

SKILL PREMIA AND PRE-MARITAL INVESTMENTS IN MARRIAGE MARKETS *

Aditya Kuvalekar

June 2026

ABSTRACT

I study a decentralized marriage market with search frictions, costly pre-marital skill investments, and non-transferable utility. Despite a fully symmetric environment, asymmetric equilibria—in which one gender systematically invests more in skills than the other—can arise. The match payoffs are microfounded through a non-cooperative household game in which spouses allocate time between labor-market work and domestic production. An asymmetric equilibrium becomes available once the high-skill wage rises past a threshold. Further, the symmetric equilibria can be fragile while the asymmetric ones are not. Thus, rising skill premia may amplify rather than narrow gender gaps in skill acquisition.

Keywords: marriage market, pre-marital investment, search frictions, skill premium, non-transferable utility

JEL Codes: C78, D83, J12, J16, J24

*Kuvalekar: University of Essex and Indian School of Business, email: a.kuvalekar@essex.ac.uk. I am grateful to Nageeb Ali, Deepal Basak, Joyee Deb, Maitreesh Ghatak, Hugo Hopenhayn, Elliot Lipnowski, Dileep Mookherjee, Andrew Newman, Sabareesh Ramachandran, Debraj Ray, Anna Sanktjohanser, and participants in ThReD 2025 for helpful comments. Refine.ink was used to check the paper for consistency and clarity.

1. Introduction

The wages commanded by highly skilled workers have been rising rapidly (Acemoglu and Autor, 2011; Goldin and Katz, 2008). How are people’s investments in skills affected by rising skill premia when their choice of skills has implications for both labor- and marriage-market outcomes? This is the central question of this paper.

From a labor-market standpoint, the implication seems clear: a higher return to skill should lead more people to incur the cost of skill acquisition. However, the answer is no longer so clear-cut when marriage-market considerations enter the picture.¹ When the members of a household divide their time, one highly skilled person can focus on market work, the other on domestic production. Thus, an agent anticipating matching with a skilled partner may be discouraged from incurring the cost of investing. This within-household specialization can occur along gender lines, causing a widening gender gap in skills.

In a canonical search-and-matching model with non-transferable utility, I show that gender disparities in skill acquisition can arise endogenously *even when no primitive asymmetry exists between the genders*. These asymmetries arise precisely when the labor-market wage for high-skilled workers exceeds a threshold. Below the threshold, the equilibrium is symmetric: both genders invest at equal rates. Above the threshold, an asymmetric equilibrium emerges alongside the symmetric one. In any asymmetric equilibrium, one gender invests significantly more in skill acquisition than the other, who specializes in domestic production. Finally, in a class of markets, whenever an asymmetric equilibrium coexists with the symmetric one, the symmetric equilibrium is *fragile* in the sense of Onuchic and Ray (2023) while the asymmetric one is not: small and asymmetric perturbations in investments get amplified by agents’ best responses to these perturbations in symmetric equilibria but die out in asymmetric ones.

The main contribution of this paper is to serve as a cautionary tale arising from the existence and robustness of asymmetric equilibria: the gender gap in skill acquisition need not close as opportunities expand. In particular, a large shock to the return to skill—of which AI is the most discussed contemporary example—can *widen* it, even though the shock is gender-neutral and the environment contains no structural discrimination.

I now briefly describe the framework. A continuum of agents from two sides, men and women, engage in undirected search to find a potential match. Prior to beginning their search, each decides whether to incur a cost to acquire high skills or to remain low skilled. An agent’s cost depends on their type, indexed

¹That pre-marital investments respond to marriage-market returns, not only labor-market ones, is well documented: see Lafortune (2013) and Chiappori et al. (2009); Goldin (2006) documents how rising labor-market opportunities for women reshaped their educational investments over the twentieth century.

by their identity in the unit interval; agents with higher types incur higher costs to acquire high skills. After making this investment choice, they begin the search phase, meeting partners from the other side randomly. When two people meet, they observe each other’s skills and decide whether to match. An individual’s payoff depends on their own skill and their partner’s skill, not on latent types. If two agents match, they immediately leave the market and are replaced by newly born agents of the same types who make fresh investment decisions before entering the unmatched pool, ensuring that the economy is perpetually in steady state.

Before describing the results, I offer some interpretations of skill in the model. The most literal reading is education, the schooling an agent acquires before entering the marriage market. An alternative interpretation could be in terms of the career one pursues after basic education. The high-skill choice can represent a demanding, competitive, time-intensive career that pays well; the low-skill choice, a less onerous occupation with shorter hours and lower pay.

I assume that everyone prefers a high-skilled partner to a low-skilled one. Therefore, negative assortative matching cannot arise, and pure-strategy equilibria fall into three classes: positive assortative matching (PAM), where high-skilled agents accept only high-skilled partners; all-match (AM), where everyone accepts everyone; and semi-assortative (SA), where high-skilled agents accept all partners but low-skilled agents reject low-skilled ones. What, then, could make two *ex ante* identical sides of the market invest differently? Not PAM: any PAM equilibrium has both genders investing at the same rate (Proposition 7, in the appendix). If asymmetry is to arise, it must come from the AM or SA classes.

And indeed, as Proposition 1 shows, if the match payoff function is submodular—an assumption maintained throughout the paper—then asymmetric equilibria can exist within the AM class.² A symmetric equilibrium always exists, of AM or SA form. Alongside it the economy may also admit an *asymmetric* equilibrium, in which one gender invests systematically more than the other; any such equilibrium takes one of two forms—*full-investment-from-one-side* (FIOS), where one gender invests fully and the other only below a cutoff type, or *no-investment-from-one-side* (NIOS), where one gender does not invest at all—and the two never coexist.

These results show asymmetric equilibria are possible; a reader may still suspect they are a knife-edge curiosity. They are not. Whether one exists depends on the match payoff function—the four payoffs, one for each own/partner skill pair, and their relative magnitudes—together with the search frictions and the investment cost. This makes the question an empirical one: Choo and Siow (2006), Galichon and Salanié (2022), and others have developed methods to identify match payoffs from observed matching patterns, and the friction and cost parameters are themselves standard objects in search-and-matching models, so the conditions can, in

²Asymmetric SA equilibria can exist, and can even be *uniquely robust*, under supermodular payoffs; their existence and robustness do not require submodularity (Appendix C works out an example). I nonetheless maintain submodularity throughout, because it is what drives the comparative statics in the skill premium that follow.

principle, be checked in data. Proposition 2 then shows asymmetric equilibria are far from exceptional: from any economy whose only equilibrium is symmetric, raising the household’s return to having a high-skilled member—the skill premium, read at the level of match payoffs—brings an asymmetric (FIOS) equilibrium into being alongside the symmetric one.

One question remains. Wherever an asymmetric equilibrium exists, a symmetric one exists beside it; is there any reason to expect the asymmetric outcome to be the one we see? Section 5 answers it through a recent refinement—*fragility*³—introduced by Onuchic and Ray (2023) in their study of gender discrimination in collaborative projects. An equilibrium is fragile if a slight upward perturbation in one gender’s investment, and downward in the other’s, is amplified rather than undone by a single round of best-response play. Proposition 4 shows that whenever an asymmetric AM equilibrium coexists with the symmetric one, it is the symmetric equilibrium that is fragile, while the asymmetric one is not. The symmetric outcome thus survives only in a world that is exactly symmetric: the small gender-correlated asymmetries any real market carries are amplified away from it but absorbed by the asymmetric equilibria. In short, rather than being pathological, the asymmetric equilibria can be the only robust equilibria in some environments.

The robustness of these asymmetric equilibria leads to the natural follow-up question: what changes in the underlying economy give rise to the changes in match payoffs that, in turn, induce asymmetric equilibria? Section 6 provides one answer. I view the match payoff function as the equilibrium outcome of a non-cooperative household game along the lines of Chiappori (1988), in which spouses allocate time between labor-market work and household production.⁴ Once match payoffs are themselves equilibrium objects rather than primitives, comparative statics can be performed on more fundamental objects of the economy—in particular, the labor-market wages paid to high- and low-skilled workers. The skill premium then has a concrete interpretation: the difference between the high-skill and low-skill wages. Proposition 5 identifies a single threshold for the high-skill wage: below the threshold, where the payoffs are not yet submodular enough to admit an asymmetric equilibrium, the only equilibrium is symmetric, and the investment cutoff rises continuously with the premium as both genders invest more in skill acquisition; above the threshold, an asymmetric (FIOS) equilibrium becomes available—with one gender fully investing and the other investing substantially less—alongside the symmetric equilibrium.

In summary, the model delivers testable predictions. First, in the low-premium

³While I adopt fragility as the refinement, the idea behind it—that an equilibrium where small perturbations in strategies get amplified by agents’ best responses should be viewed as unstable—is not exotic; related ideas appear in the refinement of equilibria in finite games (e.g., Ritzberger and Weibull, 1995).

⁴See Chiappori (2020) for a survey of various models of household labor supply decisions as games.

phase (below the threshold t^*), both genders invest at *rising* and equal rates. Second, once the premium crosses t^* , an asymmetric equilibrium becomes available alongside the symmetric one: one gender invests considerably more in skill acquisition while the other specializes in domestic production. Under the fragility refinement of Section 5, the only robust outcome is the asymmetric equilibrium. Separately, at high premia the symmetric equilibrium can itself turn selective: low-skilled agents reject one another and match only with high-skilled partners, so low–low marriages disappear. Thus moderate increases in the skill premium raise both genders’ investment equally, whereas a sufficiently large increase can lead one gender to lower its investment while the other raises its own significantly.

The rest of the paper is organized as follows. Section 2 discusses the related literature, and Section 3 lays out the model. Section 4 analyzes the post-investment matching game and the investment stage, establishing when asymmetric equilibria arise and how they respond to the skill premium. Section 5 develops the fragility refinement and shows that the symmetric equilibrium is fragile while the asymmetric ones are not. Section 6 microfound the match payoffs through a household game and recasts the comparative statics in terms of labor-market wages. Section 7 concludes.

2. Related Literature

Much of the extant literature on gender disparities in skill acquisition attributes them to primitive asymmetries between men and women. The present paper complements that line of work by showing that rising skill premia alone, in an otherwise symmetric environment, can generate the same kind of disparities.⁵

Naturally, this paper is related to the large literature on matching markets. There are two broad directions in which this literature has evolved. The initial papers, such as those of Becker (1973); Gale and Shapley (1962); Shapley and Shubik (1971), assumed frictionless matching markets. Somewhat later, Burdett and Coles (1997); Shimer and Smith (2000, 2001); Smith (2006) investigated these markets with search frictions. Across these two frameworks, another important distinction is between transferable utility (such as Becker (1973); Shapley and Shubik (1971)) versus non-transferable utility (such as Gale and Shapley (1962); Smith (2006)). In frictionless NTU environments, Legros and Newman (2007) characterize how the direction of assortative matching depends on the payoff structure in ways that diverge sharply from the TU case; the submodularity condition driving asymmetric equilibria in my paper mirrors their conditions for negative assortative matching. Eeckhout (1999) studies bilateral search with vertical heterogeneity

⁵I do not wish to suggest that structural asymmetries do not exist. My main focus, however, is on the comparative static—how pre-marital investments evolve with changing skill premia. It is *a priori* not clear whether rising skill premia exacerbate these asymmetries or mitigate them. As Proposition 5 shows, even in the absence of structural asymmetries, rising skill premia can give rise to asymmetries that would otherwise not be present.

under NTU, deriving sorting conditions that inform the post-investment matching stage of my model. The post-investment stage of my game can be viewed as a decentralized version of [Gale and Shapley \(1962\)](#) which was studied by [Adachi \(2003\)](#); [Smith \(2006\)](#). To this setting, I add an investment choice before the search phase.

A large literature has studied costly investments in marriage markets, (see [Bhaskar and Hopkins, 2016](#); [Bhaskar et al., 2023](#); [Chade and Lindenlaub, 2022](#); [Cole et al., 2001](#); [Mailath et al., 2013](#); [Nöldeke and Samuelson, 2015](#); [Peters, 2007](#), for instance). Most of these papers study the efficiency of such investments. With the exception of [Chade and Lindenlaub \(2022\)](#) and [Bhaskar and Hopkins \(2016\)](#), most other papers assume that the returns are deterministic, while these two study various issues related to how riskiness of investments affects the match outcomes as well as investments themselves. A seminal paper by [Peters and Siow \(2002\)](#) showed that, in a frictionless transferable-utility marriage market, non-cooperative premarital investments are nonetheless constrained efficient, as competition for partners provides the right investment incentives. I show that a combination of NTU and search frictions generates gender-specific asymmetries in such environments. Relatedly, [Bhaskar et al. \(2023\)](#) highlight how asymmetries in investments can occur when the two genders differ in their bargaining power. Also in a frictionless NTU environment, [Iyigun and Walsh \(2007\)](#) show how spousal allocations within the household feed back into premarital investment incentives—a mechanism my paper embeds within a search-and-matching framework.

Recently, [Atakan et al. \(2026\)](#) study a marriage market with search frictions, costly investments, and *transferable utility*. Their main finding is that the outcomes are constrained efficient and assortative in a wide variety of settings. In contrast, adopting the NTU framework, I sidestep the question of efficiency and focus solely on the disparities in skill acquisition that arise in equilibrium.

Given my emphasis on gender differences in skill acquisition, this paper is naturally related to work that highlights these issues in various contexts, e.g. [Chiappori et al. \(2009, 2017\)](#); [Low \(2024\)](#); [Siow \(1998\)](#); [Zhang \(2021\)](#). Often, unlike the current paper, these papers assume some innate differences—in preferences, fertility costs, or the technological environment—that give rise to these disparities. An exception is [Hadfield \(1999\)](#), who, like the present paper, derives gender-asymmetric specialization from symmetric primitives. Her framework, however, contains no labor market, and hence no analog of the comparative static in the skill premium that drives this paper. In a similar spirit, [Francois \(1998\)](#) is one of the early references to show how discrimination in the labor market can arise despite a fully symmetric environment, with the asymmetry emerging from the interaction between men and women within the household. [Lafortune \(2013\)](#) provides empirical support for the premise underlying this paper—that pre-marital investments respond to marriage-market returns to education, not only to labor-market returns—using variation in marriage-market competition across cohorts. On the macroeconomic side, [Greenwood et al. \(2014\)](#) document that rising assortative

mating has amplified income inequality, a consequence my mechanism suggests can be driven partly by endogenous investment asymmetries rather than match patterns alone.

3. Model

The economy consists of a unit mass of men and a unit mass of women, each indexed by a type $x \in [0, 1]$. We denote an agent’s gender by $\gamma \in \{m, w\}$, and write γ' for the opposite gender. Upon “birth,” agents choose a skill level from a finite set \mathcal{S} , which is assumed to be binary: $\mathcal{S} = \{H, L\}$, where H represents high skill and L represents low skill. The cost of acquiring skill s for an agent of type x is $C(x, s)$. Normalizing $C(x, L) = 0$, we denote $C(x, H)$ simply as $C(x)$. Throughout the paper, we adopt an affine cost specification: $C(x) = x + c$ for some $c > 0$.⁶ Agents with higher x pay more to acquire H .

Although we label the two skill levels “high” and “low,” the relevant distinction need not be educational attainment as such. The high-skill option may equally be read as entry into a demanding, high-pressure, well-remunerated career track and the low-skill option as a less time-intensive, lower-paid one; under this reading an agent’s investment decision is a choice of how ambitious a career to pursue, and the asymmetric equilibria below describe households in which one partner pursues such a career while the other does not.

If a man with skill s and a woman with skill s' match, the man’s payoff is $\phi(s, s')$, and the woman’s is $\phi(s', s)$. That is, the first argument represents an individual’s own skill and the second, their partner’s. This is a non-transferable utility (NTU) setting where payoffs depend only on skill levels, not on underlying type. Agents discount future payoffs at rate $r > 0$.

We impose the following standard ranking of payoffs:

ASSUMPTION 1: $\phi(H, H) \geq \phi(H, L) \geq \phi(L, H) \geq \phi(L, L)$.

This assumption can be microfounded through a household decision-making problem as in [Chiappori \(1988\)](#). See [Section 6](#) for an example.

I also restrict attention to equilibria in which every agent is matched eventually. Within the equilibrium classes we will consider—positive assortative matching (PAM), in which high types match only with high and low types only with low; all-match (AM), in which everyone accepts everyone; and semi-assortative (SA), in which high types accept all but low types reject low partners (all three are formally defined in [Section 3.2](#))—this rules out configurations in which some agents on one side of the market are systematically rejected by the other and remain permanently unmatched. This restriction is substantive rather than merely

⁶The possibility of asymmetric equilibria remains under general strictly increasing, strictly convex cost functions; the affine specification is adopted primarily for clean expressions that facilitate comparative statics.

definitional: it sets aside coordination-driven equilibria in which one side is collectively unwilling to match with the other’s low-skilled agents, leaving them permanently unmatched—outcomes sustained by self-fulfilling beliefs rather than by the investment incentives that are this paper’s focus.

An economy is a tuple $\langle \phi, C, \lambda, r \rangle$ consisting of the payoff functions, the cost functions for acquiring different skill levels, and the arrival rate of potential partners that we will describe momentarily.

Timing: Time is continuous and infinite. At each instant, unmatched agents engage in costless, undirected search, observing the skill distribution of the opposite pool. When two agents meet, they observe each other’s skills and decide whether to match. If they do, they leave the market and are replaced by newly born agents of the same latent type x and the same gender as those who departed. (This assumption is without loss for our purposes; see Section 3.1 for a discussion of alternative replacement rules.) These newly born agents make an irreversible skill investment decision before entering the unmatched pool.

