, especially, a National Institutes of Health and Child Development training grant to the Population Studies Center (T32 HD007339). These sources of funding released me from teaching duties throughout most of my graduate career, and thus allowed me the freedom to explore the topics that ultimately comprised my dissertation.

For camaraderie, advice on work- and non-work-related issues, blowing off steam over drinks at Alley Bar, and other acts of kindness, I thank my fellow graduate student classmates. These include Parag Mahajan, Max Risch, Dhiren Patki, Brenden Timpe, Teju Velayudhan, Meera Mahadevan, Max Gross, Matt Gross, George Fenton, Anirudh Jayanti, Avery Calkins, Pieter De Vlieger, Evan Taylor, Mike Zabek, Jessica Gillooly, Sonia Ruppcic, Jasmine Manalel, Tyler Hein, Amira Halawah, Brad Smith, Isaiah Sypher, and Lauren Brown. These individuals truly understand the effort and willpower required to complete a contemporary dissertation in the social sciences. Dhiren receives especial mention for being my tennis partner, and thus facilitating much stress-relieving yellow-ball-whacking throughout these 6 and a half years. So does Javed Basu-Kesselman, whose ping pong skills and key-card access to the Law School student lounge were greatly appreciated.

For trying their best to understand the challenges of a PhD and for unwavering support, I thank my enduring friends: Tom Kildahl, Jeff Putnam, Asad Zaman, Hubert Shang, Tony De Wolfe, Todd Schnuck, Alex Bain, Pat Jones, and Rafael Romeu. I also thank my new in-laws: Stacy Scott, Chuck Ciaccio and Gianna Ciaccio.

Perhaps the two qualities most important for completing a PhD are intellectual curiosity and self-motivation. I am fortunate to have grown up in a family with the resources and good sense to foster these qualities in me from a young age. And during times of intense self-doubt, my family did their best to remind me of their power. I thank my parents, Guyora Binder (who also read my dissertation and provided very helpful edits) and Judith Olin, and my sister, Galia Binder.

Finally, I thank my spouse, namesake and fellow population researcher, Arianna Gard. We met as graduate students at the University of Michigan, and our devotion to our work is matched only by our devotion to each other. She is a continual source of inspiration, support and advice. And we will always be proud Michigan Wolverines.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

The distribution of labor market activity across U.S. individuals has changed dramatically since 1960. While nearly all prime-age men used to participate in the labor force, men without a college education now experience substantial joblessness. At the same time, married women——especially those with college degrees——have taken up careers in increasing numbers. This dissertation explores relationships between American marriage and labor markets. It reveals new channels through which changing marriage-and-family arrangements have affected the evolution of labor market behaviors across gender and education subgroups. Its results help define the current landscape of labor and marriage inequality in the United States, and inform current debates over policies to promote job and family security.

The first chapter presents a model in which young men find employment to enhance their value as marriage partners. When the effect of employment on marital value declines, young men's employment declines as well, in preparation for a less favorable marriage market. Taking this prediction to U.S. data, I estimate that fewer young men sought employment after 2 interventions that reduced the value of gender-role-specialization within marriage: i) the adoption of unilateral divorce legislation, and ii) demand-driven improvements in women's employment opportunities. I then use a structural estimation of the model to investigate interaction between the marriage market and male labor market shocks. Simulations find that the indirect effect of a negative shock to wages on young men's employment, operating through the marriage market, is nearly as large as the direct effect that operates purely through the labor market. These findings highlight the changing marriage market as an important driver of secular decline in young men's labor market involvement.

The second chapter leverages the genealogical structure and long duration of the Panel Study of Income Dynamics to estimate intergenerational employment relationships. Previous measurements of the intergenerational transmission of women's employment status have been limited by a lack of detailed data on mothers' and daughters' employment behaviors. The intergenerational relationship is found to be strongest at the full-time employment margin for college-educated mothers, and substantially weaker at less-intensive

employment margins and for less-educated mothers. The paper also documents a stark rise in inequality in mothers' full-time employment propensities in the 21st century, and attributes roughly 36% of this trend to differential intergenerational transmission across education groups. These results suggest a disproportionate influence in high-SES families of the childhood environment on gender identity, and that family-level transmission processes deepen the long-run effects of unequal labor market opportunities on inequality in mothers' career outcomes.

The third chapter, from a work with David Lam, builds on standard marital matching models to address the question of whether it is possible to infer the existence of a "male breadwinner norm" among American families. We show that a variety of underlying social preferences about a given trait all generate positive assortative matching on that trait, and hence the same distribution of spousal trait differences in equilibrium. Applying this result to U.S. Census and administrative earnings data, we find that simple models of assortative matching can very closely replicate the observed distribution of spousal earnings differences, in which very few wives out-earn their husbands. We conclude that the distribution of spousal earnings differences in the U.S. provides little information about the existence of a male breadwinner norm or its effects on gender inequality in the labor market.

CHAPTER I

Why Bother? The Effect of Declining Marriage Prospects on Employment of Young Men

Abstract

Why have so many young men withdrawn from the U.S. labor force since 1965? This paper presents a model in which men invest time in employment to enhance their value as marriage partners. When the marriage market return on this investment declines, young men's employment declines as well, in preparation for a less favorable marriage market. Taking this prediction to data, I show that fewer young men sought employment after 2 interventions that reduced the value of gender-role-specialization within marriage: i) the adoption of unilateral divorce legislation, and ii) demand-driven improvements in women's employment opportunities. I then show, using a structural estimation, that half of the employment effect of a labor market shock to men's wages is determined by endogenous adjustment of the marriage market to the shock. These findings establish the changing marriage market as an important driver of decline in young men's labor market involvement.

JEL Codes: E24 , J12 , J21 , J22 , J24
Keywords: labor-force participation; marriage market; bargaining power; human capital investment

1.1 Introduction

Between 1965 and 2015, the share of U.S. men aged 25-34 not participating in the labor force more than tripled (Figure 1.1). Most of this aggregate change came from noncollege-educated men, for whom non-participation increased nearly seven-fold.[1] Rising joblessness of young men poses implications for human capital accumulation, family incomes, investments in children, and inequality in these outcomes.

A leading explanation for this development is that reductions in labor demand curtailed the labor market opportunities of noncollege men and induced their exit from the workforce (Juhn, Murphy and Topel, 1991; Juhn, 1992; Bound and Holzer, 2000; Moffitt, 2012; Autor, Dorn and Hanson, 2013; Acemoglu et al., 2016; CEA, 2016; Acemoglu and Restrepo, 2017; Jaimovich and Siu, 2018; Charles, Hurst and Schwartz, 2018). As shown in Panel A of Figure 1.2, declines in noncollege men's labor-force participation throughout the 1970s, 80s and 2000s occurred alongside declines in hourly earnings. These declines of quantity and price support the notion that the demand curve has been shifting inward.

Two important features of the data, however, suggest that falling labor demand is not a complete explanation. First, prominent labor demand forces were found in a recent review to account for less than half of observed decline in U.S. employment since 1999 (Abraham and Kearney, 2018). Second, it is not obvious that labor demand shifts should cause large changes to men's employment. Most estimates of the wage elasticity of male labor supply are small,[2] suggesting that male employment responds little to persistent wage changes. And in the 90 years preceding 1970, dramatic wage increases accompanied virtually no change in men's employment (Pencavel, 1986).[3] Therefore, factors beyond falling labor demand appear necessary to explain post-1965 change in noncollege men's employment.

This paper argues that changes in the labor market have interacted with changes in another market—**the marriage market**—and that such interaction provides a more complete explanation of the data. Figure 1.3 shows that noncollege men experienced a tremendous decline in marriage propensities at the same time that they withdrew from the workforce.[4] Previous reviews of this well-known "retreat from marriage" (e.g. Stevenson and Wolfers, 2007; Cherlin, 2014; Lundberg, Pollak and Stearns, 2016) have argued that changes to marriage and labor markets have reduced the attractiveness of the gender-role-specialized

[1]This is according to data from the March supplement to the Current Population Survey (CPS). Throughout this paper I consider noncollege-educated men to be men with at most one year of completed college. I often refer to this group as "noncollege" for brevity.

[2]The relevant elasticity here is the uncompensated (Marshallian) labor supply elasticity, as this describes the effect of a persistent wage change on labor supply. Coglianese (2018) surveys the literature and finds uncompensated elasticity estimates ranging from -0.02 to 0.14 (with a mean of 0.04).

[3]This relationship is evident in Panel B of Figure 1.2, which uses U.S. Census data to extend the wage and participation series back to 1940.

[4]Some of the secular decline in marriage has been offset by a rise in non-marital cohabitation and a delay in the age at first marriage. However, this is more important before age 35, as many cohabitations either turn into marriages or dissolve by this age (Mernitz, 2018). The graph presents marriage propensities for men aged 35-39 to avoid some of these complications, and still shows a steady secular decline in marriage. See the beginning of Section 1.3 for further discussion.

marriage—the predominant arrangement of working-class families in 1960. This paper argues that these changes have also reduced the value a stably-employed noncollege man could extract from the marriage market. As a consequence, young noncollege men stood to gain less from investing time in employment in 2015 than they did in 1965.

The foundation for this argument is the hypothesis that one reason young men may work hard is to improve their prospects on the marriage market. I begin by deriving such a hypothesis from a general economic framework. Male agents first choose how much time to invest in employment, taking labor market opportunities as given. Then, they participate in a competitive marriage market. The driving assumption is that the economic value of marriage depends positively on the quantity of labor men supply before the marriage market. Thus, men who have sacrificed leisure to build the most promising careers have the most to offer prospective marriage partners. I show that in equilibrium, men who expect to marry work more before entering the marriage market than men who expect to remain single. These men earn a **marriage market return** on pre-marital employment that depends on i) the effect of pre-marital employment on the economic value of marriage; and ii) the terms of marriage—that is, the share of economic value that the husband gets to claim for himself.

The framework generates two implications for secular decline in noncollege men's employment. I illustrate these with comparative statics. First, I consider a marriage market shock that lowers the value of gender-role-specialization within marriage. In response to this shock, the model predicts that i) fewer marriages involving noncollege men form; ii) the shares of marital value claimed by noncollege husbands decline; and iii) the employment propensity of young noncollege men declines. Thus, *a marriage market shock causes a shift of noncollege men's labor supply curve* at a given set of labor market opportunities. Second, I consider a labor market shock that reduces noncollege men's wage offers. The model predicts the same 3 responses. Thus, a labor market shock causes an employment response in part through its effect on marriage prospects. *This endogenous response of the marriage market to a labor market shock rotates noncollege men's labor supply curve*, magnifying men's employment sensitivity to such shocks.

I empirically assess the first implication in the context of two interventions in U.S. marriage markets. One intervention occurred during the 1960s-80s as states switched from consent-based to unilateral divorce regimes—greatly reducing legal barriers to marriage dissolution. Previous work, discussed below, has exploited differential timings of this switch across states to show that unilateral divorce led to less gender-role-specialization within marriage and less marriage formation. A second intervention occurred as technological change in the U.S. economy expanded service-related jobs and contracted manufacturing-related jobs. I exploit the fact that women and men have historically specialized in different industries to identify gender-specific employment shocks, at the local labor market level, over the period 1980-2015.[5] Holding male opportunities constant, a positive local shock to female employment opportunities should also erode the value of gender-role-specialization within marriage.

Difference-in-difference regression designs based on U.S. Census data confirm the model's predictions: following both interventions, young noncollege men married less and worked less. Single men experienced particularly large reductions in labor-force participation, consistent with the marriage market investment channel emphasized by the model. The results pass tests for internal validity and are relatively robust across an array of specifications. These interventions appear to be quantitatively important: a back-of-the-envelope calculation suggests that they were responsible for 29% of the 1965-2015 decline in labor-force participation by young men without college.

The second implication of the framework is that, in addition to the "direct" effect of reduced labor market opportunities on employment emphasized by the literature, a persistent shock to noncollege men's wage offers causes an "indirect effect" that operates through the marriage market. Because a reduced-form evaluation of a shock to noncollege men's labor market opportunities cannot distinguish the two effects, assessing this implication requires a structural approach. Guided by the theoretical framework, I specify an empirical model and estimate its parameters using the Generalized Method of Moments. I confirm that the model successfully replicates the targeted marriage and employment behavior of young men circa 1980. Then, I simulate a 10% reduction in the lifetime wages of noncollege-educated men and solve for the new equilibrium. I find that the indirect employment effect, operating through the marriage market, is roughly as large as the direct effect operating through the labor market.

These empirical results shed new light on the secular decline in employment of young, noncollege-educated men. In an era marked by young women's unprecedented access to the labor market,[6] decline in the relative labor market positions of noncollege men has caused a meaningful decline in their marriage market value. As a consequence, such men have reduced their investments in the labor market. When work is less likely to win a

[5] This gender-specific variant of the standard Bartik instrument has been applied in several recent papers (e.g. Page, Schaller and Simon, 2019).

[6] In addition to the factors I investigate here (liberalized divorce law and an increasingly favorable wage structure), factors driving the rise in women's opportunities include declining significance of gender norms and stigma associated with out-of-wedlock childbearing, increased legal protection against discrimination, increased access to contraceptive methods and abortion, and the rise of labor-saving home production technology. See Goldin (2006) and Stevenson and Wolfers (2007) for discussion, and Goldin (2006) for an argument that these changes dramatically altered the career aspirations and investments of young women.

desirable marriage contract, why bother?

The effects emphasized in this paper illuminate a contrast between the period of the late 1800s-mid 1900s and the more recent period. Throughout the former period, young wives and mothers performed almost no paid work (Goldin, 2006), while a large working class of young men decided between working for their fathers in family-owned establishments or pursuing wage labor opportunities that required little formal schooling (Ruggles, 2015). In response to rising opportunities in cities, young men could establish economic independence, although at the likely cost of not inheriting the family home and business. Legal and social enforcement of patriarchal norms also prevailed during this time: women who desired marriage had little ability to dictate its terms (Coontz, 1992). For these reasons, it is plausible that the rise in noncollege men's wage labor opportunities throughout this period exerted little impact on the economic value they could extract from the marriage market. Thus, wage changes may have had less effect on noncollege men's employment because they had less effect on marriage prospects.

1.2 Related Literature

This paper is not the first to venture outside of the labor market to explain trends in male employment. For example, much has been written about disability insurance programs.[7] The evidence suggests that rising disability insurance provision contributed importantly to the decline in employment among men older than 45 (Bound and Waidmann, 1992, 2002). Krueger (2017) hypothesized that rising opioid addiction has compromised the abilities of pain-affected individuals to work.[8] It is reasonable that displaced older workers, especially those in relatively poor health, might stand to gain little from remaining attached to the labor force—especially if they qualify for disability insurance. An important contribution of this paper, therefore, is to propose a mechanism for declining male employment that is relevant for younger men. The argument that marriage market forces have *lowered the value of employment* for young men provides a contrast to the recent argument that improved leisure technologies have *raised the value of non-employment* (Aguiar et al., 2017).

Some observers have also posited a link between the declining employment of young noncollege men since 2000 and the growing fraction of such men that have criminal records. (For evidence on the effect of incarceration on post-release employment, see

[7]See Bound and Burkhauser (1999) and Liebman (2015) for reviews.

[8]Initial tests of this hypothesis have reached different conclusions about its magnitude and significance (Currie, Jin and Schnell, 2018; Aliprantis, Fee and Schweitzer, 2019; Harris et al., 2019).

Pager, 2007 and Mueller-Smith, 2015.) Introducing the marriage market channel into the discussion helps rationalize why such men's employment began declining decades before the era of mass incarceration. However, the incarceration channel may interact in important ways with the marriage market channel, enhancing the relevance of both to the post-2000 period.

Topically, this paper relates to recent work analyzing the impact of changes in gender-specific wage structures on married men's labor supply (Knowles, 2012; Alon, Doepke and Coskun, 2018). There is also a "marriage premium" literature that assesses whether men's labor market outcomes improve upon transitioning into marriage (see Ludwig and Brüderl, 2018 for references and a discussion). In contrast to these papers, I focus on the labor supply response of *unmarried* men to anticipated changes in marriage market conditions. While only 15% of noncollege men aged 25-44 were unmarried in 1965, today over half of such men—and 65% of those aged 25-34—are unmarried. Moreover, Binder and Bound (2019) report that declining labor-force participation by husbands accounts for a small share of the total observed decline. To fully understand the decline in labor-force participation by noncollege men of prime working age, it is important to account for their decline in marriage combined with the decline in employment of unmarried men.

Substantively, this paper relates to work on marriage markets with pre-marital investments in human capital. This work has primarily focused on the education decision (Lafortune, 2013; Chiappori, Iyigun and Weiss, 2009; Chiappori, Dias and Meghir, 2018). This paper takes education as given but instead models employment behavior as a pre-marital investment choice. The framework presented in the main text assumes an efficient matching environment, as in Chiappori, Iyigun and Weiss (2009). In the Appendix, I work through an extended model wherein the arrival of unilateral divorce creates a coordination problem within marriage. This induces an inefficiently low number of men to work hard before the marriage market, consistent with what has been established in theoretical literature on investment-and-matching games (e.g Peters and Siow, 2002; Nöldeke and Samuelson, 2015).

The empirical model used in this paper builds on the transferable-utility matching model developed by Choo and Siow (2006) and extended by Chiappori, Salanié and Weiss (2017).[9] Given a distributional assumption on idiosyncratic preferences, these models apply standard discrete choice techniques to identify marital surpluses based on observed marital matching patterns. My model operates in the same manner. However, the presence of a pre-match employment decision also requires identification of the *division* of the marital surplus between spouses. This is because a man's anticipated surplus share guides his

[9]For further discussion, see the review of Chiappori and Salanié (2016).

initial employment choice, and this choice in turn affects the total marital surplus. To handle this type of interdependence (between total marital welfare and its allocation between spouses), I modify an approach developed by Galichon, Kominers and Weber (2018). With this approach, observed pre-marital employment of noncollege men, together with marital matching patterns, jointly identify marital surpluses and spousal surplus shares.

Finally, this paper furthers the literature on the long-run impacts of unilateral divorce (Rasul, 2003; Stevenson, 2007; Fernández and Wong, 2017; Reynoso, 2019). As noted by Chiappori, Dias and Meghir (2018), beyond their effects on existing families, changes to economic and family policies can have long-run consequences for marriage formation, marital matching patterns and human capital investment incentives. In revealing long-run effects of unilateral divorce on young men's marriage-market investments, the current paper contributes to this new research agenda in family economics (see also Chiappori et al., 2017; Low et al., 2018).

1.3 Young Men's Employment as Marriage Market Investment: Motivation and Theory

Economists have long recognized the importance of prior job experience and career interruptions in earnings determination (Mincer, 1958, 1962). Building up experience with a specific employer, rather than cycling between jobs, has been shown to confer additional earnings returns (Topel, 1991; Dustmann and Meghir, 2005). Young men thus reap future gains, in terms of the consumption value of future earnings, from pursuing stable employment today. In this section, I argue that there is an additional return to such behavior that is determined in the marriage market. As a result, anticipated changes in the marriage market plausibly affect young men's employment.

Before embarking, it is important to note that as marriage has declined in the United States, non-marital cohabitation and other family arrangements have risen in importance.[10] Has the "marriage" market actually changed in a meaningful way, or has marriage simply been replaced by equivalent arrangements? Some recent work on U.S. family structure has revealed the following. First, the median cohabitation duration for less-educated couples was 22-24 months in 2006-2010 data—much shorter than the median marriage duration (Copen, Daniels and Mosher, 2013). Second, a growing minority of individuals do cohabit for extended periods of time without marrying (Mernitz, 2018), but this population remains quite small. More salient is an increase in "serial cohabitation:" those whose first

[10] For a "first-generation" review, see the classic article of Bumpass (1990). For a contemporary discussion, see the recently-released review of Smock and Schwartz (2020).

cohabitation has dissolved are increasingly likely to enter a new cohabitation (Eickmeyer and Manning, 2018). Third, cohabitation has not fully replaced marriage as a childbearing arrangement in the noncollege population: single mothers account for a steadily growing share of births as well (Manning, Brown and Stykes, 2015).

As emphasized by Lundberg, Pollak and Stearns (2016), a key difference between marriage and these other family arrangements lies in marriage's function as a commitment device. To the extent that the benefits of living with a committed partner accumulate over the medium-to-long term (e.g. through investment in partnership-specific goods like children and housing), I treat the replacement of marriage by unstable cohabitation and single motherhood as a meaningful change that may impact the labor market investments of noncollege men. That is, my purpose is to model the formation of stable partnerships over the long-run. Although a dichotomous married/single framework does not adequately characterize American family structure, especially over a short-to-medium-run horizon, I submit that the dichotomous framework used by economists remains a useful abstraction for my present purpose.

Same-sex marriage has also become a more common joint living arrangement in the United States. However, nearly all marriages formed in the United States have been heterosexual. As my purpose is to consider how the marriage market may have impacted secular developments in men's employment, I restrict focus to heterosexual marriage here.

1.3.1 What men contribute to marriage

Economists think of monogamous marriage as a productive partnership between two individuals, in which each individual is better off in the partnership than either could be as single (Becker, 1973). What do men contribute to these partnerships? Gary Becker's seminal framework involves a marital surplus function that is maximized when spouses exploit their comparative advantages within the household economy (Becker, 1981). Given persistent gender gaps in labor market opportunities and the biological demands of childbirth, this tends to result in complementarities between husbands' earnings and wives' time at home in the production of marital happiness. An innovation to the specialization framework recognizes the importance of the jointly-consumed public goods (Lam, 1988). When economies of scale in public consumption outweigh specialization incentives, complementarities arise between husbands' and wives' earnings. Thus, while women's contribution to marriage is determined by competing influences that have changed with time, men who wish to form a stable marriage have had a clear dictate: form a stable career.

The hypothesis that male earning potential is an important determinant of marital value has been subject to numerous quasi-experimental tests in the last decade.[11] To the extent that marital value comes from investing in children, we might expect potential husbands' labor market positions to affect fertility behavior and children's outcomes. Black et al. (2013) found positive male earning shocks, driven by the coal boom of the 1970s, to lead to greater completed fertility. Leveraging variation in industrial business cycles combined with gender differences in employment by industry, Schaller (2016) found positive male labor demand shocks to predict increased fertility, yet positive female demand shocks to predict decreased fertility. Related work has established a similar pattern of effects on child health outcomes (Stevens and Schaller, 2011; Lindo, Schaller and Hansen, 2018; Page, Schaller and Simon, 2019)—large, positive effects of male demand yet small (and usually negative) effects of female demand. Others have associated negative shocks to fathers' lifetime earnings with decline in educational attainment and earnings of affected children (Oreopoulos, Page and Stevens, 2008).

In addition to their effects on children, the labor market positions of men appear to drive marriage formation and stability itself. Following a similar strategy to Schaller (2016), Autor, Dorn and Hanson (2019) partitioned labor market shocks from the rise in Chinese import competition since 1990 into male-specific and female-specific components. Communities receiving particularly negative male shocks experienced a relative decline in marriage and rise in single motherhood, while those receiving negative female shocks saw smaller increases in marriage. Bertrand, Kamenica and Pan (2015) found similar results and attributed them to gender identity norms—dictating that the husband out-earn his wife—as well as economic gains from specialization.

It is also plausible that the marriage market rewards men who have worked hard for reasons additional to the effects of working hard on earning potential. For example, Charles and Stephens (2004) found that divorce hazards rose following a husband getting laid off, but not from job loss events outside the husband's control (e.g. suffering a work disability or a plant closure), even though both types of events conferred similar lifetime earnings losses. Such a pattern is consistent with stable employment signaling the man's level of non-economic suitability as a partner. Lafortune and Low (2018) presented a framework in which men who enter marriages with sufficient assets are better able to solve coordination problems within marriage and reap maximum gains from specialization. Men with stable jobs may also signal to prospective wives that they are committed to investing in the marriage rather than other pursuits.

[11] See also Wilson (1996) and Cherlin (2014) for illuminating sociological narratives.

1.3.2 A simple equilibrium framework with marriage market investment

I consider a 2-period decision-making environment ($t = 0, 1$) in which men first make a human capital investment in the labor market and then match with women in a frictionless, competitive marriage market. Figure 1.4 illustrates the environment and shows the relevant payoff functions. There is no uncertainty and households do not save. In Appendix Sections A.6 and A.7, I work through a more complex framework with marital uncertainty and the possibility of unilateral divorce. I discuss additional implications of this framework in the next section.

Single household's employment choice. To start, consider a single male household making consumption and employment decisions over his two periods of life. Each man m in period t has direct utility function $U(c_{mt}, n_{mt}; \emptyset)$, where c is consumption of a market-purchased good and $n \in [0, 1]$ describes the share of the period that man spends employed.[12] (The argument after the semicolon in the utility function tracks the individual's family status: $\emptyset$ if single, and f if married to woman f.)

A standard result from static labor supply theory is that, given leisure preferences (which are assumed constant across all men of a given education status), an individual's hourly wage is sufficient to characterize his consumption, labor supply, and utility. I therefore use the indirect utility function $V(w_{m1}; \emptyset)$, where w is hourly wage, to summarize period 1. Each man faces the Mincerian wage equation:

$$\ln(w_{m1}) = \ln(w_{m0}) - \delta + (r + \delta) \cdot n_{m0} \tag{1.1}$$

where initial wages are exogenous, δ is a depreciation factor, and $r + \delta$ captures wage growth from stable employment.

Without saving, the period-0 consumption choice is pinned down by the man's period-0 employment choice: $c_{m0} = w_{m0}n_{m0} + y_{m0}$, where y is exogenous non-labor income. This yields the following employment decision problem:

$$\max_{n_{m0}} U(c_{m0}(n_{m0}), n_{m0}; \emptyset) + \beta V(w_{m1}(n_{m0}); \emptyset) \tag{1.2}$$

where β is a discount factor capturing the importance of period 1 relative to period 0. For

[12] The implicit length of each period in the model is more than one year: period 0 is meant to represent the amount of time the average non-college-educated men spends not in school before becoming stably married. In the 1979 cohort of the National Longitudinal Surveys of Youth, the average noncollege men spent 7 years unmarried between ages 20 and 32. Over such a horizon, I model employment as a continuous choice. Evidence from longitudinal administrative data presented in Appendix Table A.1 indicate that most men who spend significant time out of the labor force in a given year return to gainful employment at some point in the future: that is, men are not either always out or always in the labor force.

purposes of exposition, I assume a separable utility function:

$$U(c_{mt}, n_{mt}; \emptyset) = \text{VALUE}(c_{mt}) - \text{COST}(n_{mt}),$$

where VALUE is the value of consumption and COST is the effort cost of working. All qualitative results, however, generalize to a non-separable function.

It is straightforward to show that the solution to problem (1.2) satisfies the following first-order condition:

$$\begin{aligned} \text{MC}(n_{m0}) &= w_{m0} \cdot (\text{MV}(c_{m0}) + \beta\hat{r}V'(w_{m1})) \\ &= w_{m0} \cdot \text{MB}(n_{m0}; w_{m0}, \hat{r}) \end{aligned} \tag{1.3}$$

where MV and MC denote marginal consumption value and marginal effort cost functions, respectively, and $\hat{r} = r + \delta$. I subsume the expression inside the parentheses into the total marginal benefit function $\text{MB}(n_{m0}; w_{m0}, \hat{r})$.

The marriage market. Now, suppose that men enter a marriage market in period 1. I consider a frictionless marriage market with equal numbers of men and women. Men come in two education types: noncollege-educated (NC) and college-educated (C). In the empirical application, I also consider two education types of women. For now, I consider an arbitrary number $F \geq 1$ of female types f.

As shown in Figure 1.4, the period-1 utility payoff of a man m who marries woman f can be expressed as the sum of his utility from remaining single and the additional utility he receives from the marriage. I express this additional utility gain as $G_m(w_{m1}, n_{m0}; f)$. This gain is itself a sum of two parts: man m's share in the total economic surplus generated by the marriage, plus an idiosyncratic preference for the given partner's type. I represent idiosyncratic preferences as follows: ϵ_m^f is man m's taste for being married to female type f, ϵ_f^E is woman f's taste for marrying a man of education status E, and $\epsilon_i^{\emptyset}$ is individual i's preference for remaining single. Idiosyncratic tastes ϵ are distributed *iid* according to a mean-0, atomless distribution $H(\epsilon)$. Thus, for a given marriage between noncollege man m and woman f, individual utility gains are:

$$\begin{aligned} G_m(NC; f) &= \theta_m(NC; f) \cdot S(NC; f) + \epsilon_m^f - \epsilon_m^{\emptyset} \\ G_f(NC; f) &= \underbrace{(1 - \theta_m(NC; f)) \cdot S(NC; f)}_{\textbf{Economic}\text{ gains}} + \underbrace{\epsilon_f^{NC} - \epsilon_f^{\emptyset}}_{\textbf{Non-economic}\text{ gains}} . \end{aligned} \tag{1.4}$$

In (1.4), $S(NC; f)$ is the total economic surplus generated by the marriage and $\theta_m(NC; f)$ is man m's share in the surplus. The marriage between m and f forms if both utility gains exceed zero and if $G_f(NC; f) > G_f(C; f)$ (which is the utility gain woman f receives from marrying a college man).

It is convenient here to assume transferable utility, wherein the total economic surplus does not depend on how the surplus is shared: spouses always agree on total-surplus-maximizing behavior within marriage. This allows us to abstract from the nature of behavior of within marriage.[13] Marriage market equilibrium consists of an assignment A of individuals and a set Θ of surplus shares such that the assignment is *stable*: no two individuals would prefer to dissolve their current matches and match with each other, and no matched individual would be strictly better off remaining single. See Appendix A.6 for more details.

Lastly, I specify the economic surplus function. Consistent with the fact that virtually all college-educated men in their 20s-40s participate in the labor force, and this has not changed since the 1960s, I assume college men are employed throughout their working lives. Relative to each partner remaining single, a marriage between college-educated man m and woman f produces a surplus of $S(C; f)$. A marriage between noncollege-educated man m and woman f produces a surplus of

$$S(NC; f) = S(w_{NC,1}^{\text{married;f}}, n_{NC,0}^{\text{married;f}}; f). \tag{1.5}$$

Consistent with the class of marriage theories discussed in the preceding subsection, I assume the surplus function S is strictly increasing in its first two arguments: the earning potential and employment history man noncollege man m brings into the marriage with f.

Marriage market's impact on employment choice. Consider a noncollege man m who desires to marry woman f at the given equilibrium surplus share of $\theta_m(NC; f)$. Based on (1.4) and (1.5), the partial derivative of such a man's utility gain from marriage with respect to his employment choice is the following:

$$\frac{\partial G_m(NC; f)}{\partial n_{NC,0}^{\text{married;f}}} = \theta_m(NC; f) \cdot \underbrace{(\hat{r} \cdot w_{NC,0} \cdot S_w(w_{NC,1}^{\text{married;f}}, n_{NC,0}^{\text{married;f}}; f) + S_n(w_{NC,1}^{\text{married;f}}, n_{NC,0}^{\text{married;f}}; f))}_{\frac{dS}{dn_{NC,0}^{\text{married;f}}}}.$$

In the above expression, $S_w(\cdot)$ and $S_n(\cdot)$ are the partial derivatives of the total marital surplus with respect to period-1 wage and period-0 employment, respectively. This expression describes the **marriage market return** to period-0 employment at equilibrium prices (i.e. at the given surplus share $\theta_m(NC : f)$). In words, this is the effect of the man working

[13] The more complex environment considered in the Appendix takes a stance on behavior within marriage and results in *imperfectly* transferable utility, where the magnitude of the total surplus (in an *ex-ante* sense) depends on how it is shared. That model nests the qualitative predictions of this more restrictive model.

harder in period 0 on the total economic benefits of marriage realized in the next period, scaled by the man's share in those benefits. For ease of exposition, I shorten the portion of the expression inside the parentheses by using total derivative notation: $\frac{dS}{dn_{NC,0}^{\text{married;f}}}$.

This expression facilitates an important comparison of employment first-order conditions between a man who chooses to remain single and an otherwise identical man who chooses to marry:

$$\text{ALWAYS}-\text{SINGLE}: \quad \text{MC}(n_{NC,0}^{\text{single}}) = w_{NC,0} \cdot \text{MB}(n_{NC,0}^{\text{single}}; w_{NC,0}, \hat{r}). \tag{1.6}$$

$$\text{TO}-\text{BE}-\text{MARRIED}: \quad \text{MC}(n_{NC,0}^{\text{married;f}}) = \underbrace{w_{NC,0} \cdot \text{MB}(n_{NC,0}^{\text{married;f}}; w_{NC,0}, \hat{r})}_{\text{Marginal benefits from } \textbf{labor market}} + \underbrace{\beta\theta_m(NC; f)\frac{dS}{dn_{NC,0}^{\text{married;f}}}}_{\text{Marginal benefits from } \textbf{marriage market}}. \tag{1.7}$$

Both men share the same marginal effort cost curve, which simply depends on period-0 preferences for employment. However, relative to the always-single man, the man who will marry faces a marginal benefit curve that is shifted upward by the additional marginal benefits the man expects to receive in marriage. Recall that, because these benefits arrive in the next period, they must be scaled by the discount factor β.

Equilibrium. I consider an environment where initial male wage offers and employment preferences vary only according to education attainment. The initial wage offer for a noncollege man is $w_{NC,0}$ and for a college man is $w_{NC,0}(1+p)$ (where p is the college wage premium). The economic surplus in a marriage involving a college-educated man is some constant number S^C, while the economic surplus in a marriage involving a noncollege-educated man is given by (1.5). And, the distribution function of idiosyncratic tastes, $H(\epsilon)$, is type I extreme value. Equilibrium is defined as follows:

A rational-expectations equilibrium consists of an assignment A *of individuals in the marriage market; a set of vectors of economic surplus shares*

$$\{\Theta_f\}_{f\in F} = (\theta_m(NC; f), \theta_m(C; f), \theta_f(NC; f), \theta_f(C; f))$$

for each type of woman; and a vector of period-0 *employment choices for noncollege-educated men* $N = (n_{NC,0}^{\text{single}}, (n_{NC,0}^{\text{married;f}})_{f\in F})$*; satisfying the following conditions:*

1. $\theta_m(NC; f) + \theta_f(NC; f) = 1 = \theta_m(C; f) + \theta_f(C; f)$ *for all* $f \in F$*. That is, the divisions of marital surpluses between partners are feasible.*

2. *The assignment is stable.*

3. *Period-*0 *employment choices satisfy first-order conditions (1.6) and (1.7) for all* $f \in F$*.*

Existence and uniqueness. In the model, noncollege men make labor market investments and then move to the marriage market where they are valued in terms of these prior investments. A rational-expectations equilibrium occurs when noncollege men's investment choices are consistent with a set of prices Θ that also, at the given the investment choices, clear the marriage market (Chiappori, Iyigun and Weiss, 2009). Uniqueness of equilibrium is not generally guaranteed in these environments. (See Chiappori, Dias and Meghir (2018) for an involved discussion on this point.) However, one can show a unique equilibrium exists in this context under a reasonable boundary condition. I provide a proof in Appendix A.3.

Figure 1.5 illustrates equilibrium outcomes for noncollege men, assuming $F = 1$ for ease of exposition (i.e. women are homogeneous). The upper graph depicts the period-0 employment choices of noncollege men who will remain single (dotted MB curve) and noncollege men who will marry (solid MB curve). The steep-sloping blue line in the middle graph traces the locus of points at which the solid marginal benefit curve intersects the marginal effort cost curve. That is, it traces the points at which the period-0 employment choice of to-be married men is rational, as the surplus share $\theta_m(NC)$ varies. The left graph plots demand and supply curves for noncollege men in the marriage market, given a period-0 employment choice. The shallow-sloping orange line in the middle graph traces the locus of points at which these curves intersect. That is, it traces the points at which the marriage market assignment A is stable, as the employment choice of to-be-married men $n_{NC,0}^{\text{married}}$ varies. Equilibrium occurs where the two loci in the middle graph intersect: where the employment choice of to-be married men satisfies first order condition (1.7), given marriage market conditions; and the marriage market is stable, given the employment choice.

I can now express the central result of this section: *population employment behavior of noncollege men depends on i) the economic surplus function and ii) the equilibrium surplus shares*. A standard application of the Implicit Function Theorem to first-order conditions (1.6) and (1.7) yields the following equilibrium employment functions:

$$\begin{aligned} n_{NC,0}^{\text{single}} &= n\left(w_{NC,0}, \hat{r}\right) \\ n_{NC,0}^{\text{married;f}} &= n\left(w_{NC,0}, \hat{r}, \theta_m(NC;f), \frac{dS}{dn_{NC,0}^{\text{married;f}}}\right). \end{aligned}$$

That is, the employment rate for always-single men depends on labor market conditions only ($w_{NC,0}$ and $\hat{r}$), while the employment rate for to-be-married men depends on both labor market and equilibrium marriage market conditions $\left(\theta_m(NC;f), \frac{dS}{dn_{NC,0}^{\text{married;f}}}\right)_{f\in F}$.

With idiosyncratic preferences distributed type I extreme value, it is straightforward to show that the population employment rate for noncollege men, $n_{NC,0}$, takes the following logistic form:

$$n_{NC,0} = \frac{n^{\text{single}}_{NC,0} + \sum_f \exp\left(\theta_m(NC;f) \cdot S(w_1^{NC}, n^{\text{married;f}}_{NC,0}; f)\right) \cdot n^{\text{married;f}}_{NC,0}}{1 + \sum_f \exp\left(\theta_m(NC;f) \cdot S(w_1^{NC}, n^{\text{married;f}}_{NC,0}; f)\right)}. \tag{1.8}$$

Comparative static. For ease of exposition, return to the case where women are homogeneous: $F = 1$. Consider a shock that reduces the economic surplus by reducing noncollege men's marginal contribution $\frac{dS}{dn_{NC,0}}$. The shock of interest is a reduction in the gains from specialization, driven by an improvement in women's relative labor market opportunities; a change in the legal framework surrounding marriage; or a reduction in the strength of "male breadwinner" norms. Several changes occur. First, the reduction in $\frac{dS}{dn_{NC,0}}$ induces noncollege men to work less in period 0. Second, because the marital value of noncollege *relative* to college men has gone down, the new equilibrium must occur at a lower share for noncollege men.[14] This further erodes the marriage market return to period-0 employment. Third, the combination of the first two forces increases the number of noncollege men remaining single, leading more such men to "opt out" of the high-employment path dictated by future marriage. In the language of equation (1.8), the first two forces lead to a reduction in $n^{\text{married}}_{NC,0}$, while the third force reduces $\exp(\cdot)$, hence giving $n^{\text{single}}_{NC,0}$ more weight in the expression. The overall result is a lower population employment rate, $n_{NC,0}$, in equilibrium.

Figure 1.6 provides an illustration. The initial shock leads to a downward shift of to-be married men's marginal benefit curve. That is, for a given surplus share $\theta_m(NC)$, less premarital employment occurs in the new equilibrium, which implies a leftward shift of the blue locus in the middle graph. Because marriage is no longer as attractive at the given surplus share, both marriage market curves in the left graph shift inward, but prospective wives' demand curve shifts inward more, since college-educated men have become a better bargain relative to noncollege men. This causes a downward shift of the orange locus, as a given employment level is now associated with a lower surplus share. Equilibrium is restored at a lower value of $n^{\text{married}}_{NC,0}$, a lower share of noncollege men married, and a lower surplus share $\theta_m(NC)$.

[14] See Appendix A.3 for a proof.

1.3.3 *Prima facie* evidence from cohort data.

The theoretical framework's predictions are consistent with the patterns shown in Figure 1.3. As less-educated men entered their prime years of potential labor-force activity with diminished expectations of forming stable marriages they may also have found less to gain from forming strong attachments to the labor market. Equivalently, as secular changes reduced noncollege men's marginal contribution to the marital economic surplus, fewer marriages formed and fewer of these men built up stable attachments to the labor market while young.

Co-variation between individual employment profiles and marriage outcomes, within a given set of birth cohorts, conveys a similar message. Figure 1.7 and Table 1.1 record employment and family structure cross-tabulations based on data from the 1979 cohort of the National Longitudinal Surveys of Youth (NLSY79). The NLSY79 contains weekly labor market histories and annual information on family structure for a population-representative sample of men who were 14-22 years old when first surveyed in 1979. I divide the subsample of men with no more than one year of completed college into two groups. The "always-married" group consists of individuals who reported being married at all interviews between the ages of 32 and 40; the "always-single" group consists of those not reporting being married, or cohabiting with an unmarried partner, at any interview over the same age range.[15]

Figure 1.7 plots age profiles of participation and employment for each of the two groups. The difference in participation rates between the two groups is substantial: the average difference over the 21-31 age range is around 11 percentage points. (This is close to the cumulative decline in labor-force participation between the 1937-39 and 1982-84 cohorts observed in Figure 1.3.) The difference in employment rates between the two groups is even larger.

Table 1.1 reports the same comparison after adjustment for differences in initial labor market opportunities. Considering the original sample of less-educated men, I estimate regressions of the following form:

$$\overline{\text{LABOR}}_{i,21-31} = C_0 + C_1\overline{\text{married}}_{i,32-40} + C_2 \ln(\text{initial wage}_i) + \text{error}_i. \tag{1.9}$$

That is, I model young male employment behavior as a linear function of future marriage propensity and one's initial log hourly wage offer. Such a control is meant to proxy for

[15] I identified non-marital cohabitations in the NLSY79 using the method of Oppenheimer (2003). The "always-married" group contains 659 individuals, while the "always-single" group contains 289 individuals. See Appendix A.2 for further NLSY data processing information.

persistent differences between individuals in labor market opportunities as well as labor supply preferences (e.g. motivation to work hard). To construct this control, I average hourly wages over the first 3 years after the individual leaves school and take the log.[16] As shown in the table (which presents estimates of C_1), a 0-to-1 increase in marriage propensity over ages 32-40 is associated with a 7.3 percentage-point increase in labor-force participation and an 11.9 percentage-point increase in employment during ages 21-31 (columns 1 and 5). Controlling for wages reduces these associations to 6.7 and 10.8 percentage points—still large and statistically significant (columns 2 and 6). Adding a control for non-marital cohabitation propensity slightly strengthens the results, as to-be-cohabiters work slightly more than to-be-singles (remaining columns). Thus, even among similar initial earners, young men who end up maintaining stable marriages build stabler employment histories than young men who remain single.

These descriptions of the data suggest the empirical relevance of marriage market investment to the employment decisions of young men. However, to attribute the observed variation in employment behavior to exogenous variation in the value of marriage, all other determinants of employment behavior must be held constant. In the next section, I leverage two sources of plausibly exogenous variation in the marriage market values of less-educated men to provide causal interpretations of the marriage market investment channel.

1.4 Young Men's Employment as Marriage Market Investment: Evidence from Two Interventions in U.S. Marriage Markets

To isolate marriage market incentives to seek employment in the data, it is necessary to leverage changes in the marriage market occurring independently of changes in the labor market. Such a context does not readily arise: as discussed above, men's labor market opportunities appear to influence marriage market outcomes. In this section, I propose and empirically assess two interventions in U.S. marriage markets that plausibly avoid this identification problem.

[16]For some individuals, initial wage information was missing, due to either insufficient employment during the reference period or survey non-response during the reference period. Wage information for many of these individuals was available later in their career cycles. Using later-in-life wage information, I impute wages for these individuals based on a Mincerian wage equation. For individuals with chronically missing wage information due to insufficient employment, I impute wage offers using techniques developed by Juhn, Murphy and Topel (1991). See Appendix A.2 for further detail.

1.4.1 Intervention 1: the no-fault divorce revolution

Between the late 1960s and mid 1980s, divorce legislation liberalized dramatically across many U.S. states. In states without mutual consent requirements for divorce, the introduction of no-fault grounds for divorce—such as "irreconcilable differences" or "irremediable breakdown of the marriage"—effectively allowed one spouse to initiate a divorce without consent of the other. These laws changed on a state-by-state basis for a variety of reasons; family law experts have argued that the changes were effectively random in nature. Friedberg (1998) and Voena (2015) provide further discussion.

Theoretical predictions. The extended theoretical model presented in Appendices A.6 and A.7 offers clear predictions regarding the impact of a switch in the divorce regime on the employment behavior of less-educated men. When spouses must mutually consent to a divorce, an unhappy spouse cannot credibly exercise the threat of leaving the marriage. But in a unilateral regime, the utility a spouse could achieve in a unilateral divorce becomes a credible threat point within marriage. Thus, under a unilateral regime, a wife may choose to invest in her own career for two related reasons. The first is to insure against being poor in the event that the marriage turns out badly and the husband initiates a unilateral divorce. The second is to insure against the husband not sharing his earnings with her within marriage—if he doesn't, she can credibly exercise the threat of unilateral divorce. This action results in lower gains from specialization in the production of marital public goods, such as children's welfare and housing quality. To ensure the efficient division of household tasks, the husband must promise to transfer to the wife a larger share of household resources—a commitment that lowers his welfare in marriage and may not be enforceable. The overall effect is to lower the attractiveness of the traditionally-specialized marital arrangement.[17] This results in *ex-ante* low-surplus marriages no longer forming, and fewer young men investing in stable employment to prepare for their traditional marital role.

Previous empirical research on unilateral divorce has found evidence supporting these behavioral channels. Stevenson (2007) found evidence that marriages forming after exposure to unilateral divorce featured less specialization and a decreased willingness to invest in not-easily-divisible assets (such as having a child or supporting a spouse through school). Moreover, wives realized higher rates of labor-force participation and full-time work in marriages formed under a unilateral divorce regime. Consistent with these chan-

[17] The combination of marital uncertainty and the fact that specialization requires the wife to let her labor market skills depreciate leads to imperfect (i.e. costly) contracting. This implies a failure of the Coase theorem, which establishes that when contracting is costless, a shift in property rights should not affect the number of contracts forming.

nels lowering the attractiveness of the traditional marital arrangement, Rasul (2003) found a persistent decline in the marriage rate (new marriages per thousand adults, and new marriages per thousand single adults of marriage age) upon unilateral divorce adoption. Reynoso (2019) echoed these results and also found an increase in assortative matching on earning potential, as spouses became less able to realize specialization incentives after the shift in individual property rights.

In the language of the theoretical framework introduced above, by reducing specialization gains and raising the wife's threat point in the marriage relative to the husband's, unilateral divorce results in a combination of a reduction in noncollege men's marginal contributions to marriage $\left(\frac{dS}{dn_{NC,0}}\right)$ and a reduction in their surplus shares $(\theta_m(NC))$. This results in the additional prediction that noncollege-educated men should gain less on the marriage market from seeking stable employment while young. Using standard difference-in-difference methodology, I extend the literature on unilateral divorce by estimating the impact of this change in the marital environment on young noncollege men's participation in the labor force.

Data. Using 1960-1990 U.S. Census samples, I construct a main sample of non-institutionalized men, not currently enrolled in school, at least 18 years old, with at most one year of completed college, and with 0-20 years of potential labor market experience. When estimating labor supply regressions I focus on a slightly younger sub-sample of men with 0-15 years of potential experience. I consider *single* men as well as all men within this sample, consistent with the theory's emphasis on the pre-marital investment value of employment. When estimating marriage regressions I focus on a slightly older sub-sample of men with 5-20 years of potential experience. I adopt the coding of divorce laws presented in Appendix F of Voena (2015). See Appendix Section A.2 for further data details.

Regression specification. I adopt the following baseline linear probability model specification:

$$\begin{aligned} \mathbf{1}\{\text{outcome}_{ist}\} = \alpha \ + \ \beta \cdot \text{UD}_{st} \ &+ \ \gamma \cdot \text{demographic controls}_{it} \\ &+ \ \text{state FE} \\ &+ \ \text{time FE} \\ &+ \ \text{error}_{ist}. \end{aligned} \tag{1.10}$$

The dependent variable is a binary indicator taking the value 1 if male i, living in state s at time t, obtained the given outcome. The explanatory variable of interest, UD_{st}, is a binary indicator taking the value 1 if unilateral divorce is available in state s at time t. A standard requirement of difference-in-difference models is the inclusion of area and time fixed effects. I also control for a set of individual characteristics: race, ethnicity, education, potential experience level, nativity (foreign- or US-born) and urban/rural status.

