Risky Matching

Hector Chade* and Ilse Lindenlaub†

October 2018

Abstract

This paper analyzes the effects of changes in risk and innate heterogeneity in an equilibrium matching framework with pre-match investment. In the model, workers that differ in ability can costly invest in education before matching with heterogenous firms. At the investment stage, workers are uncertain about how skilled they will turn out (idiosyncratic risk) and also about the prevailing state of the labor market at the time of employment (aggregate risk). We derive conditions on primitives – match output function and risk attitudes – for stochastically better or more variable risk and heterogeneity to induce more investment, and also show how this affects equilibrium matching and wages. We then provide three illustrations of the economic relevance of our results, dealing with (*i*) the causes of rising household income inequality; (*ii*) the effects of social insurance policies under endogenous investment response; and with (*iii*) how this model can help disentangle changes in risk versus changes in heterogeneity in the data.

Keywords. Matching, Investment, Changes in Risk, Comparative Statics, Sorting.

*Arizona State University, `hector.chade@asu.edu`

†Yale University, `ilse.lindenlaub@yale.edu`

We are grateful to Katka Borovickova, Andreas Kleiner, Natalia Kovrijnykh, Alejandro Manelli, Costas Meghir, Giuseppe Moscarini, Larry Samuelson, and Eddie Schlee, as well as seminar participants at Collegio Carlo Alberto, 2015 SED–Warsaw, 2016 Cornell-Penn State Macroeconomic Conference, Federal Reserve Bank of St. Louis, 2017 Search and Matching Conference in Drexel University, Notre Dame, and 2017 Search and Matching Conference at Barcelona GSE Summer Forum for their helpful comments and suggestions.

Contents

1 Introduction

In many economic situations agents make investments before entering a market in order to improve their prospects. Examples abound: students engage in extra-curricular activities in high school to improve their college admissions chances, and men and women are aware that educational investment increases their value on the marriage market. Maybe the most prominent example is that of workers who invest in their skills before entering the labor market. These investments take place under multiple *risks*, both idiosyncratic (e.g., a worker's skill realization is not known a priori) and aggregate (e.g., the state of the labor market). A natural question is how changes in these risks – in the sense of "better" or "more variable" risks – affect the incentives to invest, the allocation of workers and jobs, and wages. And how these effects contrast with those of better or more spread out distributions of underlying worker and firm *heterogeneity* (e.g., workers' innate ability and firms' productivity). Understanding these effects is crucial for several economic applications: in particular, for designing effective policies that counter inequality and enhance efficiency. Yet, there is surprisingly little work on these equilibrium comparative statics effects.

This paper develops an equilibrium matching model of the labor market with two-sided heterogeneity, where risk averse workers make pre-match investments in their skills, which takes place in the presence of both aggregate and idiosyncratic risks. The model is tractable and permits a complete equilibrium analysis under a general class of primitives that includes most of the common utility and match output functions used in economic applications.

Our main contribution is the analysis of novel *equilibrium comparative statics* of how changes in the distribution of risk (aggregate and idiosyncratic) and heterogeneity (workers' ability or firms' productivity) affect educational investment, matching, and wages. We provide intuitive conditions on workers' risk preferences and technology for stochastically better or more variable distributions to induce more workers to invest. Changes in workers' investment affect the distribution of skills, thus changing the assignment of workers to firms and wages, which then feed back into investment choices. This leads to rich equilibrium effects, and we shed light on their interplay.

To illustrate the economic relevance of our framework, we explore three areas where our results can be useful. First, we use our comparative statics results to gain insight into the causes (changes in risk or innate heterogeneity) of the rise in household income inequality. Second, we shed light on the welfare impact of a variety of policies, ranging from social insurance to policies that impact initial heterogeneity or risk directly. Finally, we discuss how one can use our results to identify whether observed changes in outcomes are due to changes in risk or in innate worker/firm heterogeneity. In all cases, endogenous investment and equilibrium effects play a crucial role.