Thus, the economy is perpetually in “steady state”, i.e., the mass of each set of agents, men and women, remains to be 1 and their type distribution remains unchanged. Unmatched agents meet others drawn at random from the unmatched pool at an exponential rate λ .⁷

Strategies: We restrict attention to stationary equilibria (where strategies of any type do not vary with time). A strategy for a type x and gender γ specifies a choice of skill to acquire, and then, subsequently, an acceptance set during the search phase. This specifies the set of agents one is willing to match with should they meet them. It is without loss to restrict attention to pure strategies insofar as investment is concerned. We denote by $I^\gamma : [0, 1] \rightarrow \{0, 1\}$ the investment strategy of a type x agent of gender γ , where one indicates their decision to acquire high skills and zero does not. We assume that I^γ is measurable. Given players’ investment strategies, let $\zeta_I^\gamma := \int I^\gamma(x) dx$ denote the proportion of H types in group γ . Since the dependence on I is obvious in most cases, we will omit it. Post the agents’ investment decisions, they engage in an undirected search. Whenever they meet someone, they have to decide whether to match with that person or not. If both the agents agree to a match, then the match takes place. The acceptance strategy specifies the probability of accepting an agent of each skill type and depends on one’s own skill as well as the proportion of high types in the opposite pool. First, given Assumption 1, no agent would ever turn down a high type agent. Therefore, we merely need to specify the probability of accepting a low type for each gender and skill type. We denote by $\alpha^\gamma : [0, 1] \times \{H, L\} \rightarrow [0, 1]$ the probability that an agent of gender γ and skill s accepts a low type agent of the opposite gender given the proportion of the high types from group γ' . Often, this dependence on $\zeta^{\gamma'}$ is obvious and, therefore, we will suppress it.

⁷Since the economy is perpetually in the steady state with a unit mass of population on either side, we do not need to scale it with the population size.

3.1. Discussion of the model

The model has three key features: search frictions, costly pre-marital skill investments, and non-transferable utility. Search frictions and costly pre-marital investments are particularly pervasive in real marriage markets. Individuals—or their parents, in some contexts—make pre-marital investments anticipating labor-market outcomes that in turn shape their marriage prospects. Meanwhile, the prevalence of dating apps and marriage bureaus is clear evidence of search frictions in modern marriage markets.

The third feature, non-transferable utility (NTU), is motivated by the nature of social matches such as marriages. The seminal paper of [Becker \(1973\)](#) briefly considered such an environment, and [Smith \(2006\)](#) studied a marriage market with search frictions and NTU at length. To quote the latter:

But in defense of the NTU model, disagreements about matching in social settings are not uncommon. And whenever we observe a potential match or split desired by one party but not the other, utility is obviously not fully transferable. For the total match surplus either is positive or is not, and there can be no disagreement.

A particularly striking case of such a disagreement is the sharp discontinuity around $\frac{1}{2}$ in the share of household income earned by the wife documented by [Bertrand et al. \(2015\)](#). In the context of US couples, they show considerable aversion to situations in which the wife earns more than her husband. A world with fully transferable utility would show no such discontinuity: whatever psychological cost the husband incurs from his wife earning more could, in principle, be compensated through transfers.

Further, I have assumed that the flow payoff while unmatched is zero regardless of skill. This is a simplifying assumption. The qualitative results carry over to a setting in which an unmatched agent of skill s earns a flow payoff $u(s) \geq 0$ with $u(H) \geq u(L)$, provided $u(\cdot)$ and the gap $u(H) - u(L)$ are sufficiently small relative to the match payoffs $\phi(\cdot, \cdot)$. Intuitively, when the value of remaining unmatched is dwarfed by the value of marrying, the marriage market remains the dominant force shaping investment and acceptance decisions, and the structure of equilibria—including symmetry-breaking under submodularity and the comparative static on the skill premium—is preserved.

Let me briefly comment on the replacement rule. When two agents match and leave the market, they are replaced by newly born agents of the *same latent type and same gender as those who just departed*: a departing type- x man is replaced by a newly born type- x man, and a departing type- x woman by a newly born type- x woman. The replacement agents then make fresh investment decisions before entering the unmatched pool. In any equilibrium, since agents of the same type and gender make identical investment choices, the replacements replicate the skill composition of the departing pair exactly. As a result, the type and skill

composition of the unmatched pool is held fixed at the equilibrium investment shares throughout time, and the model remains in a stationary steady state without our having to track pool dynamics explicitly. This technology of replacement is adopted mainly for tractability: in every all-match (AM) equilibrium—one in which every agent accepts every partner they meet on the equilibrium path (formally defined in Section 3.2)—every agent matches at rate λ regardless of skill, so the steady-state pool composition would coincide with the equilibrium investment share under any reasonable alternative replacement rule. The AM class covers most of the results in Sections 4–6.

3.2. Equilibrium

Agents choose their skill levels optimally given their expected value minus the cost. Let $V^\gamma(s, \alpha^\gamma; \zeta^{\gamma'}, \alpha^{\gamma'})$ denote the expected discounted value of an unmatched agent of gender γ and skill s , given the high-skill share $\zeta^{\gamma'}$ in the opposite-gender pool and the acceptance strategies $\alpha^\gamma, \alpha^{\gamma'}$ of the two sides. These values solve standard steady-state recursive equations; the expressions are straightforward to compute, so we omit them. An agent of skill s and gender γ chooses α^γ to maximize $V^\gamma(s, \alpha^\gamma; \zeta^{\gamma'}, \alpha^{\gamma'})$ given $(\zeta^{\gamma'}, \alpha^{\gamma'})$.

An **equilibrium** consists of the investment and acceptance strategies for each gender and skill type, $\langle (I^\gamma, \alpha^\gamma)_{\gamma \in \{m, w\}} \rangle$, satisfying **(IC-invest)** and **(IC-AL)** below.

$$V^\gamma(H, \alpha^\gamma; \zeta^{\gamma'}, \alpha^{\gamma'}) - C(x) - V^\gamma(L, \alpha^\gamma; \zeta^{\gamma'}, \alpha^{\gamma'}) \begin{cases} > 0 \implies I^\gamma(x) = 1 \\ = 0 \implies I^\gamma(x) \in \{0, 1\} \\ < 0 \implies I^\gamma(x) = 0 \end{cases} \quad \text{(IC-invest)}$$

$$V^\gamma(s, \alpha^\gamma; \zeta^{\gamma'}, \alpha^{\gamma'}) \begin{cases} < \phi(s, L) \implies \alpha^\gamma(s) = 1 \\ = \phi(s, L) \implies \alpha^\gamma(s) \in [0, 1] \\ > \phi(s, L) \implies \alpha^\gamma(s) = 0 \end{cases} \quad \text{(IC-AL)}$$

(IC-invest) ensures that the agents' investment decisions are optimal, while **(IC-AL)** specifies that agents acceptance decisions are optimal. In particular, if the value to agent is higher than what (s)he would receive by matching with a low-skilled person, then (s)he must not be accepting the low-skilled person. Because the value functions do not depend on an agent's own type while the cost $C(x) = x + c$ is strictly increasing in it, the net gain to investing is strictly decreasing in x ; optimal investment is therefore a cutoff rule. In any equilibrium, the investment strategies are characterized by a cutoff type, \hat{x}_γ for each γ , such that an agent i of gender γ and type x_i invests if and only if $x_i \leq \hat{x}_\gamma$.

NOTATION 1: *In light of the above, we will denote an equilibrium by (\hat{X}, α) where*

each, $\hat{X} = (\hat{x}_m, \hat{x}_w)$ and $\alpha = (\alpha^\gamma(H), \alpha^\gamma(L))_{\gamma \in \{m, w\}}$ are vectors denoting the relevant strategies. In general, X denotes a vector (x_m, x_w) .

We turn now to the structure of sorting in equilibrium and, more importantly, the investments agents make. In light of Assumption 1, no agent will ever reject a high skilled person. Therefore, we can rule out the existence of a negatively assortative match. Thus, insofar as pure strategy equilibria are concerned, we are left with the following three possibilities:

1. All match (AM): All agents accept all partners.
2. Positive Assortative Matching (PAM): High types match only with high types, and low types only with low types.
3. Semi-assortative matching (SA): High types accept all partners, but low types reject low-skill partners and match only with high types.

A fourth possible candidate—in which one side of the market rejects all low-skill partners from the other side, leaving those low-skill agents permanently unmatched—is excluded by the restriction to equilibria in which every agent is matched eventually.

The bulk of our analysis focuses on AM and SA equilibria; PAM is analyzed separately in Appendix B. Both classes of symmetric equilibrium (AM and SA) play roles in the analysis below: which one obtains depends on whether low-skilled agents prefer to accept or reject low-skilled partners at the equilibrium cutoff. PAM, by contrast, requires the marriage market to be sufficiently thick—high types must be willing to wait through encounters with low-type partners for a high-type match—and its detailed analysis is deferred to the appendix.

4. Analysis

The analysis throughout assumes submodular payoffs ($\Delta \leq 0$). The paper’s comparative statics live in the all-match (AM) class—equilibria in which agents accept the first partner they meet—and an AM equilibrium can be asymmetric *only* under submodularity: when payoffs are supermodular ($\Delta \geq 0$), every AM equilibrium is symmetric (Proposition 6 in the appendix).⁸ Recall from Section 3.2 that pure-strategy equilibria fall into three classes: AM (all accept all), PAM (high accept only high, so low match only with low), and SA (high accept all, low reject low). PAM is discussed in Appendix B; here we analyze the AM and SA classes that drive the paper’s substantive results.

⁸Outside the AM class, asymmetric equilibria *can* exist even under supermodular payoffs—of the semi-assortative (SA) type, driven by a search channel (a low type’s pool-dependent hold-out value) rather than by household complementarity—and can themselves be *uniquely robust*. Existence and robustness therefore do not hinge on submodularity; Appendix C gives an example.

AM equilibria

In an AM equilibrium with cutoffs $\hat{X} = (\hat{x}_m, \hat{x}_w)$, every agent accepts every partner, so $\alpha^\gamma(\cdot) = 1$ for all γ . The players' values are:

$$V_A^\gamma(H; \hat{X}) = \frac{\lambda[\hat{x}_{\gamma'}\phi(H, H) + (1 - \hat{x}_{\gamma'})\phi(H, L)]}{r + \lambda}, \quad \text{and}$$

$$V_A^\gamma(L; \hat{X}) = \frac{\lambda[\hat{x}_{\gamma'}\phi(L, H) + (1 - \hat{x}_{\gamma'})\phi(L, L)]}{r + \lambda}.$$

Ignoring trivial cases where either everybody invests or nobody does, and assuming at least one of \hat{x}_m, \hat{x}_w is interior, the IC constraints are:

$$V_A^\gamma(H; \hat{X}) - C(\hat{x}_\gamma) = V_A^\gamma(L; \hat{X}) \quad \text{if } \hat{x}_\gamma \in (0, 1) \quad (\text{AM: Invest})$$

$$V_A^\gamma(H; \hat{X}) \leq \phi(H, L) \quad (\text{AM: H-L})$$

$$V_A^\gamma(L; \hat{X}) \leq \phi(L, L) \quad (\text{AM: L-L})$$

Here, **(AM: Invest)** says the cutoff type \hat{x}_γ is indifferent between investing and not. **(AM: H-L)** says that a high-skilled agent's continuation value of remaining in the search pool does not exceed the payoff $\phi(H, L)$ from accepting a low-skilled partner—so a high-skilled agent prefers to match with a low-skilled partner rather than wait. **(AM: L-L)** is the analogous statement for a low-skilled agent: their continuation value does not exceed $\phi(L, L)$, so they too accept a low-skilled partner on meeting.

SA equilibria

In a (symmetric) SA equilibrium with common cutoff $\hat{x} \in (0, 1)$, low-type agents reject low-type partners while high-type agents accept all. The pool on each side has share \hat{x} in H . Match rates differ by skill: H matches at rate λ , L matches at rate $\lambda\hat{x}$ (only with H). The values are:

$$V_S^\gamma(H; \hat{x}) = \frac{\lambda[\hat{x}\phi(H, H) + (1 - \hat{x})\phi(H, L)]}{r + \lambda} \quad (\text{same as AM})$$

$$V_S^\gamma(L; \hat{x}) = \frac{\lambda\hat{x}}{r + \lambda\hat{x}}\phi(L, H)$$

For later reference, denote the value-of-investing gap in the SA regime at common pool composition \hat{x} by

$$G(\hat{x}) := V_S^\gamma(H; \hat{x}) - V_S^\gamma(L; \hat{x}) = \frac{\lambda[\hat{x}\phi(H, H) + (1 - \hat{x})\phi(H, L)]}{r + \lambda} - \frac{\lambda\hat{x}\phi(L, H)}{r + \lambda\hat{x}}. \quad (1)$$

The IC constraints are:

$$\begin{aligned} V_S^\gamma(H; \hat{x}) - C(\hat{x}) &= V_S^\gamma(L; \hat{x}) && \text{(SA: Invest)} \\ V_S^\gamma(L; \hat{x}) &> \phi(L, L) && \text{(SA: L-L Reject)} \end{aligned}$$

The L-rejection IC (**SA: L-L Reject**) simplifies to $\lambda \hat{x}(\phi(L, H) - \phi(L, L)) > r \phi(L, L)$. It is the exact negation of the AM L-L acceptance constraint at the same cutoff: at any given \hat{x} , only one acceptance regime is optimal for the L type.

Closed-form cutoff candidates

It is convenient to summarize the payoff parameters that drive the analysis. Define

$$\begin{aligned} \Delta &:= \phi(H, H) + \phi(L, L) - \phi(H, L) - \phi(L, H) \quad (\leq 0 \text{ under submodularity}), \\ \Delta_h &:= \phi(H, L) - \phi(L, L) \quad (\geq 0), \\ \eta &:= \phi(L, H) - \phi(L, L) \quad (\geq 0). \end{aligned}$$

By Assumption 1, $\Delta_h \geq \eta \geq 0$. Throughout the paper, “SYM AM” refers to a symmetric AM equilibrium (both genders at the same cutoff under all-match acceptance), and “SYM SA” to a symmetric SA equilibrium (both genders at the same cutoff with high types accepting all and low types rejecting low partners).

The closed-form cutoff candidates for the SYM AM, FIOS, and NIOS equilibria are:

$$\begin{aligned} \hat{x}_A &:= \max \left\{ 0, \min \left\{ 1, \frac{\frac{\lambda}{r+\lambda} \Delta_h - c}{1 - \frac{\lambda}{r+\lambda} \Delta} \right\} \right\} && \text{(SYM AM),} \\ \underline{x} &:= \max \left\{ 0, \min \left\{ 1, \frac{\lambda}{r+\lambda} (\Delta + \Delta_h) - c \right\} \right\} && \text{(FIOS),} \\ \bar{x} &:= \max \left\{ 0, \min \left\{ 1, \frac{\lambda}{r+\lambda} \Delta_h - c \right\} \right\} && \text{(NIOS).} \end{aligned} \tag{2}$$

By construction, each candidate lies in $[0, 1]$; the formulae in parentheses in (2) coincide with the cutoffs whenever they are interior, and the boundary values correspond to corner equilibria where one skill type is absent. Under submodularity ($\Delta \leq 0$), $\underline{x} \leq \hat{x}_A \leq \bar{x}$. We define the FIOS profile as $(1, \underline{x})$ (all of one side invest; the other side has cutoff \underline{x}) and the NIOS profile as $(0, \bar{x})$ (no investment on one side; the other has cutoff \bar{x}); both are described up to permutation of genders. In a FIOS profile the fully-investing side has no low-skilled agents on the equilibrium path, so matching is all-match on path; the low-skilled acceptance of that side is an off-path object, pinned by the optimal behavior of an agent who deviates to low skill.⁹ The SYM SA cutoff \hat{x}_{SA} , when it arises, is defined implicitly by the SA

⁹This is what disciplines the fully-investing side’s incentive to invest. An agent of the fully-investing gender who deviated to low skill would face an opposite pool of high-skill share \underline{x} and would optimally *reject* low partners whenever $\lambda \underline{x} \eta > r \phi(L, L)$, securing the value $V_S(L; \underline{x}) =$

investment IC (**SA: Invest**) and is distinct from \hat{x}_A .

Equilibrium structure

Throughout the AM/SA analysis we maintain the condition

$$\lambda(\phi(H, H) - \phi(H, L)) < r\phi(H, L), \quad (\star)$$

which says that a high-skilled agent always strictly prefers to accept a low-skilled partner rather than wait. This is precisely what it means to focus on AM and SA equilibria: (\star) rules out PAM (Proposition 7) and ensures the symmetric equilibrium takes AM or SA form.

Before we see the formal results, I offer a quick preview of them here. Proposition 1 characterizes the equilibrium set: a symmetric equilibrium always exists, all-match (SYM AM) or semi-assortative (SYM SA), and alongside it the economy may admit asymmetric equilibria, in which one gender invests more than the other, of the FIOS or NIOS form. Proposition 2 shows these are not a curiosity: starting from any economy whose only equilibrium is SYM AM, raising the skill premium constructs an economy that admits an asymmetric equilibrium. Proposition 3 records a comparative static: as the premium rises, the investment cutoff rises with it (in the symmetric equilibrium), so more agents acquire skill. Finally, Proposition 4 shows that whenever an asymmetric equilibrium coexists with the symmetric one, the symmetric equilibrium is fragile while the asymmetric one is not, so the asymmetric outcomes, once available, are the robust ones. Section 6 then grounds the skill premium in labor-market wages through a household game.