The baseline specification also controls for the prevailing property division regime on divorce (title-based, community or equitable). The estimates proved robust to property division controls; thus, I only present estimates that include these controls.[18] For robustness, I also augment the baseline specification with region-by-year fixed effects or state-specific linear trends. Because marriage and labor market outcomes have been trending over time according to individual demographics, I also sometimes control for interactions between individual demographic characteristics and linear time trends.

1.4.2 Intervention 2: shocks to young women's employment opportunities

Gender gaps in employment and wages in the United States labor market have narrowed considerably since World War II (Goldin, 2006; Blau and Kahn, 2017). For married women and mothers with young children, relative gains in employment were particularly strong after 1960 Goldin (2006).

Young women and mothers entered the labor market as it became increasingly dominated by service-sector jobs. Previous studies have emphasized the role of technological change, which lowered the relative cost of service provision and consumption, in the growth of the service economy (Lee and Wolpin, 2006; Buera and Kaboski, 2012). It is therefore likely that secular shifts in demand for services played an important role in accommodating the secular rise in female employment. According to 1980 U.S. Census data, women greatly outnumbered men in a variety of service sectors. This was true among all prime-age workers but especially among noncollege workers. For example, women accounted for 44% of the less-educated workforce, but 63% of retail trade workers and 80% of those employed in finance, insurance and real estate sectors.

It is likely that the national rise of the service economy affected female employment differently in different local areas. Although services are on average less tradable than manufactured goods, recent work has estimated that certain service sectors face relatively low trade costs and high degrees of geographical concentration across labor market areas (Gervais and Jensen, 2019). National technological shifts impacting a local economy specializing in tradable services also raise local demand for less-tradable, personal services (Mazzolari and Ragusa, 2013). If there exists spatial variation in less-tradable services,

[18] I do not test for interactions between unilateral divorce and the prevailing property division regime, as in some previous work. This is because I focus on human capital incentives, neither of which are directly affected by the property division regime. Modeling physical capital investment (assets), on the other hand, would warrant a focus on the property division regime in interaction with unilateral divorce (Voena, 2015).

for example due to local tastes or local production advantages, a national shock to these sectors exerts differential local impacts as well (Bartik and Sotherland, 2019). Thus, in the context of technological change and growing national demand for services, areas more heavily concentrated in services should experience larger shocks to demand for female employment. Additionally, Page, Schaller and Simon (2019) report variation across areas in the extent of female specialization in services. This results in further area-level variation in female-specific employment shocks driven by national service sector expansions.[19]

In addition to services, women have also historically specialized in certain manufacturing sectors, such as textiles. The same types of technological innovations that expanded services resulted in contraction of routine-task occupations (Autor and Dorn, 2013). Thus, while changes in national demand may have facilitated female employment on average, female employment likely grew far less rapidly, or even declined, in areas concentrated in female manufacturing sectors.

Theoretical predictions. An expansion of young women's labor market opportunities reduces specialization gains and compliance with "male breadwinner norms" within the marriage market, but also raises the gains associated with joint consumption economies. To the extent that specialization gains and social norms are a dominant force in marriages involving less-educated men,[20] increased labor market demand for female workers lowers such men's contribution to marriage, both in absolute terms as well as relative to the contributions of college-educated men. As predicted by the model, these forces should reduce the equilibrium number of marriages involving noncollege-educated men, and also reduce the labor market activities of such men before the marriage market.

Data. I utilize the commuting zone concept (Tolbert and Sizer, 1996) to define local labor markets. Reliable mappings from county group identifiers in Census to 1990-defined commuting zones can be constructed starting with 1970 Census samples.[21] I begin the analysis in 1980 for two reasons. First, 5 percent (rather than 1 percent) samples are available starting in 1980, which allows for a more precise computation of base-period employment concentrations by sector. Second, starting the analysis in 1980 makes 1970 data available for use in a pre-trends test to gauge internal validity of the design. Accordingly,

[19]Appendix Table A.2 presents cross-commuting-zone (CZ) variation in the 1980 degree of female specialization in service-related sectors—both in absolute terms as well as net of male specialization. For example, among the noncollege population, the share of total female workers engaged in business services ranged from .35 to .51 between the 10th and 90th percentile CZs. The share of total female workers *net* of the share of total male workers engaged in business services ranged from .26 to .41 between the 10th and 90th percentile CZs. For all service workers, the $10-90$ range was: .61 to .85 for females, and .39 to .64 for females net of males.

[20]As reported by Bertrand, Kamenica and Pan (2015), gender identity norms, dictating that the husband should be the primary breadwinner in the household, are also stronger among noncollege-educated men.

[21]Code is available on David Dorn's website: https://www.ddorn.net/data.htm.

using Census and American Community Survey samples, I create commuting-zone-level-average data for each of the following years: 1980, 1990, 2000, and 2014-17 (referred to as 2015 in the analysis). Thus, the total number of observations is 4 time periods $\cdot 722$ commuting zones in the contiguous U.S. $= 2,888$. The underlying population of men is the same as in the unilateral divorce investigation.

Regression specification. The baseline specification regresses the male outcome level (such as the share of men currently married or in the labor force), observed in commuting zone z at time t, on an index of demand for female employment D^F_{zt} and other controls:

$$\begin{aligned} \text{outcome}_{zt} \;\; = \;\; & \alpha \; + \; \beta_F D^F_{zt} \; + \; \beta_M D^M_{zt} \\ & + \; \text{commuting zone FE} \\ & + \; \text{time FE} \\ & + \; \text{error}_{zt}. \end{aligned} \tag{1.11}$$

The baseline regression controls for an index of male labor market demand D^M_{zt}: to isolate the marriage market investment channel of interest, we require that male labor market opportunities do not change when female labor market demand is varied. Commuting zone and time fixed effects are controlled, as the underlying experiment is to compare long-run changes in male behavior within commuting zones receiving different long-run shocks to female labor market demand. As in the unilateral divorce experiment, I gauge robustness by augmenting the baseline specification with time-varying regional controls. Additionally, I sometimes control for base year demographics (cell-level averages of racial, Hispanic ethnicity, nativity, education and potential experience characteristics) interacted with linear time trends.

The demand shifters D^F_{zt} and D^M_{zt} are "Bartik instruments," constructed using shift-share methodology popularized by Bartik (1991). Considering the sample of noncollege-educated young workers, I first compute commuting-zone-level shares of the workforce employed in each of 25 different industry sectors s in the 1980 base year.[22] Then, for each year t, I interact these shares with national changes in employment of noncollege young workers by industry. I repeat this process for each gender. Thus, D^G_{zt} describes the predicted change in the employment-to-population ratio of noncollege-educated individuals of gender G in commuting zone z, between 1980 and t, due to national shifts in

[22] See Appendix A.2 for a description of these sectors, which vary slightly by gender. Consistent with the above discussion, these sectors distinguish between manufacturing and services, and among broad types of each.

industry-specific demand. Equation (1.12) illustrates:

$$D_{zt}^{G} = \sum_{s} \underbrace{\left(\frac{E_{zs,1980}^{G}}{E_{z,1980}^{G}}\right)}_{\text{Base-year share, sector } s} \cdot \underbrace{\left(\frac{E_{s,t}^{G}/P_{t}^{G}}{E_{s,1980}^{G}/P_{1980}^{G}} - 1\right)}_{\text{National shift, sector } s} \tag{1.12}$$

where $E_{zs,t}^{G}$ is the total year-t employment in commuting zone z and sector s; and P_{t}^{G} is year-t total national population. (Recall that the underlying population is less-educated individuals with 0-20 years potential experience.) By construction, the instruments take the value 0 in 1980.

Recent work has clarified the identification assumptions underlying shift-share research designs (Goldsmith-Pinkham, Sorkin and Swift, 2018; Borusyak, Hull and Jaravel, 2018). In this application, we are concerned about the exogeneity of the female Bartik instrument with respect to *residual area-level movements in male outcomes after controlling for the male Bartik instrument*. The preferred identifying assumption is that female base-year industry shares are exogenous with respect to residualized movements in male outcome variables.

1.4.3 Effects of the interventions

Main results: reductions in marriage and labor-force participation. Previous literature has found a temporary increase in divorce rates upon the passage of unilateral divorce (Friedberg, 1998; Wolfers, 2006), but also a persistent reduction in marriage rates (Rasul, 2003). Appendix Figure A.1 presents an event-study graph confirming, consistent with the theoretical model, that unilateral divorce precipitated a long-run reduction in both ever- and currently-married propensities. Appendix Table A.3 presents regressions of ever-married propensities on both interventions. It reports statistically significant reductions of close to the same magnitude as those found for currently-married propensities. Thus, both interventions reduced marriage in the long-run primarily by preventing or delaying marriage formations.

Regardless, for a young man deciding whether to maintain stable employment, both whether he will marry and how long he expects to remain married have implications for his decision. Accordingly, the main outcomes of interest are current marriage and single men's labor-force participation. Results for these outcomes are presented in Table 1.2 and Table 1.3. In each table, the top panel considers the unilateral divorce intervention, while the second panel considers the female employment shock intervention—defined as a 10 percentage-point increase in the female Bartik instrument. Each table contains 6 columns corresponding to 6 different regression specifications. All specifications include the baseline controls, while each contains a unique combination of additional control variables.

Note the fourth row in the unilateral divorce panel. This row reports results of pre-trends tests. For a given specification, this is an F-test of the hypothesis that the effect of unilateral divorce on the outcome is 0 before the passage of unilateral divorce. I accomplish this by grouping the data into 5-year bins defined by event time and estimating event-study coefficients relative to the bin containing 5-to-1 years before law passage. (See Appendix Figure A.1 for an illustration of the binning.) The pre-trends test, then, tests whether the coefficients corresponding to "10-to-6" and "at least 11" years before law passage are jointly zero. In each table, only in one case can we reject the hypothesis of no pre-trend at the 5 percent level. Most estimated p-values are above 20 percent. These results support the internal validity of the research design.

Looking at Table 1.2, estimated effects of the interventions on current marriage are negative and statistically significant. The average of the 6 point estimates is $-1.9p.p.$ for unilateral divorce and $-1.8p.p.$ for female Bartik shocks. Looking at Table 1.3, estimated effects for labor-force participation are also negative, and generally reach statistical significance at conventional levels. The average estimate is $-1.4p.p.$ for unilateral divorce and $-1.2p.p.$ for female Bartik shocks. Considering that base rates of non-participation in this population of men are $13 - 14$ percent (third row of each panel), each intervention is estimated to cause roughly a 10% increase in non-participation behavior. This pattern of results lends credence to the marriage market investment channel emphasized by the theory. As changes to the legal structure of marriage and women's labor market opportunities threatened efficient cooperation within marriage and incentivized young women to pursue their own careers, the value of marital specialization declined. This led fewer marriages to form and fewer noncollege-educated men to invest in stable employment.

One alternative interpretation of the data is that the interventions impacted the participation rate of singles simply by changing the composition of whom is single, rather than because of lower marriage market investment. This interpretation, however, does not seem plausible. The conventional wisdom is that married men are positively selected on labor market skill and motivation (Antonovics and Town, 2004): if the interventions simply led some men who would have been married to become single, such a development would presumably *raise*, not lower, the participation rate of singles. A calculation presented in Appendix A.4 shows that, for selection effects to be completely responsible for the unilateral divorce results, the men induced to remain single by unilateral divorce would require a non-participation rate of roughly 42.5 percent. This is around *triple* the pre-reform rate of non-participation among singles.[23] Moreover, if composition effects were dominant, one

[23]A similar calculation can be done for the female Bartik shock results, following the formula presented in

would expect to find a null effect on participation in the overall population of men. Appendix Table A.4 reports participation estimates for all men and continues to find negative overall effects.

Testing the identification assumption: small effects on wages. To attribute observed effects of these marriage market interventions to the marriage market channel of interest, we must be sure that men's labor market opportunities are orthogonal to intervention exposure. Failure of this condition can be interpreted as a failure of the identification assumption.

While it is impossible to directly test the identification assumption, the current context admits a reasonably informative test. This test involves investigating wages as an outcome. If the given intervention is found to "cause" substantial reductions in both employment and wage rates, then it is probable that intervention exposure is correlated with a negative shock to male labor demand. On the other hand, if male labor demand did not change with the intervention, we should not find large negative wage effects.[24]

Table 1.4 reports estimated effects of the interventions on log weekly wages of single men.[25] As in Table 1.2, the top panel considers the unilateral divorce intervention, while the second panel considers the female employment shock intervention. In the second case, log weekly wages are measured as commuting-zone-year averages, where individual observations are weighted by weeks worked. For reference, the effect of a 10 percentage-point male Bartik shock on wages is precisely estimated at 9 to 11 percent across all specifications. Against this benchmark, the estimated wage effects are small, and are always statistically insignificant. Effects of unilateral divorce range from -2.4 to 0.5 percent across specifications, and effects of a female Bartik shock range from -1.6 to 1.9 percent. These findings suggest that the interventions are uncorrelated with non-trivial changes to male labor demand.

Relevance for secular decline in male employment. Table 1.5 demonstrates the importance of the two interventions in explaining secular decline in employment of young, noncollege-educated men. Consider the first column. According to the average of the 6 estimates reported in Appendix Table A.4, unilateral divorce is associated with a $.65$

Appendix A.4.

[24]A marriage-market-caused reduction in employment at a given set of opportunities implies an inward shift of the labor supply curve: if wages are flexible and capital is inflexible, we should expect an increase in wages. But if capital also adjusts in the medium-to-long run time frame captured by the regressions (10-year periods), the initial wage impact of the supply shift may be diffused. In addition, if the marriage market effect causes less employment and hence less accumulation of labor market skills, we might expect a decline in wages. The overall expected effect of these marriage market interventions on wages is thus ambiguous. That said, a substantial negative effect likely implies a failure of the identification assumption.

[25]Results for hourly wages are similar. See Appendix A.2 for information on wage processing.

percentage-point decrease in the labor-force participation rate of all noncollege-educated men with $0-15$ years potential experience. Between 1965 and 1980, around 54 percent of this population was exposed to unilateral divorce,[26] implying that unilateral divorce caused a $0.65 \cdot 0.54 = 0.35$ percentage-point decline in participation. This amounts to 23% of the observed 1.5 percentage-point decline that occurred during the period. An analogous exercise for the female Bartik shock intervention finds that female employment shocks can account for $1.89 p.p.$, or 30%, of the observed $6.3 p.p.$ decline in labor-force participation observed between 1980 and 2015.[27]

Because the two interventions happened in non-overlapping time periods, we can compute the joint effect of the two interventions simply by adding the two individual effects together. According to rows 3 and 4 of Table 1.5, this accounting exercise yields $\frac{-0.35-1.89}{-1.5-6.3} = .287$. That is, the two interventions can account for 28.7 percent of the total observed 1965-2015 decline in labor-force participation of young noncollege men.

1.5 The Marriage Market Amplifies Men's Labor Market Shocks: Evidence from a Structural Model

The results of the last section suggest that young men's employment responds to a decline in marriage market prospects. In this section, I examine the effects on young men's employment of a decline in labor market opportunities. By reducing marriage market prospects, such a shock might reduce employment through a *marriage market effect*, in addition to the standard *labor market effect* that has been emphasized by the labor demand literature.

Figure 1.8 illustrates these two effects. Returning to the theoretical framework of Section 1.3, consider a reduction in noncollege men's earning potential. This is driven either by a reduction in the initial wage offer $w_{NC,0}$ or a reduction in returns to experience $\hat{r}$. Given a positive uncompensated employment elasticity, such a shock reduces employment via standard labor/leisure substitution.[28] This *labor market effect* is represented as a downward shift of the average man's marginal benefit curve (and a corresponding leftward shift of the steep blue locus in the bottom graph). But the shock also reduces the economic surplus ($S(w_{NC,1}^{\text{married;f}}, n_{NC,0}^{\text{married;f}}; f)$) generated by a noncollege marriage, both in absolute terms

[26] 3 states had unilateral divorce regimes in place before 1965: Alaska, Oklahoma and New Mexico. Around 30 states adopted unilateral divorce regimes between 1965 and 1980.

[27] The average commuting zone experienced a $14 p.p.$ increase in the female Bartik instrument: hence the 1.40 average exposure to a $10 p.p.$ shock reported in the table.

[28] When returns to experience $\hat{r}$ decline, it is possible to observe a reduction in employment even if the uncompensated employment elasticity is zero. See Imai and Keane (2004).

and relative to that generated by a college marriage. This leads to fewer marriages forming at a lower surplus share ($\theta_m(NC)$) for noncollege men. These dynamics shift the orange locus in the bottom graph downward, leading to a further downward shift of the marginal benefit curve. This further downward shift is the *marriage market effect*.

The goal of this section is to separately quantify these marriage market and labor market effects, and thus apprehend an additional implication of the marriage market channel for young men's employment. To do so, I specify an empirical model based on the theoretical framework and structurally estimate its parameters. Then, I impose a negative shock to noncollege men's lifetime earning potential and solve for the new model equilibrium. The structure of the model enables a decomposition of equilibrium responses into labor market versus marriage market channels.

1.5.1 Empirical setup

I consider the investment-and-matching problem laid out in Section 1.3. Agents come in two education types: noncollege-educated (NC) and college-educated (C). The shares c^M of men and c^F of women are college-educated. There are equal total numbers of men and women to be matched in period 1, after men have made their labor market investments. College men are assumed always to be in the labor force. Thus, I abstract from their initial wage offer and wage growth process: they arrive in the marriage market with wage $w_{C,1}$. Noncollege men, on the other hand, receive a period-0 wage offer of $w_{NC,0}$ and face Mincerian wage equation (1.1) when deciding how much to work in period 0. Moreover, noncollege men in period 0 receive an amount $\hat{y}$ in non-labor income. This is consistent with the fact that many out-of-work young men live with and rely on family members for income support (Binder and Bound, 2019). I do not model heterogeneity in non-labor income receipt.

Empirical specification of utilities. Single men face the following direct utility function in period t:

$$U(c_{mt}, n_{mt}; \emptyset) = \ln(c_{mt}) - \lambda \cdot \frac{n_{mt}^{1+\gamma}}{1+\gamma} \tag{1.13}$$

where λ and γ are parameters to be estimated. This function belongs to the class of separable utility functions assumed in Section 1.3. Now, define $\underline{w} = \exp(\ln(w_{NC,0}) - \delta)$. This is the wage offer in period 1 realized by a man who does not work at all in period 0. Normalizing the utility of his hypothetical man to 0, it is easy to show that the normalized indirect utility function for single men in period 1 is simply $\tilde{V}(w_{m1}) = V(w_{m1}; \emptyset) - V(\underline{w}; \emptyset) = \ln(w_{m1}) - \ln(\underline{w})$. (See Appendix A.5.)

The economic surplus for a marriage between man m and woman f, for $m, f \in \{NC, C\}$, is specified as

$$S(w_{m1}, n_{m0}; m, f) = \alpha_f \cdot (1 + \mu_m n_{m0}) \cdot (\ln(w_{m1}) - \ln(\underline{w})) - \xi(m; f) \tag{1.14}$$

where α_f, μ_m and $\xi(m; f)$ are parameters to be estimated. I assume $\mu_C = 0$ and $\mu_{NC} = \mu$; thus, μ is a signaling parameter that captures the degree to which noncollege men can compensate for their lower earning potential by working hard in period 0. To the extent that noncollege men have lower marriage propensities than college men, μ is lower than the college wage premium.

The function $\xi(m; f)$ is specified as follows:

$$\xi(m; f) = \begin{cases} 0 & m = f \\ \xi & m \neq f \end{cases} . \tag{1.15}$$

That is, the economic surplus declines by ξ in a marriage between two individuals of different education types. This captures assortative matching behavior on the marriage market.

As specified in Section 1.3, an individual's gain in a marriage between m and f equals the individual's share in the economic surplus, plus a non-economic component capturing idiosyncratic tastes. Each side of the market contains two education types of agents: thus, in equilibrium, there is a 4-vector of male surplus shares that stabilizes the marriage market. These shares, $\theta_{\mathbf{m}}(\mathbf{m}; \mathbf{f})$, are not parameters to be estimated, but rather are endogenous objects that result from the equilibrium conditions of the model. Idiosyncratic tastes (for singlehood, for marriage to a noncollege individual, and for marriage to a college individual) follow a Type I Extreme Value distribution, with scaling parameter σ_ϵ to be estimated.

It is instructive at this point to compare the period-1 marginal benefits of period-0 employment between to-be-single and to-be-married men. It is trivial to show that the marginal benefits realized by a to-be-single man are $\beta\hat{r}$. (See Appendix A.5.) A noncollege man marrying woman f receives these marginal benefits plus $\theta_m(NC; f)$ of the total marginal benefits to the marriage. This yields the following expression:

$$\begin{aligned} & \beta\hat{r} + \beta\theta_m(NC; f) \cdot \frac{dS}{dn_0^{\text{married;f}}} \\ = & \underbrace{\beta\hat{r}}_{\text{Marginal \textbf{labor market} benefits}} + \underbrace{\beta\theta_m(NC; f) \cdot \alpha_f \hat{r}(1 + 2\mu n_{m0})}_{\text{Marginal \textbf{marriage market} benefits}} . \end{aligned} \tag{1.16}$$

Given a positive surplus share and positive parameter values, it is clear that to-be-married men enjoy higher marginal benefits to period-0 employment than to-be-single men.

Preset parameters. The baseline model aims to replicate the environment of the early 1980s in the United States. A number of model parameters are preset, consistent with this baseline, based on external sources of data. First, the model abstracts from education decisions, so I preset the college shares c^M and c^F based on data reported by Autor and Wasserman (2013). I set the earning potential of college men in period 1, $w_{C,1}$, based on data reported by Binder and Bound (2019). The initial wage offer for noncollege and wage growth process—$w_{NC,0}$, δ and $\hat{r}$—are set based on estimates in the NLSY79 data reported by Braga (2018) and to ensure a college wage premium of 0.4 (its value in the early 80s inferred from Binder and Bound, 2019). Finally, nonlabor income $\hat{y}$ is inferred based on data reported by Binder and Bound (2019). Appendix Table A.5 records the preset parameters, and further information on their calibration can be found in Appendix A.5.

This leaves 7 parameters to estimate: noncollege men's preferences for leisure (λ and γ), marital value scalars for each type of woman (α_{NC} and α_C), noncollege men's signal value of working hard (μ), the penalty for marrying the opposite education type (ξ) and the idiosyncratic taste scale parameter (σ_ϵ).

Identification. I estimate these 7 parameters using the Generalized Method of Moments (Hansen, 1982). I choose 7 informative moments, based on 1980s U.S. data, for identification. Thus, the model is exactly identified: the GMM estimates are the parameter choices that make model-generated behavior exactly coincide with the 7 observed behaviors in the data.

I provide a heuristic argument that the following moments identify the model. To start, note that the only parameters that govern the employment behavior of to-be-single men are the preference parameters λ and γ. The parameter λ is identified by targeting the labor-force participation rate of "always-single" noncollege men in the NLSY79, over ages 21-31 (recall Section 1.3.3). The parameter γ is related to noncollege men's willingness to substitute between consumption and leisure. Thus, I identify γ by targeting an uncompensated labor supply elasticity. I produce this elasticity within the model by measuring the responsiveness of period-0 employment of to-be-single men to a 10% decline in $w_{NC,0}$.[29] Consistent with small uncompensated elasticities estimated in the literature, I target a value of 0.1.[30]

[29]Note that the structure of direct utility in the model forces the responsiveness of period-0 employment to $w_{NC,0}$ to be 0—that is, an exact canceling of the income and substitution effects—unless non-labor income is present. The existence of non-labor income in period 0 thus not only matches reality, but also generalizes the model.

[30]Important to the interpretation is that this 0.1 value captures the effect of a persistent wage change while holding marriage market forces constant. This is a reasonable interpretation, since microeconometric studies in the literature tend to leverage changes in the income tax code that *implicitly hold family structure*

Four remaining parameters are identified by marriage outcomes in the data. Suppose for the moment that the scale parameter on idiosyncratic tastes (σ_ϵ) is unity: as is standard in discrete choice models, the following argument identifies payoff parameters *relative* to this scale. Consider a woman of education type f. All else constant, the larger α_f is, the more marriages will form involving type-f women. Thus, I identify α_f for $f \in \{NC, C\}$ by targeting the matching statistic

$$\Pi_f = \frac{\text{marriages}^{NC,f} + \text{marriages}^{C,f}}{\sqrt{\left(\text{singleMen}^{NC} + \text{singleMen}^{C}\right) \cdot \text{singleWomen}^{f}}}.$$

This is the ratio of the number of marriages forming involving type f to the geometric average number of men and type-f women who remain unmarried. See Choo and Siow (2006) for further discussion of this statistic. The signaling parameter for noncollege men, μ, is identified by the observed *gap* in marriage propensities between college and noncollege men. The larger μ is, the smaller this gap will be in the model. The parameter ξ is identified by observed *assortative matching* behavior in the marriage market. The larger ξ is, the stronger will be the degree of positive assortative matching on education. I capture assortative matching with the correlation coefficient between married couples' college statuses.

The last parameter—σ_ϵ—is identified by the *difference in period-0 employment behavior between to-be-single and to-be-married men*. (The observed difference is reported in column 4 of Table 1.1.) Why should the dispersion of idiosyncratic tastes for marriage be at all related to employment behavior? The answer is relatively straightforward. Note that lifetime utility in this model is denominated in period-0 log consumption units. Noncollege men in period 0 convert expected period-1 marginal benefits into this scale when deciding how hard to work. But when deciding whom to marry in period 1, scale does not matter: one simply picks the option associated with the highest period-1 utility. Thus, σ_ϵ is chosen to convert the marriage market parameters that influence to-be-marrieds' period-0 employment choice—α and μ—into values consistent with this choice.

The following example is instructive. Suppose there is only one type of woman, one type of man, and μ is zero. Consider one world in which $3/4$ of men end up married and to-be-marrieds work only *slightly more* than to-be-singles. Consider another world in which $3/4$ of men marry and to-be-marrieds work *much more* than to-be-singles. Since marriage outcomes are the same in both worlds, α/σ_ϵ must be the same in both worlds.

---

constant. For example, the labor supply of married individuals in response to a change in the rate schedule facing married couples—or the labor supply of single individuals in response to an expansion of Earned Income Tax Credits for single individuals—is measured and cited as an uncompensated elasticity.

But, as stipulated by equation (1.16), σ_ϵ (and thus α) must be higher in the second world in order to generate the higher employment of to-be-married men. Thus, we see that α/σ_ϵ is identified by marriage outcomes and σ_ϵ is identified by the employment behavior of to-be-married men.

Baseline estimation. The presence of a pre-match investment decision makes the present econometric framework more complex than the standard transferable-utility matching framework (used in Choo and Siow, 2006; Chiappori, Salanié and Weiss, 2017). This is because an interdependence exists between the surplus shares that stabilize the marriage market and the overall surpluses—and hence marital matching patterns—generated in equilibrium. To see this, recall that surplus shares are inputs into the period-0 employment decision of to-be-married men (by first order condition 1.7), yet this employment decision affects the overall marital surplus and marriage formation (by equations 1.4 and 1.5). Such an interdependence does not arise in the former models. To handle this complexity, I apply the algorithm designed by Galichon, Kominers and Weber (2018) for matching models with *imperfectly* transferable utility. In such models, interdependence between surplus shares and total surpluses arises through spousal bargaining processes (e.g. Reynoso, 2019; Gayle and Shephard, 2019). Though the structure of my model differs from these, the interdependence arises for the same reason: because an equilibrium action taken by a prospective spouse affects both the overall surplus and its distribution. Following is a description of the GMM estimation procedure:

1. Guess a parameter vector.

2. Guess $n_{NC,0}^{\text{married;f}} = 1$. This is the period-$0$ employment rate of noncollege men who will marry female education type f.

3. Given the guess of $n_{NC,0}^{\text{married;f}}$, use the fixed-point algorithm of Galichon, Kominers and Weber (2018) to solve for the surplus shares $\theta_m(m;f)$ that stabilize the marriage market.

4. Use first-order condition (1.7) to solve for the values of $n_{NC,0}^{\text{married;f}}$ consistent with these surplus shares $\theta_m(NC;f)$. This becomes the new guess of $n_{NC,0}^{\text{married;f}}$.

5. Repeat steps 3 and 4 until the new guess of $n_{NC,0}^{\text{married;f}}$ is the same as the old guess.

6. The model is now at equilibrium. Compute the normed distance between the seven moments predicted by the model and those observed in the data.

7. Repeat steps 1-6 (using a computerized search algorithm) until this distance approaches zero.

Table 1.6 presents the baseline estimation results and confirms that the model replicates the 1980s environment of interest.

1.5.2 Quantifying the marriage market multiplier

With the baseline model, we can now quantify the marriage market channel associated with a shock to noncollege men's earning potential. This quantification is straightforward. Because the employment behavior of to-be-single men is not affected by marriage market forces, the response of to-be-single men to the shock pins down the labor market channel. The response of *all* men to the shock, by definition, includes both labor and marriage market channels. The size of the marriage market channel is simply the difference between the response of all men and the response of to-be-single men.

To see this, consider the following decomposition. For illustrative purposes, suppose there is just one type of woman in the marriage market. Define $\eta_{NC,0}^{\text{married}}$ as the difference in period-0 employment rates between to-be-married men and always single men. This yields the following expression for the noncollege male employment rate in period-0:

$$n_{NC,0} = n_{NC,0}^{\text{single}} + s_{NC}^{\text{married}} \cdot \eta_{NC,0}^{\text{married}},$$

where s_{NC}^{married} is the share of noncollege men who marry. Then, the change in the noncollege male employment rate due to a wage shock is:

$$\Delta n_{NC,0} = \underbrace{\Delta n_{NC,0}^{\text{single}}}_{\textbf{labor market}\text{ channel}} + \underbrace{s_{NC}^{\text{married}} \cdot \Delta \eta_{NC,0}^{\text{married}} + \Delta s_{NC}^{\text{married}} \cdot \eta_{NC,0}^{\text{married}}}_{\textbf{marriage market}\text{ channel}}. \tag{1.17}$$

Alternatively, the *marriage market multiplier* can be computed as the *ratio* of the response of all men to the response of to-be-single men: $\Delta n_{NC,0}/\Delta n_{NC,0}^{\text{single}}$. This describes how much the labor market shock is magnified by the endogenous response of the marriage market.

The model allows for two different types of labor market shocks to be simulated. One is a *downward shift* of the life-cycle wage profile, driven by a reduction in the initial wage offer $w_{NC,0}$. Another is a *flattening* of the wage profile, driven by a reduction in returns to experience $\hat{r}$. One can interpret a reduction in $\hat{r}$ as reflecting fewer opportunities for career growth among stably-employed men, or an increase in labor market imperfections that prevent stable job-holding. For example, if increased automation and import competition from China raises the probability of job displacement—and job displacement results in lost returns to tenure and involuntary depreciation of skills—then a given amount of time spent working or looking for work will result in lower wage growth. I remain agnostic about the underlying cause of a decline in $\hat{r}$.

It is important to note that such a simulation is an out-of-sample exercise. The aggregate responsiveness of noncollege men's employment and marriage outcomes to wage profile shocks was not targeted in the estimation. Hence the simulation exercise can be seen as a test of the theoretical model's predictions. In the next subsection I show that the simulated responses match up rather well with some recent empirical estimates in the literature. Both sets of responses lend credence to the marriage market channel emphasized by the theory.

Table 1.7 reports the equilibrium effects of a 10% reduction in noncollege men's earning potential, driven by a combination of a reduction in initial wage offers ($w_{NC,0}$) and a reduction in wage growth ($\hat{r}$).[31] All other agents' labor market opportunities are held constant. Looking at the first row of Panel B, the shock reduced the share of noncollege men that married in period 1 by 3.9 percentage points. This was driven by reductions both in noncollege men marrying noncollege women and in noncollege men marrying college women, although the former effect dominated. (This is not surprising given that most noncollege husbands were married to noncollege wives at baseline.) The second row shows that the shock lowered the economic surplus shares claimed by noncollege men—particularly in marriages with college women. The third row reports that the shock reduced the employment of to-be-singles by 1.3 percentage points and aggregate employment by 2.3 percentage points. Thus, had marriage prospects remained constant, the effect of the shock on employment would have been a full percentage point—or $1/2.3 = 43.5\%$—smaller. In other words, the marriage market multiplied the original employment effect of the labor market shock by a factor of $2.3/1.3 = 1.77$.

In Appendix Table A.6, I consider different 5 different shocks to noncollege men's wage profiles. All shocks reduce present-discounted lifetime wages by 10%, but they vary in their emphasis on reduced returns to experience ($\Delta\hat{r}$) versus reduced initial wage offers ($\Delta w_{NC,0}$). I find that a shock that places more emphasis on reduced returns to experience results in a smaller marriage multiplier. Specifically, as reduced returns to experience go from accounting for none of the overall wage shock to all of it, the marriage market multiplier falls from around 2.0 to 1.5. However, such a shock creates a larger period-0 employment response, as both labor market and marriage market effects rise in absolute terms.[32] Thus, labor market shocks that reduce the wage growth of noncollege men play a

[31] Specifically, I reduce initial wage offers by 9.15% and wage growth by $1.5 p.p.$. The importance of period 1 relative to period 0, β, is calibrated to 1.33 (see Appendix A.5 for details on this choice). This implies that the present-discounted lifetime wage reduction is $9.15\% \cdot \frac{1}{2.33} + (9.15 + 1.5)\% \cdot \frac{1.33}{2.33} = 10\%$.

[32] See Elsby and Shapiro (2012) for a model of the labor market effect of a reduction in the experience-

particularly important role in their falling marriage market value, and in marriage-market-induced withdrawal from the workforce.

1.5.3 Some external validation

Recent work by Autor, Dorn and Hanson (2019) has revealed economically significant relationships among labor market demand, young men's employment and marriage. These authors estimated that rising import competition from China in male-dominated sectors led to lower marriage and a greater share of men living without children in the affected communities. They additionally found an increase in the shares of young men not employed and not in the labor force. Table A3 of their study reports that a 1-unit male trade shock, over the period of a decade and holding constant female trade shocks, is associated with a $\$3,737$ loss in annual earnings for the median-earning man, a $3.06 p.p.$ (or 3.7%) decline in male employment, and a $1.97 p.p.$ (or 2.2%) decline in male labor-force participation. Under an assumption described in Appendix A.5, I convert this earnings loss into an 8.5% decline in hourly wages. The elasticities implied by these numbers are 0.41 for employment and 0.26 for participation—larger than nearly all uncompensated labor supply elasticity estimates reported in the literature.

In Table 1.8, I compare the model-simulated effects of a 10% male wage shock to the estimated effects of a 1-unit male trade shock. The effects are quite similar. Scaling the trade shock estimates by a factor of $10/8.5$, we see that a male trade shock that induces a 10% decline in wages results in a $1.97 \cdot 10/8.5 = 2.32 p.p.$ decline in male labor-force participation, while the model predicts a $2.30 p.p.$ decline. The trade shock results in a $3.57 \cdot 10/8.5 = 4.20 p.p.$ decline in marriage, while the model predicts a $3.93 p.p.$ decline. Perhaps most important, the trade shock estimates imply a young male participation elasticity of 0.26, while the simulation estimates imply an elasticity of 0.25. The parity of these results suggests that a substantial portion of the young male employment effects of rising import competition from China was generated by marriage market forces. It also suggests that the model's predictions have some degree of external validity, at least with respect to the 1990-2014 period examined by Autor, Dorn and Hanson (2019).

1.6 Conclusion

This paper identifies the changing marriage market as a partial explanation for the concerning decline in employment of young men in the United States since the 1960s. To

earnings profile.

illustrate why the marriage market might incentivize young men to seek employment, I present a marital matching framework in which men take a pre-match employment decision. When gains from traditional gender role specialization drive the economic benefits of marriage, the framework predicts that the men who work the hardest before marriage enter the highest-surplus marriages and claim the largest shares of these surpluses. Such men thus earn a **marriage market return**—in addition to the usual labor market return—from investing in stable employment. Changes in the marriage market that affect the economic benefits of marriage affect this return and, consequently, young men's equilibrium level of employment.

I develop and test two implications of this marriage market channel. First, a reduction in the gains from specialization makes noncollege men less attractive as marriage partners, and hence lowers the marriage market return they earn from investing in employment. I test this implication by leveraging two specific marriage market interventions that plausibly reduced gains from specialization: the adoption of unilateral divorce legislation, and shocks to young women's employment opportunities (holding men's opportunities constant). I estimate that these interventions generated 29% of the observed 1965-2015 decline in labor-force participation by young noncollege men. Second, a reduction in noncollege men's labor market opportunities not only reduces labor market returns to employment, but also reduces marriage market returns through its effect on gains from specialization. Simulations of a calibrated structural model find that the "indirect" effect of a 10% reduction in noncollege men's wages on employment, operating through the marriage market, is roughly as large as the "direct" labor market effect. These results establish the marriage market channel as a quantitatively important driver of secular decline young men's employment.

These findings enhance our understanding of secular decline in prime-age men's employment. Observers of this trend have tended to cite falling labor demand as its primary cause. This explanation is problematic because it requires a larger responsiveness of men's employment to wages than is revealed by the historical record (and in quasi-experimental studies of wage changes generated by tax policy). An explanation that features interaction between the labor market and the marriage market helps resolve this problem, particularly for younger men, given that older displaced workers have greater access to savings and disability benefits. In addition, the marriage market channel emphasized here likely interacts with the rising fraction of noncollege men with prison records. (See Charles and Luoh, 2010 and Schneider, Harknett and Stimpson (2018) for arguments that mass incarceration has disrupted family formation). Young men released from prison may invest less in employment in part because they face lower marriage prospects—especially if discrimination in the labor market makes the securing of stable employment harder for such men.

This paper expands a growing body of theory and evidence linking marriage market forces to human capital investments. The ground seems fertile for seeding further expansions. For example, health can be thought of as human capital (Grossman, 1972). An important narrative posits that adverse labor demand shocks have increased drug-abuse-related morbidity and mortality among noncollege men (Case and Deaton, 2015; Coile and Duggan, 2019; Autor, Dorn and Hanson, 2019). Such a narrative is potentially too crude: we may have much to learn from incorporating interactions between labor demand and marriage prospects into the study of risky health behavior.[33] Investments in housing have also recently been linked to marriage formation (Lafortune and Low, 2017), but a framework that considers the joint determination of both outcomes is yet to be seen.

Future research should consider extending the modeling framework used in this paper. The reduced-form results shown here on unilateral divorce causally demonstrate that a change in contracting technology over an uncertain future affects pre-marital investments. This suggests a fruitful extension of the matching environment to allow for uncertainty, limited commitment, and unilateral divorce. Such an extension also pairs with modeling the choices of individuals to enter non-marital cohabitation—an arrangement that features low exit costs but also low commitment. Extending econometric matching models in these ways will likely offer richer quantitative predictions on the forces determining human capital investments and the socioeconomic returns such investments receive.

[33]For example, Duncan, Wilkerson and England (2006) reported a decline in male drug and alcohol use upon transition into marriage.

1.7 Figures

Figure 1.1: Labor-force participation rates by education status: Men aged 25-34

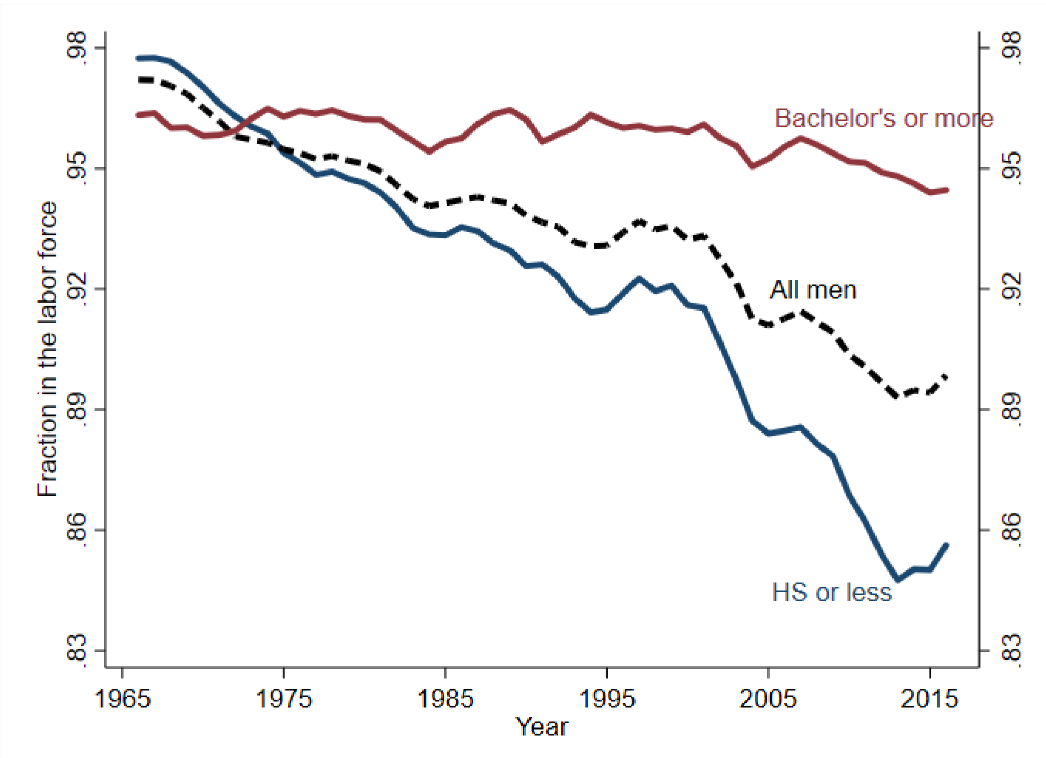

Notes: March CPS data, 1964-2017. Overall sample is white or black non-Hispanic men aged 25-34. Individuals in the “HS or less” sub-sample completed at most one year of college education. Individuals in the “BA or more” sub-sample completed at least 4 years of college education. Graph presents 3-year-centered moving averages.

Figure 1.2: Average wage and participation trends in modern U.S. history:
Men without college, ages 25-34

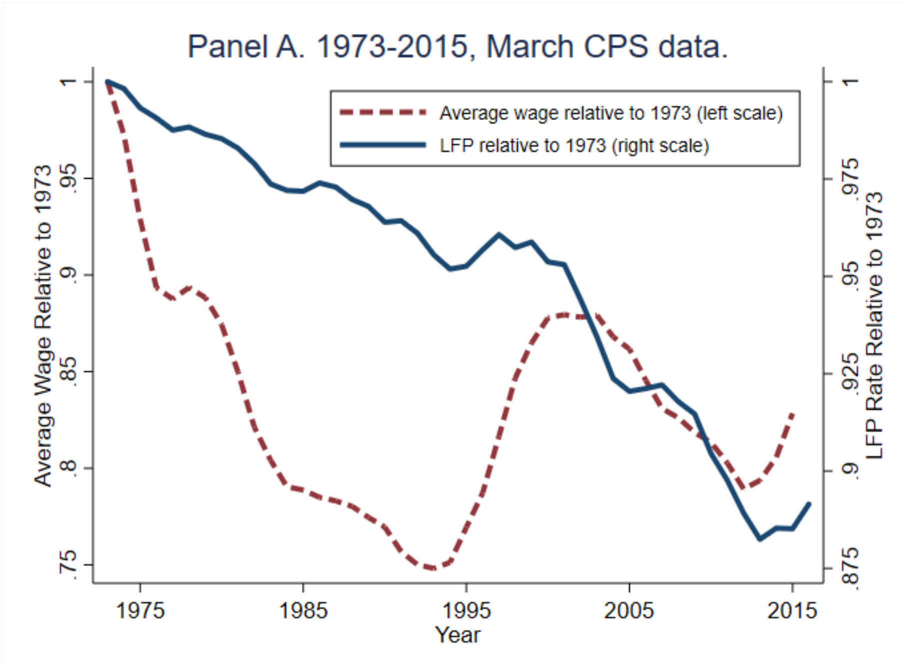

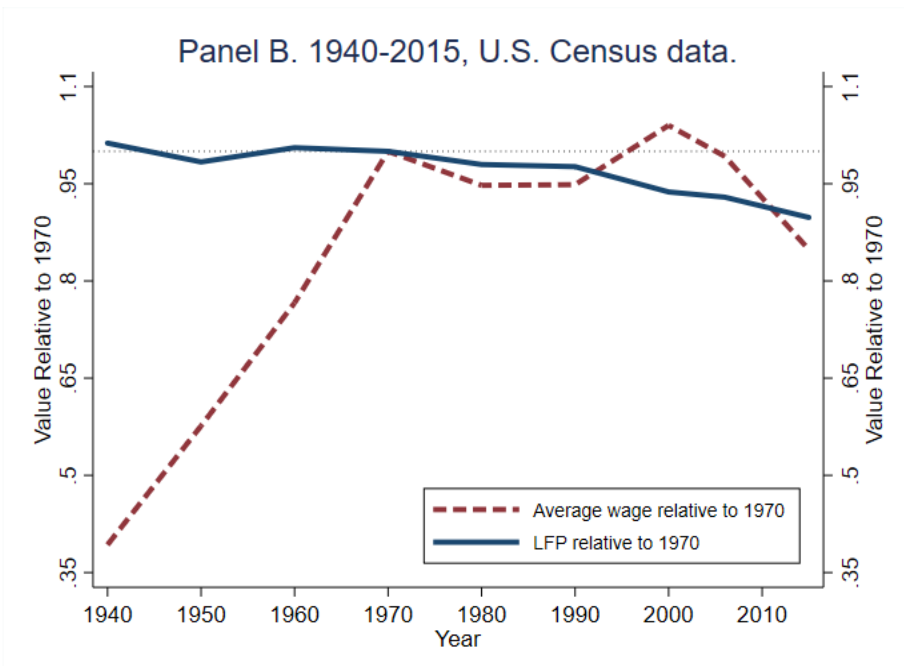

Notes: Sample is white or black non-Hispanic men, aged 25-34, with at most one year of completed college. Reported values are relative to the base year: for example, 0.9 means a 10% decline since the base year. Panel A: March CPS data, 1973-2016. Graph displays 3-year moving averages. Average hourly earnings from wage-and-salary income are constructed in a manner described in Appendix A.2. Wage offers for non-workers are imputed using the method of Juhn, Murphy and Topel (1991), described in Appendix A.2. Wage trends are not sensitive to whether non-workers are included or excluded, or to whether business income is included. Panel B: U.S. Census data, 1940-2000. American Community Survey data, 2005-2007 and 2014-2016. Average weekly earnings from wage-and-salary income are constructed in a manner described in Appendix A.2. Wage trends are not sensitive to whether business income is included.

Figure 1.3: Employment and marriage propensities across birth cohorts of noncollege-educated men

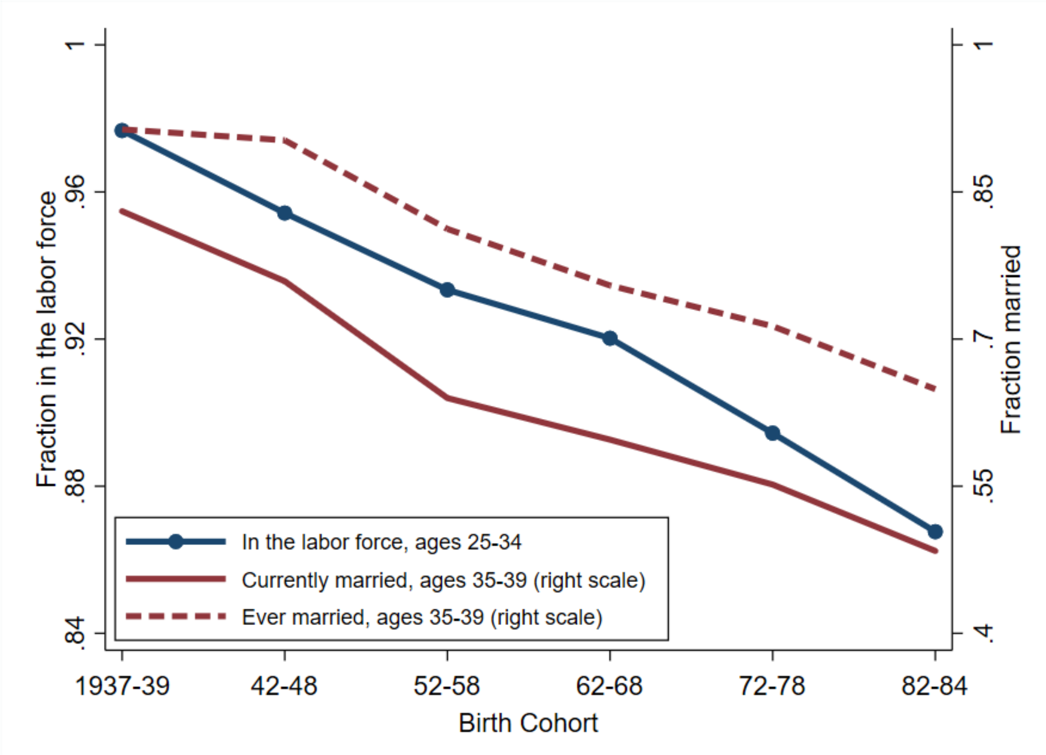

Notes: March CPS data, 1962-2017. Sample consists of men with at most one year of completed college and who are not currently enrolled in school. Age 35-39 marriage values for the 1982-84 cohort are not yet available: these values are predicted based on age 30-33 values in a manner described in Appendix A.2.