The model consists of two large populations: risk averse workers and risk neutral firms. Firms are heterogenous in productivity and, initially, workers are heterogeneous in ability with less able workers facing a higher cost (disutility) of investment. There are two stages. In the *investment stage*, workers first make a binary decision of whether or not to invest in their skills, and then they draw their skill from a distribution that is better – in first-order stochastic dominance (FOSD) sense – for those who invest. This implies that workers face idiosyncratic risk regarding their skills and, moreover, that the resulting aggregate distribution of skills is endogenous. At the beginning of the *matching stage*, an aggregate shock realizes. This is the second source of risk present in the model. This shock affects the amount of output of each worker-firm pair, which translates into uncertainty about workers' earnings when investment takes place. Following the publicly observable shock realization, both sides match pairwise in a frictionless labor market, whose Walrasian equilibrium pins down the worker-firm allocation and the market-clearing wage function.

The analysis first focuses on equilibrium existence, uniqueness, and efficiency. We show that an equilibrium exists, and provide sufficient conditions for uniqueness. Moreover, equilibrium investment is generically inefficient, with workers investing too much or too little because they fail to internalize the effect of investment on matching and thus on equilibrium wages. This type of inefficiency in matching with risky pre-match investment is new to our knowledge.

We then focus on our main analytical exercise and contribution: the equilibrium comparative statics of changes in risk and heterogeneity. We derive conditions on primitives – risk preferences and match output function – under which a stochastically better distribution (in the FOSD sense) or a riskier distribution (an increase in risk, IR, à la Rothschild and Stiglitz (1970)) of the aggregate shock, idiosyncratic skill risk, ability, or firm productivity increases the number of workers who invest. This is a natural comparative statics result and an intuitive one in applications: agents have more incentives for investment to take advantage of stochastically better shocks or better innate firm and worker heterogeneity, or to shield themselves from a riskier environment. These results are far from being foregone conclusions, due to the presence of a nonlinear equilibrium wage function whose properties need to be balanced against the properties of workers' risk attitudes. This is the essential trade off embedded in the comparative statics results.

Our first comparative statics result shows that a FOSD shift in the aggregate shock induces more workers to invest if there are sufficiently strong complementarities in production between the aggregate shock and a worker's skill, and if workers are not too risk averse. An elementary intuition relies on substitution and income effects that go in opposite directions. Complementarities between skill and shock in the match output function provide agents with more incentives to invest (substitution effect). This is because wages are increasing in the shock realization and more so for more skilled workers. In turn, a stochastically better shock leads to higher wages on average even if workers do not invest, thereby reducing the incentives for risk-averse workers to invest (income effect). Sufficiently strong complementarities and low enough absolute risk aversion cause the substitution effect to dominate the income effect, and thus *more* workers invest.

Our second comparative statics result focuses on an IR shift in the aggregate shock. The effect of an IR shift depends on workers' absolute prudence. If agents are sufficiently prudent, then a riskier shock pushes more of them to invest in their skill. It is well-known that in many risky environments prudence triggers precautionary savings to insure against bad realizations of a shock. Here, a riskier shock induces sufficiently prudent agents to engage in *precautionary investment* – a precautionary action we believe is new in the literature. How large prudence needs to be depends on curvature and complementarity properties of the match output (and thus wage) function.

We also derive similar comparative statics results when there is a FOSD or IR shift in the distributions of idiosyncratic risk (which workers draw their skills from upon investment), innate worker ability, and firm productivity. Now there is a further complication, since these shifts directly affect the equilibrium matching and wage functions (and not just indirectly via the effect on investment). Still, we provide intuitive conditions under which better and more dispersed distributions trigger more investment, which rely on the discussed trade off between risk attitude and properties of the technology. For all of our comparative statics results, we provide several natural classes of primitives – match output and utility functions – that satisfy our conditions.

The model relies on several assumptions that can be relaxed. Indeed, we show that the results are robust when we allow for two-sided risk aversion commonly used in partnership problems, continuous (instead of binary) investment, two-sided investments, and non-transferable utility. We thus provide an exhaustive analysis of this problem while maintaining our main insights.