PROPOSITION 1 (Equilibrium structure): *Maintain (\star) .*

1. Symmetric equilibrium. *There is always a unique symmetric equilibrium. It is a SYM AM equilibrium at cutoff \hat{x}_A when low-skilled agents accept low-skilled partners there ($\lambda\hat{x}_A\eta \leq r\phi(L, L)$); otherwise SYM AM does not exist and the symmetric equilibrium is SYM SA, at the unique cutoff $\hat{x}_{SA} < \hat{x}_A$.*
2. Configurations with SYM AM. *Suppose SYM AM exists, and set aside the knife-edge $1 + \frac{\lambda}{r+\lambda}\Delta = 0$.¹⁰ Any asymmetric (AM) equilibrium is then ei-*

$\lambda\underline{x}\phi(L, H)/(r + \lambda\underline{x})$ rather than the all-match value $V_A(L; \underline{x})$. The FIOS profile is therefore an equilibrium only if its cutoff type weakly prefers to invest against this *optimal* deviation value, $V_A(H; \underline{x}) - (1 + c) \geq \max\{V_A(L; \underline{x}), V_S(L; \underline{x})\}$, equivalently

$$\min\left\{\frac{\lambda}{r + \lambda}(\underline{x}\Delta + \Delta_h), G(\underline{x})\right\} \geq 1 + c.$$

When $\lambda\underline{x}\eta \leq r\phi(L, L)$ this reduces to the all-match condition $\frac{\lambda}{r+\lambda}(\underline{x}\Delta + \Delta_h) \geq 1 + c$ (FIOS is genuinely all-match); when $\lambda\underline{x}\eta > r\phi(L, L)$ it is the strictly stronger requirement $G(\underline{x}) \geq 1 + c$. NIOS raises no such subtlety: there the low-skilled side is populated on path, so its L-L acceptance constraint is imposed directly (see the proof of Proposition 1).

¹⁰At the knife-edge $1 + \frac{\lambda}{r+\lambda}\Delta = 0$ —a non-generic, measure-zero locus in the payoff parameters—the two interior investment indifference conditions become collinear, and a con-

ther *FIOS* $(1, \underline{x})$ or *NIOS* $(0, \bar{x})$, up to permutation of genders, and the two never coexist. Consequently the economy realizes exactly one of three configurations: *SYM AM* is the unique equilibrium; *SYM AM* coexists with a *NIOS* equilibrium and its mirror image; or *SYM AM* coexists with a *FIOS* equilibrium and its mirror image.

Figure 1 illustrates the configurations of Proposition 1. The proposition lays out the possible equilibrium configurations but does not say which one a given environment realizes—in particular, whether an asymmetric equilibrium is available at all. That is the role of Proposition 2 below, which shows that asymmetric (*FIOS*) equilibria emerge as the skill premium rises.

REMARK 1: *Two configurations fall outside Proposition 1. PAM arises when (\star) fails, and is treated in Appendix B. And when *SYM AM* does not exist—so the symmetric equilibrium is *SYM SA*—asymmetric equilibria in which both genders are selective, at different cutoffs, may also arise; we do not pursue these, since our interest is in the asymmetric *AM* equilibria that emerge as the skill premium rises (Proposition 2).*

REMARK 2 (Welfare): *When a symmetric and an asymmetric equilibrium coexist, they are Pareto incomparable: the fully-investing side strictly prefers the symmetric outcome, while the other side strictly prefers the asymmetric one. We omit the formal argument, which amounts to comparing the partner pool each side faces across the two equilibria.*

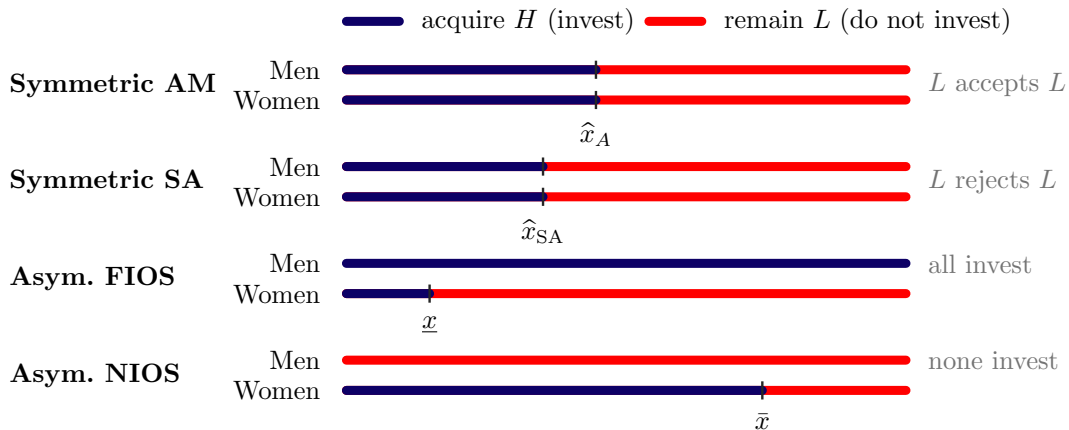


Figure 1: The equilibrium configurations of Proposition 1. Each horizontal line is the type space $[0, 1]$ for one gender; the blue segment marks the types that acquire high skill, the red segment those that remain low-skilled. The two symmetric equilibria share a common cutoff across genders and differ in acceptance behavior: under *Symmetric AM* low-skilled agents accept low-skilled partners (cutoff \hat{x}_A), while under *Symmetric SA* they reject them (cutoff $\hat{x}_{SA} < \hat{x}_A$). In *FIOS* one gender invests fully while the other invests only below \underline{x} ; in *NIOS* one gender does not invest at all while the other invests below \bar{x} . The cutoffs satisfy $\underline{x} \leq \hat{x}_A \leq \bar{x}$. Gender labels are illustrative: each asymmetric configuration has a mirror image with the roles of men and women swapped.

tinuum of interior asymmetric profiles with $\hat{x}_m + \hat{x}_w = 2\hat{x}_A$ can be sustained. We set this case aside throughout.

Proposition 1 establishes the conditions under which an asymmetric equilibrium can surface alongside the symmetric one. These conditions are stated in terms of the primitives of the surplus functions. Next, we present a simple comparative static of how changes in the underlying surplus functions can give rise to such asymmetric equilibria. Choo and Siow (2006), Galichon and Salanié (2022), and others have developed methods to estimate marriage surplus (match payoff) functions from observed matching patterns, which makes a comparative static phrased directly in terms of these surplus functions empirically meaningful. Proposition 2 establishes the pervasiveness of asymmetric equilibria as the skill premium—the increase in each household member’s payoff from having at least one highly-skilled member—rises.

DEFINITION 1: *We say that an economy $\hat{\phi} := \langle \hat{\phi}(H, H), \hat{\phi}(H, L), \hat{\phi}(L, H), \hat{\phi}(L, L) \rangle$ exhibits a higher skill premium relative to $\phi := \langle \phi(H, H), \phi(H, L), \phi(L, H), \phi(L, L) \rangle$ if $\hat{\phi}(H, H) \geq \phi(H, H)$, $\hat{\phi}(H, L) \geq \phi(H, L)$, $\hat{\phi}(L, H) \geq \phi(L, H)$ and $\hat{\phi}(L, L) = \phi(L, L)$.*

PROPOSITION 2 (Higher skill premia make asymmetric equilibria possible): *Fix r, λ, c . Suppose that the unique symmetric equilibrium of an economy ϕ with $\Delta \leq 0$ is SYM AM.¹¹ Furthermore, suppose that*

- (i) $1 + \frac{\lambda}{r+\lambda}\Delta > 0$, and
- (ii) $\frac{\lambda}{r+\lambda}\Delta_h - c \in (0, 1)$.¹²

Then there exists an economy $\hat{\phi}$ exhibiting a higher skill premium relative to ϕ such that $\hat{\phi}$ admits an asymmetric (FIOS) equilibrium.

It is straightforward to construct such examples.

EXAMPLE 1: *Suppose that $c = 2$, $r = \lambda = 1$, $\phi(L, L) = 4$. Consider the following two environments.*

- (i) $\phi(H, H) = 10, \phi(H, L) = 9, \phi(L, H) = 6$.
- (ii) $\hat{\phi}(H, H) = 13, \hat{\phi}(H, L) = 12, \hat{\phi}(L, H) = 8$.

■

Environment (ii) exhibits a higher skill premium relative to environment (i). In environment (i), the only equilibrium is symmetric: SYM AM with cutoff $\hat{x}_A = \frac{1}{3}$. In environment (ii), the symmetric equilibrium is again SYM AM, now at the higher cutoff $\hat{x}_A = \frac{4}{5}$, and an asymmetric FIOS equilibrium emerges alongside it at $\underline{x} = \frac{1}{2}$.

¹¹A simple sufficient condition for the symmetric equilibrium to be all-match is $\lambda\eta \leq r\phi(L, L)$: a low-skilled agent always prefers to marry a low-skilled partner she meets over holding out for a high-skilled one. This essentially requires the search frictions to be sufficiently large—the meeting rate λ small relative to the discount rate r —so that waiting for a better-skilled partner is not worthwhile, however valuable such a partner becomes as the premium rises. We do not impose this condition formally.

¹²Condition (ii) is not essential to the proof but simplifies the casework. It ensures the symmetric cutoff \hat{x}_A is interior.

Proposition 2 shows that asymmetric equilibria become available as the skill premium rises. Before turning to how an economy selects among the equilibria available to it, we record what happens to the symmetric equilibrium itself as the premium rises. Index a family of economies by a parameter τ , a higher value of which corresponds to a higher skill premium; concretely, τ can be taken to be the labor-market wages for high- and low-skilled workers, through which Section 6 microfounds the match payoffs $\phi(H, H), \phi(H, L), \phi(L, H)$.

PROPOSITION 3 (Investment rises with the premium): *Let $\phi(H, H)(\tau), \phi(H, L)(\tau), \phi(L, H)(\tau)$ be continuously differentiable and non-decreasing in τ , with $\phi(L, L)$ fixed, and suppose the symmetric equilibrium is SYM AM with interior cutoff $\hat{x}_A(\tau)$. If $\Delta'_h(\tau) \geq |\Delta'(\tau)|$, then $\hat{x}_A(\tau)$ is non-decreasing in τ , and strictly increasing wherever the inequality is strict.*

Observationally, then, we would find more people investing in skills as the skill premium rises, provided these private returns rise faster than the submodularity. Beyond a certain threshold, however, an asymmetric equilibrium emerges (Proposition 2). If society coordinates on it, we would instead see a divergence in the genders' investments: one gender invests more, the other less. This raises a natural question: why would society switch to an asymmetric equilibrium when a symmetric one remains available? The next section shows that whenever the asymmetric equilibria become available, the symmetric all-match equilibrium becomes fragile in the sense of Onuchic and Ray (2023)—in a sense, the asymmetric equilibria are robust to perturbations in investment that their symmetric counterpart is not. We turn to this next.

5. Equilibrium Selection

Proposition 1 establishes the possibility of an asymmetric equilibrium. However, whenever an asymmetric equilibrium exists, it is accompanied by a symmetric one. We now evaluate the robustness of these equilibria and show that the symmetric equilibria are *fragile*, where we adopt the notion of fragility due to Onuchic and Ray (2023). As mentioned before, while fragility is the notion we adopt, the idea behind it—selecting against equilibria that agents' own best responses pull apart—is standard in the refinement of equilibria in finite games. Before presenting the notion formally, we briefly describe it in words.

Roughly, an equilibrium is fragile if a small antisymmetric perturbation—one gender's investment cutoff nudged up, the other's nudged down by the same amount—gets amplified by a single round of best-response play, in the sense that both genders' cutoffs move further from the original equilibrium than the perturbation itself. The intuition is that even a small payoff-irrelevant bias across gender lines can compound into a sizable deviation, undermining the equilibrium as a focal point. Equilibria that pass this test, by contrast, absorb such perturbations rather than amplifying them.

The notion of fragility

The joint best-response iteration on the cutoff vector $(\hat{x}_m, \hat{x}_w) \in [0, 1]^2$ depends on the acceptance regime of the equilibrium under consideration. In the AM regime, define the investment best-response

$$\text{BR}_{AM}(y) := \max\{0, \min\{1, \frac{\lambda}{r + \lambda}(\Delta y + \Delta_h) - c\}\},$$

and in the SA regime,

$$\text{BR}_{SA}(y) := \max\{0, \min\{1, G(y) - c\}\}.$$

Writing BR for whichever applies, the joint iteration is

$$\Theta(\hat{x}_m, \hat{x}_w) := (\text{BR}(\hat{x}_w), \text{BR}(\hat{x}_m)),$$

and equilibria of the corresponding class are the fixed points of Θ . We adapt [Onuchic and Ray \(2023\)](#)'s definition of fragility to our setting:

DEFINITION 2 (Fragility; [Onuchic and Ray, 2023](#)): *An equilibrium $(\hat{x}_m^*, \hat{x}_w^*)$ of Θ is fragile if there exist $\delta > 0$ and $\zeta > 0$ such that for every $\epsilon \in (-\delta, \delta)$ for which the perturbation $(\hat{x}_m^* + \epsilon, \hat{x}_w^* - \epsilon)$ lies in $[0, 1]^2$,*

$$\begin{aligned} |\Theta_m(\hat{x}_m^* + \epsilon, \hat{x}_w^* - \epsilon) - \hat{x}_m^*| &\geq |\epsilon|(1 + \zeta) \quad \text{and} \\ |\Theta_w(\hat{x}_m^* + \epsilon, \hat{x}_w^* - \epsilon) - \hat{x}_w^*| &\geq |\epsilon|(1 + \zeta), \end{aligned}$$

where the subscripts on Θ refer to its component functions.

The feasibility qualifier $(\hat{x}_m^* + \epsilon, \hat{x}_w^* - \epsilon) \in [0, 1]^2$ is needed because cutoffs are bounded; at a corner equilibrium only one sign of ϵ remains feasible, and the condition is tested on that portion of $(-\delta, \delta)$.

PROPOSITION 4 (Fragility selection): *Maintain (\star) and the affine cost specification. Suppose SYM AM and an asymmetric (AM) equilibrium—FIOS or NIOS, with interior cutoff—coexist. Then SYM AM is fragile, while the coexisting asymmetric equilibrium is not.*

Thus, whenever a symmetric AM equilibrium coexists with an asymmetric (FIOS or NIOS) equilibrium, the symmetric one is fragile while the asymmetric one is not. Consequently, for a non-knife-edge set of parameters, the only equilibria that survive the fragility test are the asymmetric ones. [Section 6](#) expresses the underlying threshold in terms of observable labor-market wages.

Why the symmetric equilibrium is the fragile one. The contrast has a simple source. At the symmetric all-match equilibrium both genders' investments are interior: nudge one gender's investment cutoff up and the other's down by the same amount, and each best-responds by moving in the opposite direction.

Because the best response has slope greater than one in absolute value precisely where an asymmetric equilibrium is available ($\frac{\lambda}{r+\lambda}|\Delta| > 1$, equivalently $1 + \frac{\lambda}{r+\lambda}\Delta < 0$), the nudge is amplified rather than absorbed, and the genders' investments drift apart. At an asymmetric (FIOS) equilibrium, by contrast, one gender is at the corner—everyone invests—where its incentive to invest holds strictly; a small perturbation leaves that gender's best response pinned at the corner, so it absorbs the shock and the unraveling never begins. In each case the comparison is a single round of best response, in the sense made precise in Definition 2.

What fragility does and does not claim. The fragility result is essentially a statement about robustness, and is not an appeal to a particular adjustment dynamic—a rule specifying how agents revise their strategies over time, of the kind used to define stability in evolutionary models. The environment is exactly symmetric across genders by construction. Of course, in reality, no actual marriage market is. Small gender-correlated differences—in social norms, in expectations, in transient shocks, or simply in history—are pervasive, and the symmetric model deliberately abstracts from them. Proposition 4 asks what becomes of each equilibrium once such a perturbation is admitted: at the symmetric equilibrium an arbitrarily small antisymmetric bias is *amplified* by a single round of best responses and the genders' investments move apart, whereas at the asymmetric equilibria the same perturbations die out. The symmetric outcome therefore survives only if the environment is, and remains, perfectly symmetric, while each asymmetric outcome obtains for an open set of the small asymmetries any real market carries. Because the environment is symmetric, the surviving equilibria come in mirror-image pairs, so the argument selects *against symmetry*, not *in favour of a particular gender*; which gender invests more is determined by forces like social norms and history that the model abstracts from.

REMARK 3 (Robust asymmetry does not require submodularity): *The selection of asymmetric over symmetric equilibria is not an artifact of submodular payoffs. Outside the AM class there are supermodular economies in which the unique equilibrium surviving fragility is asymmetric—of the semi-assortative type, sustained by a distinct search channel (a low-skilled agent's pool-dependent value of holding out for a high-skilled partner) rather than by household complementarity. Submodularity is maintained throughout not for the existence or robustness of asymmetric equilibria, but because it is what drives the comparative static in the skill premium (Section 6). Appendix C works out such an example.*

6. Microfoundation from household maximization in NTU

The preceding sections established that asymmetric equilibria can arise—under conditions stated entirely in terms of the match payoff function—whenever the

payoff function is sufficiently submodular. This raises a natural question: what changes in the underlying economy give rise to the changes in match payoffs that, in turn, induce asymmetric equilibria? This section provides one answer. I view the match payoff function as the equilibrium outcome of a non-cooperative household game along the lines of [Chiappori \(1988\)](#),¹³ in which a household consisting of a man and a woman allocates two units of time between labor-market work and household production. The match payoffs $\phi(H, H), \phi(H, L), \phi(L, H), \phi(L, L)$ thus become equilibrium objects of the household game, with the labor-market wages of the two members serving as fundamental primitives. The setup serves two purposes. First, it shows that match payoff functions supporting asymmetric equilibria can arise endogenously from plausible labor-market and household primitives. Second, it allows us to perform comparative statics on more fundamental objects than the surplus function itself—most notably, the wage structure—and to ask how rising wages for high-skilled workers translate into changes in marriage-market behavior, producing what one might call “skill inequality” between genders.