Figure 1.4: 2-period model of men's pre-marital employment and marriage

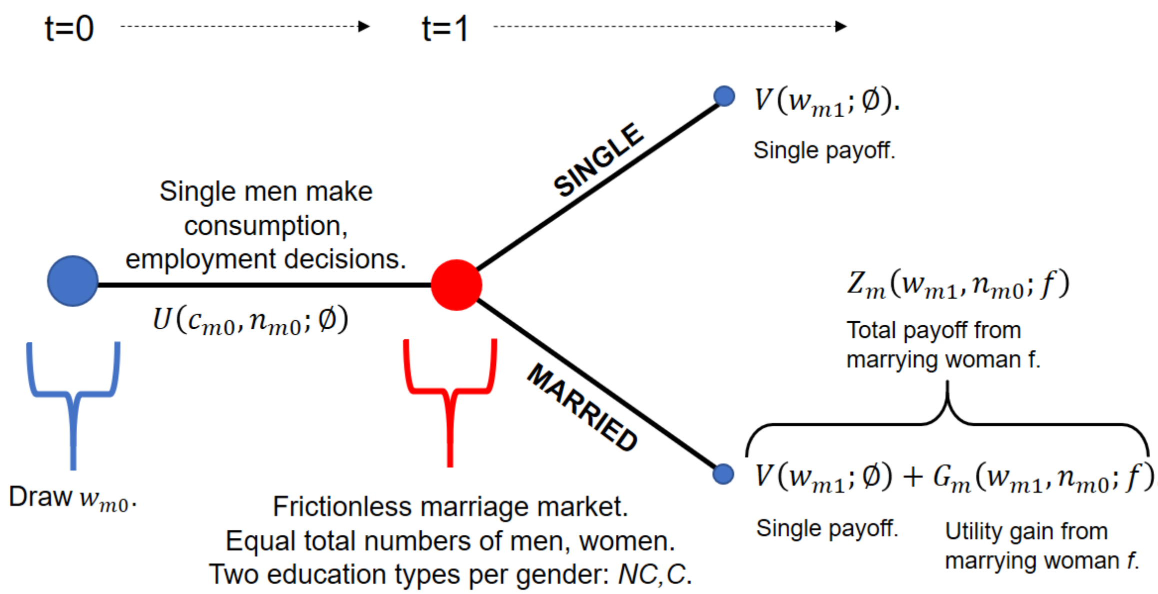

Figure 1.5: Equilibrium employment and marriage for noncollege-educated men

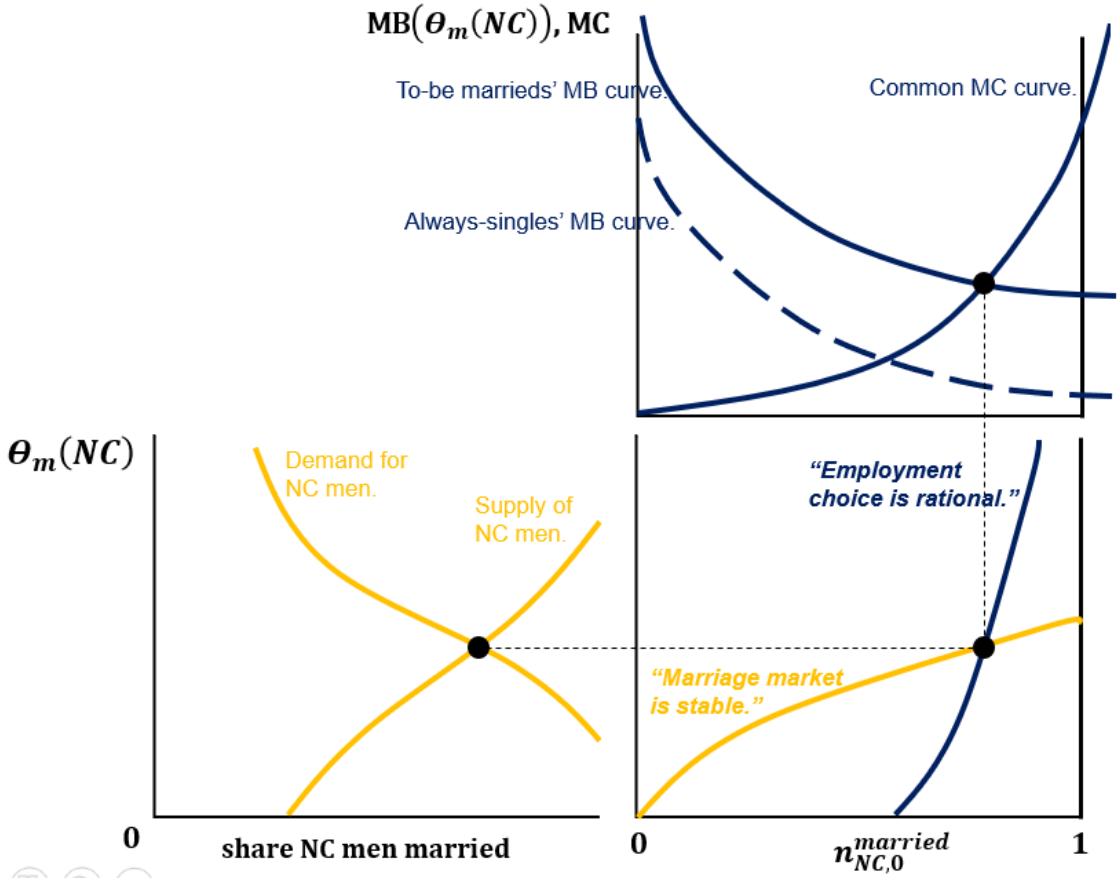

Notes: See discussion in Section 1.3.2

Figure 1.6: Comparative static: Reduction in marital surplus added by noncollege men

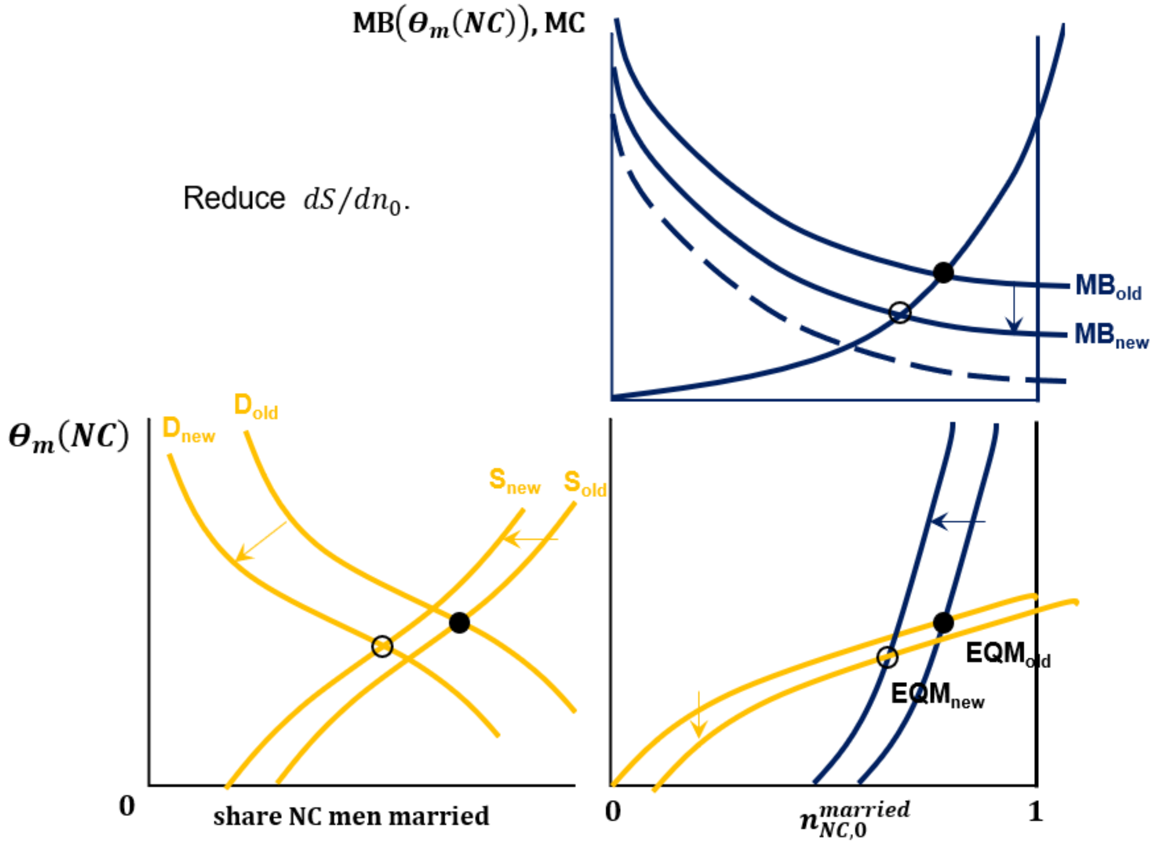

Notes: See discussion in Section 1.3.2

Figure 1.7: Life-cycle employment behavior among NLSY79 men without college

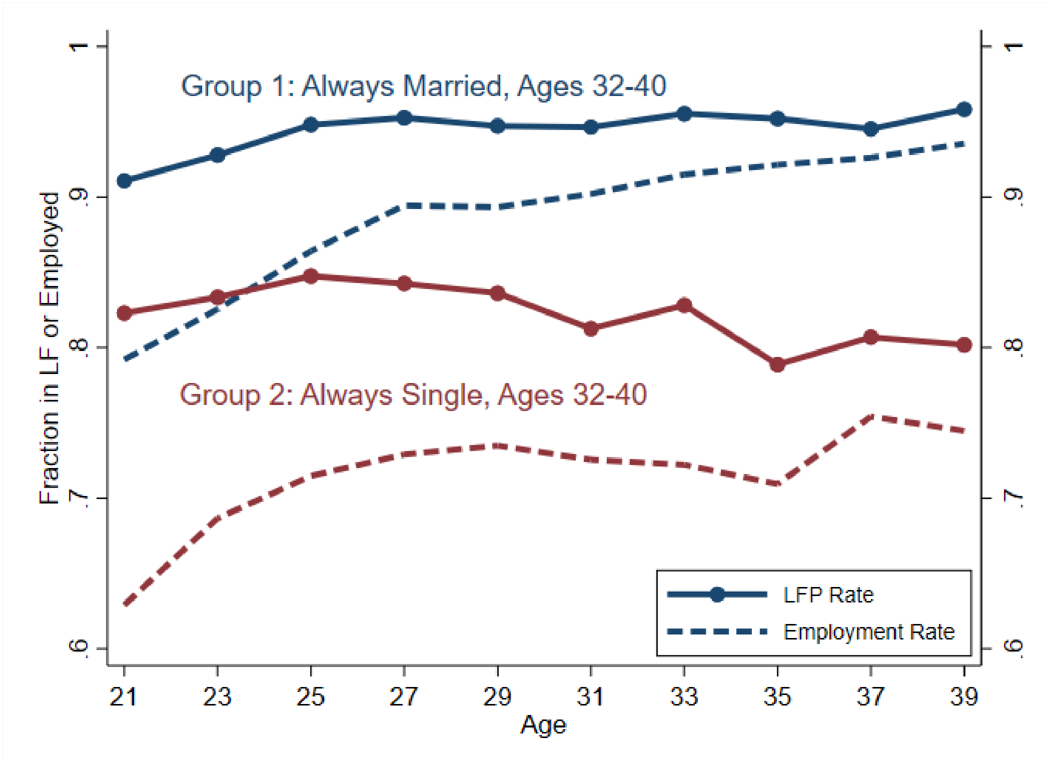

Notes: NLSY79 non-Hispanic men with high-school education or less. Sample is divided into two groups, as described in the figure. For group 2, "single" means neither married nor cohabiting with an unmarried partner. Non-marital cohabitations were identified using the method of Oppenheimer (2003). For each group, age profiles of labor-force participation and employment rates are plotted. Each point records a 2-year average of group behavior: for example, the point at age 21 records average behavior over ages 21 and 22. 659 men are in group 1 and 289 are in group 2.

Figure 1.8: Comparative static: Reduction in noncollege men's wage offers

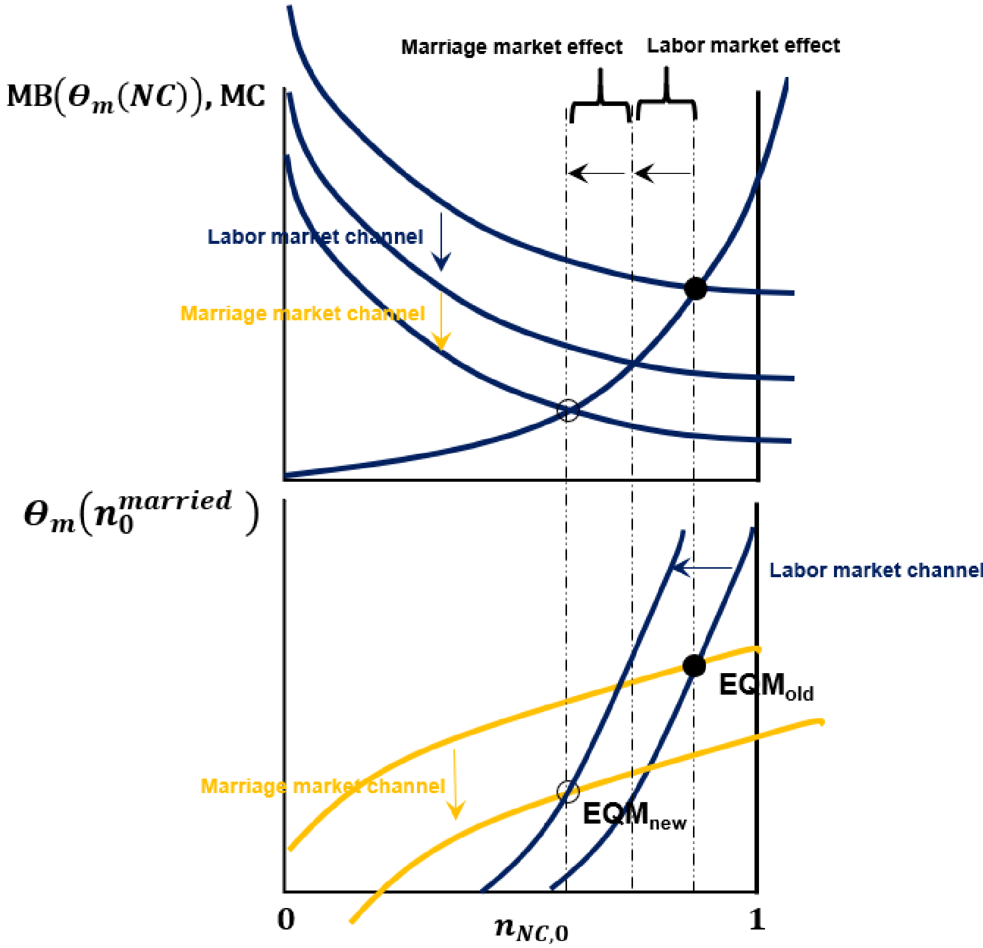

Notes: This figure graphs the equilibrium effects of a reduction in noncollege men's wage offers: some combination of a reduction in initial wages $w_{NC,0}$ and a reduction in the returns to experience $\hat{r}$. The population employment rate in period 0 is depicted by the vertical dot-dash line. The initial "labor market effect" of the shock shifts the population marginal benefit curve downward. The subsequent "marriage market effect" occurs as the marriage market adjusts to the erosion in noncollege men's earnings prospects, resulting in fewer noncollege men marrying at a lower surplus share. Such a change is captured by a downward shift of the orange locus in the bottom graph. This result in a further downward shift of the population marginal benefit curve in the top graph.

1.8 Tables

Table 1.1: Strong association between initial employment behavior and subsequent marriage propensity: NLSY79 men without college

Outcome:	% weeks in LF, ages 21-31				% weeks employed, ages 21-31			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Effect of 0-to-1 increase in marriage, ages 32-40	7.3***	6.7***	8.2***	7.6***	11.9***	10.8***	12.4***	11.3***
	(0.9)	(0.9)	(1.0)	(1.0)	(1.4)	(1.4)	(1.5)	(1.5)
Mean outcome, never-marrieds	85.8	85.8	85.8	85.8	73.5	73.5	73.5	73.5
N	1,500	1,500	1,500	1,500	1,500	1,500	1,500	1,500
Controls								
% married	✓	✓	✓	✓	✓	✓	✓	✓
% non-marital cohab.			✓	✓			✓	✓
initial wage offer		✓		✓		✓		✓

Notes: NLSY79 data. Sample consists of men with at most one year of completed college and who are not currently enrolled in school. Top row reports estimated marginal effects of a 0-to-1 increase in the share of time spent married over ages 32-40 on the share of time spent employed while ages 21-31. Even-numbered columns control for the log of the individual's hourly wage, averaged over the first 3 years since the individual left school. Hourly wage data are processed in a manner described in Appendix A.2. Non-marital cohabitations are constructed using the method of Oppenheimer (2003). Regressions are weighted by NLSY79 sampling weights multiplied by the number of employment observations in the reference period. Triple asterisks denote statistical significance at the 1% level.

Table 1.2: Effects of marriage market interventions on noncollege men's marriage propensities

	Currently married: all men					
	(1)	(2)	(3)	(4)	(5)	(6)
	Intervention 1: mutual consent → unilateral divorce					
Effect of legal change	-1.7^{***}	-2.4^{***}	-1.3^{**}	-1.8^{***}	-2.0^{***}	-2.1^{***}
	(0.6)	(0.8)	(0.6)	(0.6)	(0.6)	(0.7)
100−control mean	28.2	28.2	28.2	28.2	28.2	28.2
Pre-trends p-val	0.62	0.44	0.16	0.03	0.19	0.39
N (thousands)	2,151	2,151	2,151	2,151	2,151	2,151
	Intervention 2: increased demand for female employment					
Effect of 10 $p.p.$ Bartik shock	-1.7^{***}	-1.1^{**}	-2.5^{***}	-1.5^{***}	-2.4^{***}	-1.6^{***}
	(0.5)	(0.5)	(0.4)	(0.5)	(0.4)	(0.5)
100−control mean	31.2	31.2	31.2	31.2	31.2	31.2
N	2,888	2,888	2,888	2,888	2,888	2,888
Controls						
Baseline	✓	✓	✓	✓	✓	✓
Demos×linear trend		✓		✓		✓
Region effects			✓	✓		
State/division effects					✓	✓

Notes: Sample is single men with 5-20 years of potential experience, not enrolled in school, at least 16 years old, and with at most one year of college education. "100−control mean" reports the percentage of individuals not currently married in a mutual consent regime (Intervention 1) or in 1980 (Intervention 2). Standard statistical significance legend used. Intervention 1: OLS models based on individual data from 1960-1990 U.S. Census samples. See text for detail on baseline and additional controls. Robust standard errors are clustered on state. Regressions are weighted by Census person weights. Pre-trends p-values are p-values from an F-test of the hypothesis that the effect of unilateral divorce is 0 before the passage of unilateral divorce. Specifically, event-time is grouped into 5-year bins (due to the decadal frequency of Census data) and event-study coefficients are estimated relative to the bin containing 5-to-1 years before law passage. The F-test is whether the coefficients corresponding to "10-to-6" and "at least 11" years before law passage are jointly zero. Intervention 2: OLS models based on commuting-zone-average data at 4 time points: 1980, 1990, 2000 (U.S. Censuses) and 2015 (2014-2017 American Community Surveys). Bartik instruments are constructed using samples of men and women with 0-20 years potential experience, not enrolled in school, and with at most one year of college education. See text for detail on baseline and additional controls. Robust standard errors are clustered on commuting zone. Regressions are weighted by 1980 sample population shares for each commuting zone.

Table 1.3: Effects of marriage market interventions on noncollege men's labor-force participation

	In the labor force: single men					
	(1)	(2)	(3)	(4)	(5)	(6)
	Intervention 1: mutual consent → unilateral divorce					
Effect of legal change	-1.0^{*}	-1.4^{*}	-0.9^{*}	-1.3^{**}	-1.9^{***}	-1.9^{**}
	(0.6)	(0.8)	(0.5)	(0.5)	(0.6)	(0.9)
100−control mean	14.3	14.3	14.3	14.3	14.3	14.3
Pre-trends p-val	0.39	0.62	0.22	0.06	0.84	0.99
N (thousands)	859	859	859	859	859	859
	Intervention 2: increased demand for female employment					
Effect of 10 $p.p.$ Bartik shock	-0.8	-0.8	-1.4^{***}	-1.3^{**}	-1.5^{***}	-1.3^{**}
	(0.5)	(0.5)	(0.5)	(0.5)	(0.5)	(0.5)
100−control mean	12.9	12.9	12.9	12.9	12.9	12.9
N	2,888	2,888	2,888	2,888	2,888	2,888
Controls						
Baseline	✓	✓	✓	✓	✓	✓
Demos×linear trend		✓		✓		✓
Region effects			✓	✓		
State/division effects					✓	✓

Notes: See above table. Everything is the same except that the sample is currently single men with 0-15 years of potential experience. "100−control mean" reports the percentage of individuals not in the labor force in a mutual consent regime (Intervention 1) or in 1980 (Intervention 2). Standard statistical significance legend used.

Table 1.4: Small effects of marriage market interventions on noncollege men's wages

	Log weekly wage/100: single men					
	(1)	(2)	(3)	(4)	(5)	(6)
Intervention 1: mutual consent → unilateral divorce						
Effect of	-0.7	-0.7	0.5	0.5	-2.4	-2.2
legal change	(2.0)	(2.4)	(1.4)	(1.4)	(2.9)	(3.2)
N (thousands)	685	685	685	685	685	685
Intervention 2: increased demand for female employment						
Effect of 10 $p.p.$	1.1	-1.6	1.9	-1.0	1.7	-1.2
Bartik shock	(1.3)	(1.1)	(1.2)	(1.2)	(1.1)	(1.1)
N	2,888	2,888	2,888	2,888	2,888	2,888
Controls						
Baseline	✓	✓	✓	✓	✓	✓
Demos×linear trend		✓		✓		✓
Region effects			✓	✓		
State/division effects					✓	✓

Notes: Sample and regression specifications are the same as in the above table. See Appendix B for information on computation of log weekly wages. Standard statistical significance legend used.

Table 1.5: Contributions of marriage market interventions to secular decline in young noncollege men's LFP

Intervention:	*Unilateral divorce*	10$p.p.$ *female employment shock*
Time period:	1965-1980	1980-2015
Average effect of intervention ($p.p.$)	-0.65	-1.35
Average exposure to intervention	0.54	1.40
Predicted LFP response ($p.p.$)	-0.35	-1.89
Observed LFP response ($p.p.$)	-1.5	-6.3
% explained by intervention	**23%**	**30%**

Notes: Computations based on U.S. Census data and labor-force participation regression results. Sample is men not enrolled in school, at least 16 years old, with at most one year of completed college, and with 0-15 years potential experience. The first row reports the estimated effect of the intervention on the labor-force participation rate of this sample, averaged across the 6 specifications (see Appendix Table A.4 for regression results). The second row reports the increase in national exposure to this intervention over the specified period. The predicted change in LFP resulting from the intervention, reported in the third row, is the product of the first two rows. The contribution of the intervention to the observed decline is reported in the last row, and is equal to the third row divided by the fourth row.

Table 1.6: Baseline structural estimation results

Panel A: Parameter Estimates		
Parameter Description	Symbol	Estimated Value
Disutility of participation	λ	2.08
Substitution parameter	γ	1.14
Marital value scalar, NC women	α_{NC}	1.81
Marital value scalar, C women	α_C	1.70
Signal value of working hard, NC men	μ	0.124
Aversion to opposite education type	ξ	1.08
Marital taste scale	σ_ϵ	0.181
Panel B: Model Fit		
Moment Description	1980s U.S. data	Model
Period-0 LFP rate, to-be-singles	0.842	0.842
Period-0 LFP rate, to-be-marrieds	0.932	0.932
Uncompensated supply elasticity	0.100	0.100
Share NC women married	0.768	0.768
Share C women married	0.783	0.783
Share NC men married	0.764	0.764
Share C men married	0.790	0.790
Correlation of marrieds' C statuses	0.46	0.46

Notes: Parameters estimated by the Method of Simulated Moments. Period-0 labor-force participation for always-single and always-married groups are computed from NLSY79 data on noncollege individuals aged 21-31, with wages fixed at the noncollege sample average. See Table 1.1 and Section 1.3.3 for further detail. Baseline marriage propensities are computed on individuals aged 35-39 and born in 1942-48 from the March CPS. See Figure 1.3 notes for further detail. College is defined as at least 2 years completed college. The uncompensated supply elasticity of 0.1 is a baseline calibration discussed in Section 1.5.1. The assortative matching (correlation) parameter comes from Greenwood et al. (2016) and is based on 1980 U.S. Census data.

Table 1.7: Simulation of the effect of a 10% reduction in earning potential on noncollege men's outcomes

Group:	Singles	Married to: NC woman	Married to: C woman	All
Panel A: Baseline equilibrium				
Population share	0.238	0.631	0.132	1.000
Marr. surplus share	*N/A*	0.520	0.316	*N/A*
Period-0 LFP rate	0.842	0.932	0.932	0.911
Panel B: Changes induced by labor market shock				
Δ Pop share	.039	$-.033$	$-.006$	.000
Δ Surplus share	*N/A*	$-.013$	$-.088$	*N/A*
Δ LFP rate	$-.013$	$-.022$	$-.022$	$-.023$
Marriage market multiplier:				$\frac{.023}{.013} = 1.77$

Notes: Panel A records baseline equilibrium statistics. Panel B describes how much these equilibrium statistics change after noncollege men's wage offers are negatively shocked by 10 percent. See Section 5 of text for discussion. The marriage market multiplier for labor-force participation is ratio of the "All" response to the "Singles" response. It measures the extent to which endogenous marriage market responses amplify the direct effects of the labor market shock. See Section 1.5.2 for further discussion.

Table 1.8: Shocks to young men's labor market opportunities: comparison of modeled responses to those estimated by Autor, Dorn and Hanson (2019)

Empirical setting:	*Male trade shock identified by Autor, Dorn and Hanson (2019).*	*Reduction in wages simulated by the model.*
Effect of:	1-unit trade shock	-10% wage shock
$\%\Delta$ observed wages	-8.5%	-10.0%
Δ LFP rate	$-1.97 p.p.$	$-2.30 p.p.$
Δ marriage propensity	$-3.57 p.p.$	$-3.93 p.p.$

Notes: Male-specific labor market shocks driven by rising import competition from China were identified by Autor, Dorn and Hanson (2019) for the population of men aged 18-39. See Section 1.5.3 of the main text for further detail. Modeled responses apply to the population of noncollege men at least 18 years old and with 0-15 years of potential labor market experience. Autor, Dorn and Hanson (2019) found the China shock to disproportionately affect low-earning men who are likeliest to be noncollege-educated.

CHAPTER II

Intergenerational Transmission of Women's Employment Status: Micro Evidence and Macro Implications

Abstract

Measurement of the intergenerational transmission of women's employment status has been limited by a lack of detailed data on mothers' and daughters' employment behaviors. This paper leverages the genealogical structure and long duration of the Panel Study of Income Dynamics to construct such a data set and to estimate intergenerational employment relationships. The intergenerational relationship is found to be strongest at the full-time employment margin for college-educated mothers, and substantially weaker at less-intensive employment margins and for less-educated mothers. The paper also documents a stark rise in inequality in mothers' full-time employment propensities in the 21st century, and attributes roughly 36% of this trend to differential intergenerational transmission across education groups. These results suggest a disproportionate influence in high-SES families of the childhood environment on gender identity, and that family-level transmission processes deepen the long-run effects of unequal labor market opportunities on inequality in mothers' career outcomes.

JEL Codes: E71, J13, J22
Keywords: female labor supply, gender identity, social norms, intergenerational transmission

2.1 Introduction

Understanding the roles played by gender identity and intergenerational processes in women's labor market outcomes has become the subject of a burgeoning literature in the study of populations (Bertrand, 2011). An initial contribution by Fortin (2005) found societal gender-role attitudes to be important in explaining cross-country differences in women's employment. Subsequent works examined the persistence of these effects across generations (Fernández and Fogli, 2009; Alesina and Giuliano, 2010; Alesina, Giuliano and Nunn, 2013), finding strong correlations between the labor market behaviors of second-generation immigrant women and prevailing gender-role attitudes and labor-force participation rates in their countries of ancestry.[1] Farré and Vella (2013) found strong associations between mothers' gender-role attitudes and those voiced by their children in a U.S. sample of mother-child pairs.

These findings suggest that gender identities and, possibly, employment statuses themselves are transmitted within societies from mothers to their daughters. Yet, empirical evidence on the intergenerational transmission of women's employment status is scant in the literature. Quantifying this relationship can shed light on the possible importance of transmissions of gender identity and mother-acquired labor market skills to daughters' employment prospects. These processes, in turn, have implications for how women plan employment spells around childbirth (Rendall and Shattuck, 2019), and how this behavior affects female earnings paths and the gender gaps over the life-cycle (Kleven, Landais and Sogaard, 2019). In addition, intergenerational relationships can inform the long-run evolution of women's labor market progress, as well as the long-run effects of labor market policies enacted at one point in time on the employment outcomes of subsequent generations. In this paper, I present and interpret new estimates of intergenerational employment relationships for recent cohorts of women in the United States.

Assessing the intergenerational transmission of women's employment status requires a consideration of the diversity of choices and opportunities that inform women's labor market behavior. This consideration is increasingly important as women's labor-force experiences have become increasingly complex and unequally distributed since the 1970s (Juhn and Murphy, 1997; Goldin and Mitchell, 2017). Table 2.1, for example, shows a modest increase between 1979 and 2015 in the shares of women (and mothers) working at all, but a larger increase in the shares working full-time.[2] The table also shows a widening *education gap* in employment propensities, especially for mothers. Between 1979

[1]The related work of Boustan and Collins (2014) argued that the institution of slavery resulted in transmission of different gender norms by race. Such a process is shown to help explain the racial gap in female labor-force participation that persisted in the U.S. for most of the 20th century. Fernández, Fogli and Olivetti (2004) also found that women born in states with larger draft rates of men into World War II—a shock that pushed more mothers into the workforce—exhibited higher employment rates as adults.

[2]The table is based on March Current Population Survey data, and the sample consists of white or black women aged 25-44. Employment measures are based on reported hours worked in the reference year (i.e. the year preceding the survey year).

and 2015, the at-all employment propensity gap between college- and less-educated mothers widened from $.046$ ($= .651 - .565$) to $.123$ ($= .784 - .661$); the full-time employment propensity gap widened from $.033$ ($= .264 - .231$) to $.164$ ($= .570 - .406$). As an increasing share of daughters are raised by full-time-employed mothers—but disproportionately so for daughters of highly-educated mothers—intergenerational processes may increasingly contribute to (inequality in) contemporary women's employment outcomes.

This paper uses multiple sources of U.S. data to estimate intergenerational employment relationships between mothers and their daughters, and to roughly quantify the impact of intergenerational transmission on 21st century developments in mothers' employment. Exploiting the genealogical structure of the Panel Study of Income Dynamics (PSID), I construct variables summarizing the extent to which a first-generation mother worked throughout her daughter's upbringing. I then link these variables to the employment behavior of the second-generation daughter once she has become an adult and, in many cases, had children of her own. The resulting dataset is one of few existence in which it is possible to observe detailed employment information, covering a long time period and similar age ranges, of both mothers and their daughters. I analyze heterogeneity in intergenerational transmission by estimating logistic regressions of second-generation women's employment on their mothers' employment histories, for various employment margins (e.g. employed at all versus employed full-time) and subgroups (e.g. skilled versus less-skilled women).

Three results emerge. First, the intergenerational employment relationship is stronger at more intensive employment statuses. That is, how much a mother worked part-time while raising her daughter matters only modestly for her daughter's eventual part-time employment status. But a mother's full-time employment history matters much more for her daughter's eventual full-time employment status. Second, the intergenerational relationship strengthens when the second-generation sample of skilled women is limited to mothers with young children. Third, and perhaps most salient, the intergenerational relationship is stronger in skilled families than in less-skilled families. These results suggest that the decisions of skilled women to hold on to their full-time careers while raising young children are particularly sensitive to whether their own mothers did the same.

Motivated by the education gradient to these results, I then turn to a quantification of the role played by intergenerational transmission in the recent evolution of the *education gap* in mothers' full-time employment behavior. Using a large and nationally representative sample from the Current Population Survey, I document a marked increase between 1999 and 2016 in the share of college-educated mothers working full-time. Less-educated mothers, on the other hand, experienced virtually no increase in full-time employment.

I also document a similarly-timed divergence in gender-role attitudes in the General Social Survey. Using a regression decomposition exercise, I attribute roughly 36 percent of widening inequality in full-time employment to differential intergenerational transmission by skill level. Intergenerational forces are substantially more important in explaining this recent development than the combination of the other "economic" variables considered (e.g. wages, partners' incomes, age at childbearing and child burdens).

It is important to note that the intergenerational employment relationships estimated here lack an explicit causal interpretation, as they are determined by several mechanisms operating in tandem.[3] These include the influence of the childhood environment on gender identity, the transmission of skills acquired by a first-generation mother that facilitate an effective balance of work and family demands, and the inheritance of fixed characteristics that systematically differ between families (e.g. motivation to work hard). However, the found education gradient to the intergenerational employment relationship suggests that the transmissions of gender identity and work-family-balancing skills, or the effects of these transmissions on employment status, are stronger within higher-SES families. These heterogeneous transmission processes across families play a potentially important role in emerging inequality in mothers' career outcomes.

One implication of this paper is that not just inequality between women and men, but also inequality across women, should be considered in the design of policies to help women balance work and family demands (Blau and Kahn, 2013). This implication also raises the question of whether gender identity impacts the responsiveness of women to such policies.

2.2 Intergenerational Perspective on Women's Labor Supply

The decisions of whether and how much time to spend working in the paid labor force depend on how one's market wage compares with the values placed on alternative uses of time. These alternative uses are usually categorized as leisure or as working in the home (Gronau, 1977)—the latter of which has long been considered an important determinant of women's labor supply behavior (Killingsworth and Heckman, 1986).

2.2.1 Incorporating gender identity and family transmission into the economic framework

Akerlof and Kranton (2000) augmented the purely economic calculus by incorporating the concept of identity. In this augmented framework, the relative valuation of a woman's

[3] The intergenerational transmission of earnings status, the subject of a rich literature in economics, is also a non-causal parameter. See the classic reference of Solon (1999) for discussion.

time at home depends not just on prevailing economic conditions, but also on norms regarding how women ought to spend their time. Another influential framework of Stryker and Burke (2000) posits the interaction of two forces in the formation of individual identity: social structures and self-conception. That is, norms about women and work can be dictated "externally" by society-at-large (i.e. individuals and structures in the woman's reference network), as well as shaped "internally" from the woman's own lived experiences within her household. Societal-level norms may differ from and compete against household-level conceptions, as found by Fuwa (2004) in a cross-country sample: holding a household's own gender ideology constant, more household gender inequality was found in less-egalitarian countries.

The transmission of gender identity at the family level has long been the subject of a literature in developmental psychology. Several researchers have mounted evidence supporting a socialization process, in which children learn about prevailing gender norms from their parents (Marantz and Mansfield, 1977; Eccles, Jacobs and Harold, 1990; McHale, Crouter and Tucker, 1999; McHale, Crouter and Whiteman, 2003; Galambos, 2004). Thus, a mother's employment choice potentially determines the prescriptions associated with motherhood taken up by her daughter, and affects the daughter's employment choice when she raises a family of her own. The mother's decision to work while raising her daughter may also reflect certain predispositions or skills that are then transmitted to the daughter, such as a motivation to work hard, or practical tips about how to balance career and family.

These observations suggest that the strength of a given intergenerational employment relationship is determined by the interaction of family-level processes with norms dictated by society-at-large, and with the standard economic calculus:

Hypothesis 2.1: *The effect of a mother's particular employment choice, while raising her daughter, on her daughter's eventual choice depends on the extents to which this particular choice: i) shapes the gender identity, or reflects the employment-relevant predispositions and skills, carried by the given daughter into adulthood; and ii) is regulated by societal norms, or is made out of economic necessity.*

2.2.2 Possible heterogeneity in intergenerational transmission

To a larger extent than men, women decide between working at all, working part-time or working full-time (Killingsworth and Heckman, 1986). These margins of choice interact with family structure: most importantly, the presence of pre-school-age children in the household. As reflected in Table 2.1, women with pre-school-age children still have low propensities for full-time employment, but this group has experienced the largest increase since the 1970s. Moreover, rising inequality in women's labor market opportunities (Autor, 2014) presents skilled women with increasingly different cost-benefit analyses than less-skilled women. This is reflected in widening *education gaps*, particularly in full-time employment (Table 2.1).

Drawing on Hypothesis 2.1, there is reason to suspect that these different margins of choice, in different sub-populations, will be affected differently by intergenerational dynamics:

- A stronger intergenerational effect is likely to exist at more intensive employment margins. This is because, by the late 20th century, most women maintained some degree of attachment to the labor force (Goldin, 2006). The fact that a daughter's mother worked on an occasional basis may not constitute an informative signal. On the other hand, very few mothers worked full-time. A daughter whose mother worked full-time may enter adulthood with a substantially greater willingness to balance full-time work with the demands of being a woman in a gendered society, relative to a daughter raised by a mother without a full-time career. This relationship could reflect greater gender identity formation, or greater transmission of career-relevant skills.

- A different intergenerational effect is likely to exist for women with young children, versus in the full population of women. As suggested by the developmental psychology literature, a daughter's observation of her mother's behavior may shape her identity particularly with respect to the intersection of work and motherhood. On the other hand, the decision of a mother with young children to enter (full-time) work may be heavily regulated by societal norms, or may be economically infeasible due to child care costs (Blau, 2003). Whether the identity effect or the socioeconomic feasibility effect dominates is an empirical question for the data.

- The intergenerational effect and the above-hypothesized gradients are likely to differ by family skill level. There are three reasons to suspect a positive skill gradient. First, relative to skilled women, less-skilled women may choose to work more for reasons of economic necessity, and less to take advantage of attractive employment opportunities or egalitarian gender identities.[4] Second, a slack labor market may present less-skilled mothers with few attractive opportunities for full-time employment—whereas highly-skilled mothers have better opportunities to choose between working

[4] Less-educated women are likelier to be unmarried and to have lower-earning partners (Lundberg, Pollak and Stearns, 2016).

part-time and pursuing an attractive full-time career (Goldin, 2014). Third, inegalitarian norms appear to be stronger among less-skilled men (Bertrand, Kamenica and Pan, 2015), suggesting that less-skilled women's behavior may be more heavily regulated by societal pressures. On the other hand, a less-skilled mother's decision to pursue a full-time career, made despite (even in defiance of) these challenges, may particularly empower her daughter to behave similarly in adulthood. Once again, these trade-offs between socioeconomic feasibility and identity are theoretically ambiguous, and must be assessed in the data.

Using the Panel Study of Income Dynamics, this study assembles a two-generation dataset that links mothers' and daughters' labor supply behavior at similar points in adulthood. This permits a direct measurement of the intergenerational transmission mechanism, and an assessment of its variance across these dimensions of heterogeneity. I know of only two other studies that analyze similar data in a U.S. context: Olivetti, Patacchini and Zenou (2016), which uses AddHealth data,[5] and Galassi, Koll and Mayr (2018), which uses NLSY data. This study extends these in two ways: first, by testing for heterogeneity in the intergenerational transmission relationship; and second, by quantifying the implications of heterogeneous intergenerational transmission processes for aggregate movements in U.S. women's employment behavior.

2.3 PSID Data and Methods

The Panel Study of Income Dynamics (PSID) contains rich information on the work behavior of an individual's parents. The landmark 50-years-and-counting longitudinal study began following families in 1968, asking respondents about their work behaviors in 1967. Using the genealogical structure of the survey, which allows one to link individuals together with their family members (even if they are no longer living in the same household), one can observe the full work history of an individual's mother, for any individual born in 1967 or later to one of the multiple thousands of original PSID families.

2.3.1 Sample construction and summary statistics

Using this information I construct several variables aimed to summarize a mother's work history throughout her daughter's upbringing. The first variable, *MOM_WORK*, records

[5] One limitation of the AddHealth intergenerational data is that it is only possible to observe the first-generation mother's work behavior for one year, when the daughter is 15-18 years old. (See Olivetti, Patacchini and Zenou, 2016 for detail.) On the other hand, the PSID enables one to observe the first-generation mother's work behavior throughout the daughter's entire upbringing.

the share of years throughout her daughter's childhood that a given mother was employed at all (i.e. supplied positive hours of work during the reference year). The second variable, *MOM_PT*, records the share of years that the mother was employed part-time or greater—defined in this paper as being employed for at least 400 hours during the reference year. The third variable, *MOM_FTFY*, records the share of years that the mother was employed full-time—defined as being employed for at least 1600 hours.

I constructed these variables using 3 different daughter age ranges—0 to 17, 0 to 10, and 6 to 17—and verified that the 3 measures were highly correlated and produced similar results. The preferred specifications use the 6 to 17 age range because this results in the largest sample size. Specifically, using the 6 to 17 age range allows me to bring the 1961 through 1966 cohorts into the sample, cohorts which would have been left out if I used the 0 to 10 or 0 to 17 age ranges (since the observation period begins in 1967). I set these variables to missing if information on the mother's work history is missing for more than 6 of the 12 reference years.

Next, I construct a sample of second-generation women who are current heads of their own households, or are partners of the heads. I keep only women born between 1961 and 1987—the range of years in which it is possible to reliably observe both the first-generation mother's full work history and the second-generation woman's work behavior—and who are members of original 1968 PSID families. Because I am interested in modeling these individuals' employment behaviors relatively early in their career cycles—when they are likeliest to be raising their own children—I focus on a sample of women aged 22-49. Using the genealogical structure of the survey, I collect these women's mothers' person identifiers, and then merge information on the women to information on their mothers' work histories using these person identifiers.

Table 2.2 summarizes the resulting sample. Statistics are weighted by original PSID family weights. The sample contains $21,980$ woman-year observations, covering $2,228$ unique women. Thus, the average woman was observed around 9.9 times, with earlier birth cohorts observed more often: the average birth year of the sample is 1970, though the range of birth years covered is 1961-1987. The woman-year observations range from the 1983 interview year to the 2015 year, with the mean observation coming from 2004. The average age of the sample is 33.5 years, roughy 8 percent of the sample is black, and one-third of the sample has a 4-year college degree. 44 percent of the observations correspond to "young mothers:" women who have at least one child under the age of 6 living in the household. These young-mother households have an average of 1.43 young children. The overwhelming majority of the sample lived in a household with a gainfully employed male partner: the average male partner earned roughly $\$54,000$ (expressed at

2010 prices).

Table 2.2 also records statistics of interest on the second-generation women's mothers. 41 percent of first-generation mothers had completed some college education. The average first-generation mother in the sample was employed *at all* roughly 68 percent of the time while her daughter was aged 6-17 years old, with some mothers employed none of the time and some mothers employed all of the time. The average first-generation mother was employed *full-time, full-year* only around 31 percent of the time throughout her daughter's upbringing—once again, with some mothers fully employed none of the time and others fully employed all of the time. Thus, there exists substantial variation across second-generation women in the education and employment experiences of their mothers.

2.3.2 Regression methodology

Consistent with the definitions of first-generation mothers' employment histories, I construct three corresponding measures of second-generation women's employment statuses: binary indicators for being employed at all last year, working at least 400 hours, and working at least 1600 hours (full-time). For a given second-generation woman i observed in survey year t, I model the probability of achieving the kth employment status (where $k =$ "at all", "at least part time", or "full-time, full-year") as a logistic function of the kth measure of her mother's employment history:

$$\Pr(2nd \text{ gen emp status k}_{ist} = 1) = \frac{1}{1 + e^{-(1st \text{ gen emp status}_i \cdot \gamma_{ks} + \mathbf{X_{it}} \cdot \beta_{ks})}}. \tag{2.1}$$

By applying logistic regression to this equation, we can estimate γ, the intergenerational coefficient of interest. As hypothesized in Section 2.2, the strength of the intergenerational employment relationship likely depends on the measure of employment (k) and the nature of the sub-sample (s). Regarding the latter, we are interested in measuring how the intergenerational coefficient varies when we restrict the sample to mothers with young children, and when we vary the sample's skill level. In the results that follow, I present two alternative cuts of the sample on skill level: one according to whether the first-generation mother had obtained any college education, and one according to whether the second-generation woman had obtained a 4-year college degree.

I gauge the sensitivity of these intergenerational relationships to the inclusion of other economic and demographic characteristics $\mathbf{X_{it}}$. One important set of controls, for clarity of interpretation, consists of birth cohort fixed effects. This is because first-generation mothers' employments have trended upward throughout the sample period. If second-generation women born later also work more, for reasons unrelated to intergenerational transmission, the coincidental nature of these first- and second-generation trends will nonetheless lead us to estimate a significant intergenerational response. Controlling for birth cohort fixed effects rules out such a coincidence. I also include controls for parental education and race.

In addition to these "background" controls, I also include a set of "contemporaneous" controls, designed to capture other economic and demographic determinants of women's labor supply (Killingsworth and Heckman, 1986). These include the log real hourly wage offer, a quadratic in husband or partner's earnings, a quadratic in the number of children aged 0-5, the number of children aged 6-10, and a quadratic in age. The real hourly wage offer is simply observed wage income divided by hours employed (expressed at 2010 prices) for women who were employed at all in the given reference year. I impute hourly wage offers for women who were not employed at all, based on observed wages of demographically similar women, via a procedure described in Appendix Section B.2.4.

The preferred interpretation of the kth intergenerational transmission relationship is the effect of a 0-to-1 change in the first-generation mother's kth employment history measure, holding these background and contemporaneous variables fixed. This describes the amount by which being raised by a first-generation mother who achieved employment status k *every year* changes the second-generation woman's probability of obtaining employment status k, relative to being raised by a mother who *never* achieved employment status k. In the results that follow, I present estimates of these marginal effects, holding the other control variables constant at their sub-sample means (i.e., $\overline{\mathbf{X}}_{\mathbf{s}}$). I denote the estimate of the kth marginal effect at the sample mean, for the sth sub-sample, as $\widehat{MEM}_{ks}$:

$$\widehat{MEM}_{ks} = \frac{1}{1 + e^{-\left(\widehat{\gamma}_{ks} + \overline{\mathbf{X}} \cdot \widehat{\beta}_{ks}\right)}} - \frac{1}{1 + e^{-\left(\overline{\mathbf{X}} \cdot \widehat{\beta}_{ks}\right)}}. \tag{2.2}$$

One can also use logit odds ratio methodology to interpret the results. The estimate of the kth odds ratio at the sample mean for the sth sub-sample, or $\widehat{ORM}_{ks}$, is the following:

$$\widehat{ORM}_{ks} = \frac{1 + e^{-\left(\overline{\mathbf{X}} \cdot \widehat{\beta}_{ks}\right)}}{1 + e^{-\left(\widehat{\gamma}_{ks} + \overline{\mathbf{X}} \cdot \widehat{\beta}_{ks}\right)}}. \tag{2.3}$$

This expression describes how many times more likely it is that a typical second-generation woman of sub-sample s, whose mother *always* achieved employment status k throughout her childhood, achieved employment status k than a comparable woman whose mother *never* achieved employment status k.