There are many applications of our results and we provide three illustrations of their the economic relevance. First, something that has received much attention is the recent rise in U.S. (household) income inequality. The emerging consensus is that it is mostly due to an increase in the returns to education over time (Eika, Mogstad, and Zafar (2018)). This is an important insight, which is however based on descriptive decompositions that do not account for the equilibrium effects on educational investment. We show that considering only the direct effect of the increase in returns to education on income inequality, statistical decompositions – that keep education fixed when changing returns – are likely to overstate its impact. And this approach may lead to qualitatively wrong conclusions if the equilibrium supply-side response is strong enough.

Second, our framework is useful for studying policies affecting risk or initial heterogeneity in the presence of endogenous investment. For instance, we can analyze policies aiming to improve workers' initial conditions (e.g. through public education) or to provide insurance against shocks (think of unemployment insurance). These policies have the intuitive beneficial effects of improving the environment or helping to complete missing insurance markets. But, because investment is endogenous and inefficient in equilibrium, the planner also wants to use them to correct this inefficiency. We show that the beneficial intuitive effects of those policies can be tamed by the presence of endogenous investment – a margin often overlooked in policy design.

Third, to design policies to counter inequality or improve efficiency, it is not only important to understand their effects, but also to be able to distinguish in the data whether outcomes are driven by initial conditions or risk (see discussion in Heathcote, Storesletten, and Violante (2009)). This is a challenging task. We show that our model could help identify changes in risk versus changes in heterogeneity in the data. Our approach does not require to structurally estimate the model but relies on estimation of the reduced form wage function of our model. Endogenous educational investment is crucial for distinguishing changes in risk and heterogeneity in our setting.

While we chose to highlight the economic relevance of our model through questions surrounding the debate on risk versus heterogeneity, we stress that our analysis applies more broadly. For instance, one of our results provides a testable prediction regarding whether education is used to hedge against worse and more volatile aggregate shocks in recessions (Bloom, Floetotto, Jaimovic, Saporta-Ekstein, and Terry (2018); we model this as a simultaneous IR and reverse FOSD in the distribution of the aggregate shock). One could also analyze the effects on investment and labor market outcomes of technological change (Acemoglu and Autor (2011); e.g., a FOSD shift in firm productivity), the increased productivity dispersion across plants (Dunne, Foster, Haltiwanger, and Troske (2004); e.g. an IR shift in firm productivity), or of changes in early childhood investments (Del Boca, Flinn, and Wiswall (2014); e.g., a FOSD or IR shift in the ability distribution). Finally, one could use our results to rationalize that changes in teacher quality have lasting effects on students' outcomes in the data (Chetty, Friedman, and Rockoff (2014); e.g., a FOSD or IR shift in the distribution that workers who invest draw their skills from).

Our theory is related to the large matching literature initiated by Becker (1973), Shapley and Shubik (1972), Gale and Shapley (1962), and generalized by Legros and Newman (2007) (see Chade, Eeckhout, and Smith (2017) for a survey), which has been extensively applied to labor and marriage markets. We add a pre-match investment stage, a feature that is also in Cole, Mailath, and Postlewaite (2001), Peters and Siow (2002), Chiappori, Iyigun, and Weiss (2009), Mailath, Samuelson, and Postlewaite (2013), Nöldeke and Samuelson (2015) and Chiappori, Dias, and Meghir (2018). Bhaskar and Hopkins (2016) and Zhang (2015) also allow for risky investments. The first highlights the inefficiency of equilibrium in a marriage market, while the second shows that investments can be extreme under linear surplus. None of these papers study the equilibrium comparative statics we focus on. Indeed, to the best of our knowledge, ours is the first paper to analyze the equilibrium effects of changes in aggregate and idiosyncratic risks as well as in innate heterogeneity (FOSD and IR) in a general matching framework, highlight investment as a precautionary action, and flesh out the economic relevance of these insights for studying inequality, efficiency and policy, as well as identification.

Our paper is also related to the literature in economics of uncertainty that studies risk changes in decision problems (see Gollier (2004) for a survey). Our model deals with two complications: first, our risks affect agents' expected utility through the equilibrium wage function, while this literature focuses on decision problems where the payoff/wage function is exogenous. Second, these risks enter our wage function in a non-additive way. Both features, equilibrium and non-additive risks, not only significantly complicate the analysis but they also enrich the economics of our comparative statics results – something that becomes particularly clear in our applications.