Let us work backwards to see the idea clearly. The households’ problem is the following. A household consists of a man m and a woman w with skills $s(m)$ and $s(w)$. Wages in the market depend only on the skill and not on the gender.¹⁴ The households decide how to split their two units of time—one unit for each member—between labor market work and household work. Members’ allocation is a result of a non-cooperative game. Members of the household pool their total income and buy some consumption from it. There are also several household chores that, if not done, bring disutility to the household. Therefore, members need to allocate some fraction of their two units of time to the household chores. Neither member wishes to perform these chores and would rather work in the outside labor market. The goal of this section is to demonstrate how increases in wages for high-skilled labor can fundamentally alter the composition of skill acquisition choices in equilibrium. I will illustrate this using a simple parametric example. Consider a household (m, w) . Let the skill of member i be s_i . The market wage for skill s is t_s . Therefore, the wage of member i is $t_i := t_{s_i}$. If member i ’s allocation to the outside labor market is e_i , then the net household income is $y(s, e) := t \cdot e$, where $s := (s_m, s_w)$, $t := (t_m, t_w)$, and $e := (e_m, e_w)$ are the skill, wage, and effort vectors. The quality of the household work as a result of their combined effort at home is $h(e) := 2 - e_m - e_w$. The utility of a member i is given by,

$$U_i(s, e) = u(y(s, e), h(e)) - g(1 - e_i).$$

¹³[Chiappori \(1988\)](#) adopts a cooperative formulation to obtain the Pareto-optimal choices. While I keep the preferences of the members of the household similar, I adopt a non-cooperative approach in determining the labor supply choices. See [Lundberg and Pollak \(1993\)](#) for similar models.

¹⁴As would be clear, a gender pay gap as it exists in reality can only make it easier to construct equilibria where women invest less in skill acquisition.

Essentially, the households treat both, their income and the level of household chores as public goods, while the cost, $g(\cdot)$, of their work at home is private. We assume that $u_1 > 0, u_{11} \leq 0, u_2 > 0, u_{22} \leq 0$. Moreover, the cost $g(\cdot)$ is assumed to be differentiable and convex.

Members choose their allocations, (e_i, e_j) in a non-cooperative manner taking the other person's choice as given. Let us write the agents' first order condition (assuming interior allocation choices):

$$\begin{aligned} u_1(y(s, e), h(e))t_i - u_2(y(s, e), h(e)) + g'(1 - e_i) &= 0 \\ \implies u_1(y(s, e), h(e))(t_i - t_j) &= g'(1 - e_j) - g'(1 - e_i) \end{aligned} \quad (3)$$

Therefore, if the players' choices are interior in equilibrium, then $t_i > t_j \implies e_i > e_j$ whenever $g(\cdot)$ is strictly convex.¹⁵ Thus, if the household consists of two members with unequal skills—and therefore unequal wages—then we will have specialization in equilibrium: the more skilled member will work more in the professional market than the one with lower skills. This observation is at the core of the following numerical example that illustrates how increasing wages for high-skilled workers can give rise to asymmetric equilibria.

EXAMPLE 2: Consider a household consisting of (m, w) . The utility of member i is $U(y(s, e), h(e)) = K \cdot [\alpha \log(t \cdot e) + (1 - \alpha) \log(2 - e_m - e_w)] - \frac{1}{2}(1 - e_i)^2$, where $K = 8$ and $\alpha = 0.6$. Let $t_l = 2$. Consider the wage for high-skills going from $t_h = 3$ to $\hat{t}_h = 5$. The cost of acquiring skills for an agent with type x is $C(x) = c + x$, where $c = 0.25$. The discount and meeting rates are $r = \lambda = 1$. ■

When two members of skill s and s' match, they choose their effort choices non-cooperatively. Using (3), we solve this numerically for each pair of skills $(h, h), (h, l), (l, h), (l, l)$ to obtain $\phi(H, H), \phi(H, L), \phi(L, H), \phi(L, L)$ corresponding to $t_h = 3$ and $\hat{\phi}(H, H), \hat{\phi}(H, L), \hat{\phi}(L, H), \hat{\phi}(L, L)$ corresponding to $t_h = 5$. We provide a detailed working of this in Appendix A.6, but the parameters are as follows:

$$\begin{aligned} \phi(H, H) &= 5.3547, \phi(H, L) = 5.2114, \phi(L, H) = 4.8495, \phi(L, L) = 3.4085, \text{ and} \\ \hat{\phi}(H, H) &= 7.8067, \hat{\phi}(H, L) = 7.7253, \hat{\phi}(L, H) = 7.2253, \hat{\phi}(L, L) = 3.4085. \end{aligned}$$

In the environment with $t_h = 3$, the unique equilibrium is symmetric with $\hat{x}_A \approx 0.395$: the skill premium is too low to support any asymmetric equilibrium. When $t_h = 5$, the symmetric equilibrium is still SYM AM, now at the higher cutoff $\hat{x}_A \approx 0.666$ —as the premium rises, both genders invest more. But it is no longer alone: an asymmetric FIOS equilibrium with $\underline{x} \approx 0.041$ becomes available alongside it, in which all men invest while only about 4% of women do (the mirror image, with the genders swapped, is also an equilibrium). In the FIOS equilibrium, low-

¹⁵In fact, this also implies that if $g(\cdot)$ were affine, then the allocation choices cannot be interior if the wages are unequal.

skilled women meet only high-skilled men on path, so the acceptance constraint (AM: L-L) is vacuous for women, while it is verified to hold for the investing side. Thus, as wages rise, the economy moves from a regime in which about 40% of each gender invest symmetrically to one in which that symmetric outcome coexists with a starkly asymmetric one—one gender investing fully, the other barely at all. Despite no asymmetry in the payoff or technological features of the underlying environment, the equilibrium set can exhibit stark asymmetries between the two genders as the skill premium rises.

The pattern displayed in Example 2—a low-skill-premium regime in which the symmetric equilibrium is the only outcome, and a high-skill-premium regime in which an asymmetric (FIOS) equilibrium becomes available alongside the symmetric one—is not specific to the Cobb-Douglas example. The next definition isolates a broad class of household games for which the same comparative static holds.

DEFINITION 3: A pooled-income household (PIH) game is the simultaneous-move game in which two members $i \in \{m, w\}$ with wages (t_m, t_w) choose efforts $e_i \in [0, 1]$ to maximize

$$U_i(e_m, e_w; t_m, t_w) = \psi(t_m e_m + t_w e_w, 2 - e_m - e_w) - g(1 - e_i),$$

where

- $\psi : \mathbb{R}_{++}^2 \rightarrow \mathbb{R}$ is twice continuously differentiable, strictly increasing in both arguments, weakly concave, and satisfies $\lim_{y \rightarrow \infty} \psi(y, h) = \infty$ for every fixed $h > 0$;
- $g : [0, 1] \rightarrow \mathbb{R}$ is twice continuously differentiable, strictly increasing, and strictly convex.

We assume ψ and g are such that the game has a unique Nash equilibrium.¹⁶ The match payoff $\phi(s, s')$ at skills (s, s') is the value U_i takes (for agent i with skill s) at this equilibrium, with wages $(t_s, t_{s'})$.

The Cobb-Douglas case in Example 2 corresponds to $\psi(y, h) = K[\alpha \log y + (1 - \alpha) \log h]$ and $g(z) = \frac{1}{2}z^2$.

In what follows, we give sufficient conditions on the asymptotic behavior of the induced match payoffs $\phi(s, s'; t_h)$ for a PIH game to produce the comparative-static phenomenon: a single wage threshold t^* below which only the symmetric equilibrium exists and above which an asymmetric (FIOS) equilibrium becomes available alongside it. We refer to environments satisfying these conditions as displaying *Increasing Submodularity in High Skills* (ISHS).

DEFINITION 4 (Increasing Submodularity in High Skills): A marriage-market environment derived from a PIH game with $t_i > 0$ fixed, affine cost $C(x) = x + c$,

¹⁶Rosen (1965) provides a sufficient condition for a class of games. For the Cobb-Douglas specification used in our example, uniqueness is straightforward to verify directly.

and parameters $r, \lambda, c > 0$ displays Increasing Submodularity in High Skills (ISHS) if the induced match payoffs $\phi(s, s'; t_h)$ satisfy:

(C1) $\Delta(t_h) \rightarrow -\infty$ and $\Delta_h(t_h) \rightarrow \infty$ as $t_h \rightarrow \infty$;

(C2) $\liminf_{t_h \rightarrow \infty} \hat{x}_A(t_h) > 0$;

(C3) Assumption 1 holds for all $t_h \geq t_l$, and there exists $T < \infty$ such that for all $t_h \geq T$, $\underline{x}(t_h) \in (0, 1)$ and the FIOS profile $(1, \underline{x}(t_h))$ satisfies (AM: H-L) on path;

(C4) $t_h \mapsto 1 + \frac{\lambda}{r+\lambda}\Delta(t_h)$ is single-crossing on $[t_l, \infty)$: it changes sign at most once, and the change is from positive to negative.

In words, (C1) says that as the high-skill wage grows, the marriage payoffs become increasingly submodular: the asymmetric matches $(\phi(H, L), \phi(L, H))$ outpace the symmetric matches $(\phi(H, H), \phi(L, L))$ without bound. (C2) says that the symmetric equilibrium remains non-trivial: its investment cutoff stays bounded away from zero no matter how high the skill premium is. (C3) says that once the wage premium is large enough, the FIOS equilibrium structure is supported by the underlying payoffs—its cutoff $\underline{x}(t_h)$ lies in $(0, 1)$, and the H-skilled side’s acceptance constraint holds. (C4) is a single-crossing condition: once the payoffs become submodular enough to support FIOS ($1 + \frac{\lambda}{r+\lambda}\Delta \leq 0$), they stay so as the premium rises further. It is what makes “FIOS becomes available above a threshold” a genuine threshold statement rather than merely an eventual one; (C1) alone gives only the eventual sign of $1 + \frac{\lambda}{r+\lambda}\Delta$, not the absence of a later return to positivity.

The Cobb-Douglas environment of Example 2 displays ISHS. In particular (C4) is not vacuous there: in the corner regime (t_h large) the $\log t_h$ terms of the (H, H) , (H, L) , and (L, H) matches combine to give $\Delta(t_h) \sim -K\alpha \log t_h$, which is strictly decreasing, so $1 + \frac{\lambda}{r+\lambda}\Delta$ is eventually strictly decreasing; in the calibration of Example 2 it is monotone across its single zero at ≈ 3.51 (verified numerically). Since the remaining FIOS-feasibility conditions (C3) already hold there, the threshold t^* at which a FIOS equilibrium appears coincides with this sign change, $t^* \approx 3.51$, rather than being defined by it.

PROPOSITION 5: *Let the marriage-market environment be derived from a PIH game and assume it displays ISHS (Definition 4). Then there is a wage threshold $t^* \in (t_l, \infty)$ such that a FIOS equilibrium coexists with the symmetric equilibrium for all $t_h > t^*$; while for every t_h with $1 + \frac{\lambda}{r+\lambda}\Delta(t_h) > 0$ —in particular, for all sufficiently small t_h —no asymmetric equilibrium exists, so the symmetric equilibrium is the unique equilibrium.¹⁷ Suppose in addition that $\Delta'_h(t_h) \geq |\Delta'(t_h)|$,¹⁸*

¹⁷Both asymmetric (AM) forms require $1 + \frac{\lambda}{r+\lambda}\Delta < 0$ (the FIOS and NIOS investment ICs; see the proof of Proposition 4). In the ISHS calibration of Example 2, $1 + \frac{\lambda}{r+\lambda}\Delta$ changes sign exactly at $t^* \approx 3.51$, so below t^* the symmetric SYM AM equilibrium is the only equilibrium and above it a FIOS equilibrium coexists with it.

¹⁸This condition is satisfied for the Cobb-Douglas case in Example 2, where $\Delta'_h(t_h) = -\Delta'(t_h) = K\alpha/t_h$ once the (H, L) household fully specializes ($t_h \gtrsim 4.4$ in this calibration),

and consider the range of wages over which the symmetric equilibrium is SYM AM. Then:

1. the symmetric cutoff $\hat{x}_A(t_h)$ is continuous and non-decreasing in t_h , so as the premium rises both genders invest more; and
2. wherever the FIOS equilibrium exists, its cutoff lies strictly below the symmetric one, $\underline{x}(t_h) < \hat{x}_A(t_h)$, and is itself non-decreasing in t_h .

For the Cobb-Douglas example with $r = \lambda = 1$, $c = 0.25$, $K = 8$, $\alpha = 0.6$, $t_l = 2$, the threshold is $t^* \approx 3.51$: below it (as at $t_h = 3$) the symmetric SYM AM equilibrium is the only outcome, and above it (as at $t_h = 5$) an asymmetric FIOS equilibrium appears alongside it. The proof, in Appendix A.7, isolates the mechanism: rising t_h drives $|\Delta|$ above the level at which the FIOS incentive constraints first become satisfiable. At still higher wages the symmetric equilibrium itself eventually turns selective—low-skilled agents begin rejecting one another, so SYM AM gives way to SYM SA (here at $t_\Phi \approx 6.11$)—but our interest is the emergence of the asymmetric equilibrium within the SYM AM range $t^* < t_h < t_\Phi$.

Empirical implications of Propositions 1–5. As the wage for high-skilled workers rises, the model’s predictions divide into two phases. In the low-premium phase ($t_h < t^*$), where the payoffs are not yet submodular enough to admit an asymmetric equilibrium, the only equilibrium is symmetric and the investment cutoff rises continuously with the premium: both genders increase their investment in skill acquisition at the same rate. Empirically, this corresponds to broadly rising educational attainment, equal across genders.

In the high-premium phase, once t_h crosses t^* and while the symmetric equilibrium remains all-match, an asymmetric (FIOS) equilibrium becomes available alongside it. Under the fragility refinement of Section 5, this symmetric all-match equilibrium is fragile while the asymmetric one is not. The asymmetric equilibrium carries a sharp empirical signature: one gender invests in skill acquisition fully while the other invests substantially less and specializes in domestic production—a widening gender gap in education and labor-market participation that arises without any change in underlying preferences or technology, purely as the skill premium crosses the threshold. The model thus predicts that a rising skill premium can open up gender asymmetries exactly where none were present at lower premia.

7. Conclusion

We have studied a canonical search-and-matching model with non-transferable utility in which, despite a fully symmetric environment, asymmetries in skill investment across genders can arise. Moreover, wherever a symmetric all-match

so $\Delta'_h = |\Delta'|$ there; for smaller t_h the strict inequality $\Delta'_h > |\Delta'|$ holds (verified numerically). The hypothesis uses only the inequality.

equilibrium coexists with an asymmetric one, the former is fragile in the sense of [Onuchic and Ray \(2023\)](#) while the latter is not, suggesting that the asymmetric outcomes are the more robust ones once they become available. Finally, the underlying changes in match payoffs can be microfounded through a household problem, allowing us to perform comparative statics on more primitive objects of the economy, such as labor-market wages.

Fragility is, as [Onuchic and Ray \(2023\)](#) note, “a mathematical construct,” and I do not claim it is the literal process by which a marriage market reaches an outcome. The claim is weaker and more robust: the symmetric equilibrium survives only if the environment is, and remains, exactly gender-symmetric, because the small, payoff-irrelevant asymmetries any real market carries—and that the symmetric model assumes away—are amplified there rather than absorbed. One should therefore not expect to observe the symmetric outcome, not because a particular refinement rules it out, but because it does not survive contact with the asymmetries it ignores. The asymmetric equilibria carry no such defect.

Read alongside Propositions 1–5, the model illustrates three things: (i) in an entirely symmetric environment, asymmetric investments in skill acquisition can arise in equilibrium; (ii) such asymmetric equilibria are robust to small antisymmetric perturbations, whereas the symmetric all-match equilibria with which they coexist are fragile (Proposition 4); and (iii) the shift from symmetric to asymmetric equilibria can be triggered by rising skill premia, which in turn can be linked to readily observable economic variables such as wages in the high-skill sector. In short, somewhat paradoxically, rising skill premia may end up amplifying or triggering gaps in skill investments across genders rather than mitigating them.

References

- D. Acemoglu and D. Autor. Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter and D. Card, editors, *Handbook of Labor Economics*, volume 4B, pages 1043–1171. Elsevier, 2011.
- H. Adachi. A search model of two-sided matching under nontransferable utility. *Journal of Economic Theory*, 113(2):182–198, 2003.
- A. Atakan, M. Richter, and M. Tsur. Efficient investment, search, and sorting in matching markets. *Theoretical Economics*, 2026. Forthcoming.
- G. S. Becker. A theory of marriage: Part i. *Journal of Political economy*, 81(4): 813–846, 1973.
- M. Bertrand, E. Kamenica, and J. Pan. Gender identity and relative income within households. *The Quarterly Journal of Economics*, 130(2):571–614, 2015.