2.4 Intergenerational Regression Results

As a baseline exercise, I start by pooling all second-generation women together and examining the full-time, full-year (FTFY) employment margin. Table 2.3 reports estimated marginal effects (expressed in percentage points $/100$) of a 0-to-1 change in *MOM_FTFY*, calculated at sample means. Robust standard errors are clustered at the individual level and appear below the coefficients in parentheses.

Looking at column 1, a simple bivariate regression yields a positive and highly significant association between a first-generation mother's full-time employment history and the full-time employment behavior of her daughter in adulthood: a 0-to-1 increase in *MOM_FTFY* is associated with a 10.8 percentage-point increase in the second generation mother's probability of working full-time. This amounts to 20 percent of the second-generation sample's FTFY employment rate. Columns 2-6 successively add in the additional control variables described in the preceding sub-section. Inclusion of these controls does not largely affect the estimated intergenerational relationship: column 6, the preferred specification, reports that a 0-to-1 change in *MOM_FTFY* is associated with an 11.3 percentage-point change in a second generation mother's probability of working full-time.

To express this estimate in odds-ratio terms, we require one additional number. Recall from Table 2.2 that the average second-generation woman's mother was full-time-employed 31 percent of the time throughout her daughter's upbringing. Combining this with the second-generation sample's FTFY employment rate and the estimated marginal effect yields the following calculation:

$$\widehat{ORM}_{FTFY} = \frac{.502 + .113 \cdot (1 - .31)}{.502 + .113 \cdot (0 - .31)} = 1.24.$$

That is, a typical woman whose own mother worked full-time throughout her childhood is 1.24 times as likely to be full-time-employed in adulthood as an observationally identical woman whose own mother never worked full-time throughout her childhood.

Proceeding with the preferred specification (6), I explore heterogeneity in the intergenerational transmission relationship in Tables 2.4 and 2.5. Each table presents 12 estimates, based on 3 dimensions of heterogeneity: labor supply margin (employed at all, employed at least 400 hours, employed FTFY), motherhood nature of the sub-sample (all women, women with at least one child under 6 only), and skill nature of the sub-sample (less-skilled, skilled). Table 2.4 cuts the full sample based on the 1st-generation's education level (no college versus some college or more), while Table 2.5 cuts the full sample based on the 2nd-generation's education level (no 4-year college degree versus 4-year college degree).[6] In each table, the 4 columns correspond to the motherhood-by-skill cuts of the sample, while the 3 rows correspond to the 3 employment statuses investigated. Each cell reports the estimated intergenerational transmission relationship, its associated standard error, and the associated share of second-generation women achieving the given employment status.

Looking at the two tables together, several patterns of heterogeneity emerge. First, the intergenerational relationship tends to be stronger as the intensivity of the employment status rises—especially for the skilled sub-samples. Column 3 of Table 2.4 reports that a 0-to-1 increase in *MOM_WORK* is associated with a non-significant 5.4 percentage-point increase in the probability that a skilled second-generation woman is employed at all (or 1.09 odds ratio). However, a 0-to-1 increase in *MOM_FTFY* is associated with a highly significant 16.4 percentage-point increase in her probability of FTFY employment (or 1.35 odds ratio). Table 2.5 reveals a similar, if not slightly starker, pattern when the sample is stratified on the second generation's education status rather than the first generation's.

Second, let us compare the odd-numbered columns (full sample) to the even-numbered columns (only mothers with young children) in each table. The estimates indicate that for less-skilled women, the intergenerational relationship is slightly stronger in the full sample than in the mothers-only sub-sample. However, this pattern is reversed for skilled women: the intergenerational relationship is slightly stronger in the mothers-only sub-sample than in the full sample. These patterns hold for all 3 employment margins considered.

Third, the intergenerational relationship tends to be stronger for skilled women—especially for mothers of young children, and most especially for the FTFY employment status. Comparing column 2 with column 4 of Table 2.5, we see that a 0-to-1 increase in *MOM_FTFY* is associated with a marginally significant 6.8 percentage-point increase in a second-generation, less-educated mother's probability of full-time employment (1.20 odds ratio). This marginal effect rises to 21.9 percentage points (1.62 odds ratio) for second-generation, highly-educated mothers. This difference in marginal effects, both here and in Table 2.4, is statistically significant.

These results align with the predictions outlined in Section 2.2. First, they suggest that observing one's own mother participating in full-time employment (as opposed to occasional employment) is especially likely to lead "occupation and [full-time] employment [to] define one's fundamental identity and societal worth" (Goldin, 2006, page 1). Second, they suggest that college-educated women's employment decisions, relative less-

[6] I use a more stringent definition of skilled for the second generation than for the first generation, given that college completion rates were higher in the second generation. Nonetheless, cutting the second-generation sample based on college attendance, rather than college completion, produces qualitatively similar results.

educated women's employment decisions, are more reflective of family-transmitted identities and skills. This is particularly true for college-educated versus less-educated *mothers'* decisions to maintain employment while raising young children. Interacting these two conclusions, it appears that the intergenerational transmission mechanism is most salient for skilled women who are considering balancing a full-time career with the demands of raising young children.

2.5 Implications for Aggregate Change in Mothers' Full-Time Employment in the 21st Century

The full-time employment propensity of U.S. married women increased dramatically throughout the late 20th century (Goldin, 2006), and this aggregate increase was primarily driven by the changing behavior of skilled wives (Juhn and Murphy, 1997). These changes, combined with the intergenerational transmission relationships estimated above, suggest that intergenerational transmission has affected particularly the full-time employment behavior of U.S. mothers in the 21st century. Moreover, the skill-biased nature of these changes portend a recent rise in *skill inequality* in mothers' full-time employment.

The goal of this section is to examine these predictions. To do so, I use data from the Current Population Survey (CPS). The CPS is better equipped to examine aggregate employment trends than the PSID, as it contains a much bigger sample (roughly $60,000$ households), and has maintained national representation over time. I use the March supplement to the CPS, which elicits annual employment and earnings information for CPS households during their March interviews.[7] I begin by constructing a sample of March CPS respondents that is similar to the PSID analysis sample: women aged 22-54, white or black, and with non-missing economic and demographic information.

I use two different methods of dividing the sample by skill level. One method involves grouping all women into a high- (low-) skill sub-sample who had wage offers above (below) the median of the hourly wage offer distribution, within a cell defined by survey year and potential experience bin. To construct this stratification, I adjust wage and business incomes for top-coding and compute hourly wages as total wage and business income divided by annual hours worked,[8] trimming wage outliers from the sample. I then impute hourly wage offers for non-working women via a procedure described in Appendix Section

[7]The CPS interviews a panel of households each month for 4 months, then ignores the panel for 8 months, and then conducts 4 more monthly interviews before dropping the panel from the sample.

[8]Before the 1976 survey, information used to construct annual hours worked is missing–weeks worked information is available in the form of intervals. Using education, demographic, and weeks worked interval information, I impute annual hours worked in a manner described in Appendix Section B.2.3.

B.2.4.

Alternatively, I use the same definition of skill as used in the PSID sample: having completed 4 years of college education. This definition has two advantages over the first one. First, it is not dependent on imputation procedures. Second, it more coherently divides the population into distinct labor markets. However, this definition has a key drawback: the underlying skill composition of the college and non-college populations may have changed over time as the female college attainment rate has increased. The wage-based method of stratification groups 50 percent of the population into each skill group in each year, and thus may be more immune to changing composition. In the eventual quantitative analysis, I average the results from both definitions. The following two subsections supply some background.

2.5.1 20th century change in childhood environment, from mother-at-home to mother-at-work

To illustrate the large secular change in the environment in which daughters have been raised—from mother-at-home to mother-at-work—I manipulate the CPS sample to construct a sample of daughters aged 6-17 living with their mothers. I organize the sample by the daughter's birth year and compute mothers' FTFY employment rate for each birth cohort of daughters, separately by mother skill level. I use the high-wage-offer definition of skill. Figure 2.1 plots the resultant series of mothers' FTFY employment rates, by mother skill and by daughters' birth years. It displays a large and skill-biased change in the environment in which daughters have been raised.

Such a change is potentially relevant for 21st century patterns in mothers' FTFY employment. For example, the median-aged skilled woman with at least one young child, observed in the 1999 CPS sample, was born in 1966. Her counterpart, observed in the 2016 CPS sample, was born in 1983. According to Figure 2.1, the share of daughters in skilled households raised by full-time-employed mothers increased from 30 percent to 56 percent between these two birth cohorts: a near doubling. The share of daughters in less-skilled households raised by full-time-employed mothers also increased, but far more slowly.

2.5.2 The 21st century rise in skilled mothers' full-time employment

Accordingly, the remainder of this section focuses on the FTFY employment behavior of young mothers. Figure 2.2 presents time series of young mothers' employment propensities by skill, again using the wage-based skill classification. The left graph considers the propensity to be employed at all, while the right graph considers FTFY employment. Both graphs show a marked increase in the *skill gap* in mothers' employment in the 21st century. This is particularly true for FTFY employment: between 1999 and 2016, skilled young mothers' FTFY employment increased by roughly 12 percentage points, or over 25% on its base rate. Appendix Figure B.1 confirms that this recent increase in FTFY employment also occurred among college-educated mothers, and in U.S. Census data as well as in the CPS.

On the other hand, less-skilled young mothers' FTFY employment barely changed over this time period. By 2016, skilled young mothers' FTFY employment rate was the highest it had ever been (59 percent), and the skill gap in young mothers' FTFY employment was also the highest it had ever been (21 percentage points)—at least since the 1960s, when these statistics began to be tracked by the CPS.

2.5.3 The 21st century evolution of attitudinal support for mothers' careers

If childhood experiences shape adult perceptions, a larger share of individuals raised by working mothers might translate into greater adulthood support for mothers pursuing careers. Information collected by the General Social Survey (GSS) permits an exploration of this possibility. Of the 8 questions regarding gender attitudes asked in the GSS, 3 have been asked systematically over time and pertain to work-family balance. These questions ask respondents to rate the extent to which they agree (on a 1 to 4 scale) with the following statements: "It is much better for everyone involved if the man is the achiever outside the home and the woman takes care of the home and family" (FEFAM); "A preschool child is likely to suffer if his or her mother works" (FEPRESCH); and "A working mother can establish just as warm and secure a relationship with her children as a mother who does not work" (FECHLD).

All 3 questions were asked in the 1977 GSS and then on a regular basis since 1985. I aggregate information contained in the individual responses into a single national index, entitled $momWorkAttitude$, as follows. I begin by focusing on a sample of individuals with young children.[9] I recode responses to each of the 3 prompts into a binary scale, with 1s denoting egalitarian views and 0s denoting traditional views. For each year y I compute sample average responses to each of the 3 questions, and set $momWorkAttitude_y$ equal to the average of these 3 averages. Thus, $momWorkAttitude$ is a continuous variable ranging from 0 to 1, with higher values indicating greater social acceptance of mothers pursuing careers.

[9]Because the GSS does not report the age of the youngest child, I consider all respondents aged 22-45 with at least one child.

Figure 2.3 plots time series of $momWorkAttitude$ for each sample of skilled individuals. Skilled young mothers' FTFY employment series are overlaid: the high-wage series appears in the left graph while the college-educated series appears in the right graph. All series display substantial increases in the 1980s, slowdowns in the 1990s, and then substantial increases again after 2000. Thus, the 21st century increase in young skilled mothers' career involvement has occurred alongside an increase in households' attitudinal support for mothers' careers.[10]

2.5.4 Quantifying the contribution of intergenerational transmission

The 21st century patterns shown above are all consistent with the facts that skilled women experienced a larger change in their childhood environments than did less skilled women, and appear more responsive to these changes in the planning of their careers. Table 2.6 quantifies the contribution of these intergenerational transmission relationships to the 21st century rises in: i) skilled mothers' FTFY employment, and ii) the *skill gap* in mothers' FTFY employment.

The first row of the table reports the estimated intergenerational FTFY employment responses for young mothers by skill level. These are expressed as marginal effects at the respective sample means, and are averaged across both definitions of skill used in the PSID. For example, Table 2.4, which stratifies the sample on first-generation mother's college education, reports a marginal effect of 20.1 percentage points for skilled second-generation mothers. Table 2.5, which stratifies the sample on second-generation mother's college education, reports a marginal effect of 21.9 percentage points. Averaging these together yields the 21.0 percentage-point effect reported in the first entry of Table 2.6. An analogous procedure, conducted for less-skilled mothers, yields a 5.3 percentage-point effect.

As discussed above, between the 1966 and 1983 birth cohorts, daughters raised in skilled households experienced a 26.0 percentage-point increase in their exposure to full-time-employed mothers. On the other hand, daughters raised in less-skilled households experienced only an 11.5 percentage-point increase. (These are reported in the second row of Table 2.6.) For each skill level s, multiplying the intergenerational response together with this change in environment yields the predicted effect of intergenerational transmis-

[10]The patterns displayed in Figure 2.1, which show differential growth in the share of daughters exposed to working mothers by mother skill level, also imply differential growth in next-generation progressive attitudes by skill. Appendix Figure B.6 graphs time series of $momWorkAttitude$ by skill and confirms that such a skill divergence has occurred in the 21st century.

sion on the 1999-2016 change in FTFY employment. I denote this object as $\Delta^{IG}_{FTFY,s}$:

$$\Delta^{IG}_{FTFY,s} = \widehat{MEM}_{FTFY,s} \cdot \Delta_{MOM\_FTFY,s} \tag{2.4}$$

Applying this calculation to each skill level yields $\Delta^{IG}_{FTFY,skilled} = .055$ and $\Delta^{IG}_{FTFY,lessskilled} = .006$. The difference between these, or $\Delta^{IG}_{FTFY,skillgap} = .049$. Hence, intergenerational transmission forces raised skilled mothers' FTFY employment propensity by 5.5 percentage point between 1999 and 2016, and raised the *skill gap* in mothers' FTFY employment by 4.9 percentage points (third row of Table 2.6).

How important are these intergenerational effects relative to the observed changes in 21st century mothers' FTFY employment behavior? To answer this question, I first quantify trend behavioral changes. To do so, I keep only young mothers observed in calendar years 1999-2016, and split the CPS sample by skill level. Within each skill level s, I regress observed full-time employment behavior on a constant and a time trend τ. The parameter τ thus equals the per-year increase in the sample's FTFY employment propensity. The total observed behavioral change, or $\Delta^{OBSERVED}_{FTFY,s}$, equals the estimate of τ multiplied by 17 (the number of elapsed years):

$$\begin{aligned} FTFY_{ist} &= \alpha_s + \tau_s \cdot t + \epsilon_{ist} \\ \Delta^{OBSERVED}_{FTFY,s} &= \widehat{\tau}_s \cdot (2016 - 1999). \end{aligned} \tag{2.5}$$

I compute $\Delta^{OBSERVED}_{FTFY}$ for each skill level based on both definitions of skill (high-wage, college degree), and then average the results. As reported in the fourth row of Table 2.6, the trend increase in skilled young mothers' FTFY employment behavior in the 21st century was 13.0 percentage points; the trend increase in the *skill gap* was 13.8 percentage points.

Before computing the contribution of intergenerational transmission to these trends, I first account for the portion that can be explained by other determinants of mothers' labor supply. These are the same "contemporaneous" control variables used in the PSID analysis: the mother's own log wage offer, her partner's income, her own age, and the number and ages of her children. To the extent that these variables have been trending in the 21st century, we would expect some change in FTFY employment behavior to take place in the absence of intergenerational transmission. To accomplish this accounting exercise, I first estimate logistic regression equation (2.1), by skill level, on each 1999-2016 sample of young mothers. Based on the estimated regression coefficients, I then construct predicted probabilities of FTFY employment for each individual in the sample (denoted as $\widehat{FTFY}_{ist}$). Finally, I repeat computation (2.6), but replace the actual observed FTFY employment statuses with these predicted probabilities. The resulting quantity, $\Delta^{OTHERVARS}_{FTFY,s}$, describes the portion of the observed time trend that is attributable to these other variables:

$$\begin{aligned} \widehat{FTFY}_{ist} &= \tilde{\alpha}_s + \tilde{\tau}_s \cdot t + \tilde{\epsilon}_{ist} \\ \Delta^{OTHERVARS}_{FTFY,s} &= \hat{\tilde{\tau}}_s \cdot (2016 - 1999). \end{aligned} \tag{2.6}$$

The residual portion of the trend equals the observed trend net of this quantity.

As reported in the fifth row of Table 2.6, the residual portions of the trends are quite large. That is, 10.7 percentage-points of the 13.0 percentage-point rise in skilled mothers' FTFY employment, and 11.2 percentage-points of the 13.8 percentage-point rise in the *skill gap*, could not be explained by movements in observed determinants of mothers' labor supply. Finally, the contribution of intergenerational transmission to these residual changes, or $\text{IG explained}_{FTFY}$ can be computed as a simple ratio:

$$\text{IG explained}_{FTFY,s} = \frac{\Delta^{IG}_{FTFY,s}}{\Delta^{OBSERVED}_{FTFY,s} - \Delta^{OTHERVARS}_{FTFY,s}} \tag{2.7}$$

As reported in the second-last row of the table, intergenerational transmission accounted for 42.3% of the total observed increase in skilled young mothers' FTFY employment in the 21st century, and 35.5% of the increase in the *skill gap*. After taking into account movements in other labor supply determinants, the final row of the table reports that intergenerational transmission accounted for 51.9% of the *residual* increase in skilled young mothers' FTFY employment, and 43.8% of the *residual* increase in the *skill gap*. While nontrivial shares remain unexplained, these are sizable magnitudes. In raw terms, intergenerational transmission explains far more of the trends than the combination of the other labor supply variables. Intergenerational transmission appears to have importantly shaped the evolution of skilled mothers' career involvement in the 21st century, as well as the skill gap in career involvement.

2.6 Conclusion

Recent literature has applied social identity theory and intergenerational models of identity transmission (Bisin and Verdier, 2001) to the study of women's labor market outcomes. This literature has established important qualitative relationships, such as an intergenerational correlation of gender ideology, and a correlation between the labor market behaviors of second-generation immigrant women and those in their countries of ancestry. However, quantitative assessments of the extent to which female employment status is transmitted from one generation to the next have been virtually nonexistent. I have attempted to address this gap by assembling a data set of over $2,200$ mother-daughter pairs from the Panel Study of Income Dynamics. The data set records employment behaviors of both generations over long periods of time and at similar points in adulthood. I have also argued that the intergenerational employment relationship is likely not monolithic. Rather, the strength of the relationship plausibly varies according to the employment margin and demographic sub-group considered.

Analyses of the data reveal considerable heterogeneity in the intergenerational employment relationship. First, the relationship is stronger at more intensive employment margins—especially in the sub-samples of skilled women. Second, the relationship depends on whether the second-generation sample is limited to mothers with young children. For skilled women, this sample restriction strengthens the intergenerational relationship; for less-skilled women, it weakens the relationship. Third, the intergenerational relationship is stronger in skilled families than in less-skilled families. These findings are consistent with a labor supply theory in which intrinsic gender identity and extrinsic socioeconomic concerns interact to produce employment behaviors, and in which such interaction slants in different ways depending on the specific behavior and subgroup in consideration. Indeed, the results suggest that the decision to maintain employment while raising young children is more of an expression of gender identity for skilled mothers, but is more regulated by prevailing societal norms and economic conditions for less-skilled mothers.

This interpretation of the data sheds light on the macroeconomic evolution of mothers' employment behavior by skill. While *skill inequality* in mothers' employment is higher than it has ever been, *gender equality* in parental employment within skilled households is also at an all-time high. I calculate that 37% of widening inequality in mothers' full-time employment between 1999 and 2016 is attributable to differential intergenerational transmission of full-time employment status across skill groups. This suggests that a growth in skilled women's career opportunities (as occurred in the late 20th century) has a direct effect on the skill differential in mothers' career involvement, but also an indirect effect that operates through the family diffusion of egalitarian gender identities across subsequent generations. That is, the U.S. labor market contours of rising skill inequality and gender equality embed the combined effects of shifting economic incentives and shifting identities. Even if the growing inequality in labor market opportunities across skill groups ceases—which seems unlikelier than not—the dynamics found in this study portend continued divergence in gender identities between more- and less-skilled women.

Combining the intergenerational effects with the other economic variables considered, these factors can explain roughly 60% of the 1999-2016 rise in skilled mothers' full-time employment, and 54% of the rising *skill gap* in mothers' full-time employment. Thus, a substantial share remains unexplained. Ramey and Ramey (2010) has described an emerging “rug rat race” among skilled households in the preparation of their children for college. The rising cost of college, combined with increased willingness of parents to make expensive investments in their children (Bound, Hershbein and Long, 2009), may create greater skilled household demand for a second full-time income. Analyzing how this or other economic concerns (e.g. child care costs; Kubota, 2017) interact with the shift in gender identities emphasized in this paper is an intriguing possibility for future research. A related question left to future research is how (inherited) gender ideology impacts the responsiveness of women to “family-friendly” policies in the workplace such as paid family leave, flexible work scheduling or tax credits for child care.

2.7 Tables

Table 2.1: Overview of 1979-2015 change in employment, women aged 25-44

	All women			Mothers with young children		
	1979	1997	2015	1979	1997	2015
Panel A: Shares employed >0 hours						
Full sample	.723	.821	.803	.581	.732	.721
College grads	.827	.882	.876	.651	.778	.784
Some college or less	.699	.796	.738	.565	.712	.661
Panel B: Shares employed >1600 hours (full-time)						
Full sample	.416	.567	.614	.237	.416	.486
College grads	.521	.645	.713	.264	.445	.570
Some college or less	.390	.534	.525	.231	.403	.406

Notes: Data source: March Current Population Survey, white or black women aged 25-44. Mothers with young children have at least one child under age 6 present in the household. College grads have at least four years of completed college education. Full-time, full-year (FTFY) employment is defined as supplying at least 1600 hours of paid work in the reference year. Reported statistics for a given year correspond to a 3-year window surrounding that year (e.g. 1978-80 for 1979).

Table 2.2: PSID summary statistics

	Mean	Std deviation
Own variables (2nd generation)		
N observations	21,980	
N individuals	2,228	
Birth year	1970	7.13
Survey year	2004	7.63
Age	33.54	7.01
Black	0.08	0.26
4-year college degree	0.33	0.47
Young child present	0.44	0.50
Number of young children	0.63	0.82
Number of young children (conditional on having one)	1.43	0.62
Partner's real income (1000s)	53.76	51.26
Partner has positive income	0.94	0.23
Log real hrly wage offer	2.72	.60
Mother's variables (1st generation)		
Some college completed	0.41	0.49
When daughter was aged 6-17...		
share years employed at all	0.68	0.33
share years employed ≥ 400 hrs	0.57	0.37
share years employed FTFY	0.31	0.34

Notes: The sample consists of all PSID woman-year observations, who are members of original 1968 PSID families, meeting the following conditions: aged 22-49; born between 1961 and 1987; observed in interview years ; current head of household, or spouse or unmarried partner of household head; and whose mothers had non-missing work history information. Statistics weighted by original PSID family weights. See Section 2.3 for information on the computation of log hourly wage offers. Real income and wage figures expressed in terms of 2010 US dollars.

Table 2.3: Estimates of intergenerational transmission of women's full-time employment status

	Logit marginal effects on 2nd gen FTFY employment					
	(1)	(2)	(3)	(4)	(5)	(6)
1st gen FTFY employment	.108***	.105***	.101***	.112***	.119***	.113***
	(.032)	(.032)	(.033)	(.034)	(.035)	(.035)
2nd gen FTFY emp rate	.532	.532	.532	.503	.503	.502
Pseudo R^2	.004	.008	.010	.107	.111	.113
N individuals	2225	2225	2076	2204	2204	2063
Controls						
Birth cohort FEs		✓	✓		✓	✓
Background controls			✓			✓
Contemporaneous controls				✓	✓	✓

Notes: See Table 2.2 for sample information. Full-time employment is defined as at least 1600 hours worked in the reference year. The explanatory variable, "1st gen FTFY employment" is the share of years when the individual was aged 6-17 that her mother worked full time. The outcome variable is a binary indicator for the second-generation woman working full-time. Background controls consist of parental education fixed effects and race. Contemporaneous controls consist of the log wage offer, a quadratic in partner earnings, a quadratic in the number of children aged 0-5, the number of children aged 6-10, and a quadratic in the individual's age. The wage offer is imputed for non-working women in a manner described in Appendix section B.7. Coefficient estimates are marginal effects of a 0-to-1 increase in 1st gen FTFY employment, expressed in percentage points /100. Robust standard errors, clustered at the individual level, appear below the coefficients in parentheses. Regressions are weighted by PSID core family sampling weights. Standard statistical significance legend used.

Table 2.4: Heterogeneity in the intergenerational relationship by labor supply margin and *mother's* education level

	Logit marginal effects of 1st gen's employment status on 2nd gen's employment status			
	(1)	(2)	(3)	(4)
1st-gen mom's education	No college		Some college or more	
2nd-gen sample	All	Mom w/ young child only	All	Mom w/ young child only
Employed at all	.059*	.048	.054	.062
	(.034)	(.048)	(.041)	(.055)
2nd gen sample rate	.786	.716	.811	.743
Employed $\geq$ 400 hrs	.094***	.073*	.088**	.107*
	(.034)	(.044)	(.044)	(.058)
2nd gen sample rate	.732	.641	.761	.677
Employed FTFY	.087*	.038	.161***	.201***
	(.046)	(.047)	(.056)	(.064)
2nd gen sample rate	.498	.375	.509	.394
N individuals	1210	911	868	638

Notes: See Table 2.2 for main sample information. Working at all is defined as positive hours worked in the reference year; full-time employment is defined as at least 1600 hours worked in the reference year. Explanatory variables are measured as the share of years when the individual was aged 6-17 that her mother supplied the given amount of labor (any, at least 400 hours, or at least 1600 hours). Columns (2) and (4) restrict the second-generation sample to individual-year observations in which at least one child under the age of 6 is present in the household. Coefficient estimates are marginal effects of 0-to-1 increases in these explanatory variables on the probability that the second-generation individual supplied the given amount of labor, based on the full regression (i.e. specification 6 in the preceding Table). Robust standard errors, clustered at the individual level, appear below the coefficients in parentheses. Regressions are weighted by PSID core family sampling weights. Standard statistical significance legend used.

Table 2.5: Heterogeneity in the intergenerational relationship by labor supply margin and *daughter's* education level

	Logit marginal effects of 1st gen's employment status on 2nd gen's employment status			
	(1)	(2)	(3)	(4)
2nd-gen daughter's education	<4 years college		≥ 4 years college	
2nd-gen sample	All	Mom w/ young child only	All	Mom w/ young child only
Employed at all	.051	.032	.053	.073
	(.033)	(.045)	(.038)	(.059)
2nd gen sample rate	.772	.704	.846	.773
Employed ≥ 400 hrs	.084**	.071*	.083*	.098
	(.033)	(.043)	(.044)	(.065)
2nd gen sample rate	.718	.630	.798	.707
Employed FTFY	.086**	.068*	.187***	.219***
	(.040)	(.042)	(.068)	(.074)
2nd gen sample rate	.479	.360	.548	.427
N individuals	1445	1109	578	421

Notes: See preceding Table.

Table 2.6: The contribution of intergenerational transmission to 21st century changes in mothers' FTFY employment

	Changes in childhood environment, 1966-1983 birth cohorts	
	Skilled mothers	Less-skilled mothers
Estimated intergenerational FTFY employment response	.210	.053
Δ share exposed to full-time-employed mothers	.260	.115
	Changes in FTFY employment, 1999-2016 calendar years	
	Skilled mothers	*GAP*, skilled − less-skilled mothers
Predicted Δ due to intergenerational transmission	.055	.049
Observed Δ	.130	.138
Residual Δ, after controlling for other variables	.107	.112
% of observed Δ explained by intergenerational transmission	**42.3%**	**35.5%**
% of residual Δ explained by intergenerational transmission	**51.9%**	**43.8%**

Notes: The first row reports the intergenerational FTFY employment response coefficient for skilled mothers (column 1) or less-skilled mothers (column 2). See Section 2.5 for details on how these numbers are computed from Tables 2.4 and 2.5. The second row reports the rise in the shares of daughters exposed to full-time-working mothers between the 1966 and 1983 birth cohorts, by skill. Rows 3 thru 7 compute the contribution of these intergenerational responses to the rise in skilled mothers' full-time employment between 1999 and 2016 (column 1), or the rise in the *skill gap* in mothers' full-time employment. For skilled mothers, row 3 equals the product of the first two rows. For the *skill gap*, row 3 equals the product of the first 2 rows for skilled mothers net of the product of the first 2 rows for less-skilled mothers. The fifth row reports the residual change in skilled mothers' FTFY employment behavior between 1999 and 2016 (or the residual change in the skill gap), *after controlling for movements in other observed determinants of mother's FTFY employment*. The contribution of intergenerational transmission to the total observed change (the sixth row) equals the third row divided by the fourth row. The contribution of intergenerational transmission to the residual change (the seventh row) equals the third row divided by the fifth row.

2.8 Figures

Figure 2.1: The change in environment, from mother-at-home to mother-at-work, experienced by daughters growing up in the late 20th century

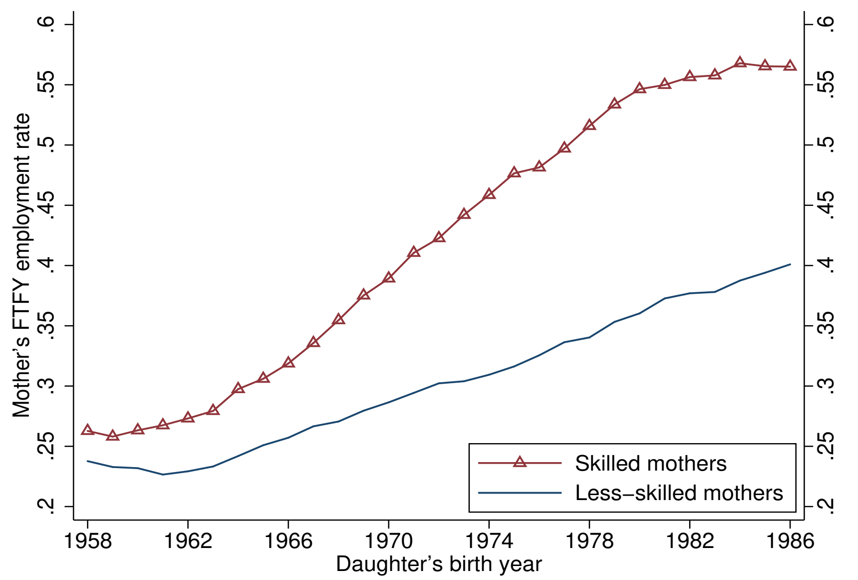

Notes: Author's computations based on March CPS data. To construct the series, I considered all girls aged 6-17, born in a given year, and with a mother of a given skill level present in the household. I computed the mothers' FTFY employment rate, and then iterated this process over the 1958-1986 birth cohorts of girls, and over both skill levels of mothers. The skill groups are constructed by imputing wage offers for non-working mothers and then classifying mothers as above or below the median of the wage offer distribution, conditional on survey year and potential experience. See Section 2.5 for further detail. Lines display 3-birth-cohort centered moving averages.

Figure 2.2: The rising skill gap in mothers' FTFY employment in the 21st century

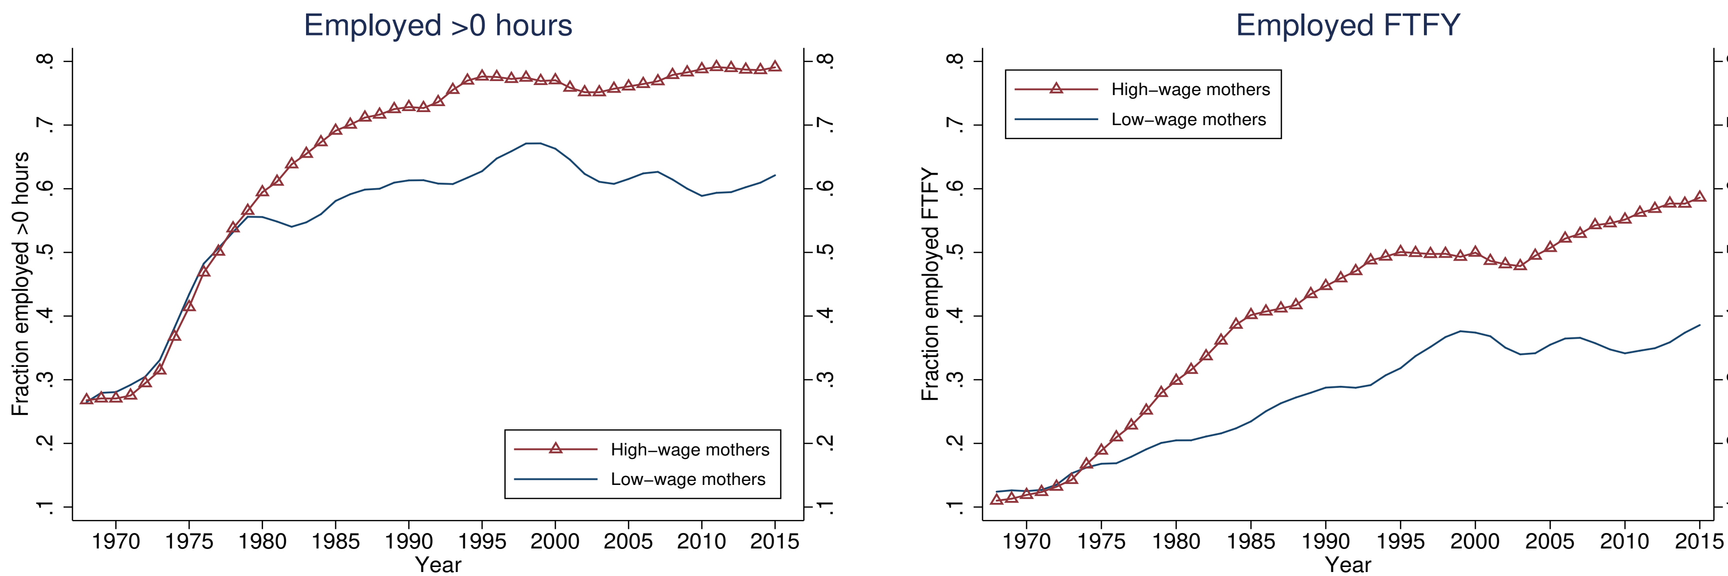

Notes: Author's computations based on March CPS data. See Section 2.5 and Appendix Section B.2 for details on sample construction and data processing. Young mothers are defined as women with at least one child under the age of 6. Full-time, full-year employment is defined as having worked at least 1600 hours in the reference year. The wage groups are constructed by imputing wage offers for non-working mothers and then classifying mothers as above or below the median of the wage offer distribution, conditional on survey year and potential experience. Graphs present 3-year centered moving averages.

Figure 2.3: The 21st century rise in attitudinal support for skilled mothers' careers

Notes on employment series: See section 2.5 of the text for main sample definition. Sample restricted to women with youngest child under 6. The left graph (right graph) further restricts the sample to high-wage (college-educated) mothers, where "high-wage" means having a wage offer above the median of the distribution (see Figures 2.2 and 2.3). Full-time, full-year (FTFY) employment is defined as working at least 1600 hours in the reference year. 3-year centered moving averages reported.

Notes on attitude series: The work-family attitude index is calculated from responses to 3 questions asked in the General Social Survey (GSS) and is coded such that higher values correspond to higher support of mothers working while raising family. See section 2.5.1 of the main text for more details. The questions comprising the index were asked in the 1977 GSS and then on a regular basis beginning in 1985. Due to small sample sizes, 5-year centered moving averages are reported (except for 1977-1985, where the index is linearly interpolated). The target sample is mothers and fathers of young children. Because the GSS does not provide age of youngest child, I considered all individuals aged 22-45 with at least one child. The left graph (right graph) further restricts the GSS sample to individuals with at least 2 (4) years completed college education.

CHAPTER III

Is There a 'Male Breadwinner Norm?' The Hazards of Inferring Preferences from Marriage Market Outcomes

From a work with David Lam

Abstract

Is there a social norm dictating that wives should not out-earn their husbands? Building on the Beckerian marriage model, this paper shows that a variety of underlying social preferences about a given trait—including indifference about the trait—all generate positive assortative matching on that trait, and hence the same distribution of spousal trait differences in equilibrium. Applying this result to U.S. Census and administrative earnings data, we find that simple empirical models of assortative matching closely replicate the observed distribution of spousal earnings differences, in which very few wives out-earn their husbands. Further, we show that the recent finding of a discontinuous drop-off in probability mass as wives start out-earning their husbands is illusory, due to the existence of a point mass of equal-earning couples that compromises econometric inference. We conclude that the distribution of spousal earnings differences in the U.S. provides little information about the existence of a male breadwinner norm or its implications for gender inequality.

JEL Codes: D10 , J12 , J16
Keywords: social norms, gender inequality, assortative matching, dual-earner couples, discontinuity test

3.1 Introduction

Do men prefer to be taller than their wives? Do women prefer to earn less than their husbands? Social scientists often use the attributes of spouses to infer individual preferences and social norms. For example, a number of studies seek to quantify the prevalence of a "male-taller" norm in marriage (Gillis and Avis, 1980; Stulp et al., 2013) and the extent to which this norm affects inter-ethnic marriage patterns (Belot and Fidrmuc, 2010). Other studies look at differences in earnings between spouses (Winkler, 1998; Brennan, Barnett and Gareis, 2001; Raley, Mattingly and Bianchi, 2006), inferring from these patterns social preferences about whether husbands should earn more than their wives as well as implications of these preferences for time allocation, the division of resources, and marital stability (Schwartz and Gonalons-Pons, 2016).

This paper demonstrates that the standard Beckerian marriage model (Becker, 1973) generates matching patterns that suggest social norms of husbands being taller and earning more than their wives, *even when individuals prefer the reverse*. Taking the example of height, we show that a broad class of loss functions for deviation from the social norm generates positive assortative matching on height in equilibrium, regardless of what the norm dictates about the ideal spousal height difference (including the absence of any norm). Positive assortative matching together with the prevailing gender gap in height results in an equilibrium in which few husbands are shorter than their wives—–even if husbands strictly prefer to be shorter than their wives. While this result is based on features of the Beckerian marriage model, we argue that it largely applies in alternative conceptualizations of the marriage market.

We apply this perspective to the context of earnings differences between spouses, a topic which has garnered significant attention following a prominent paper by Bertrand, Kamenica and Pan (2015). With data drawn from the 2000 United States Census, we use the Beckerian framework to match couples based on earnings. We consider two models: one in which observed earnings are taken as given, and a second in which an endogenous labor supply decision is made after marriage (to account for the fact that earned income is not an exogenous attribute). We assume only positive assortative matching—there is no explicit preference for wives to earn less than husbands. Even without imposing this norm, our simulations succeed in reproducing the highly skewed distribution of spousal earnings differences observed in the data. That is, there are far fewer wives out-earning than their husbands than vice versa, even though positive assortative matching is consistent with a wide class of preferences—including a preference for wives to out-earn husbands. These simulation exercises illustrate that a naive interpretation of marital matching patterns may produce incorrect inferences about underlying preferences.

The empirical strategy pursued by Bertrand, Kamenica and Pan (2015) (hereafter BKP) represents a compelling addition to the literature. Unlike other studies, BKP did not rely upon broad features of the distribution of attributes in marriage. Instead, they tested whether the distribution of the wife's share of total spousal earnings was continuous across the 50 percent threshold—–the point at which the wife goes from earning just less to just more than her husband. They found a discontinuous drop-off in probability mass across this threshold, suggesting that couples manipulate their earnings to avoid a situation in which the wife out-earns her husband. Without assuming an explicit social norm that wives should not out-earn their husbands, it is difficult to replicate this discontinuity in our simulated matching models, implying an important role for such a norm in marital matching and earnings outcomes within marriage.

Further investigation into this discontinuity result, however, suggests that it is fragile. One issue with the spousal earnings data investigated by BKP is the presence of a mass of couples earning exactly identical incomes. This generates a mass point in the distribution of the wife's share of total earned income at 50 percent. BKP tested for a discontinuity just to the *right* of 50 percent, consistent with testing for a social norm that the wife should not strictly out-earn her husband. Using the same data source,[1] we first replicate BKP's result of a sharp drop-off in probability mass across this threshold. However, when we also test for a discontinuity just to the *left* of 50 percent, we find evidence of a sharp *gain* in probability mass. This sharp gain in mass as one moves from left of 50 percent (where the wife earns less than the husband) to 50 percent (equality) could be interpreted as evidence for a social norm that the wife should earn at least as much as her husband. Thus, the data appear consistent with two nearly opposite social norms. We show that the point mass of equal-earning couples is responsible for these seemingly inconsistent results: omitting these couples eliminates the estimated discontinuities. Moreover, we argue that the point mass does not necessarily reflect social norms, but instead could be generated by (small) search frictions within the marriage market or by couples who jointly a business.

Two contemporary studies also probe the robustness of BKP's discontinuity result and reach similar conclusions. Hederos and Stenberg (2019) find, in Swedish administrative income data, that the existence of a point mass of equal-earning couples compromises the discontinuity test, and that the point mass itself is not indicative of a male-breadwinner norm. Zinovyeva and Tverdostup (2018) observe Finnish data and argue that the point mass results from a tendency of spouses to start a business together or co-work at the same establishment. Our work thus contributes to a growing body of evidence that the earnings behavior of couples who earn close to the same amount of income does not reflect male-breadwinner norms.

Our observations are also related to literature on the econometric identification of

[1] The data are administrative earnings data from the Social Security Administration. These data are linked to a household survey (the Survey of Income and Program Participation), which permits the researcher to observe earnings of matched couples. Section 3.4 provides further discussion.

matching models. A common theme within this literature is that the observation of matching patterns alone is insufficient to identify marital preferences (see, e.g., Echenique et al., 2013; Chiappori and Salanié, 2016). One strand of work has studied the identification of average marital returns to individual attributes with certain structural assumptions on idiosyncratic preferences (Choo and Siow, 2006; Galichon and Salanié, 2017). This class of models has proven useful in estimating the degree of marital complementarity between individual attributes like education (Dupuy and Galichon, 2014; Chiappori, Salanié and Weiss, 2017). Our paper studies the identification not of marital complementarities over *individual attributes*, but of societal preferences over *spousal attribute gaps*. We argue that when social norms dictate an ideal attribute gap, individual attributes will exhibit complementarity (and hence positive sorting) within marriage, *regardless of what is the ideal attribute gap*. Thus, the strong degrees of positive assortative matching on height, earnings and age that exist in the data appear particularly uninformative about male-taller, male-breadwinner or male-older norms. While there may be room for progress in the literature on this front, we believe that fruitful study of social norms in marriage likely requires the observation of more than just marital sorting patterns (e.g. divorce, labor supply, home production time, and subjective indicators of marital satisfaction).

3.2 Applying Becker's Theory of Marriage to the Study of Social Norms

We found our study of social norms on Gary Becker's economic theory of marriage (Becker, 1973, 1981). To start, consider a man m and a woman f who are considering marriage. We assume they marry if and if only if it makes both better off compared to alternatives. Denote the "output" of the marriage by Z_{mf}. Assume output can be divided

$$Z_{mf} = m_{mf} + f_{mf}, \tag{3.1}$$

where m_{mf} indicates what man m consumes when married to woman f. That is, it is possible for men to make offers to potential wives (and women to make offers to potential husbands) of some division of output.

Suppose there are multiple men and multiple women considering marriage. Drawing on results from other matching models in mathematics and economics, Becker showed that a competitive equilibrium in this marriage market will be the set of assignments that maximizes the sum of output across all marriages. The proof relies on a standard argument about the Pareto optimality of competitive markets. If an existing set of pairings does not maximize total output, then there must exist at least two couples who could switch partners and increase total output. Because output is transferable, it is possible to distribute the total output gains from the switch such that everyone is made better off.

Becker applied this result to the case of sorting on some trait A, where we will consider woman f to have a trait value A_f and man m to have trait value A_m. We characterize marital output as a function of the values of A for each partner: $Z_{mf} = Z(A_m, A_f)$. Becker showed that the marriage market equilibrium will consist of positive assortative matching on A if

$$\frac{\partial^2 Z(A_m, A_f)}{\partial A_m \partial A_f} > 0, \tag{3.2}$$

and negative assortative matching if the cross-partial in (3.2) is negative. If, for example, having a better-educated husband raises the impact of the wife's education on marital output, then we will tend to see positive assortative matching on education. We draw on this well-known result below.

3.2.1 Illustrative model of sorting on height

We build up to our key point with a very simple model of marital sorting on height. Denote female height by H_f and male height by H_m. Suppose there are two women: f_1 is $60''$ tall and f_2 is $66''$ tall. There are two men: m_1 is $66''$ tall and m_2 is $72''$ tall. Thus, there are two possible pairings: $(f_1 m_1, f_2 m_2)$, which is positive assortative matching on height, and $(f_1 m_2, f_2 m_1)$, which is negative assortative matching on height.

Assume that people get utility from their individual consumption and some bonus that comes from being married. The gains from marriage take the very simple form of some bonus K (representing, say, economies of scale in consumption or benefits of household public goods) that is offset by a quadratic loss in the deviation of the given spousal height difference from the social norm.

Consider three alternative social norms: a $6''$ male-taller norm, an equal-heights norm, and a $6''$ female-taller norm:

$$\begin{aligned}
Z(H_m, H_f) &= K - (H_m - H_f - 6)^2 && \text{[Male taller norm]} \\
Z(H_m, H_f) &= K - (H_m - H_f)^2 && \text{[Equal heights norm]} \\
Z(H_m, H_f) &= K - (H_f - H_m - 6)^2. && \text{[Female taller norm]}
\end{aligned}$$

It is trivial to show that a $6''$ male-taller norm will give rise to positive assortative matching on height in equilibrium, as this yields two marriages in perfect compliance with the norm. Notice, however, that positive assortative matching will also prevail in the case of an equal-heights norm. This is because, even though the alternate sorting yields one couple in perfect compliance with the norm (the short man and tall woman), the other couple has a height difference of $12''$ and produces a total loss of 144. The positive sorting assignment yields a lower total loss of $36 + 36 = 72$, since each couple is $6''$ from the ideal height difference. Therefore, positive sorting is the competitive equilibrium.[2] The same is true in the case of a $6''$ female-taller norm: negative sorting creates a total loss of $36 + 324 = 360$, while positive sorting creates a total loss of $144 + 144 = 288$. Hence, *all* 3 *social norms are consistent with the same competitive equilibrium: positive assortative matching.*

3.2.2 A more general model of marital matching

This non-identifiability of social norms regarding height generalizes to cases with large numbers of women and men covering a broad range of heights.

Proposition 3.1. *Consider a population with N men and N women, and assume the following marital surplus function:*

$$Z(H_m, H_f) = K - \lambda(g_{mf}), \tag{3.3}$$

where the male-female height gap $g_{mf} = H_m - H_f$. If λ is (strictly) convex in g, then strict positive sorting on height is a (the unique) marriage market equilibrium.

Proof. *Note that $\frac{\partial^2 Z(H_m,H_f)}{\partial H_m \partial H_f} = \lambda''(g) > 0$ by strict convexity of λ. Thus the payoff function satisfies condition (3.1), and so strict positive sorting on height is the unique marriage market equilibrium. If λ is merely convex, then, starting from strict positive sorting, no exchanges of partners can be made that strictly increase total marital output. Hence, such a sorting is an equilibrium.* ∎

This result establishes that positive sorting on height can be consistent with both male-taller and female-taller norms, as well without any explicit norms at all. Whatever the ideal spousal height gap may be, convex losses for deviating from the ideal implies that marital output is supermodular in spousal heights, creating a tendency for positive sorting to arise in equilibrium. The next result illustrates that if there is a gender gap in the

[2] Here is a system of transfers that would support such an equilibrium: suppose we began with the sorting in which one couple has equal heights while the other couple has a $12''$ height difference. The individuals in the mismatched couple, f_1 and m_2, see that their total marital surpluses would be higher if they could switch partners and have a $6''$ height difference instead of a $12''$ height difference. The question is whether f_1 would be able to induce m_1 to switch from f_2 to her. Her loss would decline from 72 (half of 144) to 18 (half of 36) if she changed partners. The loss for m_1 would increase from 0 to 18 if he switched partners. Clearly, f_1 can compensate m_1 for switching by making him a side payment of between 18 and 56 ($= 72 - 18$), while still leaving herself better off from the switch. The exact same story can be told for m_2 inducing f_2 to switch to him. Thus, every person will be better off after the re-sorting, even though the original sorting yielded one couple in perfect compliance with the norm.

attribute distributions, as is the case for height and for earnings, the equilibrium implied by positive sorting is highly skewed in nature.