- V. Bhaskar and E. Hopkins. Marriage as a rat race: Noisy premarital investments with assortative matching. *Journal of Political Economy*, 124(4):992–1045, 2016.
- V. Bhaskar, W. Li, and J. Yi. Multidimensional premarital investments with imperfect commitment. *Journal of Political Economy*, 131(10):2893–2919, 2023.
- K. Burdett and M. G. Coles. Marriage and class. *Quarterly Journal of Economics*, 112(1):141–168, 1997.
- H. Chade and I. Lindenlaub. Risky matching. *The Review of Economic Studies*, 89(2):626–665, 2022.
- P.-A. Chiappori. Rational household labor supply. *Econometrica*, 56(1):63–90, 1988. doi: 10.2307/1911842. URL <https://doi.org/10.2307/1911842>.
- P.-A. Chiappori. The theory and empirics of the marriage market. *Annual Review of Economics*, 12(1):547–578, 2020.
- P.-A. Chiappori, M. Iyigun, and Y. Weiss. Investment in schooling and the marriage market. *American Economic Review*, 99(5):1689–1713, 2009.
- P.-A. Chiappori, B. Salanié, and Y. Weiss. Partner choice, investment in children, and the marital college premium. *American Economic Review*, 107(8):2109–2167, 2017.
- E. Choo and A. Siow. Who marries whom and why. *Journal of Political Economy*, 114(1):175–201, 2006.
- H. L. Cole, G. J. Mailath, and A. Postlewaite. Efficient non-contractible investments in large economies. *Journal of Economic Theory*, 101(2):333–373, 2001.
- J. Eeckhout. Bilateral search and vertical heterogeneity. *Econometrica*, 67(6):1291–1311, 1999. doi: 10.1111/1468-0262.00082.
- P. Francois. Gender discrimination without gender difference: Theory and policy responses. *Journal of Public Economics*, 68(1):1–32, 1998.
- D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- A. Galichon and B. Salanié. Cupid’s invisible hand: Social surplus and identification in matching models. *Review of Economic Studies*, 89(5):2600–2629, 2022.
- C. Goldin. The quiet revolution that transformed women’s employment, education, and family. *American Economic Review*, 96(2):1–21, 2006. doi: 10.1257/000282806777212350.

- C. Goldin and L. F. Katz. *The Race Between Education and Technology*. Belknap Press of Harvard University Press, Cambridge, MA, 2008.
- J. Greenwood, N. Guner, G. Kocharkov, and C. Santos. Marry your like: Assortative mating and income inequality. *American Economic Review*, 104(5): 348–353, 2014. doi: 10.1257/aer.104.5.348.
- G. K. Hadfield. A coordination model of the sexual division of labor. *Journal of Economic Behavior and Organization*, 40(2):125–153, 1999.
- M. Iyigun and R. P. Walsh. Building the family nest: Premarital investments, marriage markets, and spousal allocations. *Review of Economic Studies*, 74(2): 507–535, 2007. doi: 10.1111/j.1467-937X.2007.00425.x.
- J. Lafortune. Making yourself attractive: Pre-marital investments and the returns to education in the marriage market. *American Economic Journal: Applied Economics*, 5(2):151–178, 2013. doi: 10.1257/app.5.2.151.
- P. Legros and A. F. Newman. Beauty is a beast, frog is a prince: Assortative matching with nontransferabilities. *Econometrica*, 75(4):1073–1102, 2007. doi: 10.1111/j.1468-0262.2007.00780.x.
- C. Low. The human capital–reproductive capital trade-off in marriage market matching. *Journal of Political Economy*, 132(2):552–576, 2024.
- S. Lundberg and R. A. Pollak. Separate spheres bargaining and the marriage market. *Journal of Political Economy*, 101(6):988–1010, 1993.
- G. J. Mailath, A. Postlewaite, and L. Samuelson. Pricing and investments in matching markets. *Theoretical Economics*, 8(2):535–590, 2013.
- G. Nöldeke and L. Samuelson. Investment and competitive matching. *Econometrica*, 83(3):835–896, 2015.
- P. Onuchic and D. Ray. Signaling and discrimination in collaborative projects. *American Economic Review*, 113(1):210–252, 2023.
- M. Peters. The pre-marital investment game. *Journal of Economic Theory*, 137(1):186–213, 2007.
- M. Peters and A. Siow. Competing premarital investments. *Journal of Political Economy*, 110(3):592–608, 2002.
- K. Ritzberger and J. W. Weibull. Evolutionary selection in normal-form games. *Econometrica*, 63(6):1371–1399, 1995.
- J. B. Rosen. Existence and uniqueness of equilibrium points for concave n -person games. *Econometrica*, 33(3):520–534, 1965.

- L. S. Shapley and M. Shubik. The assignment game i: The core. *International Journal of game theory*, 1(1):111–130, 1971.
- R. Shimer and L. Smith. Assortative matching and search. *Econometrica*, 68(2): 343–369, 2000.
- R. Shimer and L. Smith. Matching, search, and heterogeneity. *Topics in Macroeconomics*, 1(1):153460131010, 2001.
- A. Siow. Differential fecundity, markets, and gender roles. *Journal of Political Economy*, 106(2):334–354, 1998.
- L. Smith. The marriage model with search frictions. *Journal of political Economy*, 114(6):1124–1144, 2006.
- H. Zhang. An investment-and-marriage model with differential fecundity: On the college gender gap. *Journal of Political Economy*, 129(5):1464–1486, 2021.

A. Appendix: Proofs

A.1. Symmetry of AM equilibria under supermodular payoffs

The following auxiliary result is referenced in Section 4 to justify the focus on submodular payoffs. It says that under supermodular payoffs, no AM equilibrium can exhibit asymmetric investment cutoffs across genders; in particular, the FIOS and NIOS asymmetric structures are ruled out in this case. Combined with the analysis under submodular payoffs in the main text, this result confirms that asymmetric outcomes within the all-match class require submodularity of the match payoff function (Appendix C shows that the semi-assortative class admits asymmetric equilibria under supermodularity).

PROPOSITION 6: *If $\phi(\cdot, \cdot)$ is supermodular ($\Delta \geq 0$), then all AM equilibria are symmetric, i.e., $\hat{x}_m = \hat{x}_w$.*

Proof of Proposition 6. Let us split the analysis into three cases:

Case 1: $\hat{X} \in (0, 1)^2$. Suppose that $\hat{x}_m \neq \hat{x}_w$. From (AM: Invest), we have,

$$\begin{aligned} V_A^\gamma(H; \hat{X}) - V_A^\gamma(L; \hat{X}) &= \frac{\lambda}{r + \lambda} \left[\hat{x}_\gamma \underbrace{(\phi(H, H) + \phi(L, L) - \phi(H, L) - \phi(L, H))}_{\geq 0} \right. \\ &\quad \left. + \underbrace{\phi(H, L) - \phi(L, L)}_{\geq 0} \right] \\ &= C(\hat{x}_\gamma). \end{aligned}$$

Therefore,

$$\begin{aligned} [V_A^m(H; \hat{X}) - V_A^m(L; \hat{X})] - [V_A^w(H; \hat{X}) - V_A^w(L; \hat{X})] \\ &= \frac{\lambda}{r + \lambda} (\hat{x}_w - \hat{x}_m) \Delta \\ &= C(\hat{x}_m) - C(\hat{x}_w) \end{aligned}$$

Therefore, if $\hat{x}_m > \hat{x}_w$, then $LHS \leq 0 < RHS$ (with $LHS = 0$ only at the modular edge $\Delta = 0$, where the equality $LHS = RHS$ then forces $\hat{x}_m = \hat{x}_w$), a contradiction.

Case 2: $\hat{x}_m \in [0, 1)$ and $\hat{x}_w = 1$. Then,

$$\begin{aligned} V_A^m(H; \hat{X}) - V_A^m(L; \hat{X}) &= \frac{\lambda}{r + \lambda} [\hat{x}_w \Delta + \Delta_h] \leq C(\hat{x}_m) \\ V_A^w(H; \hat{X}) - V_A^w(L; \hat{X}) &= \frac{\lambda}{r + \lambda} [\hat{x}_m \Delta + \Delta_h] \geq C(\hat{x}_w) \end{aligned}$$

Therefore,

$$\begin{aligned} & [V_A^m(H; \hat{X}) - V_A^m(L; \hat{X})] - [V_A^w(H; \hat{X}) - V_A^w(L; \hat{X})] \\ &= \frac{\lambda}{r + \lambda}(\hat{x}_w - \hat{x}_m)\Delta = \frac{\lambda}{r + \lambda}(1 - \hat{x}_m)\Delta \leq C(\hat{x}_m) - C(1) \end{aligned}$$

However, $LHS \geq 0 > RHS$, a contradiction.

Case 3: $\hat{x}_m \in (0, 1]$ and $\hat{x}_w = 0$. Then,

$$\begin{aligned} V_A^m(H; \hat{X}) - V_A^m(L; \hat{X}) &= \frac{\lambda}{r + \lambda}[\hat{x}_w\Delta + \Delta_h] \geq C(\hat{x}_m) \\ V_A^w(H; \hat{X}) - V_A^w(L; \hat{X}) &= \frac{\lambda}{r + \lambda}[\hat{x}_m\Delta + \Delta_h] \leq C(\hat{x}_w) \end{aligned}$$

Therefore,

$$\begin{aligned} & [V_A^m(H; \hat{X}) - V_A^m(L; \hat{X})] - [V_A^w(H; \hat{X}) - V_A^w(L; \hat{X})] \\ &= \frac{\lambda}{r + \lambda}(\hat{x}_w - \hat{x}_m)\Delta = -\frac{\lambda}{r + \lambda}\hat{x}_m\Delta \geq C(\hat{x}_m) - C(0) \end{aligned}$$

However, $LHS \leq 0 < RHS$, a contradiction. \square

A.2. Proof of Proposition 1

Proof of Proposition 1. For the reader's convenience, recall the payoff aggregates

$$\begin{aligned} \Delta &:= \phi(H, H) + \phi(L, L) - \phi(H, L) - \phi(L, H), \\ \Delta_h &:= \phi(H, L) - \phi(L, L), \\ \eta &:= \phi(L, H) - \phi(L, L), \end{aligned}$$

and the closed-form cutoff candidates:

$$\hat{x}_A = \frac{\frac{\lambda}{r+\lambda}\Delta_h - c}{1 - \frac{\lambda}{r+\lambda}\Delta}, \quad \underline{x} = \frac{\lambda}{r + \lambda}(\Delta + \Delta_h) - c, \quad \bar{x} = \frac{\lambda}{r + \lambda}\Delta_h - c.$$

Throughout, $\Delta \leq 0$, Assumption 1 (which gives $\Delta_h \geq \eta \geq 0$ and $\phi(H, L) \geq \phi(L, L)$), and the maintained condition (\star) hold.

Proof of Part 1 (symmetric equilibrium). Throughout this proof, we treat the generic interior case in which the unclipped formula for \hat{x}_A in (2) lies in $(0, 1)$; the corner cases $\hat{x}_A \in \{0, 1\}$ correspond to the all-L or all-H profiles whose acceptance constraints are vacuous and are easily verified.

At any symmetric AM profile (\hat{X}, \hat{X}) , (AM: Invest) reduces to

$$\frac{\lambda}{r + \lambda}[\hat{X}\Delta + \Delta_h] = \hat{X} + c,$$

whose unique interior solution is $\hat{X} = \hat{x}_A$ (the LHS is non-increasing and the RHS strictly increasing in \hat{X} , since $\Delta \leq 0$). The H-L acceptance constraint (**AM: H-L**) at (\hat{x}_A, \hat{x}_A) follows from the L-L constraint (**AM: L-L**) via submodularity ($\phi(H, H) - \phi(H, L) \leq \eta$) and $\phi(H, L) \geq \phi(L, L)$:

$$\lambda \hat{x}_A (\phi(H, H) - \phi(H, L)) \leq \lambda \hat{x}_A \eta \leq r \phi(L, L) \leq r \phi(H, L).$$

Thus SYM AM's existence is determined by the L-L acceptance condition alone.

Case A: $\lambda \hat{x}_A \eta \leq r \phi(L, L)$. The profile (\hat{x}_A, \hat{x}_A) constitutes the unique SYM AM equilibrium. The SA L-rejection condition (**SA: L-L Reject**), $\lambda \hat{X} \eta > r \phi(L, L)$, is the strict negation of the L-L acceptance, so no SYM SA can coexist *at the cutoff* \hat{x}_A ; that no SYM SA exists at any other cutoff either is established at the end of this proof.

Case B: $\lambda \hat{x}_A \eta > r \phi(L, L)$. SYM AM does not exist. Define the threshold cutoff

$$x^* := \frac{r \phi(L, L)}{\lambda \eta} \in (0, \hat{x}_A),$$

at which the AM L-L acceptance just binds. At x^* , the L-type AM and SA values coincide and both equal $\phi(L, L)$:

$$\begin{aligned} V_A^\gamma(L; x^*) &= \frac{\lambda [x^* \phi(L, H) + (1 - x^*) \phi(L, L)]}{r + \lambda} = \frac{\lambda x^* \eta + \lambda \phi(L, L)}{r + \lambda} = \phi(L, L), \\ V_S^\gamma(L; x^*) &= \frac{\lambda x^* \phi(L, H)}{r + \lambda x^*} = \phi(L, L), \end{aligned}$$

where for the second line we use $r + \lambda x^* = \lambda x^* \phi(L, H) / \phi(L, L)$ (rearrangement of $r \phi(L, L) = \lambda x^* \eta$). Since H always matches at rate λ under both AM and SA, $V_S^\gamma(H; \hat{X}) = V_A^\gamma(H; \hat{X})$ for all \hat{X} . For each regime, define the net gain to investing for the cutoff type—that is, the value of being H minus the value of being L minus the investment cost—as

$$\begin{aligned} F_S(\hat{X}) &:= V_S^\gamma(H; \hat{X}) - V_S^\gamma(L; \hat{X}) - (\hat{X} + c), \\ F_A(\hat{X}) &:= V_A^\gamma(H; \hat{X}) - V_A^\gamma(L; \hat{X}) - (\hat{X} + c) = \frac{\lambda}{r + \lambda} [\hat{X} \Delta + \Delta_h] - (\hat{X} + c). \end{aligned}$$

Each $F_\sigma(\hat{X}) = 0$ is the symmetric investment indifference condition in regime $\sigma \in \{S, A\}$ at cutoff \hat{X} . Since $V_S^\gamma(L; x^*) = V_A^\gamma(L; x^*) = \phi(L, L)$ and $V_S^\gamma(H; \hat{X}) = V_A^\gamma(H; \hat{X})$, we have $F_S(x^*) = F_A(x^*)$. Since F_A is strictly decreasing in \hat{X} (as $\Delta \leq 0$) and $F_A(\hat{x}_A) = 0$ with $x^* < \hat{x}_A$,

$$F_S(x^*) = F_A(x^*) > F_A(\hat{x}_A) = 0.$$

At the upper endpoint, $F_S(1) = \lambda(\phi(H, H) - \phi(L, H)) / (r + \lambda) - 1 - c$. In the corner subcase $F_S(1) \geq 0$, the profile $(1, 1)$ is itself a symmetric equilibrium (ev-

everyone invests). Otherwise, by the intermediate value theorem, F_S has a zero $\hat{x}_{SA} \in (x^*, 1)$ satisfying (SA: Invest). At this root, $\lambda\hat{x}_{SA}\eta > \lambda x^*\eta = r\phi(L, L)$, so (SA: L-L Reject) holds strictly.

For uniqueness of \hat{x}_{SA} , note that

$$F'_S(\hat{X}) = \frac{\lambda(\phi(H, H) - \phi(H, L))}{r + \lambda} - \frac{\lambda r\phi(L, H)}{(r + \lambda\hat{X})^2} - 1$$

is strictly increasing in \hat{X} (the middle term is decreasing), so F_S is convex. A convex function with $F_S(x^*) > 0$ and $F_S(1) < 0$ has exactly one zero on $(x^*, 1)$. Hence \hat{x}_{SA} is unique.

Moreover, $\hat{x}_{SA} < \hat{x}_A$. For $\hat{X} \in (x^*, 1)$ the L-type strictly prefers to reject low partners, so $V_S^\gamma(L; \hat{X}) > V_A^\gamma(L; \hat{X})$ and hence $F_S(\hat{X}) < F_A(\hat{X})$; in particular $F_S(\hat{x}_A) < F_A(\hat{x}_A) = 0$. Since $F_S(x^*) > 0$ and \hat{x}_{SA} is the unique zero of F_S on $(x^*, 1)$, the downward crossing satisfies $\hat{x}_{SA} \in (x^*, \hat{x}_A)$.

It remains to verify that high-skilled agents accept low-skilled partners at \hat{x}_{SA} , as an SA configuration requires. A high type's acceptance constraint is $\lambda\hat{x}_{SA}(\phi(H, H) - \phi(H, L)) \leq r\phi(H, L)$; since $\hat{x}_{SA} \leq 1$, the maintained condition (\star) gives $\lambda\hat{x}_{SA}(\phi(H, H) - \phi(H, L)) \leq \lambda(\phi(H, H) - \phi(H, L)) < r\phi(H, L)$, so the constraint holds. Thus $(\hat{x}_{SA}, \hat{x}_{SA})$ satisfies all the SA constraints and is a SYM SA equilibrium.

Finally, we argue that SYM AM and SYM SA cannot coexist, even at different cutoffs. Suppose SYM AM exists at \hat{x}_A . Its L-L acceptance condition gives $\lambda\hat{x}_A\eta \leq r\phi(L, L)$, equivalently $\hat{x}_A \leq x^*$. Since F_A is strictly decreasing and vanishes at \hat{x}_A , we have $F_A(x^*) \leq F_A(\hat{x}_A) = 0$. Since L-type AM and SA values coincide at x^* (shown above), $F_S(x^*) = F_A(x^*) \leq 0$. By the convexity of F_S established above, together with the generic case $F_S(1) < 0$, we have $F_S \leq 0$ on $[x^*, 1]$. Hence F_S admits no zero on $[x^*, 1]$ at which L-rejection ($\lambda\hat{X}\eta > r\phi(L, L)$, i.e., $\hat{X} > x^*$) holds, so no SYM SA equilibrium exists. The argument in Case B above showed the symmetric converse: when SYM SA exists, SYM AM does not. The symmetric equilibrium is therefore unique in either case.

Proof of Part 2 (form of asymmetric equilibria). *At most one asymmetric AM equilibrium.* Suppose an interior asymmetric equilibrium $(\hat{x}_m, \hat{x}_w) \in (0, 1)^2$ with $\hat{x}_m \neq \hat{x}_w$ exists. The investment ICs (AM: Invest) yield

$$\begin{aligned} \frac{\lambda}{r + \lambda}[\hat{x}_w\Delta + \Delta_h] &= \hat{x}_m + c, \\ \frac{\lambda}{r + \lambda}[\hat{x}_m\Delta + \Delta_h] &= \hat{x}_w + c. \end{aligned} \tag{4}$$

Adding the two equations gives $\hat{x}_m + \hat{x}_w = 2\hat{x}_A$, while subtracting gives $(\hat{x}_m - \hat{x}_w)(1 + \frac{\lambda}{r + \lambda}\Delta) = 0$. Since $\hat{x}_m \neq \hat{x}_w$, this requires the knife-edge condition $1 + \frac{\lambda}{r + \lambda}\Delta = 0$, which we exclude. Hence no interior asymmetric AM equilibrium exists

generically.

It follows that any asymmetric AM equilibrium has at least one cutoff at $\{0, 1\}$. Suppose without loss of generality that $\hat{x}_m = 1$. The marginal woman's best response at men's cutoff 1 is the unique solution to

$$\hat{x}_w(1) = \begin{cases} 0 & \text{if } \frac{\lambda}{r+\lambda}[\Delta + \Delta_h] \leq c, \\ 1 & \text{if } \frac{\lambda}{r+\lambda}[\Delta + \Delta_h] \geq 1 + c, \\ \frac{\lambda}{r+\lambda}[\Delta + \Delta_h] - c = \underline{x} & \text{otherwise,} \end{cases}$$

uniqueness following from strict monotonicity of $\frac{\lambda}{r+\lambda}[\Delta + \Delta_h] - (y + c)$ in y . The candidate is therefore the FIOS profile $(1, \underline{x})$. For this candidate to be an equilibrium, the fully-investing gender's corner must be incentive-compatible against the deviation-optimal value of a low type (the fully-investing side carries no low types on path, so matching is all-match on path and the low type's acceptance is set by optimality): this is $V_A(H; \underline{x}) - (1 + c) \geq \max\{V_A(L; \underline{x}), V_S(L; \underline{x})\}$, i.e. $\min\{\frac{\lambda}{r+\lambda}(\underline{x}\Delta + \Delta_h), G(\underline{x})\} \geq 1 + c$, which coincides with the all-match condition $\frac{\lambda}{r+\lambda}(\underline{x}\Delta + \Delta_h) \geq 1 + c$ exactly when $\lambda\underline{x}\eta \leq r\phi(L, L)$. Symmetrically, if $\hat{x}_m = 0$, the unique candidate is the NIOS profile $(0, \bar{x})$; here the all-low side is populated on path, and its L-L acceptance $\lambda\bar{x}\eta \leq r\phi(L, L)$ is imposed directly. Up to gender permutation, FIOS and NIOS are the only candidates.

To rule out coexistence of FIOS and NIOS, suppose both $(1, \underline{x})$ and $(0, \bar{x})$ are equilibria with $\underline{x}, \bar{x} \in (0, 1)$. Note $\bar{x} = \underline{x} - \frac{\lambda}{r+\lambda}\Delta > \underline{x}$ since $\Delta < 0$. The women's ICs deliver the contradiction directly: FIOS women's indifference gives $\frac{\lambda}{r+\lambda}[\Delta + \Delta_h] - c = \underline{x} \geq 0$. Since $\bar{x} > \underline{x}$ and $\Delta < 0$,

$$\frac{\lambda}{r+\lambda}[\Delta\bar{x} + \Delta_h] - c > \frac{\lambda}{r+\lambda}[\Delta + \Delta_h] - c = \underline{x} \geq 0,$$

where the strict inequality uses $\Delta < 0$ and $\bar{x} < 1$, so $\Delta\bar{x} > \Delta$. But this contradicts NIOS men's IC $\frac{\lambda}{r+\lambda}[\bar{x}\Delta + \Delta_h] \leq c$.¹⁹ Hence at most one asymmetric (AM) equilibrium exists.

When the symmetric equilibrium is SA, the asymmetric form must be FIOS. We show that under the hypothesis $\lambda\hat{x}_A\eta > r\phi(L, L)$ (the regime in which the symmetric equilibrium is SA), no NIOS equilibrium can exist. Combined with the argument above ruling out interior asymmetric profiles, this leaves FIOS as the only possible asymmetric form.

The NIOS profile $(0, \bar{x})$ requires the L-L acceptance for L-skilled men (the all-

¹⁹At the boundaries the two profiles degenerate to mirror images of one another: $\underline{x} = 0$ makes FIOS the profile $(1, 0)$ —one side fully invests, the other not at all—and $\bar{x} = 1$ makes NIOS the profile $(0, 1)$, the same full-specialization configuration with the genders swapped. They are then not distinct coexisting equilibria, and the interior argument above is what rules out genuine coexistence.

low side) facing women's pool \bar{x} :

$$\lambda\bar{x}\eta \leq r\phi(L, L).$$

But $\bar{x} = (1 - \frac{\lambda}{r+\lambda}\Delta)\hat{x}_A \geq \hat{x}_A$, since $1 - \frac{\lambda}{r+\lambda}\Delta \geq 1$ (as $\Delta \leq 0$). Therefore,

$$\lambda\bar{x}\eta \geq \lambda\hat{x}_A\eta > r\phi(L, L),$$

contradicting the NIOS L-L constraint. Hence NIOS does not exist under failure of SYM AM. The boundary case $\bar{x} \geq 1$ (NIOS profile $(0, 1)$) yields $\lambda\eta \leq r\phi(L, L)$ for the L-L constraint at women's all-H pool; but $\lambda\eta \geq \lambda\hat{x}_A\eta > r\phi(L, L)$ (since $\hat{x}_A \leq 1$), again a contradiction. □

A.3. Proof of Proposition 2

Proof of Proposition 2. We first show that ϕ admits no asymmetric AM equilibrium. By Proposition 1 Part 2, any such equilibrium must take FIOS or NIOS form, so it suffices to rule out each.

No NIOS. The NIOS profile $(0, \bar{x})$ requires men's IC at type 0, $\frac{\lambda}{r+\lambda}[\bar{x}\Delta + \Delta_h] \leq c$. Substituting $\bar{x} = \frac{\lambda}{r+\lambda}\Delta_h - c$ and rearranging gives

$$\left(\frac{\lambda}{r+\lambda}\Delta_h - c\right) \left(1 + \frac{\lambda}{r+\lambda}\Delta\right) \leq 0.$$

By (ii), $\frac{\lambda}{r+\lambda}\Delta_h - c > 0$; by (i), $1 + \frac{\lambda}{r+\lambda}\Delta > 0$. The product is therefore strictly positive, contradicting the inequality. Hence NIOS is not an equilibrium under ϕ .

No FIOS. The FIOS investing-side IC requires $\min\{\frac{\lambda}{r+\lambda}(\underline{x}\Delta + \Delta_h), G(\underline{x})\} \geq 1 + c$; it suffices to violate the all-match term $\frac{\lambda}{r+\lambda}[\underline{x}\Delta + \Delta_h] \geq 1 + c$ (a fortiori the deviation-optimal IC, being a minimum, then also fails). Substituting $\underline{x} = \frac{\lambda}{r+\lambda}(\Delta + \Delta_h) - c$ and rearranging gives

$$\left(\frac{\lambda}{r+\lambda}\Delta_h - c\right) \left(1 + \frac{\lambda}{r+\lambda}\Delta\right) \geq \left(1 - \frac{\lambda}{r+\lambda}\Delta\right) \left(1 + \frac{\lambda}{r+\lambda}\Delta\right).$$

Under (i), $1 + \frac{\lambda}{r+\lambda}\Delta > 0$, so this reduces to $\frac{\lambda}{r+\lambda}\Delta_h - c \geq 1 - \frac{\lambda}{r+\lambda}\Delta$, equivalently $\underline{x} = \frac{\lambda}{r+\lambda}(\Delta + \Delta_h) - c \geq 1$. But submodularity ($\Delta \leq 0$) gives $\underline{x} \leq \bar{x} = \frac{\lambda}{r+\lambda}\Delta_h - c < 1$ by (ii), a contradiction. Hence FIOS is not an equilibrium under ϕ either.

Combined with the proposition's premise, the unique equilibrium of ϕ is the symmetric AM equilibrium.

Now, let us turn to constructing $\hat{\phi}$ that exhibits a higher skill premium relative to ϕ and admits a FIOS equilibrium (alongside the symmetric equilibrium guaranteed by Proposition 1). In particular, we will choose $\hat{\phi}(H, H), \hat{\phi}(H, L), \hat{\phi}(L, H)$

larger than $\phi(H, H), \phi(H, L), \phi(L, H)$ in a way that supports a FIOS equilibrium. Let us rewrite the constraints of a FIOS equilibrium.

$$\begin{aligned} \frac{\lambda}{r+\lambda}(\Delta + \Delta_h) - c &\in (0, 1). \\ \frac{\lambda}{r+\lambda}(\underline{x}\Delta + \Delta_h) &\geq 1 + c \\ \frac{\lambda}{r+\lambda}(\underline{x}\hat{\phi}(H, H) + (1 - \underline{x})\hat{\phi}(H, L)) &\leq \hat{\phi}(H, L) \end{aligned}$$

Of the three constraints above, our construction will keep $\hat{\phi}(H, H) \approx \hat{\phi}(H, L)$, so the third (H-L acceptance) constraint is automatically satisfied and we focus on the first two.

We keep $\hat{\phi}(L, L) = \phi(L, L)$ unchanged and choose $\hat{\phi}(L, H) = a$ sufficiently larger than $\phi(L, L)$, with $\hat{\phi}(H, H) \approx \hat{\phi}(H, L) = b > a$. Then $\hat{\Delta} \approx \phi(L, L) - a$ and $\hat{\Delta}_h \approx b - \phi(L, L)$. As before, (AM: Invest) for a FIOS equilibrium under $\hat{\phi}$ rearranges to

$$\left(\frac{\lambda}{r+\lambda}\hat{\Delta}_h - c\right) \left(1 + \frac{\lambda}{r+\lambda}\hat{\Delta}\right) \geq \left(1 - \frac{\lambda}{r+\lambda}\hat{\Delta}\right) \left(1 + \frac{\lambda}{r+\lambda}\hat{\Delta}\right) \quad (5)$$

This requires $1 + \frac{\lambda}{r+\lambda}\hat{\Delta} \leq 0$, equivalently $\frac{\lambda}{r+\lambda}(a - \phi(L, L)) \geq 1$. As a increases, $|\hat{\Delta}|$ increases, so the inequality is satisfied for a sufficiently larger than $\phi(L, L)$. We also need $\hat{x} = \frac{\lambda}{r+\lambda}(\hat{\Delta} + \hat{\Delta}_h) - c \approx \frac{\lambda}{r+\lambda}(b - a) - c$ to lie in $(0, 1)$. Given any admissible a , $\hat{x} < 1$ puts an upper bound on b and $\hat{x} > 0$ puts a lower bound. Once these hold, (5) reduces, under $1 + \frac{\lambda}{r+\lambda}\hat{\Delta} \leq 0$, to $\hat{x} \leq 1$, which is exactly the upper bound on b . Finally, since the fully-investing side carries no low types on path, its corner IC is the deviation-optimal condition $\min\{\frac{\lambda}{r+\lambda}(\hat{x}\hat{\Delta} + \hat{\Delta}_h), G(\hat{x})\} \geq 1 + c$ (see the definition of the FIOS profile). The all-match term holds by the reduction just given; the SA term $G(\hat{x}) = \frac{\lambda}{r+\lambda}[\hat{x}\hat{\phi}(H, H) + (1 - \hat{x})\hat{\phi}(H, L)] - \lambda\hat{x}\hat{\phi}(L, H)/(r + \lambda\hat{x}) \approx \frac{\lambda}{r+\lambda}b - \lambda\hat{x}a/(r + \lambda\hat{x})$ exceeds $1 + c$ once b is taken large enough (with the gap $b - a$, hence \hat{x} , held fixed), which is compatible with the bounds above. Therefore, choosing a large enough, then b large enough with $b - a$ set so that $\hat{x} \in (0, 1)$, guarantees that a FIOS equilibrium exists under $\hat{\phi}$. This completes the construction. \square

A.4. Proof of Proposition 3

Proof of Proposition 3. Write $\hat{x}_A = (\frac{\lambda}{r+\lambda}\Delta_h - c)/(1 - \frac{\lambda}{r+\lambda}\Delta) =: N/D$, with $D = 1 - \frac{\lambda}{r+\lambda}\Delta > 0$ and $N = \frac{\lambda}{r+\lambda}\Delta_h - c = \hat{x}_A D > 0$. Differentiating, $\hat{x}'_A = \frac{\lambda}{r+\lambda}(D\Delta'_h + N\Delta')/D^2$, so $\hat{x}'_A \geq 0$ if and only if $\Delta'_h \geq -\hat{x}_A\Delta'$. Since $\hat{x}_A \in (0, 1)$ and $\Delta'_h \geq 0$: if $\Delta' \geq 0$ the right-hand side is non-positive; if $\Delta' < 0$ then $-\hat{x}_A\Delta' = \hat{x}_A|\Delta'| \leq |\Delta'| \leq \Delta'_h$. In either case $\hat{x}'_A \geq 0$. \square

A.5. Proof of Proposition 4

Proof of Proposition 4. Coexistence of FIOS at $(1, \underline{x})$ with interior $\underline{x} \in (0, 1)$ requires the men's investment IC at cutoff 1. The deviation-optimal form of this IC is $\min\{\frac{\lambda}{r+\lambda}(\Delta \underline{x} + \Delta_h), G(\underline{x})\} \geq 1 + c$ (see the definition of the FIOS profile); in particular it implies the all-match term $\frac{\lambda}{r+\lambda}(\Delta \underline{x} + \Delta_h) - c \geq 1$. Substituting the FIOS women's IC $\frac{\lambda}{r+\lambda}\Delta_h - c = \underline{x} - \frac{\lambda}{r+\lambda}\Delta$ into the latter gives $(1 + \frac{\lambda}{r+\lambda}\Delta)(\underline{x} - 1) \geq 0$; with $\underline{x} < 1$, this forces $1 + \frac{\lambda}{r+\lambda}\Delta \leq 0$, generically strict.²⁰ The analogous derivation for NIOS at $(0, \bar{x})$ with interior $\bar{x} > 0$ yields $\bar{x}(1 + \frac{\lambda}{r+\lambda}\Delta) \leq 0$, again forcing $1 + \frac{\lambda}{r+\lambda}\Delta < 0$.

SYM AM is fragile. At the interior fixed point (\hat{x}_A, \hat{x}_A) of Θ with $\text{BR} = \text{BR}_{AM}$, the symmetric IC gives $\frac{\lambda}{r+\lambda}(\Delta \hat{x}_A + \Delta_h) - c = \hat{x}_A$, so $\Theta_m(\hat{x}_A + \epsilon, \hat{x}_A - \epsilon) = \text{BR}_{AM}(\hat{x}_A - \epsilon) = \hat{x}_A - \frac{\lambda}{r+\lambda}\Delta \epsilon$. Symmetrically, $\Theta_w(\hat{x}_A + \epsilon, \hat{x}_A - \epsilon) = \hat{x}_A + \frac{\lambda}{r+\lambda}\Delta \epsilon$. Both deviations have magnitude $\frac{\lambda}{r+\lambda}|\Delta||\epsilon|$. With $\frac{\lambda}{r+\lambda}|\Delta| > 1$ (equivalent under $\Delta < 0$ to $1 + \frac{\lambda}{r+\lambda}\Delta < 0$), Definition 2 holds with $\zeta = \frac{\lambda}{r+\lambda}|\Delta| - 1 > 0$ and any δ small enough that the perturbations stay in $(0, 1)^2$.