Proposition 3.2. *Consider a population with N men and N women, and assume the marital surplus function is given by equation (3.5). If the male height distribution exhibits first order stochastic dominance (FOSD) over the female distribution,*[3] *then there exists a marriage market equilibrium in which no wives are taller than their husbands. Moreover, if the loss function exhibits strict convexity, this equilibrium is unique.*

Proof. *By Proposition 3.1, a marriage market equilibrium characterized by strict positive sorting on height exists, and is unique if the loss function is strictly convex in the height gap. In a strict positive sorting equilibrium, the spouses of each couple have heights of identical rank in their respective distributions. Therefore, by the FOSD assumption, the husband is taller than the wife in each couple.* ∎

These results suggest important implications for the study of social norms in marriage. To the extent that society has preferences about the ideal height, age or earnings gap between spouses, competitive forces operating within the marriage market will push the equilibrium toward perfect rank-order sorting on these traits. And if men in the marriage market tend to be taller, older or higher-earning than comparable women, positive assortative matching leads to a situation in which the large majority of wives are shorter, younger or lower-earning than their husband—suggesting a male-taller, male-older or male-breadwinner norm when one may not exist.[4]

It is important to note that the results depend on the convexity of the loss function. This creates a situation in which the marginal loss from further deviation from the social norm is larger when the current deviation is larger. With this structure, the competitive equilibrium is not generally the one which maximizes the number of couples in perfect compliance with the social norm—–it is instead one in which many couples may deviate from the norm by small amounts, but few couples deviate by large amounts. Note that increasing marginal losses from deviating from the social norm is isomorphic to diminishing marginal returns from greater adherence to the social norm. Social welfare is thus maximized when this commodity (level of adherence to the norm), is distributed as equally as possible in the population. Such a situation is realized by positive assortative matching.

Note also that the convex structure nests "kinked" loss functions. For example, a struc-

[3] That is, at any common rank in the distributions, the male attribute is larger than the female attribute. Although this may sound like a strong assumption, it is quite realistic in the cases of both height and earnings. For example, FOSD holds for the earnings distributions of husbands and wives in the 2000 US Census, the data used in our empirical investigation of income differences between spouses (see section 3.3).

[4] This conclusion generalizes the findings of Belot and Francesconi (2013), in British speed dating data, that the pool of potential partners appears to be more important than underlying preferences in the determination of who matches with whom.

ture in which households are indifferent when the husband is taller (shorter) than the wife, but face a loss increasing in the amount by which the husband is shorter (taller) than the wife, is also consistent with strict positive assortative matching on height in equilibrium (though because kinked loss functions are not necessarily strictly convex, such an equilibrium may not be unique). Fisman et al. (2006) found evidence of kinked preference structures in a speed dating experiment implemented with Columbia University graduate students: men appeared to value female intelligence and ambition only when it did not exceed their own.

3.2.3 Extensions of the model

Propositions 3.1 and 3.2 predict an equilibrium in which no wives are taller than their husbands (or, analogously, no wives earn more than their husbands). This stark result is clearly counterfactual. Several factors are presumably at work in actual marriage markets—couples do not match on a single trait, there are search frictions, there is not perfectly transferable utility, etc. We consider some of these issues below.

Sorting on multiple attributes. It is important to consider matching on multiple attributes in our current context: if the economic gains to marriage depend on attributes other than height, then the distribution of height gaps in marriage will clearly depend on how these attributes are correlated with height in the population.

As a simple example, suppose there is an additional attribute X which enters the marital output function, such that the economic gains from marriage unrelated to the height gap are no longer constant:

$$Z_{mf} = Z(X_m, H_m, X_f, H_f) = K(X_m, X_f) - \lambda(g_{mf}). \tag{3.4}$$

Suppose that X is positively correlated with H, that X positively affects marital output, and that $\frac{\partial^2 K(X_m,X_f)}{\partial K_m \partial K_f} > 0$. Thus, K satisfies sufficient condition (3.1) for positive assortative matching on X in equilibrium, while λ generates positive assortative matching on H by Proposition 3.1. It is impossible to know without further assumptions whether the prevailing equilibrium will consist of positive sorting on X, on H, or on some function of X and H. However, given that X and H are positively correlated in the population, some degree of positive sorting on H must exist in equilibrium. Therefore, given a significant gender gap in H, this model still predicts that an equilibrium in which few wives are taller than their husbands is consistent with a variety of social preferences over the spousal height gap. There could also be no social preferences regarding height whatsoever—λ could be constant—yet the positive correlation between X and H would still lead to an equilibrium making it look as if a male-taller norm exists.[5]

As an example, the work of Mansour and McKinnish (2014) illustrates subtle sorting patterns when individuals care about multiple attributes. Mansour and McKinnish observed that couples in which the husband is significantly older than the wife tend to be negatively selected on education and earning potential. They argued that this pattern results from the fact that higher-earning individuals tend to locate in marriage markets with more similarly-aged individuals. That is, the preponderance of "husband-significantly-older" couples among relatively low-earning individuals may have little to do with a husband-significantly-older norm in this population.

Non-transferable marital surplus. The Becker model assumes that the gains from marriage are fully transferable between spouses via monetary payments. In this setup, the allocation of marriages and the transfers are determined in equilibrium as prospective partners make binding agreements in the marriage market. If the division of the marital surplus cannot be negotiated in the marriage market—that is, bargaining over the marital surplus occurs after marriage, or bargaining weights are fixed ex ante——the market clears on the basis of what prospective partners expect to obtain from bargaining within marriage.[6] Pollak (2019) notes that such a setup is consistent with using the Gale-Shapley framework (Gale and Shapley, 1962) rather than Becker's framework to analyze the marital equilibrium.

When the marital surplus is fully non-transferable, equilibrium marital outcomes may have the potential to offer some identifying information about the underlying social norm. For example, if the ideal is for husbands to be two inches taller than wives, we might expect to see a point mass at two inches in the height gap distribution. Though such an allocation would likely result in some very tall men matching with some very short women, the impossibility of credible side payments prevents these individuals from attracting more suitable partners.

This situation strikes us as unrealistic. For example, consider the man who is just taller than the male partner of the tallest woman. This man would be forced to match with one of the shortest women and realize a very low marital surplus. He stands to gain a tremendous amount from matching with the tallest woman, and the total marital surplus of her marriage would shrink only very slightly from matching with him. Thus, even a

[5]A systematic analysis of which characteristics significant affect the marital surplus was performed by Dupuy and Galichon (2014). Applying the analysis to rich Dutch data on married couples, the authors uncovered significant complementarities between spousal education, height, health and personality traits. These results underscore the points that multiple characteristics influence marriage formation, and that the patterns in the data support the prevalence of assortative matching on these characteristics.

[6]For a survey of the implications of household bargaining models for distribution of resources within marriage, see Lundberg and Pollak (1996).

modest degree of credible divisibility of the marital surplus might inspire a reshuffling of partners in this context.[7] The resulting equilibrium will therefore blend elements of positive sorting with a cluster of couples with height gaps near 2 inches. Its exact nature will depend not just on preferences but on the prevailing height distributions and market size (i.e. the availability of close partner substitutes), and on the efficacy of the transfer technology.

It is also worth noting that in the case of fully non-transferable utility, the equilibrium allocation of partners is generally not unique (Roth and Sotomayor, 1992). Thus, even if one is willing to assume fully non-transferable utility, additional structural assumptions are likely necessary for identification of preferences.

3.2.4 Application to analysis of spousal height differences

Two recent studies of spousal height differences helps illustrate our point about the difficulty of inferring preferences from equilibrium matches. Stulp et al. (2013) analyzed the distribution of height differences among couples in the United Kingdom's Millennium Cohort Study. They compared the actual distribution of height differences to hypothetical distributions based on random matching, drawing several inferences based on this comparison. Table 3.1 presents their data, divided into bins of 5 cm (2 inch) height differences. A key observation is that the actual distribution has fewer women who are taller than their husbands than would occur through random matching. The authors argue that this is consistent with a "male-taller" norm. They also interpret the data as supporting a "male-not-too-tall" norm, since there are fewer men who are more than 25 cm taller than their wives than would occur through random matching.

It is easy to see that the data are consistent with other social norms as well. These include what might be called a "wife-not-too-short" norm or a "heights-not-too-different" norm. In fact, a better way to describe the norm implied by Table 3.1 might be a norm to keep the difference in heights between husbands and wives close to the overall average difference in heights between men and women in the population. The three bins closest to the actual average height difference of 14.1 cm (5.5 inches) are the bins that occur more frequently in the actual distribution than in the random matching distribution. The bins with the height differences farthest from 14.1 cm are the bins that occur with the lowest frequency relative to random matching. Sohn (2015) found similar sorting patterns in Indonesia: spousal height gaps closest to the mean gap occurred more frequently in the observed data than in hypothetical random sortings, while those furthest away from the

[7] For example, for every \$10 the man promises to transfer to her, she believes she will only actually receive \$2 in marriage. This is a situation of *imperfectly transferable utility*.

mean occurred less frequently. Notice that this is exactly what will happen if there is a tendency for positive assortative matching on height, as this pushes the equilibrium toward an outcome in which the height gap is uniform across all marriages. Thus, it is likely that a wide variety of underlying preferences could produce these observed distributions of spousal height gaps.

3.3 The Skewed Distribution of Spousal Earnings Differences in the U.S. does not Imply a Male Breadwinner Norm.

We apply the above insights to an empirical investigation of earnings differences between spouses, where, like height, a persistent gender gap also exists. As gender gaps in wages and hours worked have shrunk throughout the late 20th century, the proportion of wives earning similar to or more than their husbands has risen (Winkler, 1998; Winkler, McBride and Andrews, 2005). However, this proportion remains small (Bertrand, Kamenica and Pan, 2015), and the gender gap in earnings remains substantial (Blau and Kahn, 2017). A tendency for positive sorting combined with this gender gap would lead to a skewed marriage market equilibrium in which most husbands out-earn their wives—— even if there is no social norm dictating this outcome. Whether the observed gender gap in earnings is influenced by gender norms operating in the home, independent of the labor market, is an important question, with implications for family and labor market policy. Our analysis indicates that this question cannot readily be answered by analyzing the distribution of spousal earnings differences.

Our approach is to simulate marriage market equilibria using observed earnings in U.S. Census data and simple matching processes. We start by matching men and women randomly. Next, we match women and men assuming positive assortative matching on observed earnings perturbed with noise (to approximate characteristics other than observed earnings influencing the marriage market). Finally, recognizing that earnings is not an exogenous attribute but is affected via a labor supply decision, we assume positive sorting on unobserved potential earnings and endogenize labor supply choices made after marriage. The goal of these exercises is to illustrate that simple matching models, which do not include any specific preference about the husband earning more than the wife, can reproduce the highly skewed observed distribution of spousal earnings differences.

Following BKP, we summarize spousal earnings differences by plotting the distribution of *the share of the couple's total earnings that was earned by the wife*. Thus, 0.01 indicates that the wife earned 1 percent of the couple's total earnings, and 1.0 indicates that she earned all of it. 0.50 represents a couple in which wife and husband earned equal amounts.

3.3.1 Empirical distributions of spousal earnings differences

We begin with a sample of men and women drawn from the 5 percent sample of the 2000 U.S. Census (Ruggles et al., 2015). Following BKP, we restrict the sample to couples ages 18-65 and process earned income variables following the procedure outlined in that paper's main text and appendix. We keep only couples in which both spouses report positive earnings. Figure 3.1 displays two 20-bin histograms of the distribution of the share of total earnings earned by the wife: the one published in BKP and our replication. As in BKP, we apply a local linear smoother to the histogram bins, allowing for a break in the smoothed distribution at 0.50. The two distributions are almost identical, and both display a substantial reduction in probability mass to the right of 0.50.

For our simulation exercises, we further restrict the sample to relatively young couples (aged 18-40) without children. Our final sample consists of $109,569$ dual-earning couples. Figure 3.2 plots the sample distribution of the wife's share of total earnings. The main difference between this distribution and that in Figure 3.1 is that there is less mass below 0.25, which likely reflects the impact of specialization after childbearing.[8] Our simple simulations are not set up to handle the dynamic considerations of fertility and its effect on the wife's labor supply and earning potential. Nonetheless, imposing this sample restriction does not change the fact that most of the distribution lies to the left of 50 percent (where the wife earns less than the husband), and the probability mass drops sharply as one moves to the right of 50 percent. These are the stylized facts we will attempt to replicate in the following exercises.

3.3.2 Simulated distributions

Random matching of couples. Figure 3.3 displays a smoothed distribution of the wife's share of total earnings based on random matching, again allowing for a break at 0.50, overlaid on the observed distribution. Like the observed distribution, the distribution generated by random matching contains a mode around 0.42 and a drop-off in mass to the right of that point. Moreover, significantly fewer wives slightly out-earn their husbands than vice versa; the point of equal earnings (0.50) corresponds to the $70th$ percentile of the distribution. This benchmark exercise demonstrates that the prevailing male and female earnings distributions exert a strong influence on spousal earnings differences.

[8]This additional restriction is motivated by the well-known fact that women disproportionately reduce their working hours or exit the labor force to raise young children and later re-enter the workforce with lower earnings potential (Mincer and Ofek, 1982; Hotchkiss and Pitts, 2007; Attanasio, Low and Sánchez-Marcos, 2008; Bertrand, Goldin and Katz, 2010). We abstract from this endogenous specialization decision after childbearing. BKP's Appendix Figures A.1 and A.2 show similar effects of children and marital tenure on the observed distribution of the wife's share of total earnings.

Notice that Figure 3.3 follows a similar pattern to the distribution of height differences shown in Table 3.1. The bins in Figure 3.3 that occur more frequently in the actual distribution than in the distribution with random matching are those closest to 0.42, the average wife's share of total earnings. (Although Figure 3.3 is in shares rather than differences, the pattern would look similar if plotted in absolute or proportional income differences.) A key feature is that the actual distribution is pushed toward the mean earnings difference and away from extremes, exactly as in our simple theoretical examples.

Positive assortative matching on potential earnings. We take male and female earnings as observed in our sample (denoted as Y_i^m for males and Y_i^f for females). We create couples by rank-order-matching individuals not according to observed earnings, but rather observed earnings perturbed with noise. That is, for each individual i of gender g we assign

$$W_i^g = Y_i^g + u_i, \tag{3.5}$$

where u is normally distributed white noise, and pair up males and females according to their ranks of W. This is consistent with at least two interpretations. One interpretation is that couples are perfectly sorted on permanent earning potential and the white noise represents transitory earnings shocks realized after marriage. A second is that men and women care about other characteristics as well as earnings, or that assortative matching on earnings is imperfect, for example due to the presence of search frictions. Under the latter interpretation, equilibrium sorting on observed earnings plus noise is the reduced form of a more complicated matching process.

Figure 3.4 displays the distribution of the wife's share of total earnings, simulated from this simple model, overlaid on the actual distribution.[9] The simulated distribution is very similar to the actual distribution: it exhibits a sharp drop in mass across the 50 percent threshold and contains few couples in which the wife out-earns her husband.[10] Thus, given the gender gap in earnings distributions, the observed distribution of spousal earnings differences is largely consistent with positive assortative matching on earnings. As the previous section indicates, this matching is consistent with a wide variety of underlying

[9] The standard deviation of u is set to 16,000 for this simulation and is chosen to match the observed data. This choice is slightly larger than the standard deviation of transitory earnings for males in 2000 implied by the numbers reported in Gottschalk and Moffitt (2009). Thus, one might prefer to interpret the simulation as reflecting both elements of transitory earnings variance and imperfect positive sorting on potential earnings.

[10] The successful fit of this simulation is striking when we consider the fact that the simulation assumes one national marriage market. If we instead considered separate marriage markets defined by state and age (or, state, age and ethnicity), and allowed the prevailing earnings distributions and choice of noise term to vary by marriage market, we would (by greater modeling flexibility) be able to replicate the aggregate distribution of spousal earnings differences even more closely. The point remains that a simple matching model with no explicit social norm broadly succeeds in replicating the data.

preferences. It could be based on a desire for equality in spousal earnings, a preference for wives to earn more than their husbands, or economic gains from marriage related to household public goods (i.e. with no explicit preference for equal or unequal spousal earnings).

Positive assortative matching on potential earnings with endogenous labor supply. One shortcoming of the previous exercise is that it treats the observed distributions of men's and women's earnings as fixed attributes, determined outside of the household. Yet, a literature dating back to Becker (1981) argues that household incentives, such as gains from specialization, influence spousal labor supply choices. Moreover, BKP argue that social norms themselves may influence how many hours a wife chooses to work in the market: if she is at risk of out-earning her husband in a full-time job, she may work fewer hours. To address this shortcoming, we endogenize the wife's earnings via a simple labor supply model—that does not assume an explicit male-breadwinner norm—and explore the model's predictions about the distribution of spousal earnings differences.

We assume that, for a given male m and female f, the match output function is given by

$$Z_{mf} = Z(Y_m, Y_f, P_f) = \frac{C_{mf}^{1-\gamma}}{1-\gamma} - \psi P_f, \tag{3.6}$$

with $C_{mf} = 0.61(Y_m + Y_f P_f)$; where C is consumption of a composite good, Y_m and Y_f denote each spouse's permanent income, P_f is the wife's labor supply decision (constrained to be in the unit interval), γ is the coefficient of relative risk aversion, and ψ is the disutility incurred by the household if the wife works.[11] This specification of household utility has been used in recent work investigating determinants of wives' labor supply (e.g. Attanasio, Low and Sánchez-Marcos, 2008 and subsequent papers). It assumes household consumption is a public good with congestion; the 0.61 is a McClements scale calibration capturing consumption economies of scale. In this setup, where spouses consume an indivisible public good, positive sorting on permanent earnings occurs in marriage market equilibrium so long as each member's permanent earnings positively affects match output (Becker, 1973, 1981). It is trivial to show that this holds here (regardless of the wife's eventual labor supply decision).

After marriage, the household takes household potential income as given and chooses the wife's labor supply $P_f \in [0, 1]$ to maximize the above utility function.[12] With an interior

[11]This parameter could capture specialization incentives or social norms. Notice, however, that the disutility faced by the household is continuous in the wife's labor supply decision—–it does not change discontinuously if the wife supplies enough labor to out-earn her husband. Thus there is no discontinuous incentive for the wife to earn less than her husband.

[12]We assume the household acts as a unitary decision-maker, committing to equation (3.9) at the time of

solution, the household will choose

$$P_f^* = \frac{\frac{1}{0.61} \cdot \left(\frac{\psi}{0.61 \cdot Y^f}\right)^{-\frac{1}{\gamma}} - Y^m}{Y^f}. \tag{3.7}$$

If P_f^* lies outside of the unit interval, the appropriate corner solution applies.

We calibrate the model by imposing $\gamma = 1.5$, a standard value estimated in the macro literature. We assume log-normally distributed potential earnings and allow the work disutility parameter, ψ, to be heterogeneous in the population and negatively correlated with Y_f. In total there are 8 remaining parameters, which we calibrate by targeting 8 moments in our observed data: the means and standard deviations of male and female log observed income, the observed mean gender earnings ratio conditional on earning positive income ($P_f^* > 0$), the observed mean gender earnings ratio conditional on full-time work (defined in the data as at least 1600 hours worked in the last calendar year; defined in the model as $P_f^* > 0.95$), the female employment rate (defined in the data as the share of wives working positive hours in the last calendar year), and the female full-time employment rate. We do not target any moment related to marital matching or spousal earnings differences, as doing so would threaten the external validity of our inferences.

Table 3.2 summarizes the calibration. Overall the model does a good job of replicating the targets in the data. With the calibrated model we simulate the distribution of the wife's share of total spousal earnings (Figure 3.5). The simulated distribution again matches the actual distribution quite closely, delivering a sharp drop in probability mass at the 0.50 threshold.

3.4 Discontinuity or Point Mass? Assessing Alternative Evidence for a Male Breadwinner Norm

The above evidence suggests that social scientists wishing to test the importance of social norms need to find strategies beyond interpreting skewed distributions of spousal attributes. The challenge in doing so makes the discontinuity found by Bertrand, Kamenica and Pan (2015) at the equal-earnings threshold a compelling addition to the literature. The logic behind BKP's discontinuity test runs as follows. Suppose we observe the distribution of the share of total spousal earnings that was earned by the wife in the neighborhood of 0.50. Suppose we find that this distribution exhibits a sharp change in probability mass at the 0.50 threshold—that is, there are far fewer wives barely out-earning

marriage and then choosing P_f^* after observing the earnings shocks.

their husbands than husbands barely out-earning their wives. Because standard models of the marriage market, involving agents optimizing continuous utility functions, should not generate discontinuous equilibrium distributions, such a finding suggests the existence of a utility penalty that applies if and only if the wife out-earns the husband.

BKP estimated a discontinuous drop-off in probability mass across the equal-earnings threshold in a variety of Census samples. However, inference is complicated by the fact that earnings are not precisely measured in Census survey data. Mis-measurement occurs for several reasons. First, earnings are reported, rather than measured directly. (Moreover, earnings for both spouses are typically reported by one household member.) Second, earnings are imputed for individuals who do not answer earnings questions, and the earnings of high-earning individuals are top-coded at a common value. Third, reported earnings are rounded by the U.S. Census Bureau in public-use samples, usually to the nearest thousand. These issues create a large point mass of couples with identical earnings. Even after employing several procedures to adjust the data, BKP still found that around 3 percent of dual-earning Census couples have identical earnings. (We corroborate this finding.)

To overcome these limitations, BKP also assembled a sample of administrative earnings records from the Social Security Administration (SSA). These data have been linked to a household survey (the Survey of Income and Program Participation, or SIPP) which allows couples to be identified.[13] In this administrative data sample, the point mass of equal-earning couples still exists but is much smaller: only around one quarter of one percent of all dual-earning couples earn identical incomes. BKP obtained a similar discontinuity result in this sample.

Without the point mass, the straightforward way to implement BKP's procedure would be to test for a discontinuity in the distribution exactly at 0.50, and interpret the finding of a significant drop-off in the density function as evidence for a social norm that the wife should not out-earn her husband. The presence of the point mass presents a challenge, which BKP acknowledge in footnote 7 of their paper. To circumvent this problem, they tested for a discontinuity just to the right of 0.50. One might interpret this treatment of the data as a test for whether there is a social norm dictating that a wife should not *strictly* out-earn her husband. Their finding of a significant drop-off in the density function to the right of 0.50, combined with the presence of the point mass of equal earners, might suggest that couples manipulate their earnings on the margin to comply with such a social

[13]The data come from a pre-linked and cleaned Census Bureau data product called the Gold Standard File (GSF). Users work with synthetic versions of the data remotely and then have Census run final programs internally on the actual GSF, subject the output to a disclosure review, and then release the output. More information can be found in Benedetto, Stinson and Abowd (2013) and here: http://www.census.gov/programs-surveys/sipp/guidance/sipp-synthetic-beta-data-product.html.

norm.

This treatment of the data seems sensible *a priori*, but the existence of the mass point violates one of the assumptions required by the discontinuity test, namely, that the distribution is continuous everywhere except possibly at the supposed breakpoint (McCrary, 2008). Like a non-parametric regression discontinuity design, the test involves local linear smoothing of a finely-binned histogram on either side of the supposed breakpoint, and asymptotic inference is based on the size of the bins shrinking to zero at the correct rate as the number of observations increases to infinity. In BKP's application of the test, for a small bin size, the bin immediately before the breakpoint will (by containing the point mass) be taller than the bin immediately after the breakpoint. This could exert undue influence on the discontinuity estimate, especially if a small bin size and bandwidth is used to perform the test.

3.4.1 Gauging the robustness of BKP's discotninuity test results

To investigate the sensitivity of the discontinuity test to the presence of the point mass, we replicate BKP's SIPP-SSA data sample and analysis. BKP constructed a sample of earnings data for all dual-earning couples aged 18 to 65 observed in the first year they were in the SIPP panel. They considered SIPP panels 1990 through 2004. We construct a sample according to the same conditions but include the 1984 and 2008 SIPP panels as well, which are available in the most recent version of the SIPP-SSA data product. We obtain a sample of around $83,000$ couples—about $9,500$ more than in BKP's sample. Despite using a slightly different sample, the resultant distribution of the wife's share of total spousal earnings is virtually identical to BKP's, as illustrated in Figure 3.6.

In our replicated sample, 0.21 percent of all dual-earning couples earn identical incomes, compared to 0.26 percent in BKP's sample. To see the impact of this mass point on the distribution, Figure 3.7 zooms in on the portion of the distribution between 45 and 55 percent, displaying histograms with a very small bin size of 0.001 (about the size used in the discontinuity tests). The top histogram retains the mass point, while the bottom histogram removes it. The two histograms look very different: the top one exhibits a large spike right at 0.50, while the bottom one does not. Moreover, though the data are noisy for such a small bin size, the histogram on the right does not look particularly discontinuous at 0.50. These illustrations suggest that the point mass may exert an undue influence on the discontinuity estimates.

Using our sample we perform 3 different versions of the McCrary test for a discontinuity in the distribution at 50 percent, based on three different treatments of the point mass: keeping the point mass and testing for a discontinuity at $.500001$, keeping the point mass and testing for a discontinuity at $.499999$, and deleting the point mass and testing for a discontinuity exactly at 0.50. For each version we use 4 different sets of tuning parameters. McCrary's test procedure involves an algorithm which automatically chooses a bin size for the histogram and a bandwidth within which to apply the local linear smoother to the histogram. McCrary (2008) recommends using a smaller bandwidth than the automatically-selected one (around half the size) to conduct robust asymptotic inference. We consider the automatically selected bandwidth, which in this case is around $.084$; and then bandwidths of $.045$, $.023$, and $.011$. The last bandwidth may be too narrow for optimal statistical inference, but using successively smaller bandwidths allows us to gauge the sensitivity of the test to the presence of the point mass (which becomes increasingly dominant as the bandwidth shrinks).

Table 3.3 reports the discontinuity estimates, which equal the estimated log increase in the height of the density function as one travels from just to the left of the supposed breakpoint to just to the right. A negative number thus indicates a sharp drop and a positive number indicates a sharp gain. The first version of the test replicates BKP's choice of retaining the point mass of couples and testing for a discontinuity just to the right of 50 percent ($.500001$). With the standard bandwidth and bin size, we estimate that the density function drops by a statistically significant 12.4 percent across the threshold. This is very similar to BKP's reported estimate of a 12.3 percent drop in their very similar sample (reported on p. 576). Observe that as the bandwidth shrinks, the estimate of the sharp drop rises in magnitude, such that with the smallest bandwidth we estimate a 57.5 percent drop—over 4 times as large as the first estimate. This suggests that the point estimates are sensitive to the existence of the point mass.

When we retain the point mass and test for a discontinuity just to the *left* of 50 percent, we find the exact opposite result: the density function jumps discontinuously *upward*. Once again, the estimate starts out reasonably small (6.4 percent) and becomes very large (45.1 percent) as the bandwidth shrinks. The finding of a sharp increase in the distribution at 50 percent suggests that couples manipulate earnings to avoid a situation in which the wife earns strictly less than her husband. Put another way, the data appear consistent with a social norm dictating that *a wife should earn at least as much as her husband*. This is nearly opposite to the social norm dictating that *a wife should not earn strictly more than her husband*, which is supported by the first version of the results.

The third column of results derives from deleting the point mass and testing for a discontinuity exactly at 50 percent. Two features stand out. First, while the estimates are negative, they are no longer statistically significant—–moreover, the estimate based on the standard bandwidth matches closely estimates generated by performing the test with the standard bandwidth on our simulated data from the above section. Second, the estimates do not rise appreciably in magnitude or statistical significance as the bandwidth shrinks, likely because the point mass is no longer present. Therefore, if we ignore the one quarter of one percent of couples earning identical incomes, the conclusion that the observed distribution of spousal earnings differences could be consistent with a variety of underlying social preferences (including no explicit social norm) is supported by the data. A related conclusion is that while BKP's discontinuity test eschews the theoretical critique of the literature we levied in section 2, it does not produce robust empirical results, given the point mass of couples earning identical incomes.

3.4.2 A further inquiry into the point mass

Considering these conclusions, it is worth exploring why the point mass exists in the first place, and what it means to remove it from the sample. For example, the existence of the point mass could indicate a social preference, in the population or a certain subpopulation, for strict equality of spousal earnings. Further exploration of the 2000 Census data reveals the following facts about the couples who report identical earnings in comparison to the full sample.[14] [15] First, couples who report identical earnings are almost six times more likely to both be self-employed than couples who report different earnings (13.0 percent versus 2.3 percent). Among couples in which husband and wife indicate being self-employed in the same occupation and industry (a likely indicator of running a family business), 34 percent report identical incomes. (These couples represent 0.18 percent of the full sample of couples.) Since income from a family business can be allocated in any way between husband and wife on tax returns, this suggests that one source of identical incomes is couples choosing to divide family business income equally for income tax purposes.[16]

In addition, there are couples in which the husband and wife do appear to earn identical salary incomes. Couples reporting that husband and wife both earn wages (i.e., are not self-employed) and report identical earnings, occupations, and industries (suggesting that

[14]The Gold Standard File provides very little occupational information about the couples, which is why we use the Census for this exploration. It is important to keep in mind that the point mass of couples with identical earnings is over 10 times as large in the Census data, due to rounding of reported earnings as well as possible reporting biases. That is, many couples who report identical earnings in the Census data do not have identical administrative earnings records. However, it is reasonable to assume that couples who report identical earnings are (much) likelier than those who do not to have identical administrative records.

[15]All of these facts are based on the sample of couples in the 2000 Census 5 percent sample in which both husband and wife are age 18 to 65 with positive earnings.

[16]For couples filing jointly there will generally be no tax implications from the way family business income is allocated between husband and wife on Schedule C tax forms, though there might be implications for Social Security.

they are likely to have identical jobs) constitute 0.34 percent of the sample. Elementary, middle school, and secondary teachers make up 18.9 percent of this group, by far the largest occupation. Taken together, the group of self-employed and salaried couples with identical incomes, occupations, and industries constitute 0.52 $(= 0.18 + 0.34)$ percent of all couples. Some of these are presumably "false positives," given the fact that Census data are self-reported and rounded. But this suggests that it is not difficult to account for the 0.2-0.3 percent of couples with identical earnings in the administrative data. (Zinovyeva and Tverdostup (2018) found that co-working spouses could account for the entire point mass of equal-earning couples observed in Finnish data.)

Our interpretation of these cases (couples with family businesses reporting identical incomes and couples with identical earnings in occupations such as school teachers) is that they do not provide much information about a social norm related to husbands earning more than wives. They could constitute evidence for an equal-earning norm in a subset of the population, but they could also indicate frictions in the marriage market which lead a disproportionate share of equal-earning individuals to marry, for example, because they met through work. That is, there could be a small utility loss for the husband not out-earning his wife which is outweighed by the search cost of finding a more suitable partner. For example, McKinnish (2007) and Mansour and McKinnish (2018) have presented evidence suggesting that the workplace plays an important role in marital sorting and dissolution, consistent with the search-friction paradigm.

Whatever the cause of the point mass of equal earners, we have shown that its presence compromises the validity and robustness of BKP's discontinuity test at the equal-earning threshold. It remains unclear whether observed distributions of spousal earnings differences offer identifying information about underlying social norms.

3.5 Conclusion

Our theoretical and empirical results demonstrate that it is misleading to infer preferences about spousal attribute differences from their observed distribution in marriage market equilibrium. Marriage market outcomes are affected by preferences as well as the underlying distributions of attributes. If men are taller or higher-earning than women on average, preferences which lead to positive assortative matching will produce equilibria in which it is unusual for women to be taller or higher-earning than their husbands. Even a preference for men to be shorter than their wives can lead to positive assortative matching and, consequently, an equilibrium in which men tend to be taller than their wives.

Our simulations produce distributions of spousal earnings shares which closely resemble the observed distribution using very simple models of assortative matching—without making any assumptions about preferences regarding husbands earning more than wives. The one feature we cannot reproduce with our simulations is the discontinuous drop-off in probability mass to the right of the equal-earning threshold, reported by Bertrand, Kamenica and Pan (2015). However, we show that this discontinuity is less informative than it first appears, since it is the result of a point mass of equal-earning couples. This mass causes a sharp drop to the right of 50 percent in the distribution of the wife's share of total earned income, which is consistent with a social norm that wives should not earn more than their husbands. But it also causes a sharp drop to the left of 50 percent, a result that is consistent with a social norm that husbands should not earn more than their wives. When we remove the point mass we do not see any evidence of a discontinuity at the equal-earnings threshold. Whether these individuals are retained or removed from the sample, their presence compromises the robustness of BKP's strategy of using a discontinuity test at the equal-earning threshold to infer the presence of a husband-breadwinner norm.

To be clear, our results do not imply that male breadwinner norms do not exist. The literature includes other types of analysis, with BKP providing other pieces of evidence in their paper that are not based on inferences drawn from the distribution of spousal earnings differences. These include analyses of marriage rates, divorce rates, labor force participation, work hours, and housework time as a function of the actual or predicted probability that the wife out-earns the husband.[17] We are particularly intrigued by the release of a study that finds that husbands tend to inflate, and wives deflate, reported earnings on surveys when the wife's "true" administrative earnings exceed her husband's (Murray-Close and Heggeness, 2018). It is outside the scope of this paper to analyze these other tests of the social norm hypothesis. Our argument is simply that observed differences in spousal attributes are not, in and of themselves, good evidence for social norms related to these attributes.

It is also interesting to consider whether social norms may themselves be driven by the underlying distributions of traits. In the analysis of Stulp et al. (2013), there is a tendency for spouses to be pushed toward the actual mean difference in heights of 14 cm. We showed that this tendency can be explained as the result of positive assortative matching, with no need for a social norm related to height differences. But if there were a fundamental social preference for husbands to be taller than their wives, it would seem surprising if such a preference coincidentally matched the actual difference in mean heights between men and women. If there is such a norm, it presumably was influenced by the actual differences in heights between men and women. A plausible explanation is that positive assortative

[17]For a further analysis of housework time in German longitudinal data, see Wieber and Holst (2015).

matching produced distributions like those we observe, which in turn led individuals to perceive that there must be some normative reason for husbands to be taller than their wives.

This explanation is relevant to the case of earnings differences as well. Women's labor market opportunities in the United States have increased dramatically in the last 50 years, yet substantial gender career and earnings gaps remain, especially in marriage. It is possible that labor market change has outpaced social change, and slow-moving gender norms play a key role in generating these extant gender gaps in marriage. Inquiries into the existence and potential consequences of these norms are likely to continue to be an active area of research. We believe this research will be stronger and more convincing if researchers are sensitive to the challenges involved in drawing inferences about social norms from observed marriage market outcomes.

3.6 Tables

Table 3.1: Height differences between husbands and wives, UK Millennium Cohort Study

Husband height minus wife height (cm)	Proportion in: actual distribution	random distribution	Ratio of actual to random
< -10	0.6%	1.3%	0.47
-10 to -5	1.5%	2.6%	0.58
-5 to 0	1.9%	2.5%	0.77
0 to 5	8.5%	8.7%	0.97
5 to 10	16.3%	14.5%	1.12
10 to 15	21.3%	19.2%	1.11
15 to 20	20.7%	19.7%	1.05
20 to 25	15.3%	15.8%	0.97
25 to 30	8.8%	9.4%	0.94
30 to 35	3.7%	4.2%	0.87
> 35	1.4%	2.1%	0.66

Notes: Data taken from Table 1 of Stulp et al. (2013).

Table 3.2: Model calibration

Parameter	Symbol	Calibrated value
Mean, male log earnings	μ^m	10.350
Std dev, male log earnings	σ^m	0.750
Mean, female log potential earnings	μ^f	10.160
Std dev, female log potential earnings	σ^f	0.700
Mean, disutility of work	ψ	0.002
Std dev, disutility of work	σ^ψ	0.001
Earnings-disutility correlation, females	ρ	-0.400
Std dev, transitory earnings shock (1000s)	σ^u	13.000
Targets in the data	**Data**	**Model**
Mean, male log earnings	10.350	10.350
Std dev, male log earnings	0.750	0.750
Mean, female log earnings	10.000	9.980
Std dev, female log earnings	0.870	0.870
Gender earnings ratio, all	0.740	0.71
Gender earnings ratio, full-timers only	0.80	0.79
Labor force-participation rate, females	0.88	0.91
Full-time employment rate, females	0.67	0.67

Notes: Calibration of model discussed in Section 3.3.

Table 3.3: Different treatments of the point mass produce different discontinuity estimates.

		Hypothesized breakpoint:		
Bandwidth	Bin size	.500001	.499999	$.5\overline{0}$ (omit point mass)
.084	.0016	$-.124^{***}$	$.064^{**}$	$-.034$
		(.031)	(.031)	(.032)
.045	.0016	$-.184^{***}$	$.129^{***}$	$-.031$
		(.040)	(.040)	(.043)
.023	.0016	$-.310^{***}$	$.240^{***}$	$-.040$
		(.055)	(.055)	(.061)
.011	.0005	$-.575^{***}$	$.451^{***}$	$-.078$
		(.078)	(.081)	(.091)

Notes: Sample of spousal earnings data taken from the SIPP-SSA Gold Standard Files. See Section 3.4 for discussion of the sample. The first reported bandwidth and bin size correspond to those automatically selected by the McCrary (2008) test algorithm. McCrary (2008) recommends using a small bandwidth than the automatically selected one, as is done in the second thru fourth rows. Point estimates report the change in log height of the density function as one travels from just left of the hypothesized breakpoint to just right of it. Asymptotic standard errors reported below coefficient estimates in parentheses; standard statistical significance legend used.

3.7 Figures

Figure 3.1: Distributions of gender relative income in the 2000 U.S. Census

A. BKP Figure III

B. Replication of BKP Figure III

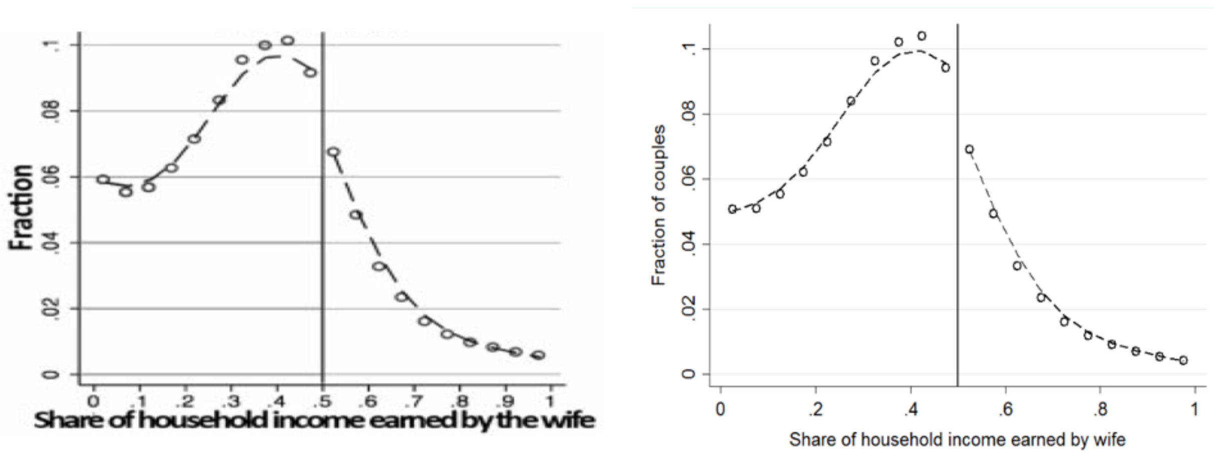

Notes: Graph A is a screenshot of Figure III of Bertrand, Kamenica and Pan (2015). Graph B is our replication. Each graph is based on a sample of dual-earning married couples in which both husband and wife are between 18 and 65 years of age. Each graph plots a 20-bin histogram of the distribution, across couples, of the share of total household income that was earned by the wife. The dashed lines depict the lowess smoother applied to each histogram, allowing for a break at 0.5.

Figure 3.2: Distribution of gender relative income in the 2000 U.S. Census: Couples aged 18-40 without children

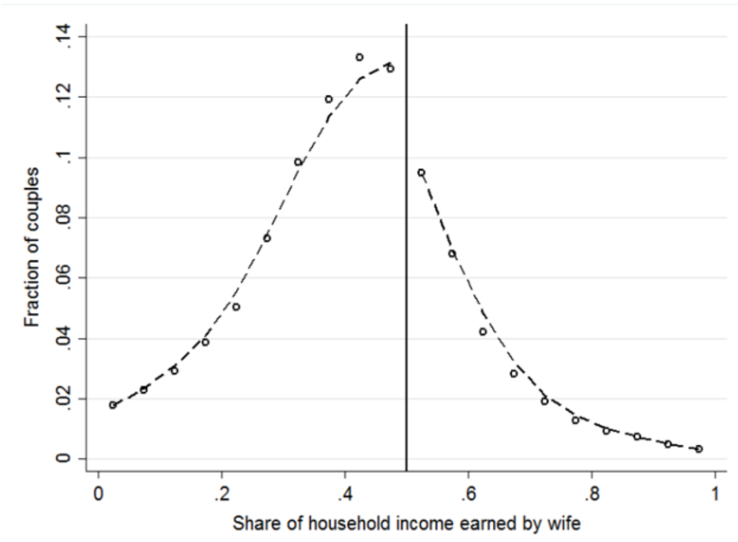

Notes: The sample includes dual-earning married couples who do not have children and where both husband and wife are between 18 and 40 years of age. The figure plots a 20-bin histogram of the distribution, across couples, of the share of total household income that was earned by the wife. The dashed lines depict the lowess smoother applied to the histogram, allowing for a break at 0.5.

Figure 3.3: Distributions of gender relative income in the 2000 U.S. Census: Actual and random sorting

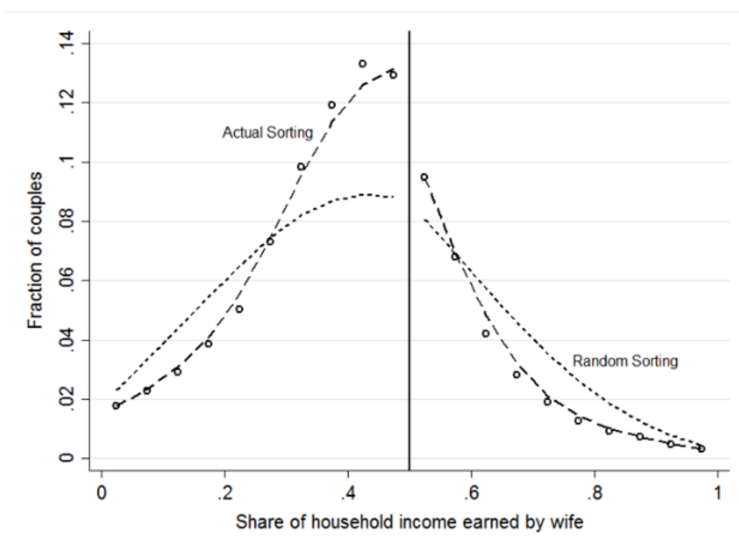

Notes: The sample is the same as in Figure 3.2. The figure plots a 20-bin histogram of the actual distribution of the wife's share of total spousal earnings ("Actual Sorting"), and a 20-bin histogram of a simulated distribution based on random sorting of couples ("Random Sorting"). The dashed lines represent the lowess smoother applied to each histogram on either side of 0.5.

Figure 3.4: Distributions of gender relative income in the 2000 U.S. Census:
Actual sorting and simulated sorting with exogenous earnings

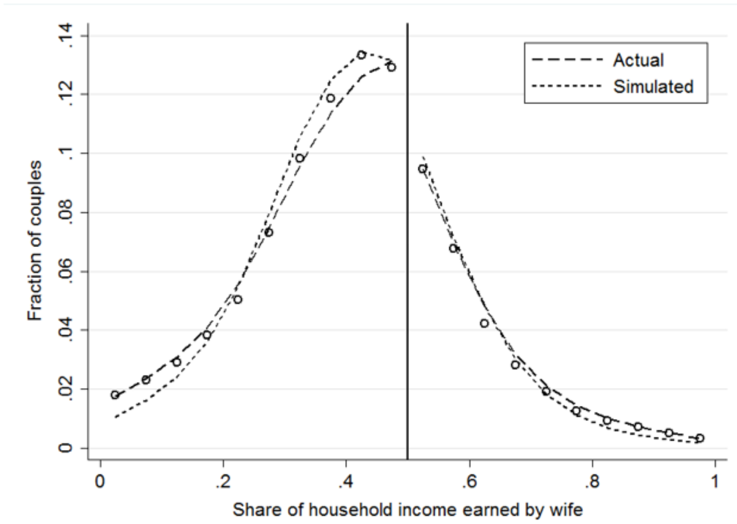

Notes: The sample is the same as in Figure 3.2. The figure plots a 20-bin histogram of the actual distribution of the wife's share of total spousal earnings ("Actual Sorting"), and a 20-bin histogram of a simulated distribution based on assortative matching of couples on observed income plus noise ("Simulated Sorting"). See Section 3.3 for further detail on the simulation. The dashed lines represent the lowess smoother applied to each histogram on either side of 0.5.

Figure 3.5: Distributions of gender relative income in the 2000 U.S. Census:
Actual sorting and simulated sorting with endogenous earnings

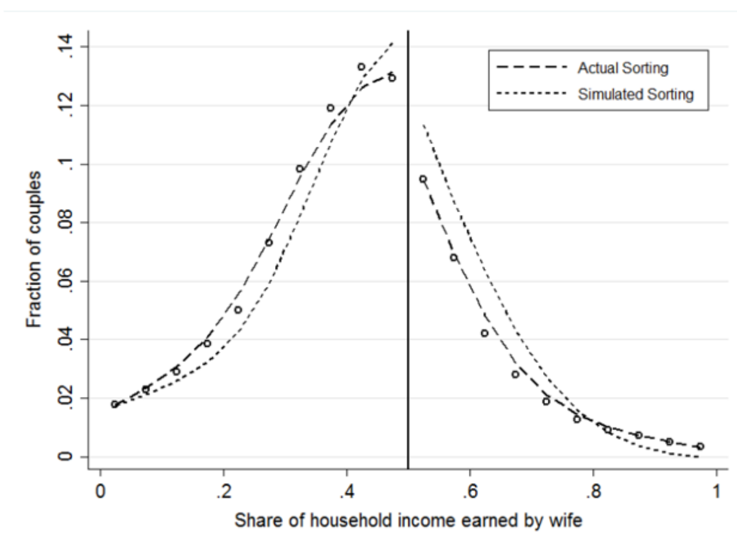

Notes: The sample is the same as in Figure 3.2. The figure plots a 20-bin histogram of the actual distribution of the wife's share of total spousal earnings ("Actual Sorting"), and a 20-bin histogram of a simulated distribution based on assortative matching of couples on potential income plus noise ("Simulated Sorting")—and in which the wife's observed earnings are endogenized via a neoclassical labor supply decision. See Section 3.3 for further detail on the simulation. The dashed lines represent the lowess smoother applied to each histogram on either side of 0.5.