FIOS is not fragile. At $(1, \underline{x})$, feasibility requires $\epsilon \leq 0$. The men's unclipped best response to $\underline{x} - \epsilon$ is

$$M(\epsilon) = \frac{\lambda}{r+\lambda}(\Delta(\underline{x} - \epsilon) + \Delta_h) - c = \underline{x}(1 + \frac{\lambda}{r+\lambda}\Delta) - \frac{\lambda}{r+\lambda}\Delta - \frac{\lambda}{r+\lambda}\Delta \epsilon,$$

using the FIOS women's IC. At $\epsilon = 0$, $M(0) - 1 = (1 + \frac{\lambda}{r+\lambda}\Delta)(\underline{x} - 1) = |1 + \frac{\lambda}{r+\lambda}\Delta|(1 - \underline{x}) > 0$. By continuity, $M(\epsilon) > 1$ for $|\epsilon| < \delta_F := |1 + \frac{\lambda}{r+\lambda}\Delta|(1 - \underline{x})/(\frac{\lambda}{r+\lambda}|\Delta|)$, so $\Theta_m = 1$ throughout this neighborhood. Hence $|\Theta_m(1 + \epsilon, \underline{x} - \epsilon) - 1| = 0$ for every $\epsilon \in (-\delta_F, 0]$, violating the first inequality of Definition 2 for any positive ζ . FIOS is not fragile.

NIOS is not fragile. The symmetric argument at $(0, \bar{x})$, with $\epsilon \geq 0$ the feasible direction. Substituting the NIOS women's IC $\frac{\lambda}{r+\lambda}\Delta_h - c = \bar{x}$ gives the unclipped men's best response $m(\epsilon) = \bar{x}(1 + \frac{\lambda}{r+\lambda}\Delta) - \frac{\lambda}{r+\lambda}\Delta \epsilon$. At $\epsilon = 0$, $m(0) = \bar{x}(1 + \frac{\lambda}{r+\lambda}\Delta) < 0$ with slack $\bar{x}|1 + \frac{\lambda}{r+\lambda}\Delta|$; by continuity $m(\epsilon) < 0$ for $\epsilon \in [0, \delta_N]$ where $\delta_N := \bar{x}|1 + \frac{\lambda}{r+\lambda}\Delta|/(\frac{\lambda}{r+\lambda}|\Delta|)$, so $\Theta_m = 0$ throughout, the first inequality fails, and NIOS is not fragile. \square

A.6. Derivation of the payoffs in Example 2

This subsection derives the four match payoffs $\phi(H, H)$, $\phi(H, L)$, $\phi(L, H)$, $\phi(L, L)$ reported in Section 6 from the household game. Recall the parameters: $K = 8$, $\alpha = 0.6$, $t_l = 2$, $t_h \in \{3, 5\}$, and the utility of member i in a match with effort

²⁰The not-fragile conclusion below is likewise unaffected by the correction: it requires only that the fully-investing side's best response stay pinned at the corner 1 in a neighborhood of the perturbation, which holds because the deviation-optimal IC holds strictly at \underline{x} (and hence, by continuity, nearby).

vector (e_m, e_w) is

$$U_i(e_m, e_w) = K[\alpha \log(t \cdot e) + (1 - \alpha) \log(2 - e_m - e_w)] - \frac{1}{2}(1 - e_i)^2,$$

where $t \cdot e = t_m e_m + t_w e_w$ is household income.

First-order conditions. Differentiating U_i with respect to e_i ,

$$\frac{\partial U_i}{\partial e_i} = K\alpha \frac{t_i}{t \cdot e} - K(1 - \alpha) \frac{1}{2 - e_m - e_w} + (1 - e_i).$$

A non-cooperative interior equilibrium requires this to vanish for both $i \in \{m, w\}$.

Symmetric matches: $(s, s) \in \{(L, L), (H, H)\}$. When $t_m = t_w =: t$, the model is symmetric in the two members, so $e_m = e_w =: e^*$ and the FOC reduces to

$$\frac{K\alpha}{2e^*} - \frac{K(1 - \alpha)}{2(1 - e^*)} + (1 - e^*) = 0.$$

The wage t cancels, so e^* is the same for (L, L) and (H, H) . With $K = 8$ and $\alpha = 0.6$, the equation becomes $\frac{2.4}{e^*} - \frac{1.6}{1 - e^*} + (1 - e^*) = 0$, whose unique root in $(0, 1)$ is $e^* \approx 0.6222$. Substituting,

$$\begin{aligned} \phi(s, s) &= K[\alpha \log(2te^*) + (1 - \alpha) \log(2 - 2e^*)] - \frac{1}{2}(1 - e^*)^2 \\ &= 8[0.6 \log(2t \cdot 0.6222) + 0.4 \log(0.7556)] - \frac{1}{2}(0.3778)^2. \end{aligned}$$

Numerically, $\phi(L, L) = 3.4085$, and $\phi(H, H) = 5.3547$ at $t_h = 3$ and $\phi(H, H) = 7.8067$ at $t_h = 5$.

Asymmetric matches: $(s_m, s_w) = (H, L)$. With $t_m = t_h$ and $t_w = t_l$, the wage gap pushes the high-wage member to the boundary $e_m = 1$. The KKT condition for member m at $e_m = 1$ requires $\partial U_m / \partial e_m|_{(1, e_w)} \geq 0$, which at any candidate $(1, e_w)$ reduces to $K\left(\alpha \frac{t_h}{t_h + t_l e_w} - (1 - \alpha) \frac{1}{1 - e_w}\right) \geq 0$; for the parameters considered this holds whenever $e_w \leq 0.23$ (at $t_h = 3$) or $e_w \leq 0.26$ (at $t_h = 5$)—in particular at the Nash value of e_w identified below ($e_w \approx 0.149$ and $e_w = 0$, respectively), so $e_m = 1$ is member m 's best response there.

The Nash structure then depends on member w 's FOC at $e_m = 1$.

At $t_h = 5$: the parameters satisfy $\alpha t_l / t_h < (1 - \alpha) - 1/K$ (here $\alpha t_l / t_h = 0.24 < 0.275$), so at $(1, 0)$ member w 's marginal payoff $K[\alpha t_l / t_h - (1 - \alpha)] + 1 = -0.28 < 0$. Member w is also at the boundary, and the Nash equilibrium is the complete corner $(e_m, e_w) = (1, 0)$. Substituting,

$$\begin{aligned} \phi(H, L) &= K\alpha \log(t_h) - \frac{1}{2}(1 - 1)^2 = 4.8 \log(5) = 7.7253, \\ \phi(L, H) &= K\alpha \log(t_h) - \frac{1}{2}(1 - 0)^2 = 4.8 \log(5) - 0.5 = 7.2253. \end{aligned}$$

At $t_h = 3$: the parameters satisfy $\alpha t_l / t_h > (1 - \alpha) - 1/K$ (since $\alpha t_l / t_h = 0.4 >$

0.275), so at $(1, 0)$ member w 's marginal payoff $K[\alpha t_l/t_h - (1 - \alpha)] + 1 = +1 > 0$. Member w strictly prefers to raise e_w above zero, so $(1, 0)$ is *not* Nash. The Nash equilibrium is a partial corner: $e_m = 1$ and e_w solves member w 's interior FOC,

$$K\alpha \frac{t_l}{t_h + t_l e_w} - K(1 - \alpha) \frac{1}{1 - e_w} + (1 - e_w) = 0.$$

Substituting $K = 8$, $\alpha = 0.6$, $t_h = 3$, $t_l = 2$ and solving numerically yields $e_w \approx 0.1492$. The resulting match payoffs are

$$\begin{aligned} \phi(H, L) &= 8[0.6 \log(3 + 2 \cdot 0.1492) + 0.4 \log(1 - 0.1492)] \\ &\quad - \frac{1}{2}(1 - 1)^2 \approx 5.2114, \\ \phi(L, H) &= 8[0.6 \log(3 + 2 \cdot 0.1492) + 0.4 \log(1 - 0.1492)] \\ &\quad - \frac{1}{2}(1 - 0.1492)^2 \approx 4.8495. \end{aligned}$$

The household therefore specializes only partially at $t_h = 3$: the high-wage member supplies all of his time to the market, while the low-wage member supplies a small but positive amount of market labor. At $t_h = 5$, by contrast, complete specialization $(1, 0)$ is the Nash outcome.

Verification of the equilibrium claims. With $r = \lambda = 1$ (so $\frac{\lambda}{r+\lambda} = 1/2$) and $c = 0.25$:

For $t_h = 3$: $\Delta = -1.2977$, $\Delta_h = 1.8029$, giving

$$\hat{x}_A = \frac{\frac{\lambda}{r+\lambda} \Delta_h - c}{1 - \frac{\lambda}{r+\lambda} \Delta} = \frac{0.5 \cdot 1.8029 - 0.25}{1 + 0.5 \cdot 1.2977} \approx 0.3951.$$

The L-L acceptance constraint at \hat{x}_A is $\lambda \hat{x}_A (\phi(L, H) - \phi(L, L)) \approx 0.5693 < 3.4085 = r\phi(L, L)$, so (\hat{x}_A, \hat{x}_A) is the unique equilibrium.

For $t_h = 5$: $\Delta = -3.7354$, $\Delta_h = 4.3168$, giving $\hat{x}_A = 0.6655$ from the same formula. Here $1 + \frac{\lambda}{r+\lambda} \Delta = -0.8677 < 0$, so the FIOS investment IC is feasible; meanwhile the L-L acceptance constraint at \hat{x}_A reads $\lambda \hat{x}_A (\phi(L, H) - \phi(L, L)) = 0.6655 \cdot 3.8168 = 2.5401 < 3.4085 = r\phi(L, L)$, so (\hat{x}_A, \hat{x}_A) remains a SYM AM equilibrium. Alongside it, the FIOS cutoff is

$$\underline{x} = \frac{\lambda}{r + \lambda} (\Delta + \Delta_h) - c = 0.5 \cdot 0.5814 - 0.25 = 0.0407,$$

interior and far below \hat{x}_A . The investing side's incentive constraint holds, $\min\{\frac{\lambda}{r+\lambda}(\underline{x}\Delta + \Delta_h), G(\underline{x})\} = \min\{2.08, 3.58\} \geq 1 + c$, so the FIOS profile $(1, \underline{x})$ coexists with the SYM AM equilibrium (\hat{x}_A, \hat{x}_A) .

A.7. Proof of Proposition 5

Proof of Proposition 5. Throughout, let $\phi(s, s'; t_h)$ denote the match payoff in the PIH game with skill pair (s, s') , where $t_l > 0$ is fixed and $t_h \geq t_l$ varies. We abbreviate $\phi(H, H)(t_h) = \phi(H, H; t_h)$, $\phi(H, L)(t_h) = \phi(H, L; t_h)$, $\phi(L, H)(t_h) = \phi(L, H; t_h)$, and use $\phi(L, L) = \phi(L, L; t_h)$ which is independent of t_h (the (L, L) match has both members at wage t_l , which does not involve t_h). Define

$$\begin{aligned}\Delta(t_h) &:= \phi(H, H)(t_h) + \phi(L, L) - \phi(H, L)(t_h) - \phi(L, H)(t_h), \\ \Delta_h(t_h) &:= \phi(H, L)(t_h) - \phi(L, L), \\ \hat{x}_A(t_h) &:= \frac{\frac{\lambda}{r+\lambda}\Delta_h(t_h) - c}{1 - \frac{\lambda}{r+\lambda}\Delta(t_h)}, \\ \underline{x}(t_h) &:= \frac{\lambda}{r + \lambda}[\Delta(t_h) + \Delta_h(t_h)] - c, \\ \bar{x}(t_h) &:= \frac{\lambda}{r + \lambda}\Delta_h(t_h) - c.\end{aligned}$$

Step 1: Continuity of the household game's Nash equilibrium.

Fix any $(t_m, t_w) \in \mathbb{R}_+^2$. By Definition 3, the game has a unique Nash equilibrium $e^*(t_m, t_w) \in [0, 1]^2$. Because each U_i is continuous in (e, t_m, t_w) and the strategy space $[0, 1]^2$ is compact, the set of Nash equilibria has closed graph in (t_m, t_w) ; being single-valued by uniqueness, the equilibrium $e^*(\cdot)$ is therefore continuous in (t_m, t_w) . Hence $\phi(s, s'; t_h)$ is continuous in t_h , and so are Δ , Δ_h , \hat{x}_A , \underline{x} , \bar{x} .

Step 2: Asymptotic divergence of Φ .

Define

$$\Phi(t_h) := \lambda \hat{x}_A(t_h) (\phi(L, H)(t_h) - \phi(L, L)) - r \phi(L, L).$$

The symmetric L-L acceptance (**AM: L-L**) holds at the symmetric profile iff $\Phi(t_h) \leq 0$. We claim $\Phi(t_h) \rightarrow +\infty$ as $t_h \rightarrow \infty$.

By assumption (C2), $\liminf_{t_h \rightarrow \infty} \hat{x}_A(t_h) =: x^* > 0$, so $\hat{x}_A(t_h) \geq x^*/2$ for t_h sufficiently large. Furthermore, $\phi(L, H)(t_h) \rightarrow \infty$: at the household Nash equilibrium (e_m^*, e_w^*) at wages (t_h, t_l) , the best-response inequality for member m (using $e_m = 1/2$ as a feasible deviation) gives

$$U_m(e_m^*, e_w^*; t_h, t_l) \geq \psi\left(\frac{t_h}{2} + t_l e_w^*, \frac{3}{2} - e_w^*\right) - g(1/2) \geq \psi(t_h/2, 1/2) - g(1/2),$$

which tends to ∞ by the unboundedness of ψ in its first argument. Since $\phi(L, H)(t_h) = U_w(e_m^*, e_w^*) = U_m(\cdot) + [g(1 - e_m^*) - g(1 - e_w^*)]$ and the bracketed cost difference is bounded (g is bounded on $[0, 1]$), $\phi(L, H)(t_h) \rightarrow \infty$. Combining with $\hat{x}_A(t_h) \geq x^*/2$ at large t_h , $\Phi(t_h) \rightarrow \infty$.

Since $\hat{x}_A(t_l) \leq 0$ (the unclipped formula gives $-c < 0$) so $\Phi(t_l) \leq -r\phi(L, L) <$

0, and Φ is continuous (Step 1), the threshold

$$t_\Phi := \inf\{t_h \geq t_l : \Phi(t_h) \geq 0\}$$

is finite. Under the additional hypothesis $\Delta'_h \geq |\Delta'|$, $\hat{x}_A(t_h)$ is non-decreasing (Step 5) and $\eta(t_h) = \phi(L, H)(t_h) - \phi(L, L)$ is non-decreasing (household income at the (L, H) match rises in t_h), so $\Phi(t_h) = \lambda \hat{x}_A(t_h) \eta(t_h) - r\phi(L, L)$ is non-decreasing and crosses zero exactly once. Thus t_Φ marks a single AM-to-SA transition of the symmetric equilibrium: for $t_h < t_\Phi$ it is SYM AM, and for $t_h > t_\Phi$ it is SYM SA. (The coexistence claim does not rely on this: a symmetric equilibrium exists at every t_h by Proposition 1, whether in AM or SA form.)

Step 3: Defining t^* .

By (C1), $\Delta(t_h) \rightarrow -\infty$, so $1 + \frac{\lambda}{r+\lambda}\Delta(t_h) \leq 0$ for sufficiently large t_h . Define $t_h^\dagger := \inf\{t_h \geq t_l : 1 + \frac{\lambda}{r+\lambda}\Delta(t_h) \leq 0\}$; this is finite by (C1) and Step 1. By the single-crossing condition (C4), $1 + \frac{\lambda}{r+\lambda}\Delta(t_h) \leq 0$ holds for *all* $t_h \geq t_h^\dagger$, not merely at the infimum—so the FIOS investment feasibility, once attained, persists. By (C3), there exists $T_2 < \infty$ such that for all $t_h \geq T_2$, $\underline{x}(t_h) \in (0, 1)$ and the FIOS profile $(1, \underline{x}(t_h))$ satisfies (AM: H-L) on path. Define

$$t^* := \max\{t_h^\dagger, T_2\} \in (t_l, \infty).$$

That is, t^* is a threshold above which the three FIOS-existence conditions hold simultaneously: feasibility of the FIOS investment IC ($1 + \frac{\lambda}{r+\lambda}\Delta \leq 0$), interior FIOS cutoff ($\underline{x} \in (0, 1)$), and FIOS H-L acceptance on path. t^* should not be confused with t_Φ from Step 2: t^* marks the threshold above which FIOS becomes feasible, while t_Φ marks the AM-to-SA transition of the symmetric equilibrium. The two may sit on either side of each other in general.

Step 4: At $t_h > t^*$, a FIOS equilibrium exists alongside the symmetric one.

Fix $t_h > t^*$. The FIOS profile is $(1, \underline{x}(t_h))$. By the definition of t^* , $\underline{x}(t_h) \in (0, 1)$, so the cutoff is interior. The women's investment IC at \underline{x} is the indifference relation $\frac{\lambda}{r+\lambda}[\Delta + \Delta_h] = \underline{x} + c$, satisfied by definition of \underline{x} . The men's investment IC at type 1 is the deviation-optimal condition $\min\{\frac{\lambda}{r+\lambda}(\underline{x}\Delta + \Delta_h), G(\underline{x})\} \geq 1 + c$ (the fully-investing side carries no low types on path, so matching is all-match on path and the low type's acceptance is set by optimality; see the definition of the FIOS profile). The all-match term $\frac{\lambda}{r+\lambda}[\underline{x}\Delta + \Delta_h] \geq 1 + c$ reduces, using $1 + \frac{\lambda}{r+\lambda}\Delta \leq 0$ for all $t_h > t^*$ (Step 3, by (C4)), to $\underline{x} \leq 1$, which holds. For the SA term, $G(\underline{x}) = \frac{\lambda}{r+\lambda}[\underline{x}\phi(H, H) + (1 - \underline{x})\phi(H, L)] - \lambda \underline{x} \phi(L, H)/(r + \lambda \underline{x}) \rightarrow \infty$ as $t_h \rightarrow \infty$ (the high-type value diverges by the argument of Step 2 while the bracketed coefficient on the divergent payoffs is positive for $\underline{x} < 1$), so $G(\underline{x}) \geq 1 + c$ for t_h large; raising T_2 if necessary, both terms hold for $t_h > t^*$. The H-L acceptance constraint (AM: H-L) is satisfied on path by the definition of t^* , and low-skill women meet

only high-skill men on path. Hence FIOS is an equilibrium. By Proposition 1 Part 1, a symmetric equilibrium also exists (SYM AM at cutoff \hat{x}_A if $\Phi(t_h) \leq 0$, i.e., $t_h \leq t_\Phi$; SYM SA at cutoff \hat{x}_{SA} otherwise). The two coexist.

Step 5: Within-regime comparative statics.

Monotonicity of $\hat{x}_A(t_h)$ in the symmetric AM regime $[t_l, t_\Phi]$ (under the additional hypothesis $\Delta'_h \geq |\Delta'|$). On this regime, the cutoff is given by the closed form

$$\hat{x}_A(t_h) = \frac{\frac{\lambda}{r+\lambda}\Delta_h(t_h) - c}{1 - \frac{\lambda}{r+\lambda}\Delta(t_h)} = \frac{N(t_h)}{D(t_h)},$$

where $N(t_h) := \frac{\lambda}{r+\lambda}\Delta_h(t_h) - c$ and $D(t_h) := 1 - \frac{\lambda}{r+\lambda}\Delta(t_h)$. Under the smoothness of the PIH game (Step 1), both N and D are continuously differentiable in t_h . By the quotient rule,

$$\frac{d\hat{x}_A}{dt_h} = \frac{N'(t_h)D(t_h) - N(t_h)D'(t_h)}{D(t_h)^2} = \frac{\frac{\lambda}{r+\lambda}[\Delta'_h(t_h)D(t_h) + \Delta'(t_h)N(t_h)]}{D(t_h)^2}.$$

Within the symmetric regime, $\hat{x}_A(t_h) = N(t_h)/D(t_h) \in [0, 1)$, which (combined with $D > 0$) gives $0 \leq N(t_h) \leq D(t_h)$. Under the hypothesis $\Delta'_h(t_h) \geq |\Delta'(t_h)|$ (for the Cobb-Douglas case this holds with equality at the full corner $t_h \gtrsim 4.4$, where $\Delta'_h(t_h) = -\Delta'(t_h) = K\alpha/t_h$, and strictly for smaller t_h), and using $\Delta'(t_h) \geq -|\Delta'(t_h)|$,

$$\begin{aligned} \Delta'_h(t_h)D(t_h) + \Delta'(t_h)N(t_h) &\geq \Delta'_h(t_h)D(t_h) - |\Delta'(t_h)|N(t_h) \\ &\geq |\Delta'(t_h)|D(t_h) - |\Delta'(t_h)|N(t_h) \\ &= |\Delta'(t_h)|[D(t_h) - N(t_h)] \geq 0. \end{aligned}$$

Hence $d\hat{x}_A/dt_h \geq 0$, and \hat{x}_A is non-decreasing in t_h on the AM regime $[t_l, t_\Phi]$.

The FIOS cutoff lies below the symmetric one. The SYM AM indifference condition $\frac{\lambda}{r+\lambda}(\Delta\hat{x}_A + \Delta_h) - c = \hat{x}_A$ and the FIOS cutoff $\underline{x} = \frac{\lambda}{r+\lambda}(\Delta + \Delta_h) - c$ together give $\hat{x}_A - \underline{x} = \frac{\lambda}{r+\lambda}\Delta(\hat{x}_A - 1) = \frac{\lambda}{r+\lambda}|\Delta|(1 - \hat{x}_A)$, strictly positive whenever $\Delta < 0$ and $\hat{x}_A < 1$. Hence, wherever a FIOS equilibrium exists alongside SYM AM, its cutoff lies strictly below the symmetric cutoff. Finally $\underline{x}(t_h) = \frac{\lambda}{r+\lambda}[\Delta(t_h) + \Delta_h(t_h)] - c$ is non-decreasing in t_h , since $d\underline{x}/dt_h = \frac{\lambda}{r+\lambda}(\Delta'_h + \Delta') \geq \frac{\lambda}{r+\lambda}(\Delta'_h - |\Delta'|) \geq 0$ under the slope hypothesis.

Discontinuous appearance of FIOS at $t_h = t^$.* The FIOS cutoff is $\underline{x}(t_h) = \frac{\lambda}{r+\lambda}[\Delta(t_h) + \Delta_h(t_h)] - c = \frac{\lambda}{r+\lambda}[\phi(H, H)(t_h) - \phi(L, H)(t_h)] - c$, which stays small as the premium rises—in the Cobb-Douglas case $\phi(H, H) - \phi(L, H)$ becomes independent of t_h once the (H, L) household fully specializes ($t_h \gtrsim 4.4$), where the $\log t_h$ terms cancel, so \underline{x} levels off near 0.04. At $t_h = t^* \approx 3.51$, the symmetric cutoff is $\hat{x}_A(t^*) \approx 0.51$, far above $\underline{x} \approx 0.03$. The discontinuity is not in any single equilibrium cutoff (the symmetric trajectory is continuous, as shown above) but

in the equilibrium *set*: a new FIOS equilibrium emerges at t^* , far from the symmetric one. If the population coordinates on the FIOS equilibrium for t_h slightly above t^* , the investment profile jumps discretely from $(\hat{x}_A, \hat{x}_A) \approx (0.51, 0.51)$ to $(1, \underline{x}) \approx (1, 0.03)$: one gender goes from \hat{x}_A to full investment, the other drops from \hat{x}_A to $\underline{x} \ll \hat{x}_A$.

Steps 3–4 establish that a FIOS equilibrium coexists with the symmetric equilibrium for all $t_h > t^*$; and since both asymmetric forms require $1 + \frac{\lambda}{r+\lambda}\Delta < 0$ (the FIOS and NIOS investment ICs), no asymmetric equilibrium exists while $1 + \frac{\lambda}{r+\lambda}\Delta(t_h) > 0$. Step 5 establishes the comparative statics on the SYM AM range: the symmetric cutoff \hat{x}_A is continuous and non-decreasing in t_h , the FIOS cutoff satisfies $\underline{x} < \hat{x}_A$ wherever it exists, and $\underline{x}(t_h) = \frac{\lambda}{r+\lambda}(\Delta + \Delta_h) - c$ is itself non-decreasing (since $d\underline{x}/dt_h = \frac{\lambda}{r+\lambda}(\Delta'_h + \Delta') \geq 0$ whenever $\Delta'_h \geq |\Delta'|$). The threshold t^* satisfies $t_l < t^* < \infty$. \square

Online Appendix

B. PAM equilibrium

This section presents the formal analysis of the PAM (positive assortative matching) equilibrium class, referenced from the Remark following Proposition 1.

Setup

A PAM equilibrium is one in which high-skilled agents reject low-skilled partners ($\alpha^\gamma(H) = 0$) while low-skilled agents accept low-skilled partners ($\alpha^\gamma(L) = 1$); since the high types reject them, low types are matched only with low types. The players' payoffs (with some abuse of notation), denoted by $V_P^\gamma(\cdot)$, are:

$$\begin{aligned} V_P^\gamma(H; \hat{X}) &= \frac{\lambda \hat{x}_{\gamma'} \phi(H, H)}{r + \lambda \hat{x}_{\gamma'}}, & (\text{PAM: Values}) \\ V_P^\gamma(L; \hat{X}) &= \frac{\lambda(1 - \hat{x}_{\gamma'}) \phi(L, L)}{r + \lambda(1 - \hat{x}_{\gamma'})}. \end{aligned}$$

For such an equilibrium to exist with an interior \hat{X} , the IC constraints are:

$$\begin{aligned} V_P^\gamma(H; \hat{X}) - C(\hat{x}_\gamma) &= V_P^\gamma(L; \hat{X}) & (\text{PAM: Invest}) \\ V_P^\gamma(H; \hat{X}) &\geq \phi(H, L) & (\text{PAM: Assortativity}) \end{aligned}$$

Here, the investment IC checks for the indifference of the cutoff type \hat{x}_γ , while the assortativity constraint ensures that a high-skill person prefers to wait for another high-skill partner than to match with a low-skill one.

PROPOSITION 7: *No PAM exists if $\lambda[\phi(H, H) - \phi(H, L)] < r\phi(H, L)$. Whenever a PAM exists, it is symmetric, i.e., $\hat{x}_m = \hat{x}_w$.*

Proof of Proposition 7. From (PAM: Assortativity), we have that

$$\frac{\lambda \hat{x}_{\gamma'} \phi(H, H)}{r + \lambda \hat{x}_{\gamma'}} \geq \phi(H, L)$$

The LHS is increasing in $\hat{x}_{\gamma'}$, and hence, is maximized at 1. Substituting $\hat{x}_{\gamma'} = 1$, we get that $V^\gamma(H; \hat{X})$ is maximized at $\frac{\lambda \phi(H, H)}{r + \lambda}$. Thus, if $\lambda[\phi(H, H) - \phi(H, L)] < r\phi(H, L)$, then (PAM: Assortativity) cannot be satisfied for any \hat{X} establishing the first part of the Proposition.

For the symmetry of a PAM whenever it exists, suppose $\hat{X} = (\hat{x}_m, \hat{x}_w)$ is a PAM equilibrium with $\hat{x}_m \neq \hat{x}_w$. Without loss of generality, $\hat{x}_m > \hat{x}_w$, which

forces $\hat{x}_m > 0$ and $\hat{x}_w < 1$. The investment IC (**PAM: Invest**) at the cutoffs (in its boundary-inclusive form) gives,

$$\begin{aligned} V_P^m(H; \hat{X}) - V_P^m(L; \hat{X}) &\geq C(\hat{x}_m) \quad (\text{type } \hat{x}_m > 0 \text{ invests}), \\ V_P^w(H; \hat{X}) - V_P^w(L; \hat{X}) &\leq C(\hat{x}_w) \quad (\text{type just above } \hat{x}_w < 1 \text{ does not}). \end{aligned}$$

Subtracting the second from the first,

$$[V_P^m(H; \hat{X}) - V_P^m(L; \hat{X})] - [V_P^w(H; \hat{X}) - V_P^w(L; \hat{X})] \geq C(\hat{x}_m) - C(\hat{x}_w) > 0,$$

where the final inequality uses C strictly increasing and $\hat{x}_m > \hat{x}_w$. On the other hand, from (**PAM: Values**),

$$\begin{aligned} & \left[V_P^m(H; \hat{X}) - V_P^m(L; \hat{X}) \right] - \left[V_P^w(H; \hat{X}) - V_P^w(L; \hat{X}) \right] \\ &= \left[\frac{\lambda \hat{x}_w \phi(H, H)}{r + \lambda \hat{x}_w} - \frac{\lambda(1 - \hat{x}_w) \phi(L, L)}{r + \lambda(1 - \hat{x}_w)} \right] \\ & \quad - \left[\frac{\lambda \hat{x}_m \phi(H, H)}{r + \lambda \hat{x}_m} - \frac{\lambda(1 - \hat{x}_m) \phi(L, L)}{r + \lambda(1 - \hat{x}_m)} \right] = h(\hat{x}_w) - h(\hat{x}_m), \end{aligned}$$

where $h(x) := \frac{\lambda x \phi(H, H)}{r + \lambda x} - \frac{\lambda(1-x) \phi(L, L)}{r + \lambda(1-x)}$ is strictly increasing. Since $\hat{x}_m > \hat{x}_w$, $h(\hat{x}_w) - h(\hat{x}_m) < 0$, contradicting the previous inequality. \square

C. Asymmetric equilibria under supermodular payoffs

The main analysis maintains submodularity ($\Delta \leq 0$), which is what ties the emergence of asymmetric equilibria to a rising skill premium (Section 6). It is *not*, however, what makes asymmetric equilibria possible or robust. We show here, by example, that a *supermodular* economy ($\Delta > 0$) can admit an asymmetric equilibrium that is the unique survivor of the fragility test of Definition 2. The asymmetry is of the semi-assortative (SA) type and is sustained by a *search channel*—a low-skilled agent’s pool-dependent value of holding out for a high-skilled partner—rather than by household complementarity, which is absent when payoffs are supermodular.

Recall from Section 5 the SA best response $\text{BR}_{SA}(y) = \max\{0, \min\{1, G(y) - c\}\}$, with G the value-of-investing gap (1). A *symmetric* SA equilibrium is a fixed point $\hat{x} = \text{BR}_{SA}(\hat{x})$; an *asymmetric* SA equilibrium is a two-cycle: a pair (\hat{x}_m, \hat{x}_w) with $\hat{x}_m = \text{BR}_{SA}(\hat{x}_w)$ and $\hat{x}_w = \text{BR}_{SA}(\hat{x}_m)$, $\hat{x}_m \neq \hat{x}_w$, both cutoffs interior and in the SA region ($\lambda \hat{X} \eta > r \phi(L, L)$). Differentiating (1),

$$G'(y) = \frac{\lambda}{r + \lambda} (\phi(H, H) - \phi(H, L)) - \frac{\lambda r \phi(L, H)}{(r + \lambda y)^2}.$$

The second term is the search channel: a low type matches only with high-skilled partners, at rate λy , so the value of remaining low, $\lambda y \phi(L, H)/(r + \lambda y)$, climbs steeply with the opposite-side pool y . Where this channel is strong, G is steeply decreasing—*whatever the sign of Δ* .

Applying Definition 2 to Θ with $\text{BR} = \text{BR}_{SA}$ yields a slope criterion. At a symmetric SA equilibrium (\hat{x}, \hat{x}) , the perturbation $(\hat{x} + \epsilon, \hat{x} - \epsilon)$ maps, to first order, to $(\hat{x} - G'(\hat{x})\epsilon, \hat{x} + G'(\hat{x})\epsilon)$, so the equilibrium is fragile iff $|G'(\hat{x})| > 1$. At an asymmetric SA equilibrium (\hat{x}_m, \hat{x}_w) the same perturbation maps to $(\hat{x}_m - G'(\hat{x}_w)\epsilon, \hat{x}_w + G'(\hat{x}_m)\epsilon)$, so it is fragile iff *both* $|G'(\hat{x}_m)| > 1$ and $|G'(\hat{x}_w)| > 1$, and robust as soon as one cutoff has $G' > -1$ (equivalently $|G'| < 1$), where that side absorbs the perturbation.

EXAMPLE 3: Let $r = \lambda = 1$, $c = 2.875$, $\phi(H, H) = 13$, $\phi(H, L) = 7$, $\phi(L, H) = 6$, and $\phi(L, L) = \frac{1}{4}$. These payoffs are supermodular, $\Delta = 13 + \frac{1}{4} - 7 - 6 = \frac{1}{4} > 0$. The economy has a unique symmetric equilibrium, which is SA and fragile, together with an asymmetric SA equilibrium (and its mirror image) that is not fragile; the asymmetric SA equilibrium is the unique survivor of the fragility test. ■

The payoffs satisfy Assumption 1 and the no-PAM condition (\star), as $\lambda(\phi(H, H) - \phi(H, L)) = 6 < 7 = r\phi(H, L)$. Because $1 + \frac{\lambda}{r+\lambda}\Delta = \frac{9}{8} > 0$, no FIOS or NIOS equilibrium exists (Proposition 6 rules out asymmetric AM equilibria under supermodular payoffs). The symmetric equilibrium is therefore SA: the all-match candidate cutoff $\hat{x}_A = \frac{4}{7}$ would require low types to accept low partners there, $\lambda\hat{x}_A\eta \leq r\phi(L, L)$, which fails ($\frac{4}{7} \cdot \frac{23}{4} = \frac{23}{7} \approx 3.29 > \frac{1}{4}$). The unique symmetric SA cutoff is $\hat{x}_{SA} \approx 0.212$, where $G'(\hat{x}_{SA}) \approx -1.09$, so by the criterion above it is *fragile*.

Alongside it lies an asymmetric SA equilibrium at $(\hat{x}_m, \hat{x}_w) = (\frac{1}{3}, \frac{1}{8})$ and its mirror image. Indeed $G(\frac{1}{3}) = 3$, so $\text{BR}_{SA}(\frac{1}{3}) = 3 - 2.875 = \frac{1}{8}$, and $G(\frac{1}{8}) = \frac{77}{24}$, so $\text{BR}_{SA}(\frac{1}{8}) = \frac{77}{24} - \frac{23}{8} = \frac{1}{3}$; both cutoffs exceed $x^* = \frac{1}{23}$, so on each side low types reject low partners, as an SA configuration requires. Here the slope $G'(\frac{1}{3}) = -\frac{3}{8} > -1$, so this equilibrium is robust. As the symmetric SA and the asymmetric SA (with its mirror) exhaust the equilibrium set, the asymmetric SA equilibrium is the unique survivor of the fragility test—in a supermodular economy.

The asymmetric equilibrium arises from the benefit to a low type of holding out for a high-skilled partner. The example thus shows that neither the existence nor the robustness of asymmetric equilibria requires submodularity. I focus on submodular payoffs primarily because they connect most readily with the comparative statics in the skill premium, and because the microfoundation with Cobb-Douglas preferences delivers submodular payoffs.