Figure 3.6: Distributions of gender relative income in U.S. administrative record data

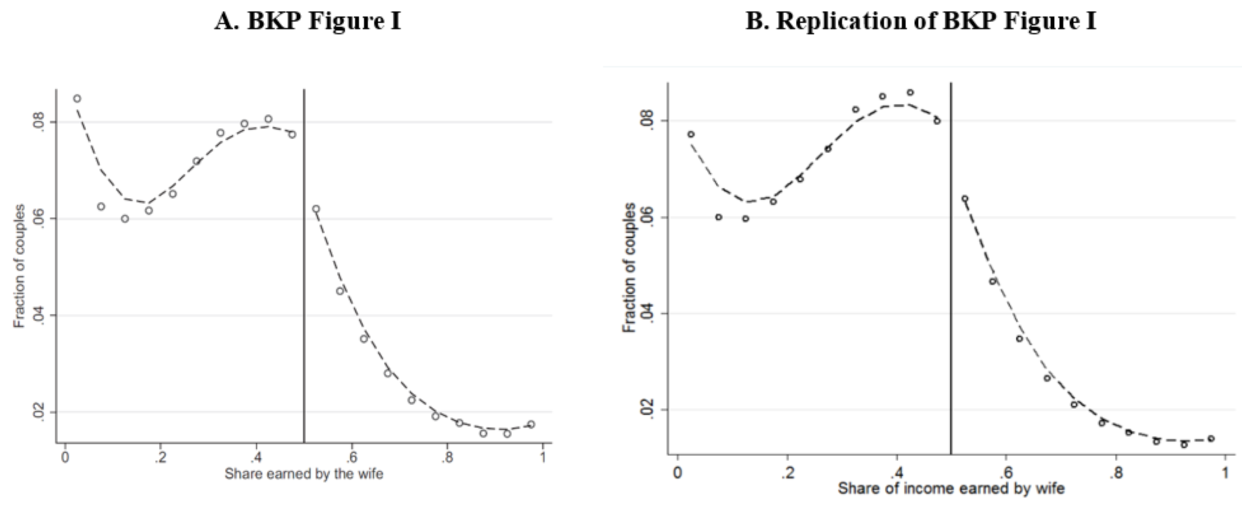

Notes: Graph A is a screenshot of Figure I of Bertrand, Kamenica and Pan (2015). The data underlying this graph are administrative income data from the SIPP-SSA Gold Standard File covering the 1990 to 2004 SIPP panels. Graph B is our replication. We use the latest version of the Gold Standard File, which includes the 1984 and 2008 SIPP panels as well. The sample in each graph includes all dual-earning couples aged 18 to 65, with income information taken from the first year the couple was observed in the SIPP panel. See Section 3.4 for further discussion. Each graph plots 20-bin histograms of the observed distribution of the wife's share of total spousal earnings. The dashed lines represent the lowess smoother applied to each histogram on either side of 0.5.

Figure 3.7: Distributions of gender relative income in U.S. administrative record data: Couples who earn close to the same incomes

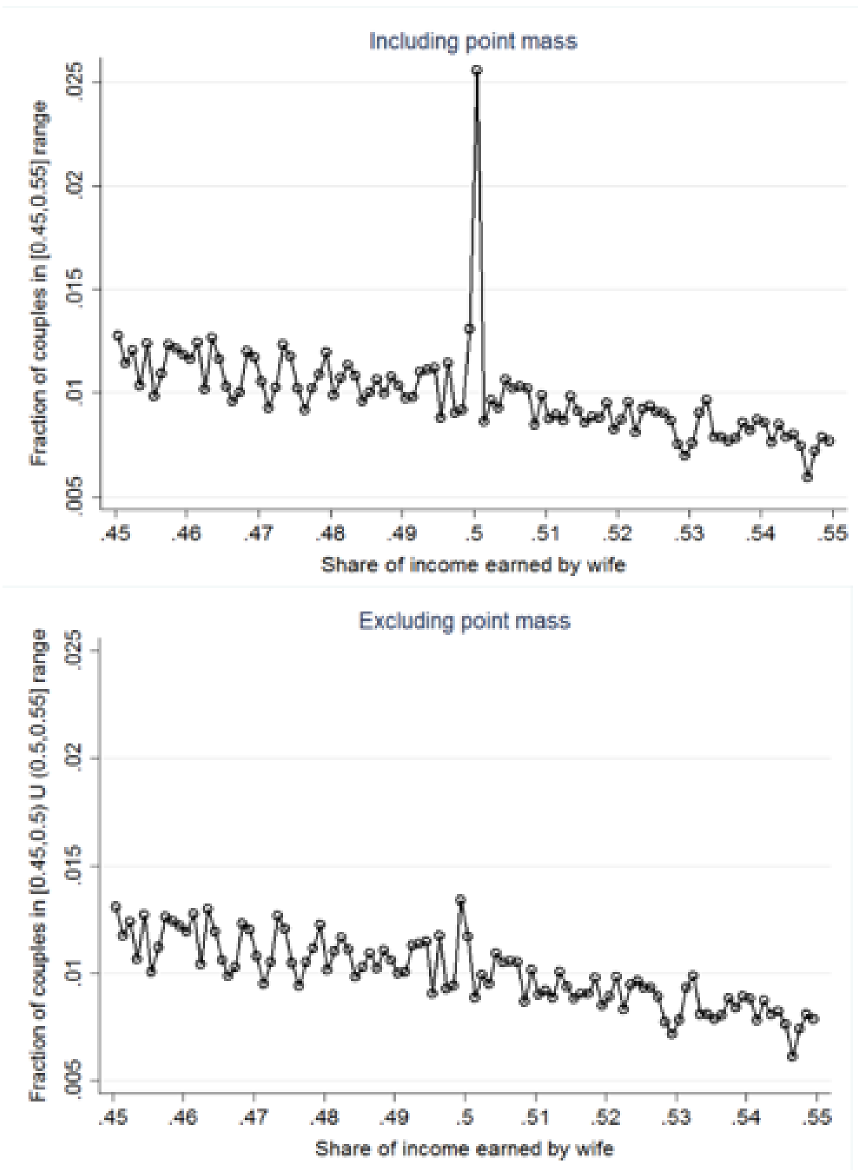

Notes: The sample is the same as in graph B of Figure 3.6, but is restricted to couples in which the wife earns between 45 and 55 percent of total income. The graph in the top panel retains the point mass of couples earning identical incomes; the graph in the bottom panel excludes it. The bin size used in both graphs is $.001$; each graph contains 100 bins.

APPENDICES

APPENDIX A

Chapter I Supporting Material

A.1 Appendix Figures and Tables

Figure A.1: The effect of unilateral divorce on marriage propensities was long-lasting, not temporary.

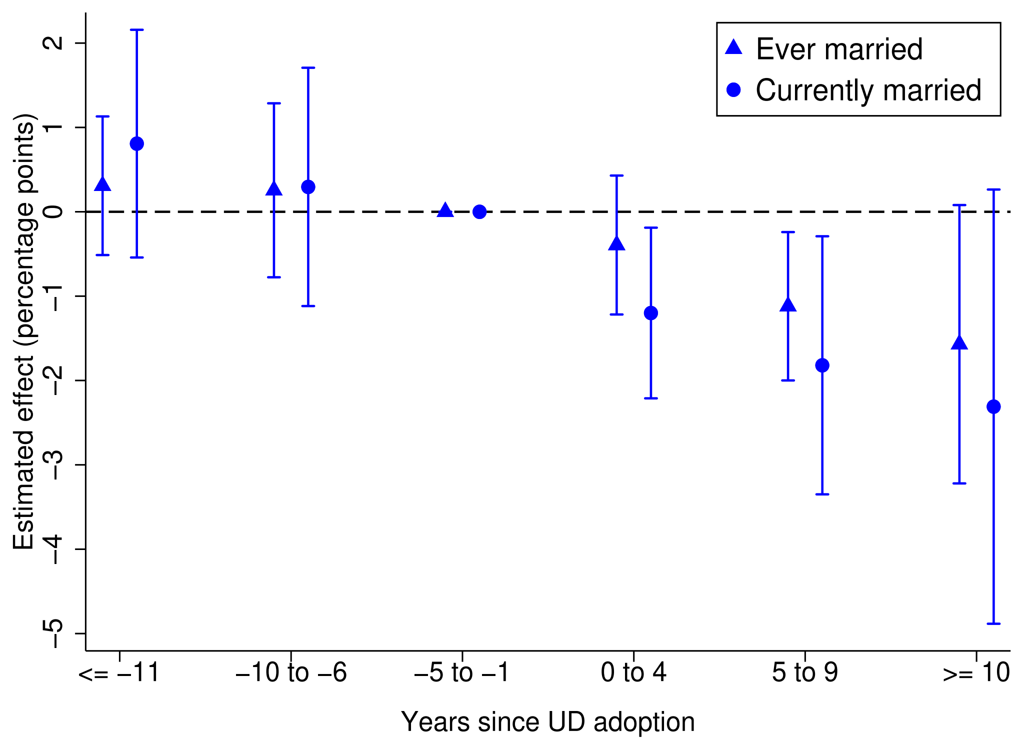

Notes: Event-study regression estimates from 1960-1990 U.S. Census data. Sample consists of less-educated men at least 18 years old, no longer enrolled in school, and with 6-21 years of potential labor market experience. Markers display point estimates (in percentage points) for associated outcome; whiskers display 95 percent confidence intervals. Standard errors are clustered on state. Event studies are estimated in 5-year bins because Census data exist only once every ten years.

Table A.1: Probabilities of transitioning back to yearly labor-force participation (P) for current-year nonparticipants (N): Men with high school education only, 1978-2011

Group	Race	$N \to P$ transition probabilities: 1 year	2 years	5 years	10 years
Experienced a $P \to N$ transition in first 15 years observed	Whites	0.50	0.66	0.83	0.89
	Blacks	0.40	0.57	0.76	0.83
Non-participant in first year observed	Whites	0.43	0.54	0.72	0.81
	Blacks	0.34	0.47	0.63	0.74

Notes: Taken from Binder and Bound (2019). Data source: Social Security Administration earnings records (1978-2011) linked to all Survey of Income and Program Participation panels from 1984-2008. Sample consists of men with exactly a high school degree, at least 19 years old, and with 0-15 years of potential labor market experience in the first year observed. Non-participation is defined as having total administrative earnings for the year less than a minimum threshold of 0.5 times the federal minimum wage times 40 hours times 13 weeks. Two groups of initial non-participants are considered: men who experienced a transition to yearly non-participation within the first 15 years of observation, and men who were yearly non-participants in the first year observed. Columns report the share of each group that had transitioned back to yearly participation by the given number of years—1, 2, 5 or 10—after the initial experience of non-participation.

Table A.2: 1980 variation across commuting zones in extent of female specialization in service sectors

	All Workers			*Noncollege Workers*		
Service sector $\to$	Business^{-}	Business	All	Business^{-}	Business	All
Percentile	*Panel A: Share of total female workers employed in given sector group*					
10th	.17	.33	.67	.18	.35	.61
25th	.19	.36	.76	.21	.38	.71
50th	.21	.39	.81	.23	.42	.77
75th	.23	.42	.86	.26	.46	.83
90th	.25	.46	.88	.28	.51	.85
Percentile	*Panel B: Share of total female workers − share of total male workers employed in given sector group*					
10th	.09	.21	.34	.13	.26	.39
25th	.10	.25	.39	.14	.30	.48
50th	.12	.27	.46	.17	.33	.53
75th	.14	.30	.51	.19	.37	.60
90th	.17	.33	.55	.22	.41	.64

Notes: 1980 U.S. Census data. See Section 1.4 for definition of commuting zones (of which 722 exist in the mainland U.S.). Sample consists of workers at least 18 years old and with 0-18 years potential experience. Noncollege workers are those with at most one year of completed college. Sector groups are constructed using occupation groups defined below in Appendix Table A.8. Business^{-} consists of groups 9-11; Business consists of Business^{-} plus administrative support, which is group 12; All services consist of groups 1, 3-5, 7-16. There are 24 total occupation groups.

Table A.3: Effects of the marriage market interventions on ever-married propensities

	Propensity to be ever-married					
	(1)	(2)	(3)	(4)	(5)	(6)
Intervention 1: mutual consent $\rightarrow$ unilateral divorce						
Effect of	-1.2^{*}	-1.9^{**}	-0.9	-1.5^{**}	-1.3^{***}	-1.4^{***}
legal change	(0.6)	(0.9)	(0.6)	(0.7)	(0.4)	(0.5)
$100-$control mean	21.0	21.0	21.0	21.0	21.0	21.0
Pre-trends p-val	0.40	0.69	0.11	0.03	0.10	0.45
N (thousands)	2,151	2,151	2,151	2,151	2,151	2,151
Intervention 2: increased demand for female employment						
Effect of 10 $p.p.$	-1.2^{***}	-1.1^{**}	-1.7^{***}	-1.4^{***}	-1.7^{***}	-1.2^{**}
Bartik shock	(0.4)	(0.5)	(0.4)	(0.5)	(0.4)	(0.5)
$100-$control mean	23.0	23.0	23.0	23.0	23.0	23.0
N	2,888	2,888	2,888	2,888	2,888	2,888
Controls						
Baseline	✓	✓	✓	✓	✓	✓
Demos$\times$linear trend		✓		✓		✓
Region effects			✓	✓		
State/division effects					✓	✓

Notes: U.S. Census and ACS data. Sample consists of men with at most one year of completed college, at least 18 years old, and with 5-20 years potential labor market experience. See notes to Table 1.4 for further detail.

Table A.4: Effects of the marriage market interventions on labor-force participation: all young men

	Propensity to be in the labor force					
	(1)	(2)	(3)	(4)	(5)	(6)
Intervention 1: mutual consent → unilateral divorce						
Effect of legal change	-0.4	-0.8^{*}	-0.3	-0.6^{**}	-0.8^{**}	-1.0^{*}
	(0.3)	(0.4)	(0.3)	(0.3)	(0.4)	(0.6)
$100-$control mean	6.8	6.8	6.8	6.8	6.8	6.8
Pre-trends p-val	0.50	0.46	0.33	0.06	0.80	0.94
N (thousands)	$2,151$	$2,151$	$2,151$	$2,151$	$2,151$	$2,151$
Intervention 2: increased demand for female employment						
Effect of 10 $p.p.$ Bartik shock	-0.8^{*}	-1.1^{**}	-1.5^{***}	-1.6^{***}	-1.6^{***}	-1.6^{***}
	(0.4)	(0.5)	(0.4)	(0.5)	(0.5)	(0.5)
$100-$control mean	7.5	7.5	7.5	7.5	7.5	7.5
N	$2,888$	$2,888$	$2,888$	$2,888$	$2,888$	$2,888$
Controls						
Baseline	✓	✓	✓	✓	✓	✓
Demos×linear trend		✓		✓		✓
Region effects			✓	✓		
State/division effects					✓	✓

Notes: U.S. Census and ACS data. Sample consists of men with at most one year of completed college, at least 18 years old, and with 0-15 years potential labor market experience. See notes to Table 1.4 for further detail.

Table A.5: Preset parameters in the structural estimation

Parameter Description	Symbol	Value	Source
Importance of period 1 relative to period 0	β	1.33	See notes.
Period-1 weekly wage, college men	w_1^C	$\exp(6.98)$	Binder and Bound (2019), Braga (2018)
Period-0 weekly wage offer, noncollege men	w_0^{NC}	$\exp(6.25)$	Binder and Bound (2019), Braga (2018)
Share women college-educated	c^F	0.36	Autor and Wasserman (2013)
Share men college-educated	c^M	0.42	Autor and Wasserman (2013)
Wage depreciation	δ	0.253	Braga (2018)
Returns to experience	$\hat{r}$	0.636	Braga (2018)
Non-labor income's share in period-0 full income	$\hat{y}$	0.24	Binder and Bound (2019)

Notes: External parameters in the model. The calibration of $\beta = 1.33$ is driven by the following assumptions: there are 6 years in period 0, there are 24 years in period 1, and the annual discount factor is 0.92. This leads the male agent, at the beginning of his life, to value period 1 by 1.33 times more than period 0. The wage depreciation and returns to experience parameters are cumulative over the 6 years of period 0. Appendix Section A.5 contains further detail on the choices of these parameters.

Table A.6: Downward shift of the wage profile versus flattening of the wage profile: equilibrium employment responses to different 10% wage shocks for noncollege men

	(1)	(2)	(3)	(4)	(5)
Δ *PDV lifetime wages*	-10%	-10%	-10%	-10%	-10%
$\Delta \ln \left(w_0^{NC}\right)$ (downward shift of wage profile)	$-.100$	$-.075$	$-.050$	$-.025$	0.00
$\Delta \hat{r}$ (flattening of wage profile)	0.00	$-.044$	$-.088$	$-.131$	$-.173$
LFP responses					
$\Delta n_0^{singles}$	$-.008$	$-.025$	$-.042$	$-.059$	$-.074$
Δn_0^{all}	$-.016$	$-.041$	$-.065$	$-.088$	$-.109$
Marr mkt multiplier	1.93	1.62	1.55	1.50	1.47

Notes: Results are computed using the empirical model developed in Section 1.5. For each simulation, the baseline wage processes for noncollege men are changed by the amount specified in the table, and the new equilibrium is computed. Each simulation imposes a -10% shock to noncollege men's present-discounted lifetime earning potential, but varies the degree to which the shock is generated by a *downward shift* of the wage profile (caused by lowering w_0^{NC}) versus a *flattening* of the wage profile (caused by lowering $\hat{r}$). The marriage market multiplier is defined as the ratio of the total employment response to the response for the always-single men. See Section 1.5.2 of the main text for details.

A.2 Data details

This section describes the processing of the various sources of data used throughout the paper.

A.2.1 Census and March CPS data samples.

The main samples used throughout this paper consist of civilian, non-institutionalized men at least 18 years old, with at most one year of completed college education, and with non-missing race, ethnicity, nativity, and employment information (i.e. current labor force status along with weeks spent employed in the reference year).[1] Samples for Figures 1.2 and 1.3 stratify on age (25-34 for labor-force participation series; 35-39 for marriage series), while samples used in regressions stratified on years of potential experience (0-15 for labor-force participation; 5-20 for marriage). Individuals with imputed values for any of these demographic were excluded from all samples. Individuals with imputed labor force variables were excluded from labor-force samples, while individuals with imputed marital status were excluded from marriage sample. When constructing hourly wages, I additionally excluded individuals with imputed weekly hours worked information (subject to such information being reported in the given survey year). After verifying that wage results were not sensitive to whether I excluded individuals with imputed wage incomes, I elected to keep such individuals in all samples.

For the gender-specific, shift-share demand shock analysis of Section 1.4, I adapted David Dorn's code to construct 1990-defined commuting zones for 1980-2015 Census data. His code is available here: `https://www.ddorn.net/data.htm`. To calculate the gender-specific shifts and shares that go into the shift-share instruments, I used the noncollege population of women (or of men) with 0-20 years of potential experience. See Section 1.4 of the main text for the formula used to construct these instruments.

A.2.2 Construction of hourly wage series appearing in Figure 1.2.

Constructing hourly wages in the March CPS (shown in Panel A) involved several steps. First, I adjusted observed wage-and-salary incomes for topcoding. Before the 1995 survey, these incomes were top-coded at a common value. I replaced these cases with the top-code multiplied by 1.5. From 1996-2010, wage incomes above top-code thresholds were

[1]Hispanic ethnicity was not tracked by the CPS until 1976 and nativity was not tracked until 1994. Thus, only individuals with missing values for these variables after these years were excluded from the CPS samples.

replaced by means of incomes above the top-code, conditional on certain observed characteristics. After 2010, wage incomes above the state-specific threshold were systematically swapped with other reported values within a bounded interval. I elected not to implement any top-coding adjustments after 1995.

Second, I used the Personal Consumption Expenditures Deflator to convert nominal values into 2017 dollars. Third, I computed annual hours worked. After the 1976 survey, this simply involved multiplying weeks worked by usual hours worked per week. Before 1976, weeks worked information is available only in the form of intervals, and usual hours worked per week is not available. However, hours worked last week is available. Thus, before 1976, I separately imputed weeks worked and usual hours worked per week using observed demographic information interacted with the observed weeks worked bin and with hours worked last week. I used 1976-1981 data to condition these imputation regressions.

Fourth, I divided real annual earnings by annual hours worked to compute real hourly wages. I trimmed calculated wage observations of below $\$2.50$ or above $\$175$ from the sample. (Trimming wage outliers did not much affect the overall wage trends reported in the figure.) Finally, I computed the average wage in a given year by applying the exponential function to the average log hourly wage. I divided the result by the average wage in 1973. Thus, the series reported in Panel A can be interpreted as the ratio of the geometric average hourly wage in the given year to that in 1973. (A value of 0.9 in year t then indicates that the geometric average hourly wage fell by 10% between 1973 and t.)

I experimented with one important adjustment to the above series, which reports geometric average hourly earnings for those employed in the reference year. Conceptually, we might prefer the geometric average hourly wage a man might expect (that is, including those who ended up not employed the whole year). Inspired by the procedure developed in Juhn, Murphy and Topel (1991), I imputed wage offers to these full-year non-workers. This involved regressing observed wages on demographic variables, *limiting the sample to men who were employed 13 weeks or less in the reference year*, and then using the regression output to predict average wage offers for the non-working population. I performed this imputation procedure in 6-year bins across time. Panel A of Figure 1.2 includes these imputations, which mildly affects the wage series but does not alter the broad pattern across time.

The construction of geometric average weekly wages in the U.S. Census data (shown in Panel B) proceeded similarly. I adjusted wage-and-salary income data for top-coding by replacing top-coded values with the top-code multiplied by 1.5 for 1940-2000 data. Thereafter, I did not implement any top-coding adjustments. Weeks worked information is available only in the form of intervals for survey years 1960, 1970, and 2008 onward. For 1960 and 1970, I used the same imputation procedure as above on 1950 and 1980 data to impute weeks worked. For 2014-2016 data, I used the same imputation procedure as above on 2005-2007 data to impute weeks worked. Because usual hours worked per week information is not available for 1940-1970, I decided to proceed with weekly wages as my concept of interest. I computed them by dividing real annual wage income by total weeks worked. Given the long time coverage of the data and the dramatic wage structure changes that took place during this time, I elected not to trim weekly wage outliers from the sample.

Because the 1950 earnings data are particularly sparse, adjusting the wage series to include the wage offers of full-year non-workers was not feasible.[2] However, due to large changes in wage-and-salary versus self-employment that transpired during this time (especially before 1980—Ruggles, 2015), I found it important to experiment with including business income in the annual earnings concept. This presented an additional challenge, since non-wage income is not tracked in 1940: only an indicator for whether the individual had over $\$50$ in non-wage income is reported. Using 1950 data on total business-and-farm income, I constructed a similar indicator (adjusting for inflation) for 1950 respondents. Then, using demographic information in interaction with this indicator, I imputed business-and-farm income for 1940 respondents based on 1950 data. Ultimately, inclusion of business-and-farm income exerted surprisingly little impact on the resultant series. For this reason, and to preserve consistency with the March CPS series, the reported series in Panel B just consider wage-and-salary income.

A.2.3 Construction of predicted marriage values appearing in Figure 1.3.

This figure reports marriage propensities for men aged 35-39, based on birth cohort group. Because the last cohort group, 1982-84, had not yet completed its 30s as of 2017 (the last year of data), this statistic is not yet observable. Instead, the figure reports predicted marriage propensities during ages 35-39 based on observed marriage propensities during ages 30-33. I estimated

$$\text{marrPropensity}^c_{35-39} = C_1 \text{marrPropensity}^c_{30-33} + C_2 t + \text{error}^c$$

using observed data from the 1952-58, 62-68 and 72-78 cohort groups—where t is the average birth year of the cohort group (1955, 1965 or 1975). Using the resultant estimates

[2]Excluding 1950, the weekly wage series were largely not sensitive to the inclusion of imputed wages for full-year non-workers.

of C_1 and C_2 and the observed marriage propensity during ages 30-33 for the 1982-84 birth cohort, I then predicted the marriage propensity for ages 35-39. I repeated this procedure for each of ever-married and currently-married propensities.

A.2.4 NLSY79 sample.

I used data on noncollege men from the 1979 cohort of the National Longitduinal Surveys of Youth in Figure 1.7 and in Table 1.1. The NLSY79 contains weekly labor market histories and annual information on family structure for a population-representative sample of individuals that were 14-22 years old when first surveyed in 1979. I constructed a sample of noncollege men, aged 16-40, according to the same specifications as those used for the Census and March CPS. I used these data to compare the employment behavior during the ages of 21-31 of "always-married" to "always-single" men. The always-married group always reported being married during the ages of 32 to 40, while the always-single group never reported being married or cohabiting with an unmarried partner during these ages. I identified non-marital cohabitation in the NLSY using a method employed by Oppenheimer (2003).

Table 1.1 required the measurement of real hourly wages. Because annual earnings and annual hours worked are reported each year, calculating hourly wages was straightforward relative to the Census and March CPS. However, I implemented some adjustments for missing wage data. The regressions reported in Table 1.1 require $\ln(\text{initial wage}_i)$ for each individual i, defined as the log of average hourly wages observed over the first 3 years after i left school. Many individuals for whom initial wage information was missing had wage information later in their career cycles. I imputed initial wages for these individuals based on later-in-life wages and other variables. Specifically, I took the sample of individuals with complete wage data and regressed initial wages on later-in-life wages (defined as the log of average hourly wages during ages 30-31), education and race dummies, and education dummies interacted with a quadratic in share of weeks spent employed between the ages 21 and 31. I then predicted initial wages based on this regression.

A few individuals had chronically missing wages due to very low attachment to the labor force. I imputed initial wage offers for these individuals by applying a variant of the Juhn, Murphy and Topel (1991) procedure. Specifically, after applying the above adjustments, I assigned these individuals the average observed wage of individuals in the *bottom quartile* of the employment distribution, conditional on race and education. (The employment distribution is the distribution across individuals of the share of time spent employed between the ages of 21-31.)

After making the above adjustments, I estimated regression specification (1.9) and generated Table 1.1. When estimating the regression, I weighted the individual data by NLSY79 sampling weights multiplied by the number of employment observations observed in the 11-year reference period (i.e. ages 21-31).

A.2.5 Coding of divorce legislation dates.

I adopted the coding presented in Appendix Table 8 of Voena (2015). See her paper and the literature she cites for more information on coding of these laws. The recent literature has, more or less, followed the original coding of Friedberg (1998), with almost no discrepancies between papers.

A.2.6 Industry and occupation groups used in Bartik shock analysis.

Table A.7: Crosswalk between sector groupings used to construct male Bartik shocks and Census industry codes

Group #	Description	*ind1990* codes
1	agriculture	10-32
2	mining	40-50
3	construction	60
4	low-tech manufacturing	100-152, 230-262, 220-222
5	basic manufacturing	270-351, 360-370, 390-392, 160-172, 210-212
6	high-tech manufacturing	352, 371-381, 189-201
7	transportation	400-432
8	telecommunications	433-442
9	utilities	450-472
10	wholesale trade	500-571
11	retail trade	580-691
12	finance, insurance, real estate	700-712
13	business services	721-760
14	personal services	761-791
15	recreation services	800-810
16	health services	812-840
17	other prof. services	841-893
18	public administration	900-932

The *ind1990* variable, which captures consistently-coded industries across time, comes from IPUMS.

Table A.8: Crosswalk between sector groupings used to construct female Bartik shocks and Census occupation codes

Group #	Description	*occ1990* codes
1	managerial	3-22, 26-34
2	scientists	43-79
3	medical	80-106
4	teachers	113-163, 187
5	law & social science	164-179, 234
6	artists	183-199
7	medical support	203-208, 445-447
8	science technicians	213-233, 235
9	finance, insurance, real estate	23-25, 243-256
10	sales	258-283
11	financial clerks	337-344, 375-383, 386
12	other admin support	37, residual 300s codes
13	housekeepers, child care	405, 468
14	protective services	415-427
15	food prep and service	434-444, 686-688
16	other personal services	448-465, 469
17	agriculture	473-498
18	mechanics	503-549, 36
19	construction	558-599, 35
20	mining	614-617
21	precision production	634-684
22	system operators	693-699, 796-799, 628
23	assembly workers	703-789
24	transportation	803-889

The *occ1990* variable, which captures consistently-coded occupations across time, comes from IPUMS.

A.2.7 Details on administrative earnings data used to construct participation transition probabilities

Table A.1 is based data from the SIPP Synthetic data product produced by the U.S. Census Bureau. See Benedetto, Stinson, and Abowd (2013) for further information. The data contain 4 administrative sources of earnings: total nondeferred earnings from FICA-covered jobs; total deferred earnings from FICA-covered jobs; total non-deferred earnings from jobs not covered by the FICA tax; and total deferred earnings from jobs not covered by the FICA tax. I sum all 4 sources together to come up with a measure of total yearly earnings. I computed yearly labor-force participation statuses based on whether total earnings for the year were above a certain minimum threshold. Following Coglianese (2018), I used a threshold of one-half of the federal minimum wage times 40 hours per week times 13 weeks per year.

A.3 Mathematics omitted from Section 1.3

A.3.1 Proof of equilibrium existence and uniqueness

See Chiappori, Iyigun and Weiss (2009). $\blacksquare$

A.4 Demonstration: the effect of unilateral divorce on labor-force participation in the population of single men is not driven by selection

I require 4 numbers for this demonstration. The first 3 numbers are drawn from Section 4.3 of the main text. First, I estimate that unilateral divorce reduced labor-force participation by noncollege single men by 1.4 percentage points. Second, the pre-reform rate of non-participation for less-educated men is listed at 14.3 percent. Third, the pre-reform rate of singlehood for less-educated men is 28.2 percent. Finally, unilateral divorce increased the rate of singlehood by roughly 1.4 percentage points (Appendix Table A.3).

Suppose now that the entire 1.4 percentage-point estimate owes to the fact that unilateral divorce changed the composition of whom is single. Using these numbers, we can solve for the rate of non-participation n required for this to be the case. Specifically,

$$
\begin{aligned}
1.4 &= \text{new nonparticipation rate} - \text{old nonparticipation rate} \\
&= \text{new rate} - 14.3 \\
&= \frac{14.3 \cdot 28.2 + n \cdot 1.4}{28.2} - 14.3 \\
&= (n - 14.3) \cdot \frac{1.4}{28.2}.
\end{aligned}
$$

This yields $n = 14.3 + 1.4 \cdot 28.2 / 1.4 = 42.5$. Thus, the previously-married individuals induced to become single as a result of unilateral divorce require a very large non-participation rate—42.5 percent—for changing selection to be responsible for the estimated effect of unilateral divorce on the labor-force participation of singles.

A.5 Section 1.5 details

A.5.1 Derivation of single men's indirect utility

Note that the utility function in period 1 is equivalent to

$$U(c_{m1}, n_{m1}; \emptyset) = \ln(w_{m1}) + \ln(n_{m1}) - \lambda \cdot \frac{n_{m1}^{1+\gamma}}{1+\gamma},$$

since $c_{m1} = w_{m1} n_{m1}$. I assume single men do not have access to non-labor income in period 1. The allowance for non-labor income in period 1, such as disability benefits, would make the algebra not reduce as nicely but would not qualitatively change the analysis.

This separability implies that the period-1 wage offer does not affect period-1 employment: the first-order condition for period-1 employment is simply

$$\frac{1}{n_{m1}} - \lambda n_{m1}^{\gamma} = 0.$$

That is, men with the same work preferences λ, γ work the same amount in period 1 and suffer the same amount of work-induced disutility, regardless of wage offer. Now, define $\underline{w} = \exp(\ln(w_0^{NC}) - \delta)$. This is the wage offer in period 1 realized by a man who does not work at all in period 0. Normalizing the utility of his hypothetical man to 0, the normalized indirect utility function in period 1, $\tilde{V}(w_{m1})$, reduces to:

$$\begin{aligned} \tilde{V}(w_{m1}) &= V(w_{m1}) - V(\underline{w}) \\ &= \ln(w_{m1}) - \ln(\underline{w}) \\ &= \hat{r} n_{m0}. \end{aligned} \tag{A.1}$$

Thus, the period-1 marginal benefit of period-0 employment for a single man is simply $\beta\hat{r}$, where β is the discount factor.

A.5.2 Calibration of preset parameters in the model

As reported in Appendix Table A.5, I preset 8 of the empirical model's parameters outside of the estimation. The notes to that table describe the calibration of the discount factor β. Here I describe how I calibrated the other 7 parameters.

I start with the wage parameters: the model requires an initial (period-0) wage offer of noncollege men (w_0^{NC}), a period-1 wage offer of college men (w_1^C), and wage growth ($\hat{r}$) and depreciation (δ) processes for noncollege men. According to Figure 1 of Binder and Bound (2019), the geometric average wage of college graduates aged 25-54 in the early 1980s was roughly $\$26.75$ (expressed in 2017 US dollars). Assuming a 40-hour work week returns the weekly wage reported in the table. The wage growth and depreciation processes are calibrated based on Table 11 of Braga (2018). This table reports Mincerian male wage regressions by education status, including estimates of the effect of an additional year spent unemployed and an additional year spent out of the labor force on wages. To calibrate the wage depreciation parameter for noncollege men, I average these two estimates within education group. Then, I take a $3/4$ to $1/4$ weighted average of the estimate for high school graduates and the estimate for high school dropouts. Finally, I multiply the resultant number by 6, as there are 6 years in period 0. The result is δ. Table 11 of Braga (2018) also reports log wage returns to a cubic in cumulative years of prior work experience. Once again, I take a $3/4$ to $1/4$ weighted average of the high school graduate and high school dropout estimates. Using these coefficients, I then predict the log wage change that occurs when work experience increases by 6 years. Finally, I add the resulting prediction to the δ. The result is $\hat{r}$.

3 of the 4 wage parameters have now been calibrated: we still require a value for w_0^{NC}. I calibrate this value as follows. Figure 1 of Binder and Bound (2019) reports a geometric average hourly wage among high school graduates aged 25-54 of roughly $\$19.00$ in the early 1980s. This implies a college wage premium of $26.75/19 = 1.41$. This yields the following equation for w_0^{NC}:

$$w_0^{NC} = \frac{w_1^C}{1.41 \cdot \exp(\hat{r} - \delta)}.$$

I calibrate the education distribution—the shares of men (c^M) and women (c^F) with college education—based on statistics reported by Autor and Wasserman (2013). Using U.S. Census data, they report the shares of adults with some completed college and with a 4-year college degree by birth cohort. I set c^G, for each gender G, equal to a $3/4$ to $1/4$ weighted average of the "some college" and "college degree" attainment rates in the 1945 birth cohort.

Finally, I calibrate the amount of nonlabor income that noncollege men have access to in period 0 using statistics reported in Table 2 of Binder and Bound (2019). This table uses 1992-2017 March CPS data to determine an average breakdown of income sources in households containing a noncollege man who worked less than 13 weeks in the reference year. I consider men aged 25-34. Taking a $3/4$ to $1/4$ weighted average across high school graduate and high school dropout statuses, followed by a $6/7$ to $1/7$ weighted average across whites and blacks, I calculate that the average man in the sample received roughly $\$4,300$ in benefits. It is likely that such men also received transfers from household members: for example, 25-34-year-old white high school graduates received $\$4,700$ in own benefits, but an additional $\$38,300$ of income was received by other household members. In absence of specific data on intra-family transfers, I decided to double the own-benefit figure and use this number—$\$8,600$—as the the average noncollege man's total non-labor income. Dividing this number by the noncollege man's full income (i.e. non-labor income plus income from full-time work) yields the figure $\hat{y}$ reported in the table. Implicit in this calibration is the assumption that non-labor income is exogenous: noncollege men are guaranteed $\hat{y}$ of their full income in transfers/benefits in period 0, regardless of how much they work.

A.5.3 Setting up the comparison of the simulation results to those estimated by Autor, Dorn and Hanson (2019)

In Section 1.5.3 I compare the effects of the 10% reduction in noncollege men's wages, simulated by the model, to the effects of a China trade shock to men's labor market opportunities (holding female opportunities constant), estimated by Autor, Dorn and Hanson (2019) (ADH).

To do this comparison, it is necessary to determine the effect of the trade shock on men's hourly wages—a response that is not reported by ADH. However, the authors do report effects on annual earnings: according to Table A3 of their study, the median-earning man in their sample lost $\$3,737$ in earnings, on a base of $\$26,000$, as a result of a 1-unit trade shock. This amounts to a 14.3% loss. In Table A3 they also report that a 1-unit trade shock caused a $3.06p.p.$ decline in male employment. A similar study found that a contemporary Bartik shock to the male manufacturing sector caused a $4.6p.p.$ decline in employment and a 7.9% decline in hours worked of noncollege men. These numbers suggest that a 1-unit trade shock caused a

$$3.06p.p.\text{ employment decline} \cdot \frac{7.9\%\text{ hours decline}}{4.6p.p.\text{ employment decline}} = 5.3\%\text{ hours decline.}$$

To first-order approximation, the percent change in hourly wages caused by the trade shock is simply the percent loss in earnings net of the percent loss in hours, or $14.3 - 5.3 = 9.0\%$.

This number, combined with ADH's estimated effects of the trade shock on male labor-force participation (reported in Table A3) and marriage (reported in Table 3), underlies the comparison undertaken in Table 1.8 of the main text.

A.6 A richer equilibrium model of labor supply, marriage, child investment and divorce

Here I specify and characterize the solution to a male employment and marriage problem with uncertainty, the possibility of divorce, and imperfectly transferable utility. This section lays out the model, while the next section derives comparative static results on the young male labor supply effects of i) a shift in the divorce regime from mutual consent to unilateral, and ii) an incremental improvement in women's labor market opportunities. This model provides further motivation for the reduced-form test undertaken in Section 1.4 of the main text—especially the unilateral divorce test, since the stylized model of Section 1.3 abstracts from divorce.

As argued in the main text, for a young man, working full-time today plausibly yields not only higher future labor market returns, but also higher marital surpluses. Competition over mates in a frictionless, competitive marriage market determines marital assignments (who marries and who remains single), intra-household division of resources, and as a byproduct, the marriage market return to male human capital investment. This return, combined with the known profile of market wages, guides males' decisions to invest and to marry. These dynamics give rise to a rational-expectations equilibrium, as in the schooling-and-marriage framework of Chiappori, Iyigun and Weiss (2009).

A.6.1 Preliminaries

The starting point for the model is an economy populated by an equal number of males m and females f. Agents live for 3 periods: $t = 0$, $t = 1$ and $t = 2$. Life is structured as follows. In period 0, men decide whether to invest in their labor market skills. Investment implies working full-time and potentially losing out on utility from leisure, while not investing implies working part-time.[3] Investment in period 0 raises one's wage in subsequent periods.

At the beginning of period 1, males and females match in a competitive marriage market. Singlehood is an absorbing state: individuals who choose to remain single make unitary labor supply and consumption decisions over $t = 1, 2$. Married couples decide on a child investment strategy in $t = 1$ (described below), which returns a payoff of child welfare in $t = 2$. The model does not rule out the production and raising of children by unmarried mothers, but it does impose that child welfare is higher on average when children are raised in married households.

[3] I abstract from benefit or other non-labor income receipt. Allowing for non-labor income weakens incentives to supply labor and plausibly strengthens all subsequent results.

Marriage quality is subject to shocks. At the beginning of period 2, a match quality shock is realized that causes couples to evaluate whether to remain married. If they do so, they make joint labor supply and consumption decisions as a couple in $t = 2$. If they divorce, they behave as singles in period 2 and additionally enjoy a modified amount of child welfare.

Each male is endowed with either of two labor market skill levels: $s_{m,0} \in \{L, H\}$. We can think of these as low or high education statuses. Investing men improve their skill levels, resulting in four possible male skill levels in the marriage market: $s_{m,1} \in \{L, L_+, H, H_+\}$. Females are homogeneous in labor market skill.[4] Individuals also vary in fixed tastes for marriage: each individual i has marital taste θ_i, where θ_i is distributed with bounded support $[0, \theta^{max}]$ and mean $\theta/2$ according to an atomless distribution function $G(\theta)$. Marital tastes are independently and identically distributed across gender and skill level.

Marital taste is additive separable from the "material" payoff to marriage. For a given individual i who has married an individual j, utility in marriage in period 2 is of the form

$$u_i(j) = \text{material utility}_i(j) + \theta_i + \psi_{ij}. \tag{A.2}$$

The variable ψ_{ij} is a random match quality shock that arrives at the beginning of period 2 and takes the value 0 with probability $1/2$ and $-\theta$ with probability $1/2$, where $\theta \leq \theta^{max}$.[5] As in Chiappori, Iyigun and Weiss (2009), this specification imposes a structure in which interactions between agents in the marriage market depend on their skill levels only: total marital surplus for any given marriage is simply additive in individual bliss payoffs.

The life-cycle is depicted in Figure A.2.

A.6.2 The marriage market and household planning problem

Individuals are utility price-takers in a frictionless, competitive, heterosexual marriage market. In the marriage market, prospective couples commit to a feasible, incentive-compatible, ψ-contingent plan that includes child investment, labor supply, consumption and divorce decisions and delivers each individual his or her utility price in expectation. This is akin to the frameworks of Chiappori et al. (2017); Reynoso (2017); Chiappori, Dias

[4]All insights generalize to a situation in which females are heterogeneous in labor market skill, although the analysis becomes more complicated. While this assumption means the model is not suited to understand changes in assortative matching on skill, it will deliver relevant theoretical predictions on male marriage and participation rates by skill. The structural modeling analysis, presented in Section 1.5, involves multiple types of females.

[5]Thus, *ex-ante* marital bliss follows a mean-0 distribution; that is, the average individual does not systematically prefer married life to unmarried life in an *ex-ante* sense.

Figure A.2: Male life-cycle in the model.

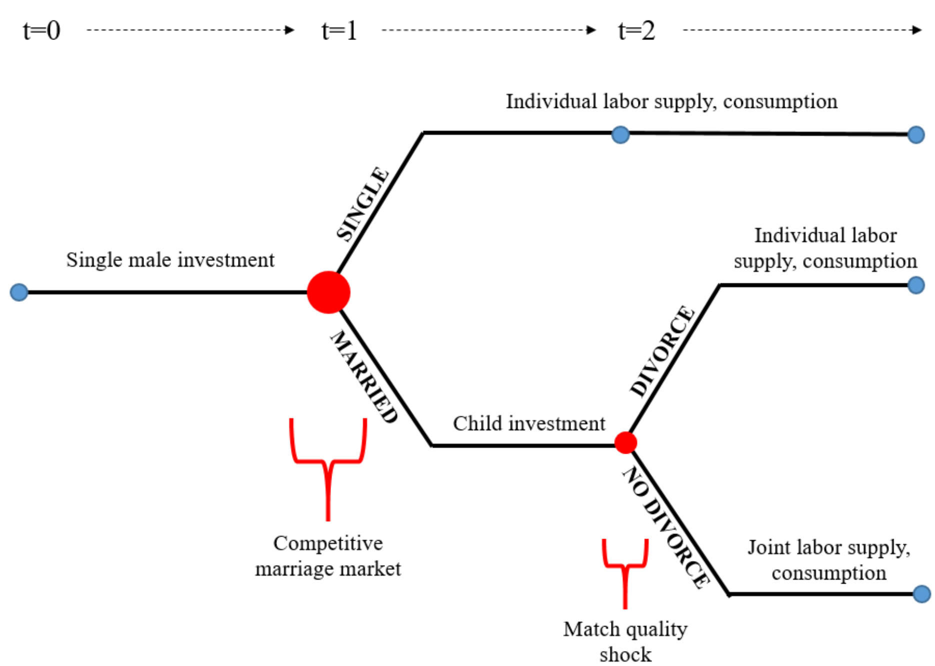

and Meghir (2018).

For two individuals i and j of opposite genders, denote the expected material payoff of i given j's utility price (material utility + marital bliss) as $\mathsf{v}_i(j;\mathsf{V}_j)$. Equilibrium in the marriage market consists of an assignment matrix A and material price vector V that yields a *stable matching*.

Definition 1. *A* stable matching $\{\mathsf{A},\mathsf{V}\}$ *satisfies the following conditions:*

$$\mathsf{v}_i(\mathsf{A}(i),\mathsf{V}_{\mathsf{A}(i)}) = \mathsf{V}_i \ \forall\ i \tag{A.3}$$

and no (i,j) *pairing exists in which*

$$\begin{aligned} & \mathsf{A}(i) \neq j \\ \textit{with} \quad & \mathsf{v}_i(j,\mathsf{V}_j) \geq \mathsf{V}_i \\ \textit{and} \quad & \mathsf{v}_j(i,\mathsf{V}_i) \geq \mathsf{V}_j, \end{aligned} \tag{A.4}$$

and in which at least one of the inequalities is strict.

Condition (A.3) links assignments to prices by stipulating that the assignments guarantee each individual his or her expected utility price in equilibrium. Condition (A.4) requires that the assignment A is a *core* allocation (Shapley and Shubik, 1971; Roth and Sotomayor, 1992). That is, in equilibrium, there does not exist a set of individuals who would be better off dissolving their assigned matches and matching with each other.

Let $\mathsf{A}(i) = \emptyset$ correspond to a situation in which individual i remains single. Note that $\mathsf{V}_\emptyset = 0$; that is, each individual matched to "single" pays "single" a utility price of 0 and retains the total output from remaining single. Thus, by setting $j = \emptyset$ in (A.4), we see that stability also requires the equilibrium assignment to deliver everyone who marries a (weakly) larger expected utility price than that from remaining single.

I now describe household behavior and its relationship with marriage-market utility prices. Single individuals derive material utility from consumption of two commodities: a market-purchased commodity c and a home-produced commodity, which is formed by combining monetary inputs x with time inputs L. Divorcees have the same structure of material utility as singles. In married couples, the home-produced commodity is a public good, creating consumption economies relative to divorced and single households.

Assumptions 1: child welfare. *Without loss of generality, normalize the utility from being single in period* 1 *to* 0*. In period* 1*, married couples invest in their children. Child investment requires at least* 0.5 *units of time from the mother and returns the additive separable value* $\alpha w_M + k$ *in period 2 (where* w_M *is the father's wage). Child welfare is produced according to the following cost function:*

$$\chi(w_M, L_{F,1}) = \begin{cases} 0 & L_{F,1} = 1 \\ \alpha w_M & L_{F,1} = 0.5 \end{cases} \tag{A.5}$$

Net child welfare, denoted as $\omega(w_M, L_{F,1}) = \alpha w_M + k - \chi(w_M, L_{F,1})$*, is a public good in both married households. Divorced wives continue to enjoy* ω*, while divorced husbands only enjoy* $(1-d)\omega$*.*[6]*. I do not assume any loss of child welfare on divorce; only that the ex-husband has less access to the child. However, all results generalize to a situation in which overall child welfare decreases on divorce, as suggested by (Gruber, 2004).*

For a couple (m, f) that decides to marry, I represent a potential household plan as follows (m, f arguments dropped after the first line for convenience):

$$\begin{aligned} \rho(\psi; m, f) &= \{\ \rho_1(m,f),\ \rho_2(\psi; \rho_1(m,f), m, f)\ \} \\ &= \{\ L_{f,1},\ \{\mathsf{D}(\psi, L_{f,1}), L_{m,2}(\mathsf{D}), L_{f,2}(\mathsf{D}), c_{m,2}(\mathsf{D}), c_{f,2}(\mathsf{D})\}\ \} \\ &\in \ \left\{\ \{0.5, 1\} \times \{\{0,1\} \times [0,1]^2 \times R_+^2\}\ \right\}. \end{aligned}$$

[6]The parameter d is a "distance" parameter reflecting a standard custodial arrangement (Chiappori, Iyigun and Weiss, 2015)

That is, a married couple in period 1 chooses whether the wife f should stay home full-time while raising the child. In period 2 the couple chooses whether or not to divorce, conditional on the realization of ψ and the child investment strategy from period 1. Conditional on the divorce decision, labor supply and consumption choices occur.

An individual i deciding to remain single simply chooses labor supply and consumption in period 2. I represent such a plan as follows:

$$\begin{aligned} \sigma(i) &= \{\ L_{i,2}, c_{i,2}\ \} \\ &\in \{\ [0,1] \times R_+\ \}. \end{aligned}$$

Cooperation in period 2—I make no assumption about bargaining over the material surplus: as in Chiappori (1988, 1992) and myriad subsequent work, I only require Pareto efficiency. Given a wife's labor supply decision from period 1 and her associated utility price, this property admits a general expression of the husband's expected payoff in marriage. For a male m marrying a female f,

$$\begin{aligned} \mathsf{v}_m(f; L_{f,1}, \widetilde{\mathsf{V}}_f(L_{f,1}; m)) \quad &= \quad \max_{\rho_2(\psi; m,f)} E_\psi[u_m(\rho_2(\psi; L_{f,1}, m, f))] \\ \text{s.t.} \quad & \quad [\text{ICC}]\ \ E_\psi[u_f(\rho_2(\cdot))] \geq \widetilde{\mathsf{V}}_f \\ & \quad [\text{F}_{\text{LS}}]\ \ L_{f,2}, L_{m,2} \in [0,1] \\ & \quad [\text{F}_{\text{BC}^{\text{M}}}]\ \ c_{m,2} + c_{f,2} + x_2 = w_m(1 - L_{m,2}) + w_f(1 - L_{f,2}) \\ & \quad \quad \text{if } \mathsf{D}(\cdot) = 0 \\ & \quad [\text{F}_{\text{BC}^{\text{D}}}]\ \ c_{m,2} + x_{m,2} + c_{f,2} + x_{f,2} = w_m(1 - L_{m,2}) + w_f(1 - L_{f,2}) \\ & \quad \quad \text{if } \mathsf{D}(\cdot) = 1 \end{aligned} \tag{A.6}$$

where $u_i(\rho_2(\cdot))$ denotes a given individual i's total period-2 payoff (material utility $+$ possible marital bliss) from following plan ρ_2. The first constraint is an incentive compatibility constraint, guaranteeing that female f is paid her utility price $\widetilde{\mathsf{V}}_f$ from participating in a marriage with m—conditional on her labor supply choice in period 1. The last 3 constraints are feasibility constraints of time and consumption allocations in period 2.

If an individual i chooses to remain single, he or she receives

$$\begin{aligned} \mathsf{v}_i(\emptyset) \quad &= \quad \max_{\sigma(i)} u_i(\sigma(i)) \\ \text{s.t.} \quad & \quad [\text{F}_{\text{LS}}]\ \ L_{i,2} \in [0,1] \\ & \quad [\text{F}_{\text{BC}}]\ \ c_{i,2} + x_{i,2} = w_i(1 - L_{i,2}) \end{aligned} \tag{A.7}$$

in period 2.

Possibility of non-cooperation in period 1—When expected utility is transferable between spouses at a one-for-one rate, the wife and husband will always agree on the child investment strategy in period 1 that maximizes joint material output in period 2. If contracting between spouses is imperfect, however, spouses may not agree on how the wife spends her time.

Assumption 2. *The wife retains control over how she spends her time in period* 1.

Thus, when married to male m, female f's labor supply decision solves the following program

$$\max_{L_{f,1}\in\{0.5,1\}} \widetilde{\mathsf{V}}_f(L_{f,1}; m). \tag{A.8}$$

By assumption 1, splitting time between working in the market and in the home lowers child welfare by αw_m in period 2. Despite this loss, the wife may decide to work part-time in the market if doing so appreciably raises the marital surplus from other goods, or raises the share of the marital surplus she can claim in period 2. Disagreement between wife and husband over the wife's time allocation arises when $\mathsf{v}_m(f, 1, \widetilde{\mathsf{V}}_f(1; m)) > \mathsf{v}_m(f, 0.5, \widetilde{\mathsf{V}}_f(0.5; m))$ but $\widetilde{\mathsf{V}}_f(1; m) < \widetilde{\mathsf{V}}_f(0.5; m)$. This will become relevant when the divorce regime is unilateral.

It is instructive to characterize each individual's marriage decision at equilibrium prices. I operate from the perspective of males making proposals to females, although an equilibrium stable matching can be constructed with either side as the proposer. For a given male m, the marriage decision simply involves choosing the assignment—including possibly remaining single—that maximizes his expected payoff. That is:

$$\mathsf{V}_m = \max_f \mathsf{v}_m(f; L^*_{f,1}, \widetilde{\mathsf{V}}_f(L^*_{f,1}; m)) \tag{A.9}$$

where $L^*_{f,1}$ is the wife's solution to program (A.8) given a marriage to m. Thus, males take prospective wives' period-1 labor supply decision rules as given when deciding whether and whom to marry.

Given equilibrium prices, each female f remains single if no male desires her, or chooses whether to marry the given male m who wishes to match with her. Thus:

$$\mathsf{V}_f = \begin{cases} \mathsf{v}_f(\emptyset) & \text{if no proposer} \\ \max\left\{ \mathsf{v}_f(\emptyset),\ \widetilde{\mathsf{V}}_f(L^*_{f,1}; m) \right\} & \text{otherwise.} \end{cases} \tag{A.10}$$

In this way, the assignment problem in the marriage market is *decentralized*, with each individual choosing the mate—including possibly remaining single—that yields him or her the highest payoff, subject to being desired by the corresponding mate.

A.6.3 Male investment decisions

We are now in a position to describe male investment behavior in period $t = 0$. In a rational-expectations equilibrium, each male is aware of the marriage market price vector $\mathsf{V}_M\ \forall M$. This knowledge, together with the known costs of investing, is sufficient to pin down equilibrium investment behavior.

For a given male of type M, represent the investment decision as I_M. The period-0 "effort" cost of investing is

$$(u_{M,0}|I_M = 0) - (u_{M,0}|I_M = 1) = e_M.$$

Male M will invest in period 0 if the effort cost of doing so is expected to be returned in the marriage market. Thus,

$$I_M = \mathbf{1}\{V_{M^+} - V_M \geq e_M\} \tag{A.11}$$

for $M \in \{L, H\} \times \{0, 1\}$.

Definitions 2. *The* labor market return *to investing is the period-2 return enjoyed by a single male. Thus, for a given male M:*

$$\lambda_M = \mathsf{v}_{M^+}(\emptyset) - \mathsf{v}_M(\emptyset).$$

The marriage market return *to investing is the additional expected return in period 2 received from participating in the marriage market:*

$$\mu_M = \mathsf{V}_{M^+} - \mathsf{V}_M - \lambda_M.$$

Thus, for males M who choose to remain single, $\mu_M = 0$.

With these definitions, we can conveniently re-express the decision to invest as occurring exactly when the sum of the labor market and marriage market returns exceeds the period-0 cost:[7]

$$I_M = \mathbf{1}\{\lambda_M + \mu_M \geq e_M\}. \tag{A.12}$$

[7]It is worth noting that this specification of investment behavior generally holds, regardless of the utility specifications underlying the marriage market prices. The only implicit assumption is that each agent i acts to maximize V_i, which is itself an expected utility outcome. This is equivalent to assuming agents obey *von-Neumann-Morgenstern rationality*: ex-post utility is separable across states of the world and expected utility is a probability-weighted linear combination of ex-post utilities. Equivalently, the happiness an individual expects to obtain when married is not contingent on the actions he might take were he to remain single.

A.6.4 Equilibrium

Definition 3. *A rational-expectations equilibrium consists of a set of investment strategies, marital assignments, and ex-ante and ex-post utilities—*$\{I, \{\mathsf{A}, \mathsf{V}\}, \mathsf{v}, u\}$*—satisfying the following properties:*

1. *Period-*0 *male investment decisions* I *maximize expected utility, as given by (*A*.12).*

2. *The marital assignment* A *is stable, as given by (*A*.3) and (*A*.4). Decentralized marital choices satisfy (*A*.9) and (*A*.10). Moreover, the utility prices* V *clear the market.*

3. *The ex-ante utilities* v *and ex-post period utilities* u *result from solutions to the problems described by (*A*.6), (*A*.7) and (*A*.8).*

As noted by Chiappori, Iyigun and Weiss (2009), the rational-expectations structure of the model transforms male investment and marriage decisions into a multinomial discrete choice problem. That is, given the investment cost structure and the marriage market utility prices, each male makes the choice in equilibrium yielding the highest expected utility from the given choice set: $\{\{\text{invest}, \text{marry}\}, \{\text{invest}, \text{single}\}, \{\text{no invest}, \text{marry}\}, \{\text{no invest}, \text{single}\}\}$. Section A.8 of this Appendix provides explicit utility and human capital specifications. With these, I prove 2 propositions that conveniently shrink the set of possible equilibria.

Proposition 1. *Highly skilled men* $(s_{M,0} = H)$ *always invest. Less-skilled men* $(s_{M,0} = L)$ *only invest if they also expect to marry.*

Proof. *See Appendix Section A.8.* ∎

Proposition 2. *All men who choose to marry are better off investing.*

Proof. *Denote the cost of investing for less-skilled men, net of the labor market return, as* $\hat{e}_L = e_L - \lambda_L$*. The proof, contained in Appendix Section A.8, shows that the marriage market return to investing exceeds* $\hat{e}_L$*.* ∎

Though these propositions are not necessary, they facilitate the claim that a rational-expectations equilibrium exists. This is proved in Section A.8 of this Appendix as Proposition 3. The proof relies on the adaptation of the salary-adjustment algorithm (Crawford and Knoer, 1981; Kelso and Crawford, 1982) to the case of a one-to-one matching market with a finite set of types, suggested by Reynoso (2017). This algorithm guarantees the construction of a market-clearing stable matching. It is shown that the equilibrium is characterized by a triple of threshold marital taste values $(\underline{\theta}_L, \underline{\theta}_H, \underline{\theta}_F)$, with each individual i of type I (less-skilled male, highly skilled male, female) marrying exactly when $\theta_i \geq \underline{\theta}_I$. It is shown that these threshold values solve the following system of equations:

$$\begin{aligned} \underline{U}_{L^+} - \mathsf{V}_{\underline{F}} + \underline{\theta}_L &= \mathsf{v}_L(\emptyset) \\ \underline{U}_{H^+} - \mathsf{V}_{\underline{F}} + \underline{\theta}_H &= \mathsf{v}_H(\emptyset) \\ \mathsf{V}_{\underline{F}} &\geq \mathsf{v}_F(\emptyset) \\ 1 - \frac{1}{2}\left(G(\underline{\theta}_L) + G(\underline{\theta}_H)\right) &= 1 - G(\underline{\theta}_F). \end{aligned} \tag{A.13}$$

$\underline{F}$ represents the female with marital taste $\underline{\theta}_F$ and $\underline{U}_M$ $(M \in \{L^+, H^+\})$ is the level of joint marital output, exclusive of the husband's marital taste, that is produced in a marriage involving female $\underline{F}$. Thus, in equilibrium, it is exactly the males with marital taste above the threshold value who would prefer to marry the threshold-taste female. The last equation is the market-clearing condition: in an equilibrium stable matching, the number of males who marry must equal the number of females who marry.

Even though the marriage market is competitive, the sequential nature of male investment and marriage decisions generally raises the possibility of multiple equilibria (Nöldeke and Samuelson, 2015). For example, we could have a possible equilibrium in which the marriage market return to investing is low and males do not invest, and another possible equilibrium in which the marriage market return is high and males invest. The structure of the current model rules out this multiplicity problem (as proven in propositions 1 and 2).

A.7 Equilibria under different divorce regimes and gender gaps

In this section, I analyze equilibria that arise under different divorce law regimes and women's relative labor market opportunities. I start from a baseline equilibrium in which the divorce regime is mutual consent and there is a significant gender gap in incentives to work in the labor market. I show that a switch to unilateral divorce, under plausible conditions, leads to a new equilibrium with lower marriage and participation rates for less-skilled males. Furthermore, I show that regardless of the prevailing divorce regime, an increase in females' labor market opportunities also produces an equilibrium decline in marriage and male labor-force participation.

A.7.1 Base case: mutual consent divorce

Constructing the ex-ante *Pareto frontier.*—To begin, I make the following simplifying assumption.

Assumption 3: Divorce allocations. *The Pareto frontier in divorce does not depend on the wife's labor supply decision in period* 1*.*

If the wife works in the first period, she prevents her own human capital depreciation and thus increases the pool of resources to be shared on divorce. But at the same time, this action decreases child welfare, lowering total resources in divorce. For convenience, we assume that these effects cancel, and thus that there is one stable Pareto frontier in divorce.[8] This facilitates, although is not necessary, for the following result.

Proposition 4. *When the divorce regime is mutual consent, spouses will always agree on the efficient child investment strategy, which is to have the wife stay home full-time in period* 1*.*

Proof. *See next section of this Appendix.* ∎

A consequence of proposition 4 is that to construct the ex-ante Pareto frontier we need only consider the Pareto frontier in divorce and the ex-post Pareto frontiers in marriage conditional on the wife staying at home in period 1. In a mutual consent divorce, an allocation must be found that makes both spouses better off than in the current allocation in marriage. For purposes of exposition, we consider a possible marriage involving a less-skilled male who has invested, and in which spouses have sufficiently low marital taste that divorce is sometimes a possibility. It is straightforward to show that the *ex-post* Pareto frontier in marriage is the following hyperplane:

$$\mathsf{V}_m + \mathsf{V}_f = \overline{\mathsf{V}}_{mf} + \psi_{mf} \tag{A.14}$$

where $\overline{\mathsf{V}}_{mf} = \beta w_{L^+} + 2\omega + \theta_m + \theta_f + \psi_{mf}$ and $\omega = \alpha w_{L^+} + k$. As a consequence of the TU property, total resources to be shared do not depend on how they are shared. This hyperplane defines a straight line with slope -1 in the $(\mathsf{V}_m, \mathsf{V}_f)$ Euclidean space.

The Pareto frontier in divorce is defined by the following equation:

$$\mathsf{V}_m(\hat{y}) + \mathsf{V}_f(\hat{y}) = \mathsf{V}^D_{mf}(\hat{y}) \tag{A.15}$$

where $\mathsf{V}^D_{mf} = \beta\left(\frac{w_{L^+}+w_F}{4} + \frac{\hat{y}^2}{4w_{L^+}} + \frac{\hat{y}^2}{4w_F}\right) + (2-d)\omega$ (see Appendix calculations). Because the home-produced commodity becomes private in divorce, utility is not transferable in the same manner in divorce as it is in marriage. In particular, the total level of resources to be shared now does depend on the monetary transfer $\hat{y}$. This equation defines a downward-sloping, convex line in the $(\mathsf{V}_m, \mathsf{V}_f)$ Euclidean space. Moreover, whenever $\theta_m + \theta_f + \psi_{mf} \geq$

[8] Leveraging some calculations in the next section, one can show that this assumption is equivalent to assuming $(2-d)\alpha w_m = \frac{\beta\delta w_F}{4} + \frac{\beta(\underline{y}^2 - \overline{y}^2(1+\delta))}{4w_F(1+\delta)}$.

$-d\omega$, this frontier is completely contained within the marriage frontier.

Figure A.3 depicts the construction of the *ex-ante* Pareto frontier when the divorce regime is mutual consent. First, note that each spouse must receive a preferred allocation (in expectation) to the allocation from remaining single—otherwise the marriage would not form. The single allocations are given by the axes of the graph. The straight lines correspond to the *ex-post* Pareto frontiers in marriage, conditional on each possible realization of the taste shock ψ_{mf}. The convex line is the Pareto frontier in divorce. Next, note that the bad taste shock happens with a 50 percent probability. Thus, in the region of the Euclidean space where both straight lines lie to the northeast of the convex line, the *ex-ante* Pareto frontier traces out the locus of mid-way points of 45 degree lines drawn between the two straight lines. Essentially, this is the "average" of the two straight lines.

In the regions of the Euclidean space where the convex line lies to the northeast of the inner straight line, the *ex-ante* Pareto frontier traces out the locus of mid-way points of 45 degree lines drawn between the outer straight line and the convex line (by assumption 3). Essentially, this is the "average" of the outer straight line and the convex line. The result is the 3-segment, thick red line shown in the figure.

Figure A.3: Construction of Pareto frontier under a mutual consent divorce regime

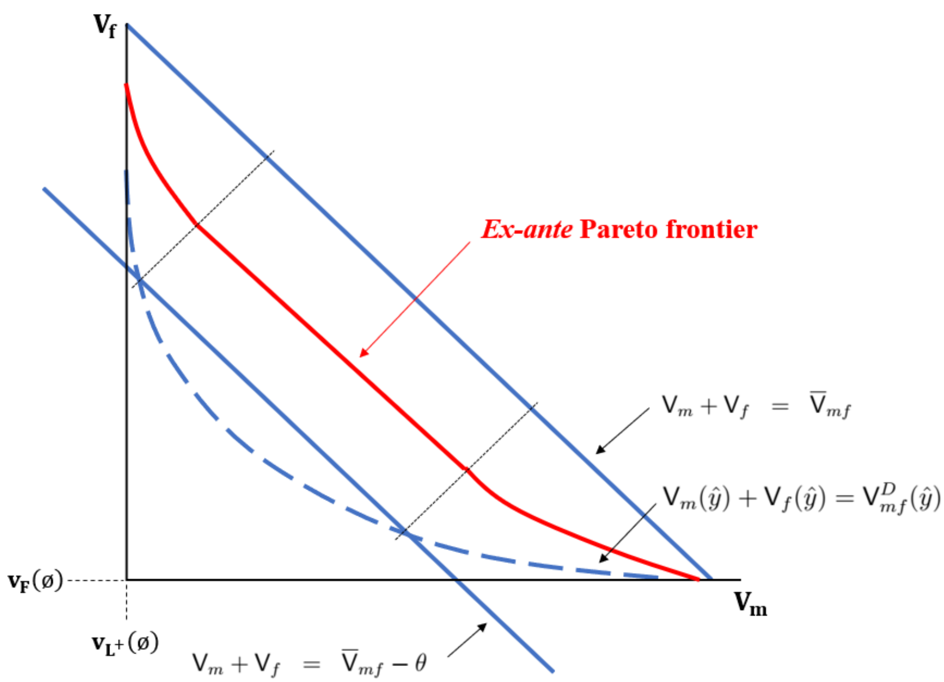

As described above.
The figure considers a hypothetical marriage involving a less-skilled male, and in which the two partners have low enough joint marital taste that divorce is a possibility for some allocations.

Baseline equilibrium.—Figure A.4 depicts allocations in a baseline equilibrium for a given set of individuals with relatively low marital taste. The left-hand Pareto diagram considers a marriage involving a less-skilled man and the right-hand diagram considers one involving a highly-skilled man.

Let us unpack Figure A.4. First, less-skilled men require at least $\hat{e}_L$ of the ex-ante marital surplus to participate in the marriage—otherwise, their investment cost is not covered and they would prefer not to marry (and not to invest). This constraint is represented by the vertical dotted line. Thus, the segment of the Pareto frontier between the two small black dots describes the range of feasible and incentive-compatible expected utility outcomes at the moment of marriage for these individuals. The large red dots denote potential equilibrium allocations resulting from the salary-adjustment algorithm. So long as at least one male of both skill types desires to marry in equilibrium, reciprocating female partners must be indifferent about who they marry. Otherwise, the market would not clear: prices would adjust until indifference was established or until one type no longer desired marriage. This condition is represented by the horizontal dotted line connecting the two red dots.

Figure A.4: Hypothetical couple in baseline equilibrium

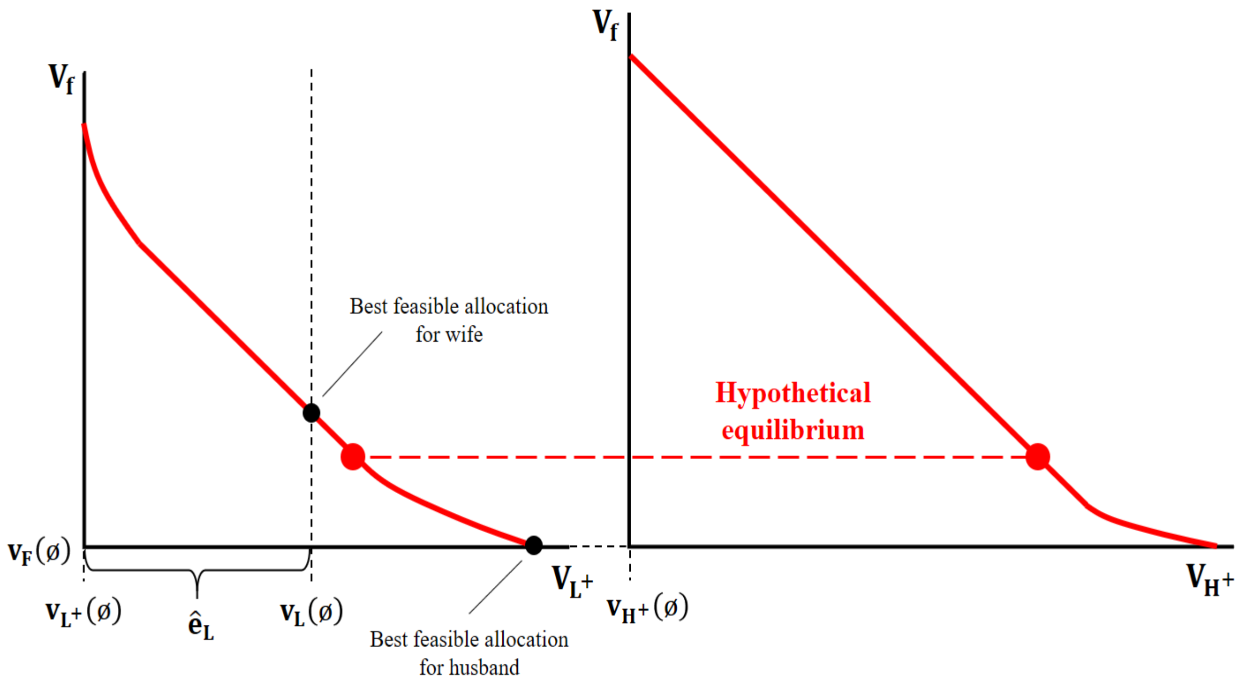

Depiction of equilibrium allocations in the baseline scenario (mutual consent divorce, gender gap in labor market opportunities) for a given set of individuals (less-skilled male, highly-skilled male, female). Each individual has marital taste above the threshold required for marriage but distinctly below the maximum taste value.

A.7.2 Introducing unilateral divorce

I now introduce the first of two novel theoretical results: in the absence of perfect contracting technology in marriage, laws governing the terms of divorce may change the *ex-ante* gains from marriage enough to impact less-skilled males' pre-marital labor supply decisions.

Proposition 5. *When the divorce regime changes from mutual consent to unilateral, and:*

- *the returns to female experience ($\beta w_F \delta$) are sufficiently high relative to the loss of child welfare from the wife working in the market (αw_M);*
- *the marital taste shock θ is sufficiently large;*
- *and the ex-husband's access to child welfare $(1-d)$ is sufficiently small,*

spouses in relatively-low-taste marriages will disagree on the child investment strategy. As a consequence, such marriages no longer form in equilibrium, and the less-skilled men who no longer marry also no longer invest in period 0.

The last claim of the proposition follows from propositions 1 and 2: less-skilled males who no longer marry will also no longer invest. Thus, the claim to establish is that unilateral divorce reduces formation of certain marriages involving less-skilled males. The proof of this claim is contained in Figure A.5.

The conditions stated in the proposition give rise to a situation where the maximum allocation the husband can guarantee the wife within marriage while still respecting his participation constraint (point A) is inferior to her autarky value (i.e. her value in divorce in the absence of a divorce settlement) from working in period 1 (point B). Thus, if the two marry, the wife will work in period 1 regardless of the planned allocation of resources—which, in turn, violates the husband's *ex-ante* participation constraint. (The best he can do in this situation is represented by point C). Therefore, such a marriage—which was incentive-compatible in a mutual consent divorce regime—no longer forms in a unilateral divorce regime.

This result depends on the incomplete-contracts nature of marriage. As shown in Figure A.5, there exist feasible *ex-ante* allocations in marriage that make both spouses better off compared with being single. However, to attain such an allocation, the wife would have to forego the opportunity to increase her bargaining position in the marriage (by working in period 1). In the absence of binding agreements that can be made in the marriage market (e.g. costlessly enforceable prenuptial contracts), there is nothing to prevent the wife from entering the workforce in period 1. This is an example of Pareto inefficiencies that may arise in bargaining models of marriage (Lundberg and Pollak, 1994, 2003). While these models originally conceived of Pareto inefficient allocations arising *within* marriage, I show that bargaining after marriage can lead to an inefficient number of marriages forming.[9]

In sum, by strengthening a woman's property right over her earning potential in the labor market, unilateral divorce reduces the share of the gains from specialization within marriage that husbands can claim. This makes preparing for marriage by sinking human capital investment no longer worth it for men with low initial labor market skill and relatively low taste for marriage. While such men reap enough material benefits from specialization gains in a mutual consent environment to coordinate on a path of investment and marriage, unilateral divorce lowers the marriage market returns to their investment enough to discourage this path.

Figure A.5: Hypothetical marriage no longer forms under unilateral divorce

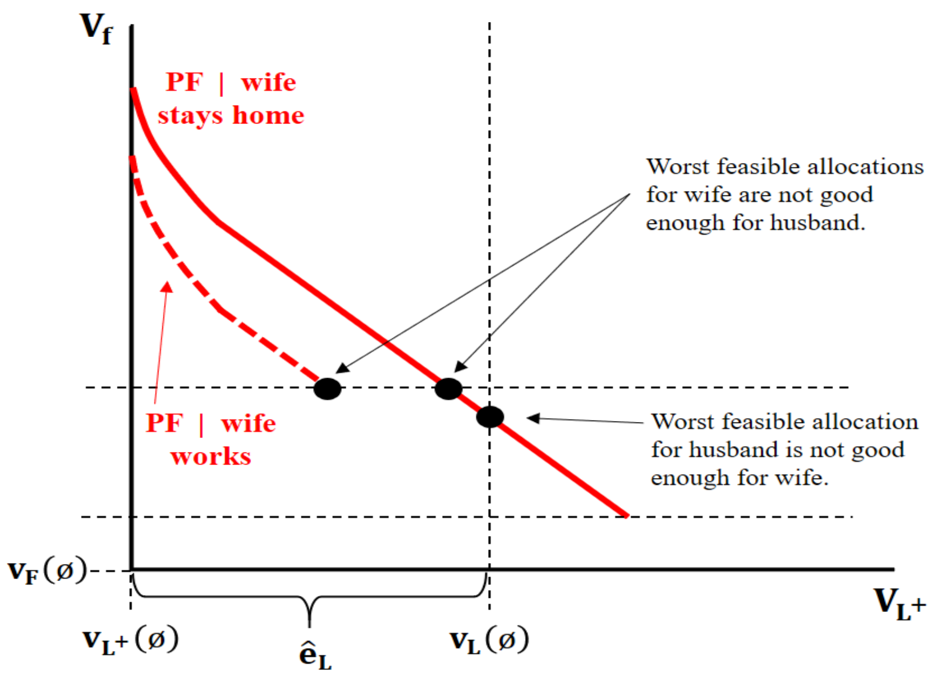

As described in the text. The upper horizontal dotted line is the wife's autarky value conditional on working in period 1. She requires this value in marriage not to divorce. The vertical dotted line denotes the less-skilled husband's participation constraint. He requires this value to invest and marry. The figure considers the same individuals as represented in the left-hand graph of Figure A.4.

[9] This is consistent with the recent observations of Pollak (2019).

A.7.3 Improving women's relative labor market opportunities

I now introduce the second novel result.

Proposition 6. *Regardless of prevailing divorce regime, when women's labor market opportunities improve relative to those of men—that is, w_F increases and/or α falls—less-skilled male labor-force participation decreases.*

The proof of this proposition is contained in Figure A.6. We consider a scenario in which women's initial wage offers rise from their starting level of w_F to a level less than or equal to w_{L^+}—the wage of less-skilled husbands. Additionally, the value of the mother's time at home in the production of child welfare (α) declines. So long as $\alpha \geq 0$ and $w_F \leq w_{L^+}$, it remains efficient for wives to stay home in period 1 and continue to specialize in home production in period 2. Thus, wives generally do not access their increased earning potential in marriage. The only effect of an increase in wives' labor market opportunities relative to husbands' is a reduction in gains from specialization in period 1. This is represented as an inward shift of the *ex-ante* Pareto frontiers in Figure A.6.

On the other hand, an improvement in women's wages raises the value of single life, since single women participate in the labor force.[10] It also raises wives' autarky value from participating in the labor force in period 1, since a unilateral divorce regime grants a wife property right over her earning potential. Thus, regardless of divorce regime, the level of utility a prospective wife must be guaranteed to participate in a marriage rises while the total marital surplus falls. As a result, certain marriages that were incentive-compatible under a large gender gap in labor market opportunities become incompatible as women's opportunities rise. This effect spills over and impacts the pre-marital labor supply behavior of less-educated men.

It is important to note that this result depends on the assumption that marriage is driven primarily by gains from task specialization in the production of household public goods. If consumption complementarities are more important than task specialization—which happens in the model when each partner has a wage above β—an increase in women's labor market opportunities *raises* the gains from marriage and wives' corresponding demand for husbands' earning potential. In the main text I empirically estimate the effect of an improvement in women's labor market opportunities on less-educated male labor supply, using 1980-2015 U.S. data, and find the effect to be negative.

[10] Refer to the optimal labor supply decisions calculated in the Appendix.

Figure A.6: Improvements in women's relative labor market opportunities reduce marriage formation and less-educated male labor supply.

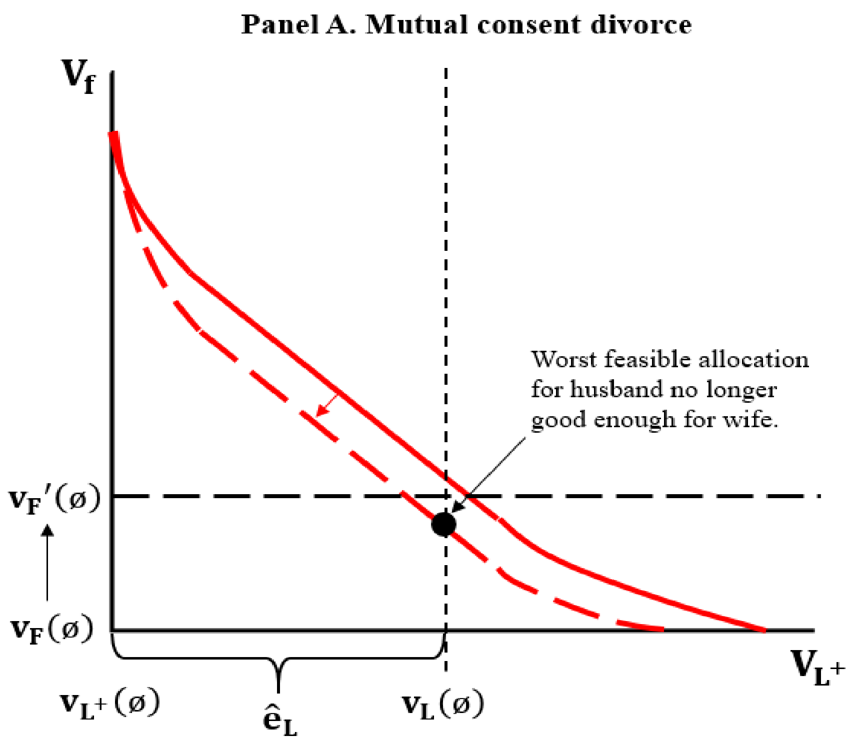

As described in the text.

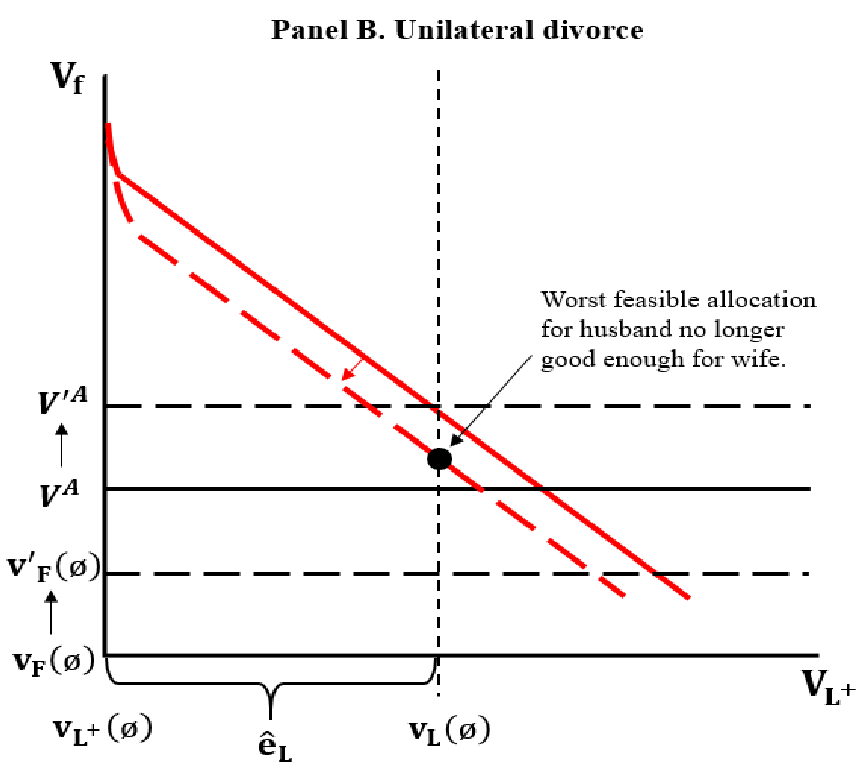

These panels consider a combination of an improvement in women's relative wages and a decline in the non-pecuniary returns to the wife staying at home while raising the child in period 1.

A.8 Omitted calculations and proofs

A.8.1 Utility and wage specifications

Singles' utility $(t = 0, 2)$—Single men (women) have the following utility function for periods 0 and 2 (period 2 only):

$$u(x, L, c) = (x + \beta L) \cdot c. \tag{A.16}$$

Marrieds' utility $(t = 2)$—Married couple (m, f) has the following *material* utility functions for period 2. Utility is conditional on period-2 allocations as well as the child investment strategy from period 1:

$$\begin{aligned}
&u_m(x, L_m, L_f, c_m, w_m, L_{f,1}) \\
&\quad = (x + \beta \cdot [L_m + L_f]) \cdot c_m + \alpha \cdot w_m \mathbf{1}\{L_{f,1} = 1\} + k \\
&\quad = (x + \beta \cdot [L_m + L_f]) \cdot c_m + \omega \\
&u_f(x, L_m, L_f, c_f, w_m, L_{f,1}) \\
&\quad = (x + \beta \cdot [L_m + L_f]) \cdot c_f + \alpha \cdot w_m \mathbf{1}\{L_{f,1} = 1\} + k. \\
&\quad = (x + \beta \cdot [L_m + L_f]) \cdot c_f + \omega
\end{aligned} \tag{A.17}$$

where $\omega = \omega(w_m, L_{f,1})$ is defined by (A.5). These utilities belong to the quasi-linear class described by Bergstrom and Cornes (1983), and thus satisfy the property of transferable utility in the presence of public goods. Married spouses in period 2 thus allocate time and consumption to maximize total material utility, and then use private consumptions c_m and c_f to transfer to each spouse his and her requisite shares of the pie.

Divorcees' utility $(t = 2)$—Divorced couple (m, f) has the following utility functions for period 2:

$$\begin{aligned}
u_m(x_m, L_m, c_m, \omega) &= (x_m + \beta L_m) \cdot c_m + (1 - d)\omega \\
u_f(x_f, L_f, c_f, \omega) &= (x_f + \beta L_f) \cdot c_f + \omega.
\end{aligned} \tag{A.18}$$

Thus, there are two sources of material utility loss on divorce: less enjoyment of child welfare and a loss of efficiency in home commodity production and consumption (because the home commodity is no longer public). These losses are weighed against the innovation in marital bliss (the match quality shock ψ) in the decision to divorce.

The wage structure—Initial wages satisfy the following inequalities:

$$w_{F,1} < w_{L^+} < \beta \le w_H.$$

Next-period wages increase by a factor of $R > 0$ if the individual works full-time in the current period, stay the same if the individual works half-time and depreciate by a factor of δ if the individual does not work at all in the market. That is:

$$\begin{aligned} w_{L^+} &= w_L(1+R) \\ w_{F,2} &= \begin{cases} w_{F,1}(1-\delta) & L_{F,1}=0 \\ w_{F,1} & L_{F,1}=0.5 \end{cases} \end{aligned}$$

Optimal labor supplies in $t=2$—Given these specifications, it is straightforward to show the following (full derivations in the next subsection):

- Married couples in $t=2$ fully specialize, with the wife staying at home and the husband working full-time. Couples involving a less-skilled husband achieve a joint material utility of $\beta w_{L,2} + 2\omega$, while couples involving a highly-skilled husband achieve a joint utility of $(\beta + w_{H,2})^2/4 + 2\omega$.[11]

- Divorced wives and less-skilled husbands, in the absence of a divorce settlement, work half-time in $t=2$, while divorced highly-skilled husbands continue to work full-time. In divorces without settlements, less-skilled ex-husbands obtain $\beta w_{L,2}/4 + (1-d)\omega$; highly-skilled ex-husbands obtain $w_{H,2}^2/4 + (1-d)\omega$; and ex-wives obtain $\beta w_{F,2}/4 + \omega$.

- Single individuals choose the same labor supplies in $t=2$ as divorced individuals who do not negotiate a settlement. Thus, single and divorced less-skilled males generally work less than less-skilled husbands.

A.8.2 Solutions to static labor supply problems in period 2

Married couples—Because utility within marriage is transferable, married couples in period 2 behave as a unitary decision-maker in choosing efficient time and aggregate resource allocations (Mazzocco, 2007; Chiappori, Dias and Meghir, 2018). Thus, a married

[11]Recall that the husband's wage in period 2 depends on his investment decision from period 0.

couple (M, F) in period 2 solves

$$\begin{aligned} \max_{x, L_M, L_F, C} \ & (x + \beta[L_M + L_F]) \cdot C \\ \text{s.t.} \ \ & w_M(1 - L_M) + w_F(1 - L_F) = x + C \end{aligned}$$

and feasibility constraints on time allocation decisions (L_M, L_F).

By assumption 4, the wife has a lower market wage than the husband and so has a comparative advantage in home production. Therefore, she supplies $L_F = 1$. The couple then solves

$$\begin{aligned} \max_{x, L_M, C} \ & (x + \beta[L_M + 1]) \cdot C \\ \text{s.t.} \ \ & w_M(1 - L_M) = x + C. \end{aligned} \tag{A.19}$$

Notice that the marginal utility of x is C, while the marginal utility of C is $x + \beta[L_M + 1]$. Hence, optimization dictates spending the first $\beta[L_M+1]$ of income on C and then equating marginal spending on C and x thereafter. Recall that for a less-skilled husband M, full-income w_M is below β. Therefore, regardless of the less-skilled husband's time allocation decision, the efficient demand for x is 0. These considerations imply the following problem

$$\begin{aligned} & \max_{L_M} \ \beta(L_M + 1) w_M (1 - L_M) \\ = & \max_{L_M} \ \beta w_M (1 - L_M^2) \end{aligned}$$

which is clearly solved by setting $L_M = 0$. Thus, total joint from consumption in a marriage with a less-skilled husband is βw_M and private consumptions satisfy $c_M + c_F = C = w_M$.

For marriages involving a highly-skilled husband M, full specialization prevails as well, although the spending allocations differ. Because the highly-skilled husband's wage exceeds β, some money will be spent on x. Specifically, the first order conditions to program (A.19) dictate the spending of $(w_M - \beta)/2$ on x and allocating the rest to C. Total joint utility from consumption is thus $(\beta + w_M)^2/4$ and individual private consumptions satisfy $c_M + c_F = C = (\beta + w_M)/2$.

Divorcees with a settlement—I start by considering a divorce involving a less-skilled husband. In addition to the indirect utilities from consumption derived below, divorcees also enjoy child welfare. Ex-husbands enjoy $(1 - d)\omega$ on top of consumption utility and ex-wives enjoy ω.

Suppose a court orders, or the ex-couple agrees, for the ex-husband to pay the ex-wife a transfer of $\hat{y}$ (which can be negative). The ex-husband, then, solves

$$\max_{x_M, L_M, c_M} (x_M + \beta L_M) c_M$$
$$\text{s.t. } w_M(1 - L_M) = c_M + x_M + \hat{y}.$$

Because a less-skilled husband's wage does not exceed β, the marginal utility of time at home exceeds the marginal utility of x. Accordingly, he will not work full-time and will devote all monetary resources toward c. Thus the above problem is equivalent to

$$\max_{L_M} \beta L_M (w_M(1 - L_M) - \hat{y})$$

subject to the time constraint on L_M. For ease of exposition, suppose that the space of possible transfers $\hat{y}$ is such that an interior time allocation always prevails for less-skilled ex-husbands and ex-wives. (This will imply $|\hat{y}| \le w_F$.) This problem has first-order condition

$$\beta(w_M(1 - L_M) - \hat{y}) + \beta L_M(-w_M)$$

and second order condition

$$-2\beta w_M < 0.$$

Setting the first order condition to 0 and solving for L_M returns the Marshallian leisure demand

$$L_M(w_M, \hat{y}) = \frac{w_M - \hat{y}}{2w_M}$$

and indirect utility

$$\begin{aligned} U_M(w_M, \hat{y}) &= \beta L_M(w_M, \hat{y}) \cdot (w_M(1 - L_M(w_M, \hat{y})) - \hat{y}) \\ &= \beta \frac{w_M - \hat{y}}{2w_M} \left(\frac{w_M + \hat{y}}{2} - \hat{y} \right) \\ &= \beta \frac{(w_M - \hat{y})^2}{4w_M}. \end{aligned} \tag{A.20}$$

Running through an analogous set of derivations for the wife returns a Marshallian leisure demand of

$$L_F(w_F, \hat{y}) = \frac{w_F + \hat{y}}{2w_F}$$

and indirect utility of

$$\begin{aligned} U_F(w_F, \hat{y}) &= \beta L_F(w_F, \hat{y}) \cdot (w_F(1 - L_F(w_F, \hat{y})) + \hat{y}) \\ &= \beta \frac{w_F + \hat{y}}{2w_F} \left(\frac{w_F - \hat{y}}{2} + \hat{y} \right) \\ &= \beta \frac{(w_F + \hat{y})^2}{4w_F}. \end{aligned} \tag{A.21}$$

Highly-skilled ex-husbands have a wage exceeding β. Thus, the marginal utility of x exceeds the marginal utility of time at home, and so these individuals work full-time regardless of the divorce settlement. Standard optimization dictates that these men equally divide total income, $w_M - \hat{y}$, between x and c, reaching an indirect utility of

$$U_M(w_M, \hat{y}) = \frac{(w_M - \hat{y})^2}{4}. \tag{A.22}$$

Singles, divorcees without a settlement—Divorces without settlements are equivalent to divorces with settlements in which $\hat{y} = 0$. In these "autarky allocations," less-skilled ex-husbands and ex-wives each work half-time and obtain consumption utilities of $\beta w_M/4$ and $\beta w_F/4$, while highly-skilled ex-husbands continue to work full-time and obtain a consumption utility of $w_M^2/4$. By definition, singles in period 2 behave the same way as divorcees who do not negotiate a settlement. The only difference between singles and divorcees without a settlement is that the latter group enjoys additional utility from child welfare.

Hence, less-skilled single and divorced men generally do not work full-time, but less-skilled husbands do.

A.8.3 Omitted proofs

Proof of Proposition 1.

The first claim is that highly skilled men always invest in period 0. Recall the investment decision rule:

$$I_M = \mathbf{1}\{\lambda_M + \mu_M \geq e_M\}$$

First, note that the marriage market return to investing, μ_M, is weakly positive for all men M. This is because men always have the option of remaining single and simply enjoying the labor market returns to investment. Because the wage return to investing, R, is strictly positive, the labor market return to investing is strictly positive. From the above calculations, we can show that the effort cost of investing for skilled men

$$\begin{aligned} e_H &= u_{H,0}(I_H = 0) - u_{H,0}(I_H = 1) \\ &= \left(\frac{\beta + w_H}{2}\right)^2 - \left(\frac{w_H}{2}\right)^2, \end{aligned}$$

is weakly negative, since $w_H \geq \beta$. Thus, $\lambda_H + \mu_H > e_H$, so highly skilled men always invest.

The second claim is that less-skilled men who do not marry choose not to invest. This would happen if $\lambda_L < e_L$. To prove this claim, we introduce the following assumption.

Assumption 4. $R < \frac{\beta - w_L}{\beta}$.

From the above calculation,

$$\begin{aligned} e_L &= u_{L,0}(I_L = 0) - u_{L,0}(I_L = 1) \\ &= \frac{\beta w_L}{4} - \frac{w_L^2}{4} \\ &= \frac{w_L}{4}(\beta - w_L) \end{aligned}$$

and

$$\begin{aligned} \lambda_L &= \frac{\beta w_L(1+R)}{4} - \frac{\beta w_L}{4} \\ &= \frac{w_L}{4} R\beta \end{aligned}$$

By assumption 4, $R < \frac{\beta - w_L}{\beta}$ and so $\lambda_L < e_L$. Therefore, less-skilled men who choose to remain single also do not invest. ▮

Proof of Proposition 2. The claim is that all men who choose to marry are better off investing. To prove this, it is convenient to introduce the following assumption and lemma.

Assumption 5. $w_L > \beta\left(1 - \frac{7-\xi}{2} R\right)$, *where* $\xi = \frac{(w_F/w_L)^2}{1+R}$.

Lemma 1. *No marriage forms simply as a route to divorce.*

Proof of lemma. Consider a marriage involving a less-skilled man. According to the above derivations, total utility in marriage in period 2, exclusive of marital bliss, is $\beta w_M + 2\omega$. By summing up the indirect utilities derived above, total consumption utility in divorce is

$$
\begin{aligned}
&\beta\frac{(w_M-\hat{y})^2}{4w_M}+\beta\frac{(w_F+\hat{y})^2}{4w_F} \\
&=\left(\frac{\beta w_M}{4}-\frac{\beta\hat{y}}{2}+\frac{\beta\hat{y}^2}{4w_M}\right)+\left(\frac{\beta w_F}{4}+\frac{\beta\hat{y}}{2}+\frac{\beta\hat{y}^2}{4w_F}\right) \\
&=\frac{\beta w_M}{4}+\frac{\beta w_F}{4}+\frac{\beta\hat{y}^2}{4w_M}+\frac{\beta\hat{y}^2}{4w_F} \\
&<\frac{\beta w_M}{2}+\frac{\beta w_F}{2}
\end{aligned}
$$

given the feasible set of transfers $|\hat{y}| \leq w_F$. Notice that this quantity, in turn, is strictly less than total consumption utility in marriage (βw_M). Moreover, total child welfare in divorce $(2-d)\omega$ is less than that in marriage (2ω), so total material utility is strictly lower in divorce. Because utility in marriage is transferable, it is possible to guarantee each spouse a higher material utility in marriage than in divorce, for any given feasible divorce allocation. Thus, regardless of the divorce regime, it is better for the marriage to continue when the match quality shock $\psi = 0$; only when $\psi = -\theta$ is divorce a possibility.

As a consequence, any marriage that forms in period 1 involving a less-skilled man dissolves in period 2 with probability at most $1/2$—the chance of drawing a bad match quality.

The same logic holds for marriages involving a highly-skilled man. ∎

Proof of proposition. Denote the cost of a less-skilled man investing, net of the labor market return, as $\hat{e}_L = e_L - \lambda_L$. What we must show is that this net cost is more than returned in the marriage market; that is, that $\mu_L \geq \hat{e}_L$. From the proof to proposition 1, we have

$$
\begin{aligned}
\hat{e}_L &= e_L - \lambda_L \\
&= \frac{w_L}{4}(\beta - w_L) - \frac{w_L}{4}R\beta \\
&= \frac{w_L}{4}(\beta(1-R) - w_L).
\end{aligned}
$$

By the lemma, a less-skilled husband divorces with probability at most $1/2$. Suppose for the moment that the couple has zero marital taste and that $\alpha = 0$, so child welfare does not depend on the husband's wage. This yields

$$\begin{aligned}
\mu_L &= \mathsf{V}_{L^+} - \mathsf{V}_L \\
&= (\mathsf{V}_{L^+} + \mathsf{V}_F) - (\mathsf{V}_L + \mathsf{V}_F) \\
&\geq \frac{1}{2}\left((U_{L^+,2} + U_{F,2}) - (U_{L,2} + U_{F,2}) \mid \text{stay married}\right) \\
&\quad + \frac{1}{2}\left((U_{L^+,2} + U_{F,2}) - (U_{L,2} + U_{F,2}) \mid \text{divorce}\right) \\
&= \frac{1}{2}(\beta w_L(1+R) - \beta w_L) \\
&\quad + \frac{1}{2}\left(\frac{\beta w_L(1+R)}{4} + \frac{\beta \hat{y}^2}{4w_L(1+R)} - \frac{\beta w_L}{4} - \frac{\beta \hat{y}^2}{4w_L}\right) \\
&= \frac{1}{2}R\beta w_L + \frac{1}{2}R\beta\left(\frac{w_L}{4} - \frac{\hat{y}^2}{4w_L(1+R)}\right) \\
&\geq \frac{1}{2}R\beta w_L\left(1 + \frac{1-\xi}{4}\right) \\
&= \frac{1}{8}R\beta w_L(5-\xi)
\end{aligned} \tag{A.23}$$

where the second-last line is derived from factoring a w_L out of the last term and then applying the definition of ξ under the assumption that $\hat{y} = w_F$. By this assumption, the equality in the second-last line becomes an inequality, since the feasible set of transfers $|\hat{y}| \leq w_F$.

Given these calculations, observe that the expression $\mu_L \geq \hat{e}_L$ is equivalent to

$$\begin{aligned}
\frac{1}{8}R\beta w_L(5-\xi) - \frac{w_L}{4}(\beta(1-R) - w_L) &\geq 0 \\
\frac{w_L}{4}\left(\frac{1}{2}R\beta(5-\xi)\right) - \frac{w_L}{4}(\beta(1-R) - w_L) &\geq 0 \\
\frac{1}{2}R\beta(5-\xi) - \beta(1-R) + w_L &\geq 0 \\
w_L &\geq \beta\left(1 - \frac{7-\xi}{2}R\right),
\end{aligned}$$

which is true by assumption 5.

Notice that the above scenario is a "worst case." In reality, $\alpha \geq 0$, the divorce probability is weakly decreasing in the male's wage (since the material marital surplus rises in the male's wage), the marital taste may be positive, and the divorce settlement $|\hat{y}| \leq w_F$. It is easy to show that each of these possibilities raises the expected marriage market returns to investing further.[12] Thus, if investment conditional on marriage is profitable in this worst

[12] Calculations available upon request.

case scenario, it is profitable always. ∎

Notice that assumption 5, then, is sufficient but not necessary for this proposition to hold.

Proof of Proposition 3. The claim is that a rational-expectations equilibrium exists. To begin, it is convenient to introduce the following lemma.

Lemma 2. *Take two individuals of the same skill type, i and j, with $\theta_i > \theta_j$. No equilibria exist in which j marries while i remains single.*

Proof of lemma. Suppose such a marriage market allocation arose in equilibrium. Suppose that j preferred to be single at the given price vector. Then, by definition, such an allocation is not an equilibrium. Suppose instead j was indeed best off marrying and promising his/her partner her equilibrium marriage market price of P. Then i would prefer to be married as well, and can promise j's partner a price of $\mathsf{P} + \theta_i - \theta_j - \epsilon$, where $\epsilon \in (0, \theta_i - \theta_j)$. Both i and j's initial partner will be better off in the new allocation. Therefore, the initial allocation is not an equilibrium. ∎

Proof of proposition. By the lemma, the algorithm converges toward a triple of marital taste threshold values $(\underline{\theta}_L, \underline{\theta}_H, \underline{\theta}_F)$. For the market to clear, we require $1 - G(\underline{\theta}_F) = 1 - \frac{1}{2}\left(G(\underline{\theta}_L) + G(\underline{\theta}_H)\right)$.

Then, define conditions on the $\underline{\theta}$ values. Link the output of the algorithm to male investment strategy. Skilled males always invest. Less-skilled males invest only if they can get at least $\hat{e}_L$ of the marital surplus from marrying.

Investing when expected marital gain is less than $\hat{e}_L$ is not rational. Not investing and marrying is ruled out by proposition 2. Hence this investment rule and associated stable matching is an equilibrium. ∎

Proof of Proposition 4. The claim is that in a mutual consent divorce regime, and where the initial female wage is less than β, spouses will always agree on the efficient child investment strategy, which is to have the wife stay home full-time in period 1.

At the initial wage structure, regardless of the wife's labor supply decision in period 1, a married household will fully specialize in period 2—with the wife staying at home full-time (see above calculations). Thus, the only impact in marriage of the wife working in the market in period 1 is to decrease total enjoyment of child welfare. Because utility within marriage is transferable, this implies that for any marital allocation in which the wife works in the first period, there exists an allocation in which the wife does not work that both partners prefer. That is, the *ex-post* Pareto frontier in marriage conditional on the wife staying at home is dominated by the *ex-post* frontier conditional on her working.

By assumption 3, the Pareto frontier in divorce is unaffected by the wife's labor supply choice. Moreover, by Lemma 1, no marriage forms simply as a route to divorce: regardless of male and female types, any marriage that forms has at most a 50 percent chance of dissolving in period 2: i.e., the probability of drawing a bad match quality shock. Thus the *ex-ante* Pareto frontier is a convex combination of the *ex-post* frontiers in marriage and the frontier in divorce. As a consequence, the *ex-ante* Pareto frontier conditional on the wife staying at home dominates the frontier conditional on the wife working. That is, given any *ex-ante* allocation in which the wife works in the first period, there exists a feasible *ex-ante* allocation in which the wife stays home that both partners weakly prefer. $\blacksquare$

APPENDIX B

Chapter II Supporting Material

B.1 Appendix Figures and Tables

Figure B.1: Confirming the 21st century increase in skilled young mothers' full-time employment across measure of skill and data source

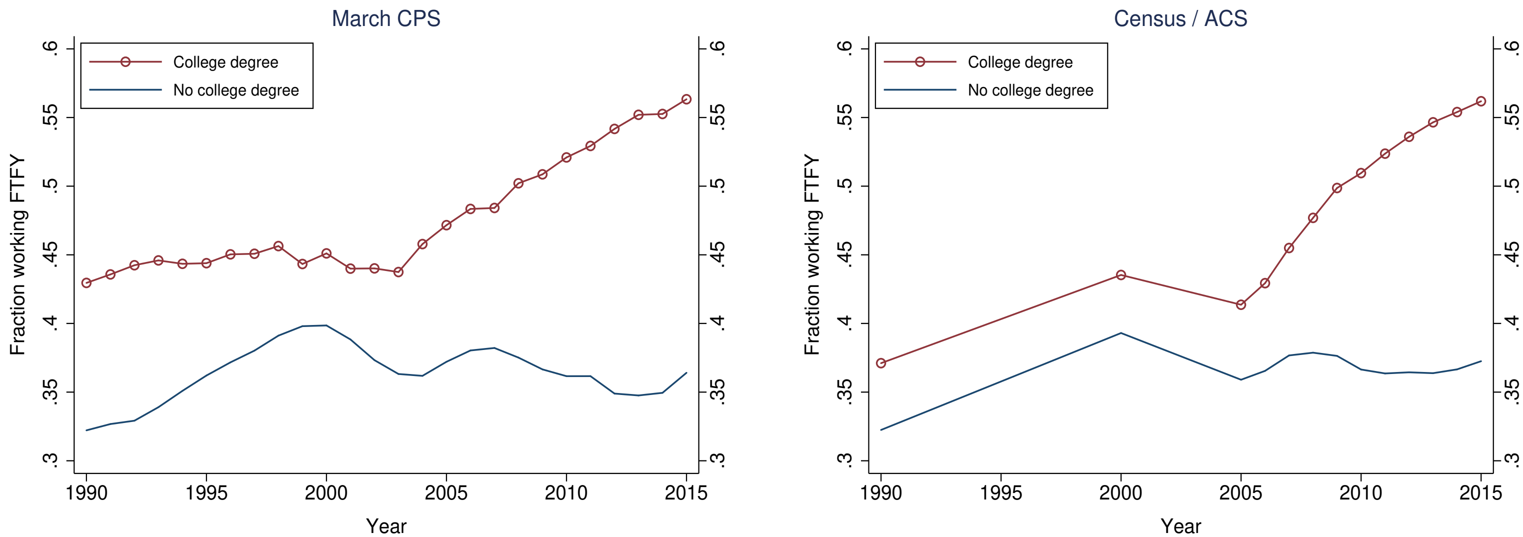

Notes: See Section 2.2 for main sample definition. These graphs restrict the main sample to the sub-sample of individuals with at least one child under the age of 6, and stratifies this sub-sample by 4-year college attainment status. The March CPS series consists of 3-year centered moving averages. The Census / ACS series consists of 1990, 2000, and 2005 data followed by 3-year centered moving averages for 2006-2015. The CPS series defines full-time, full-year (FTFY) employment as at least 1600 hours worked in the reference year. The Census series defines FTFY employment as having been employed for at least 40 weeks and working at least 35 hours per week in the reference year.

Figure B.2: Evolution of egalitarian gender attitudes by skill, 1985-2014

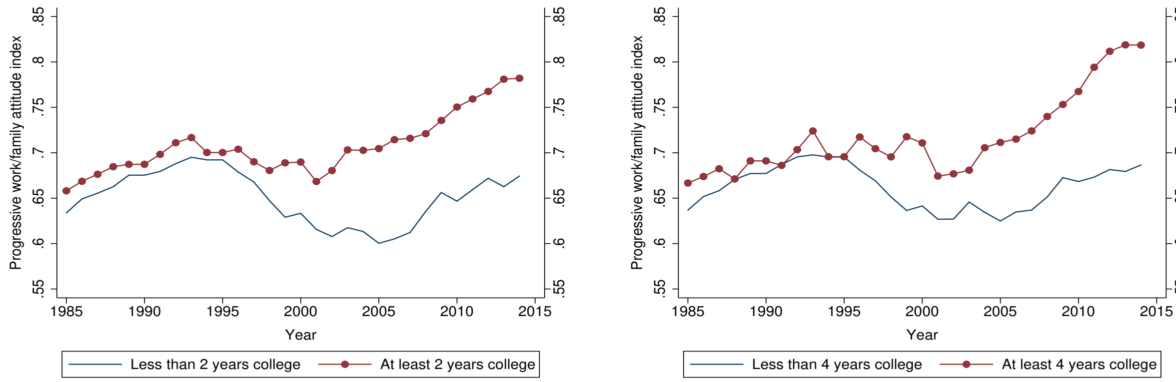

Notes: The progressive work-family attitude index is calculated from responses to 3 questions in the General Social Survey and is coded such that higher values correspond to higher support of mothers working while raising family. See section 2.4 for further detail. Due to small sample sizes, 5-year centered moving averages are reported. The main sample considered is individuals aged 22-45 with at least one child.

B.2 Data Details

This section describes in detail the processing of the various sources of data used throughout the paper.

B.2.1 Adjusting for top-codes

Around 3 percent of women and 5 percent of men in the data have their earnings top-coded by the CPS. I followed a procedure similar to Autor, Katz and Kearney (2008) adjust for top-coded values. The CPS separately top-codes the primary earnings source (ERN-VAL) and "other" earnings sources (WS-VAL). Before 1996, both sources of earnings were top-coded at 99,999. I replaced top-coded values with the top-code multiplied by 1.5. Between 1996 and 2010, primary earnings above the top-code threshold are assigned the average value of other top-coded earners.[1] I did not further adjust these values. Secondary earnings are top-coded at 25,000 between 1996 and 2002, and at 35,000 between 2003 and 2010. Once again I replaced top-coded secondary earnings values with the top-code multiplied by 1.5. Since 2011, and for both sources of earnings, reported earnings above the top-code threshold are systematically swapped with other values above the top-code, within a bounded interval. I elected not to adjust these values.

B.2.2 Imputing annual hours worked before the 1976 survey

The March CPS sample used throughout the paper begins in 1964. Before the 1976 survey, only binned weeks worked per year and the number of hours worked last week are reported, instead of the actual number of weeks worked and usual hours worked per week. For these years, annual hours worked were in the following manner. I consider a sample of women from the 1976-1981 survey years, where all 4 items are observed (binned weeks, number of weeks, hours worked last week, and usual hours worked per week).[2] On this sample I regress the number of weeks worked last year on a race dummy, region dummies, 6 educational attainment dummies (0-8 years, 9-11 years, high school degree, some college, BA, post-BA), a quadratic in potential experience, and the education dummies interacted with the experience quadratic. I perform this regression within each weeks worked bin (1-13, 14-26, 27-39, 40-47, 48-49, 50-52). I then impute the number

[1] The top-code threshold and replacement value is conditional on race, sex, and FTFY employment status.

[2] I use only 1976-1981 to impute the data instead of all years post-1976 in case the relationship between actual weeks and binned weeks, or between usual hours and hours worked last week, has changed over time.

of weeks worked, for the 1964-1975 sample of women not in the 0 weeks worked bin, as fitted values from these regressions.[3]

Next, I do a similar imputation for usual hours worked per week. On the 1976-1981 sample I regress usual hours worked per week on a race dummy, regional dummies, education dummies, all of the above interacted with the number of hours worked last week, the number of hours worked last week, and weeks worked interval dummies. I impute usual hours worked per week, for the 1964-1975 sample of women not in the 0 weeks worked bin, as fitted values from these regressions.

Imputed annual hours worked is the product of imputed weeks and imputed usual hours worked per week.

B.2.3 Imputing wage offers for non-working women in both the PSID and CPS

Hourly wages are calculated as total wage and business income divided by annual hours worked. Wages are converted to 2010 US dollars via the Consumer Price Index. To reduce measurement error, I trim wages below $2.50 or above $150 from the sample.

I use a Heckman two-step procedure to impute wages for non-working women. First I perform a probit regression of labor-force participation status on education dummies, a quadratic in potential experience, their interaction, education dummies interacted with a race dummy, a race dummy, education dummies interacted with metropolitan central city status dummies, metro central city status dummies, education dummies interacted with region dummies, region dummies, education dummies interacted with survey year dummies, survey year dummies, a 3-level education variable (high school or less, some college, BA or more) interacted with marital status dummies (married, separated or divorced, never married), marital status dummies, and marital status dummies interacted with the number of children under age 5. (The geographical variables are not used in the PSID, due to lack of sufficient sample size.)

Next I perform a Mincerian regression, on the sample of working women, of the observed log hourly wage on all of the above controls, excluding the interaction between marital status dummies and the number of young children, and including an inverse Mills ratio control function. Finally I impute log wage offers for non-working women using fitted values from this regression. I iterate this imputation procedure over 6-year bins to allow the wage and selection equations to change over time.

[3] In a few cases the fitted value was larger than 52. I replaced these fitted values with 52.

B.2.4 Constructing the high-wage-offer / low-wage-offer classification

As discussed in the main text, the main sample contains individuals with 0-40 years of potential labor market experience. Following a procedure used by Juhn and Murphy (1997), I partition the main sample into cells defined by survey year and potential experience group. I consider 16 potential experience groups: $0-2, 3-4, 5-6, 7-8, \ldots, 27-28, 29-30, 31-40$. Within each cell, I compute the median of the population-weighted sample and assign as high-skill (low-skill) each woman with a wage offer above (below) the median. Note that the grouping of those with 31-40 years of potential experience into one bin is done in recognition of the fact that the main sample only includes individuals aged 22-54. Because of this age ceiling, variation in potential experience above 30 is almost certainly due to variation in educational attainment rather than variation in age conditional on educational attainment. Thus, using a finer binning would group too many of the oldest individuals in the sample into the high-wage bin, and not enough into the low-wage bin.

BIBLIOGRAPHY

Abraham, Katharine G, and Melissa S Kearney. 2018. "Explaining the Decline in the U.S. Employment-to-Population Ratio: A Review of the Evidence." National Bureau of Economic Research Working Paper 24333.

Acemoglu, Daron, and Pascual Restrepo. 2017. "Robots and Jobs: Evidence from US Labor Markets." National Bureau of Economic Research Working Paper 23285.

Acemoglu, Daron, David H. Autor, David Dorn, Gordon H. Hanson, and Brendan Price. 2016. "Import Competition and the Great US Employment Sag of the 2000s." *Journal of Labor Economics*, 34(1): 141–198.

Aguiar, Mark, Mark Bils, Kerwin Kofi Charles, and Erik Hurst. 2017. "Leisure Luxuries and the Labor Supply of Young Men." National Bureau of Economic Research Working Paper 23552.

Akerlof, George A., and Rachel E. Kranton. 2000. "Economics and Identity." *The Quarterly Journal of Economics*, 115(3): 715–753.

Alesina, Alberto, and Paola Giuliano. 2010. "The Power of the Family." *Journal of Economic Growth*, 15(2): 93–125.

Alesina, Alberto, Paola Giuliano, and Nathan Nunn. 2013. "On the Origins of Gender Roles: Women and the Plough." *The Quarterly Journal of Economics*, 128(2): 469–530.

Aliprantis, Dionissi, Kyle Fee, and Mark E. Schweitzer. 2019. "Opioids and the Labor Market." Accessed at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3179068.

Alon, Titan, Matthias Doepke, and Sena Ekin Coskun. 2018. "Trends in Work and Leisure: It's a Family Affair." Society for Economic Dynamics 2018 Meeting Papers 552.

Antonovics, Kate, and Robert Town. 2004. "Are All the Good Men Married? Uncovering the Sources of the Marital Wage Premium." *American Economic Review*, 94(2): 317–321.

Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos. 2008. "Explaining Changes in Female Labor Supply in a Life-Cycle Model." *The American Economic Review*, 98(4): 1517–1552.

Autor, David H. 2014. "Skills, Education, and the Rise of Earnings Inequality among the 'Other 99 Percent'." *Science*, 344(6186): 843–851.

Autor, David H., and David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review*, 103(5): 1553–97.

Autor, David H., and Melanie Wasserman. 2013. "Wayward Sons: The Emerging Gender Gap in Labor Markets and Education." Third Way Next Initative, Third Way. accessed Aug 11, 2017 at: http://www.thirdway.org/report/wayward-sons-the-emerging-gender-gap-in-labor-markets-and-education.

Autor, David H., David Dorn, and Gordon Hanson. 2019. "When Work Disappears: Manufacturing Decline and the Falling Marriage-Market Value of Men." *American Economic Review: Insights*.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *The American Economic Review*, 103(6): 2121–2168.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." *The Review of Economics and Statistics*, 90(2): 300–323.

Bartik, Timothy J. 1991. *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute for Employment Research.

Bartik, Timothy J., and Nathan Sotherland. 2019. "Local Job Multipliers in the United States: Variation with Local Characteristics and High-Tech Shocks." Upjohn Institute Working Paper 19-301.

Becker, Gary S. 1973. "A Theory of Marriage: Part I." *Journal of Political Economy*, 81(4): 813–846.

Becker, Gary S. 1981. *A Treatise on the Family.* Harvard University Press.

Belot, Michéle, and Jan Fidrmuc. 2010. "Anthropometry of Love: Height and Gender Asymmetries in Interethnic Marriages." *Economics and Human Biology*, 8(3): 361 – 372.

Belot, Michéle, and Marco Francesconi. 2013. "Dating Preferences and Meeting Opportunities in Mate Choice Decisions." *Journal of Human Resources*, 48(2): 474–508.

Benedetto, Gary, Martha Stinson, and John M. Abowd. 2013. "The Creation and Use of the SIPP Synthetic Beta." Accessed at https://census.gov/content/dam/Census/programs-surveys/sipp/methodology/SSBdescribe_nontechnical.pdf.

Bergstrom, Theodore C., and Richard C. Cornes. 1983. "Independence of Allocative Efficiency from Distribution in the Theory of Public Goods." *Econometrica*, 51(6): 1753–1765.

Bertrand, Marianne. 2011. "New Perspectives on Gender." *Handbook of Labor Economics*, 4: 1543 – 1590.

Bertrand, Marianne, Claudia Goldin, and Lawrence F. Katz. 2010. "Dynamics of the Gender Gap for Young Professionals in Corporate Sectors." *American Economic Journal: Applied Economics*, 2(3): 228–255.

Bertrand, Marianne, Emir Kamenica, and Jessica Pan. 2015. "Gender Identity and Relative Income within Households." *The Quarterly Journal of Economics*, 130(2): 571–614.

Binder, Ariel J., and John Bound. 2019. "The Declining Labor Market Prospects of Less-Educated Men." *Journal of Economic Perspectives*, 33(2): 163–90.

Bisin, Alberto, and Thierry Verdier. 2001. "The Economics of Cultural Transmission and the Dynamics of Preferences." *Journal of Economic Theory*, 97(2): 298 – 319.

Black, Dan A., Natalia Kolesnikova, Seth G. Sanders, and Lowell J. Taylor. 2013. "Are Children "Normal"?" *The Review of Economics and Statistics*, 95(1): 21–33.

Blau, David M. 2003. "Child Care Subsidy Programs." In *Means-Tested Transfer Programs*. Vol. 76, , ed. Robert Moffitt, Chapter 7, 455–469. Chicago:University of Chicago Press.

Blau, Francine D., and Lawrence M. Kahn. 2013. "Female Labor Supply: Why Is the United States Falling Behind?" *American Economic Review: Papers and Proceedings*, 103(3): 251–256.

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

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2018. "Quasi-Experimental Shift-Share Research Designs." National Bureau of Economic Research Working Paper 24997.

Bound, John, and Harry J. Holzer. 2000. "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s." *Journal of Labor Economics*, 18(1): 20–54.

Bound, John, and Richard V. Burkhauser. 1999. "Economic Analysis of Transfer Programs Targeted on People with Disabilities." In . Vol. 3 of *Handbook of Labor Economics*, 3417 – 3528. Elsevier.

Bound, John, and Timothy Waidmann. 1992. "Disability Transfers, Self-Reported Health, and the Labor Force Attachment of Older Men: Evidence from the Historical Record." *The Quarterly Journal of Economics*, 107(4): 1393–1419.

Bound, John, and Timothy Waidmann. 2002. "Accounting for Recent Declines in Employment Rates among Working-Aged Men and Women with Disabilities." *The Journal of Human Resources*, 37(2): 231–250.

Bound, John, Brad Hershbein, and Bridget Terry Long. 2009. "Playing the Admissions Game: Student Reactions to Increasing College Competition." *The Journal of Economic Perspectives*, 23(4): 119–146.

Boustan, Leah Platt, and William J. Collins. 2014. "The Origin and Persistence of Black-White Differences in Women's Labor Force Participation." In *Human Capital in History: The American Record.* , ed. Leah Platt Boustan, Carola Frydman and Robert A. Margo, 205–240. University of Chicago Press. Chapter 6.

Braga, Breno. 2018. "Earnings Dynamics: The role of Education throughout a Worker's Career." *Labour Economics*, 52: 83 – 97.

Brennan, Robert T., Rosalind Chait Barnett, and Karen C. Gareis. 2001. "When She Earns More Than He Does: A Longitudinal Study of Dual-Earner Couples." *Journal of Marriage and Family*, 63(1): 168–182.

Buera, Francisco J., and Joseph P. Kaboski. 2012. "The Rise of the Service Economy." *American Economic Review*, 102(6): 2540–69.

Bumpass, Larry L. 1990. "What's Happening to the Family? Interactions between Demographic and Institutional Change." *Demography*, 27(4): 483–498.

Case, Anne, and Angus Deaton. 2015. "Rising morbidity and mortality in midlife among white non-Hispanic Americans in the 21st century." *Proceedings of the National Academy of Sciences*, 112(49): 15078–15083.

CEA. 2016. "The Long-Term Decline in Prime-Age Male Labor Force Participation." Accessed at: https://www.whitehouse.gov/sites/default/files/page/files/20160620_cea_primeage_male_lfp.pdf.

Charles, Kerwin Kofi, and Melvin Stephens, Jr. 2004. "Job Displacement, Disability, and Divorce." *Journal of Labor Economics*, 22(2): 489–522.

Charles, Kerwin Kofi, and Ming Ching Luoh. 2010. "Male Incarceration, the Marriage Market, and Female Outcomes." *The Review of Economics and Statistics*, 92(3): 614–627.

Charles, Kerwin Kofi, Erik Hurst, and Mariel Schwartz. 2018. "The Transformation of Manufacturing and the Decline in U.S. Employment." In . *NBER Macroeconomics Annual: Volume 33*, , ed. Martin Eichenbaum and Jonathan A. Parker. University of Chicago Press.

Cherlin, Andrew. 2014. *Labor's Love Lost: The Rise and Fall of the Working-Class Family in America.* Russell Sage Foundation.

Chiappori, Pierre-André. 1988. "Rational Household Labor Supply." *Econometrica*, 56(1): 63–90.

Chiappori, Pierre-André. 1992. "Collective Labor Supply and Welfare." *Journal of Political Economy*, 100(3): 437–467.

Chiappori, Pierre-André, and Bernard Salanié. 2016. "The Econometrics of Matching Models." *Journal of Economic Literature*, 54(3): 832–61.

Chiappori, Pierre-André, Bernard Salanié, and Yoram Weiss. 2017. "Partner Choice, Investment in Children, and the Marital College Premium." *American Economic Review*, 107(8): 2109–67.

Chiappori, Pierre-André, Monica Costa Dias, and Costas Meghir. 2018. "The Marriage Market, Labor Supply, and Education Choice." *Journal of Political Economy*, 126(S1): S26–S72.

Chiappori, Pierre-André, Murat Iyigun, and Yoram Weiss. 2009. "Investment in Schooling and the Marriage Market." *The American Economic Review*, 99(5): 1689–1713.

Chiappori, Pierre-André, Murat Iyigun, and Yoram Weiss. 2015. "The Becker-Coase Theorem Reconsidered." *Journal of Demographic Economics*, 81(2): 157–177.

Chiappori, Pierre-André, Murat Iyigun, Jeanne Lafortune, and Yoram Weiss. 2017. "Changing the Rules Midway: The Impact of Granting Alimony Rights on Existing and Newly Formed Partnerships." *The Economic Journal*, 127(604): 1874–1905.

Choo, Eugene, and Aloysius Siow. 2006. "Who Marries Whom and Why." *Journal of Political Economy*, 114(1): 175–201.

Coglianese, John. 2018. "The Rise of In-and-Outs: Declining Labor Force Participation of Prime Age Men." Accessed at `https://www.johncoglianese.com/publication/in-and-outs/in-and-outs.pdf`.

Coile, Courtney C., and Mark G. Duggan. 2019. "When Labor's Lost: Health, Family Life, Incarceration, and Education in a Time of Declining Economic Opportunity for Low-Skilled Men." *Journal of Economic Perspectives*, 33(2): 191–210.

Coontz, Stephanie. 1992. *The Way We Never Were: American Families and the Nostaliga Trap.* Basic Books.

Copen, Casey E., Kimberly Daniels, and William D. Mosher. 2013. "First Premarital Cohabitation in the United States: 2006-2010 National Survey of Family Growth." *National Health Statistics Reports*, 64(64): 1–15.

Crawford, Vincent P., and Elsie Marie Knoer. 1981. "Job Matching with Heterogeneous Firms and Workers." *Econometrica*, 49(2): 437–450.

Currie, Janet, Jonas Y Jin, and Molly Schnell. 2018. "U.S. Employment and Opioids: Is There a Connection?" National Bureau of Economic Research Working Paper 24440.

Duncan, Greg J., Bessie Wilkerson, and Paula England. 2006. "Cleaning up Their Act: The Effects of Marriage and Cohabitation on Licit and Illicit Drug Use." *Demography*, 43(4): 691–710.

Dupuy, Arnaud, and Alfred Galichon. 2014. "Personality Traits and the Marriage Market." *Journal of Political Economy*, 122(6): 1271–1319.

Dustmann, Christian, and Costas Meghir. 2005. "Wages, Experience and Seniority." *The Review of Economic Studies*, 72(1): 77–108.

Eccles, Jacquelynne S., Janis E. Jacobs, and Rena D. Harold. 1990. "Gender Role Stereotypes, Expectancy Effects, and Parents' Socialization of Gender Differences." *Journal of Social Issues*, 46(2): 183–201.

Echenique, Federico, Sangmok Lee, Matthew Shum, and M. Bumin Yenmez. 2013. "The Revealed Preference Theory of Stable and Extremal Stable Matchings." *Econometrica*, 81(1): 153–171.

Eickmeyer, Kasey J., and Wendy D. Manning. 2018. "Serial Cohabitation in Young Adulthood: Baby Boomers to Millennials." *Journal of Marriage and Family*, 80(4): 826–840.

Elsby, Michael W. L., and Matthew D. Shapiro. 2012. "Why Does Trend Growth Affect Equilibrium Employment? A New Explanation of an Old Puzzle." *The American Economic Review*, 102(4): 1378–1413.

Farré, Lídia, and Francis Vella. 2013. "The Intergenerational Transmission of Gender Role Attitudes and its Implications for Female Labour Force Participation." *Economica*, 80(318): 219–247.

Fernández, Raquel, Alessandra Fogli, and Claudia Olivetti. 2004. "Mothers and Sons: Preference Formation and Female Labor Force Dynamics." *The Quarterly Journal of Economics*, 119(4): 1249–1299.

Fernández, Raquel, and Alessandra Fogli. 2009. "Culture: An Empirical Investigation of Beliefs, Work, and Fertility." *American Economic Journal: Macroeconomics*, 1(1): 146–77.

Fernández, Raquel, and Joyce Cheng Wong. 2017. "Free to Leave? A Welfare Analysis of Divorce Regimes." *American Economic Journal: Macroeconomics*, 9(3): 72–115.

Fisman, Raymond, Sheena S. Iyengar, Emir Kamenica, and Itamar Simonson. 2006. "Gender Differences in Mate Selection: Evidence From a Speed Dating Experiment." *The Quarterly Journal of Economics*, 121(2): 673–697.

Fortin, Nicole M. 2005. "Gender Role Attitudes and the Labour-market Outcomes of Women across OECD Countries." *Oxford Review of Economic Policy*, 21(3): 416–438.

Friedberg, Leora. 1998. "Did Unilateral Divorce Raise Divorce Rates? Evidence from Panel Data." *The American Economic Review*, 88(3): 608–627.

Fuwa, Makiko. 2004. "Macro-level Gender Inequality and the Division of Household Labor in 22 Countries." *American Sociological Review*, 69(6): 751–767.

Galambos, Nancy L. 2004. "Gender and Gender Role Development in Adolescence." In *Handbook of Adolescent Psychology*. Vol. 2, , ed. Richard M. Lerner and Laurence Steinberg, 233–262. John Wiley and Sons, Inc.

Galassi, Gabriela, David Koll, and Lukas Mayr. 2018. "Intergenerational Correlation of Employment: Is there a Role for Work Culture?" European University Institute, working paper.

Gale, D., and L. S. Shapley. 1962. "College Admissions and the Stability of Marriage." *The American Mathematical Monthly*, 69(1): 9–15.

Galichon, Alfred, and Bernard Salanié. 2017. "The Econometrics and Some Properties of Separable Matching Models." *American Economic Review*, 107(5): 251–55.

Galichon, Alfred, Scott Duke Kominers, and Simon Weber. 2018. "Costly Concessions: An Empirical Framework for Matching with Imperfectly Transferable Utility." Available: https://www.journals.uchicago.edu/doi/pdfplus/10.1086/702020.

Gayle, George-Levi, and Andrew Shephard. 2019. "Optimal Taxation, Marriage, Home Production, and Family Labor Supply." *Econometrica*, 87(1): 291–326.

Gervais, Antoine, and J. Bradford Jensen. 2019. "The Tradability of Services: Geographic Concentration and Trade Costs." *Journal of International Economics*, 118: 331 – 350.

Gillis, John Stuart, and Walter E. Avis. 1980. "The Male-Taller Norm in Mate Selection." *Personality and Social Psychology Bulletin*, 6(3): 396–401.

Goldin, Claudia. 2006. "The Quiet Revolution That Transformed Women's Employment, Education, and Family." *The American Economic Review*, 96(2): 1–21.

Goldin, Claudia. 2014. "A Grand Gender Convergence: Its Last Chapter." *The American Economic Review*, 104(4): 1091–1119.

Goldin, Claudia, and Joshua Mitchell. 2017. "The New Life Cycle of Women's Employment: Disappearing Humps, Sagging Middles, Expanding Tops." *Journal of Economic Perspectives*, 31(1): 161–82.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2018. "Bartik Instruments: What, When, Why, and How." National Bureau of Economic Research Working Paper 24408.

Gottschalk, Peter, and Robert Moffitt. 2009. "The Rising Instability of U.S. Earnings." *Journal of Economic Perspectives*, 23(4): 3–24.

Greenwood, Jeremy, Nezih Guner, Georgi Kocharkov, and Cezar Santos. 2016. "Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment, and Married Female Labor-Force Participation." *American Economic Journal: Macroeconomics*, 8(1): 1–41.

Gronau, Reuben. 1977. "Leisure, Home Production, and Work–the Theory of the Allocation of Time Revisited." *Journal of Political Economy*, 85(6): 1099–1123.

Grossman, Michael. 1972. "On the Concept of Health Capital and the Demand for Health." *Journal of Political Economy*, 80(2): 223–255.

Gruber, Jonathan. 2004. "Is Making Divorce Easier Bad for Children? The Long-Run Implications of Unilateral Divorce." *Journal of Labor Economics*, 22(4): 799–833.

Hansen, Lars Peter. 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica*, 50(4): 1029–1054.

Harris, Matthew C., Lawrence M. Kessler, Matthew N. Murray, and Beth Glenn. 2019. "Prescription Opioids and Labor Market Pains: The Effect of Schedule II Opioids on Labor Force Participation and Unemployment." *Journal of Human Resources*.

Hederos, Karin, and Anders Stenberg. 2019. "Gender Identity and Relative Income within Households: Evidence from Sweden." Swedish Institute for Social Research Working Paper 3. accessed at: http://www.diva-portal.org/smash/get/diva2:1332276/FULLTEXT01.pdf.

Hotchkiss, Julie L., and M. Melinda Pitts. 2007. "The Role of Labor Market Intermittency in Explaining Gender Wage Differentials." *American Economic Review: Papers and Proceedings*, 97(2): 417–421.

Imai, Susumu, and Michael P. Keane. 2004. "Intertemporal Labor Supply and Human Capital Accumulation." *International Economic Review*, 45(2): 601–641.

Jaimovich, Nir, and Henry E Siu. 2018. "Job Polarization and Jobless Recoveries." National Bureau of Economic Research Working Paper 18334.

Juhn, Chinhui. 1992. "Decline of Male Labor Market Participation: The Role of Declining Market Opportunities." *The Quarterly Journal of Economics*, 107(1): 79–121.

Juhn, Chinhui, and Kevin M. Murphy. 1997. "Wage Inequality and Family Labor Supply." *Journal of Labor Economics*, 15(1): 72–97.

Juhn, Chinhui, Kevin M. Murphy, and Robert H. Topel. 1991. "Why has the Natural Rate of Unemployment Increased over Time?" *Brookings Papers on Economic Activity*, 1991(2): 75–142.

Kelso, Alexander S., and Vincent P. Crawford. 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica*, 50(6): 1483–1504.

Killingsworth, Mark R., and James J. Heckman. 1986. "Female Labor Supply: A Survey." In *Handbook of Labor Economics*. Vol. 1-2,, , ed. Orley Ashenfelter and Richard Layard.

Kleven, Henrik, Camille Landais, and Jakob Egholt Sogaard. 2019. "Children and Gender Inequality: Evidence from Denmark." *American Economic Journal: Applied Economics*, 11(4): 181–209.

Knowles, John A. 2012. "Why are Married Men Working so Much? An Aggregate Analysis of Intra-Household Bargaining and Labor Supply." *The Review of Economic Studies*, 80(3): 1055–1085.

Krueger, Alan B. 2017. "Where Have all the Workers Gone? An Inquiry into the Decline of the U.S. Labor Force Participation Rate." *Brookings Papers on Economic Activity*, 2017(Fall): 1–87.

Kubota, So. 2017. "Child Care Costs and Stagnating Female Labor Force Participation in the US." Princeton University dissertation.

Lafortune, Jeanne. 2013. "Making Yourself Attractive: Pre-marital Investments and the Returns to Education in the Marriage Market." *American Economic Journal: Applied Economics*, 5(2): 151–78.

Lafortune, Jeanne, and Corinne Low. 2017. "Tying the Double-Knot: The Role of Assets in Marriage Commitment." *American Economic Review*, 107(5): 163–67.

Lafortune, Jeanne, and Corinne Low. 2018. "Betting the House: The Role of Homeownership in Marital Commitment and Child Investments." https://delia-furtado.uconn.edu/wp-content/uploads/sites/1822/2019/03/Lafortune_house.pdf.

Lam, David. 1988. "Marriage Markets and Assortative Mating with Household Public Goods: Theoretical Results and Empirical Implications." *The Journal of Human Resources*, 23(4): 462–487.

Lee, Donghoon, and Kenneth I. Wolpin. 2006. "Intersectoral Labor Mobility and the Growth of the Service Sector." *Econometrica*, 74(1): 1–46.

Liebman, Jeffrey B. 2015. "Understanding the Increase in Disability Insurance Benefit Receipt in the United States." *Journal of Economic Perspectives*, 29(2): 123–50.

Lindo, Jason M., Jessamyn Schaller, and Benjamin Hansen. 2018. "Caution! Men *not* at Work: Gender-specific Labor Market Conditions and Child Maltreatment." *Journal of Public Economics*, 163: 77 – 98.

Low, Hamish, Costas Meghir, Luigi Pistaferri, and Alessandra Voena. 2018. "Marriage, Labor Supply and the Dynamics of the Social Safety Net." National Bureau of Economic Research Working Paper 24356.

Ludwig, Volker, and Josef Brüderl. 2018. "Is There a Male Marital Wage Premium? New Evidence from the United States." *American Sociological Review*, 83(4): 744–770.

Lundberg, Shelly, and Robert A. Pollak. 1994. "Noncooperative Bargaining Models of Marriage." *The American Economic Review*, 84(2): 132–137.

Lundberg, Shelly, and Robert A. Pollak. 1996. "Bargaining and Distribution in Marriage." *Journal of Economic Perspectives*, 10(4): 139–158.

Lundberg, Shelly, and Robert A. Pollak. 2003. “Efficiency in Marriage.” *Review of Economics of the Household*, 1(3): 153–167.

Lundberg, Shelly, Robert A. Pollak, and Jenna Stearns. 2016. “Family Inequality: Diverging Patterns in Marriage, Cohabitation, and Childbearing.” *Journal of Economic Perspectives*, 30(2): 79–102.

Manning, Wendy D., Susan L. Brown, and Bart Stykes. 2015. “Trends in Births to Single and Cohabiting Mothers, 1980-2013.” Family Profiles FP-15-03, National Center for Marriage & Family Research.

Mansour, Hani, and Terra McKinnish. 2014. “Who Marries Differently Aged Spouses? Ability, Education, Occupation, Earnings, and Appearance.” *The Review of Economics and Statistics*, 96(3): 577–580.

Mansour, Hani, and Terra McKinnish. 2018. “Same-Occupation Spouses: Preferences or Search Costs?” *Journal of Population Economics*, 31(4): 1005–1033.

Marantz, Sonia A., and Annick F. Mansfield. 1977. “Maternal Employment and the Development of Sex-Role Stereotyping in Five- to Eleven-Year-Old Girls.” *Child Development*, 48(2): 668–673.

Mazzocco, Maurizio. 2007. “Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment.” *The Review of Economic Studies*, 74(3): 857–895.

Mazzolari, Francesca, and Giuseppe Ragusa. 2013. “Spillovers from High-Skill Consumption to Low-Skill Labor Markets.” *The Review of Economics and Statistics*, 95(1): 74–86.

McCrary, Justin. 2008. “Manipulation of the Running variable in the Regression Discontinuity Design: A Density Test.” *Journal of Econometrics*, 142(2): 698 – 714. The regression discontinuity design: Theory and applications.

McHale, Susan M., Ann C. Crouter, and Corinna J. Tucker. 1999. “Family Context and Gender Role Socialization in Middle Childhood: Comparing Girls to Boys and Sisters to Brothers.” *Child Development*, 70(4): 990–1004.

McHale, Susan M., Ann C. Crouter, and Shawn D. Whiteman. 2003. “The Family Contexts of Gender Development in Childhood and Adolescence.” *Social Development*, 12(1): 125–148.

McKinnish, Terra G. 2007. “Sexually Integrated Workplaces and Divorce: Another Form of On-the-Job Search.” *Journal of Human Resources*, XLII(2): 331–352.

Mernitz, Sara E. 2018. “A Cohort Comparison of Trends in First Cohabitation Duration in the United States.” *Demographic Research*, 38(66): 2073–2086.

Mincer, Jacob. 1958. “Investment in Human Capital and Personal Income Distribution.” *Journal of Political Economy*, 66(4): 281–302.

Mincer, Jacob. 1962. “On-the-Job Training: Costs, Returns, and Some Implications.” *Journal of Political Economy*, 70(5, Part 2): 50–79.

Mincer, Jacob, and Haim Ofek. 1982. “Interrupted Work Careers: Depreciation and Restoration of Human Capital.” *The Journal of Human Resources*, 17(1): 3–24.

Moffitt, Robert A. 2012. “The Reversal of the Employment-Population Ratio in the 2000s: Facts and Explanations.” *Brookings Papers on Economic Activity*, Fall(2): 201–264.

Mueller-Smith, Michael. 2015. “The Criminal and Labor Market Impacts of Incarceration.” http://www.antoniocasella.eu/nume/Mueller-Smith_2014.pdf.

Murray-Close, Marta, and Misty L. Heggeness. 2018. “Manning Up and Womaning Down: How Husbands and Wives Report their Earnings when She Earns More.” U.S. Census Bureau Social, Economic and Housing Statistics Division Working Paper 2018-20.

Nöldeke, Georg, and Larry Samuelson. 2015. “Investment and Competitive Matching.” *Econometrica*, 83(3): 835–896.

Olivetti, Claudia, Eleonora Patacchini, and Yves Zenou. 2016. “Mothers, Peers and Gender Identity.” Boston College, unpublished working paper.

Oppenheimer, Valerie Kincade. 2003. “Cohabiting and Marriage during Young Men’s Career-Development Process.” *Demography*, 40(1): 127–149.

Oreopoulos, Philip, Marianne Page, and Ann Huff Stevens. 2008. “The Intergenerational Effects of Worker Displacement.” *Journal of Labor Economics*, 26(3): 455–000.

Page, Marianne, Jessamyn Schaller, and David Simon. 2019. “The Effects of Aggregate and Gender-Specific Labor Demand Shocks on Child Health.” *Journal of Human Resources*, 54(1): 37–78.

Pager, Devah. 2007. *Marked: Race, Crime, and Finding Work in an Era of Mass Incarceration.* The University of Chicago Press.

Pencavel, John. 1986. “Chapter 1: Labor supply of men: A survey.” In *Handbook of Labor Economics*. Vol. 1, 3 – 102. Elsevier.

Peters, Michael, and Aloysius Siow. 2002. “Competing Premarital Investments.” *Journal of Political Economy*, 110(3): 592–608.

Pollak, Robert A. 2019. “How Bargaining in Marriage Drives Marriage Market Equilibrium.” *Journal of Labor Economics*, 37(1): 297–321.

Raley, Sara B., Marybeth J. Mattingly, and Suzanne M. Bianchi. 2006. “How Dual Are Dual-Income Couples? Documenting Change From 1970 to 2001.” *Journal of Marriage and Family*, 68(1): 11–28.

Ramey, Gary, and Valerie A. Ramey. 2010. “The Rug Rat Race.” *Brookings Papers on Economic Activity*, Spring(1): 129–199.

Rasul, Imran. 2003. “The Impact of Divorce Laws on Marriage.” Accessed at: https://www.researchgate.net/profile/Imran_Rasul/publication/228558333_The_Impact_of_Divorce_Laws_on_Marriage/links/0deec522ff75c5aac4000000.pdf.

Rendall, Michael S., and Rachel M. Shattuck. 2019. “First Birth before First Stable Employment and Subsequent Single-Mother “Disconnection” before and after the Welfare Reform and Great Recession.” *Journal of Poverty*, 23(2): 83–104.

Reynoso, Ana. 2017. “Marriage, Marital Investments, and Divorce: Theory and Evidence on Policy Non Neutrality.” Accessed at: https://sites.lsa.umich.edu/ana-reynoso/wp-content/uploads/sites/704/2019/05/AR_MIDth_latest.pdf.

Reynoso, Ana. 2019. “The Impact of Divorce Laws on the Equilibrium in the Marriage Market.” Accessed at: http://www.anamreynoso.com/assets/AR_JMP_latest.pdf?no_cache=0.4412435197042406.

Roth, Alvin E., and Marilda Sotomayor. 1992. “Chapter 16: Two-Sided Matching.” In *Handbook of Game Theory with Economic Applications*. Vol. 1, 485 – 541. Elsevier.

Ruggles, Steven. 2015. “Patriarchy, Power, and Pay: The Transformation of American Families, 1800-2015.” *Demography*, 52(6): 1797–1823.

Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, , and Matthew Sobek. 2015. Integrated Public Use Microdata Series: Version 6.0. [dataset]. Minneapolis: University of Minnesota. http://doi.org/10.18128/D010.V6.0.

Schaller, Jessamyn. 2016. “Booms, Busts, and Fertility: Testing the Becker Model Using Gender-Specific Labor Demand.” *Journal of Human Resources*, 51(1): 1–29.

Schneider, Daniel, Kristen Harknett, and Matthew Stimpson. 2018. “What Explains the Decline in First Marriage in the United States? Evidence from the Panel Study of Income Dynamics, 1969 to 2013.” *Journal of Marriage and Family*, 80(4): 791–811.

Schwartz, Christine R., and Pilar Gonalons-Pons. 2016. “Trends in Relative Earnings and Marital Dissolution: Are Wives Who Outearn Their Husbands Still More Likely to Divorce?” *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2(4): 218–236.

Shapley, L. S., and M. Shubik. 1971. “The Assignment Game I: The Core.” *International Journal of Game Theory*, 1(1): 111–130.

Smock, Pamela J., and Christine R. Schwartz. 2020. “The Demography of Families: A Review of Patterns and Change.” *Journal of Marriage and Family*, 82(1): 9–34.

Sohn, Kitae. 2015. “The Male-Taller Norm: Lack of Evidence from a Developing Country.” *Journal of Comparative Human Biology*, 66(4): 369 – 378.

Solon, Gary. 1999. "Chapter 29: Intergenerational Mobility in the Labor Market." In . Vol. 3 of *Handbook of Labor Economics*, , ed. Orley C. Ashenfelter and David Card, 1761 – 1800. Elsevier.

Stevens, Ann Huff, and Jessamyn Schaller. 2011. "Short-Run Effects of Parental Job Loss on Children's Academic Achievement." *Economics of Education Review*, 30(2): 289 – 299.

Stevenson, Betsey. 2007. "The Impact of Divorce Laws on Marriage-Specific Capital." *Journal of Labor Economics*, 25(1): 75–94.

Stevenson, Betsey, and Justin Wolfers. 2007. "Marriage and Divorce: Changes and Their Driving Forces." *The Journal of Economic Perspectives*, 21(2): 27–52.

Stryker, Sheldon, and Peter J. Burke. 2000. "The Past, Present, and Future of an Identity Theory." *Social Psychology Quarterly*, 63(4): 284–297.

Stulp, Gert, Abraham P. Bunk, Thomas V. Pollet, Daniel Nettle, and Simon Verhulst. 2013. "Are Human Mating Preferences with Respect to Height Reflected in Actual Pairings?" *PLoS One*, 8(1): 1–7.

Tolbert, Charles M., and Molly Sizer. 1996. "U.S. Commuting Zones and Labor Market Areas: A 1990 Update." Available: http://ageconsearch.umn.edu/record/278812.

Topel, Robert. 1991. "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority." *Journal of Political Economy*, 99(1): 145–176.

Voena, Alessandra. 2015. "Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?" *American Economic Review*, 105(8): 2295–2332.

Wieber, Anna, and Elke Holst. 2015. "Gender Identity and Women's Supply of Labor and Non-Market Work: Panel Data Evidence from Germany." 1517. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2688965#.

Wilson, William J. 1996. *When Work Disappears: The World of the New Urban Poor.* Alfred A. Knopf.

Winkler, Anne E. 1998. "Earnings of Husbands and Wives in Dual-Earner Families." *Monthly Labor Review*, 121(4): 42–48.

Winkler, Anne E., Timothy D. McBride, and Courtney Andrews. 2005. "Wives who Outearn Their Husbands: A Transitory or Persistent Phenomenon for Couples?" *Demography*, 42(3): 523–535.

Wolfers, Justin. 2006. "Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results." *American Economic Review*, 96(5): 1802–1820.

Zinovyeva, Natalia, and Maryna Tverdostup. 2018. "Gender Identity, Co-Working Spouses and Relative Income within Households." 11757. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3249871.