2 The Model

There is a unit measure of heterogeneous workers and the same measure of heterogeneous firms.[1]

Workers are risk averse with von Neumann-Morgestern utility function for income given by $u : \mathbb{R}_+ \to \mathbb{R}$, which is continuous on $\mathbb{R}_+$ and three times continuously differentiable on $\mathbb{R}_{++}$, with $u' > 0$, $u'' \leq 0$, and $u''' \geq 0$.[2] Initially, all workers have the same skill level that we normalize to zero, and they differ in a characteristic θ, their ability, distributed according to a continuously differentiable cumulative distribution function (cdf) Q on $[0,1]$, with positive density q.

Before entering the labor market workers make a binary investment choice $a \in \{0,1\}$. If $a = 1$, the worker draws a skill $x \in [0,1]$ from a cdf H_1; if $a = 0$, the worker draws $x \in [0,1]$ from H_0. These distributions are ordered by strict FOSD: $H_1(x) \leq H_0(x)$ for all x, with strict inequality over some interval.[3] Moreover, H_i is continuously differentiable with positive density h_i, $i = 0, 1$. Investment is costly, but less so for those with higher ability θ. A worker's investment cost when his ability is θ by $c(\theta) \geq 0$, and c is continuously differentiable on $(0,1]$, and strictly decreasing in θ, with $c(1) = 0$. To rule out uninteresting cases we assume that $\lim_{\theta \to 0} c(\theta) = +\infty$. If a worker with θ invests and obtains income w, then his payoff is $u(w) - c(\theta)$; it is $u(w)$ if he does not invest.

Firms are risk neutral, profit-maximizers, and heterogeneous in a productivity attribute $y \in [0,1]$, distributed according to a continuously differentiable cdf G, with positive density g.

After workers' investment decision, firms and workers match pairwise in a competitive labor market. Before matching, they observe the realization of an aggregate shock $\alpha \in [0,1]$, drawn from a continuous cdf L.[4] The output produced by a worker-firm pair is given by a positive function f on $[0,1]^3$, so if a worker with skill x matches with a firm of productivity y, and the shock realization is α, then they produce $f(x,y,\alpha)$. The function f is increasing and supermodular, three times continuously differentiable, with positive derivatives f_x, f_y, f_α, f_{xy}, $f_{x\alpha}$, and $f_{y\alpha}$, strictly so on $(0,1]^3$. Also, $f(0,0,\alpha) = 0$ for all α: a pair with the lowest attributes is unproductive. An example of such technology is $f(x,y,\alpha) = \alpha xy$. If a worker is unmatched, then his income is zero.[5]

The timing then is as follows: workers decide whether to invest; their attributes are drawn; the aggregate shock realizes; and workers and firms match in a competitive market.

[1] Although the two sides of the market can be men and women, buyers and sellers, students and colleges, etc., for clarity we cast it as a labor market where workers can make investments in their skills before matching with firms.

[2] We use the following conventions. Given a function $z : X \to \mathbb{R}$, $X \subset \mathbb{R}^n$, we write $z \geq 0$ (≤ 0) if the inequality holds for all values in the domain of z. Also, we use increasing, decreasing, concave, convex, positive, etc., in the weak sense, adding "strictly" whenever needed. When z is differentiable on $[a,b]$, the derivatives at the endpoints are the one-sided derivatives. All random variables have support $[0,1]$, and we use $\int$ instead of $\int_0^1$. Finally, whenever there is no risk of confusion, we suppress the arguments of functions.

[3] A random variable $X \in [0,1]$ dominates $Y \in [0,1]$ in FOSD if $F_X(s) \leq F_Y(s)$ for all s, where F_X and F_Y are the cdf's of X and Y. Equivalently, X FOSD Y if $\mathbb{E}[z(X)] \geq \mathbb{E}[z(Y)]$ for all increasing functions z.

[4] Even though we interpret α as an aggregate shock before matching, with some minor modifications we could alternatively interpret it as a worker's (uninsurable) idiosyncratic shock realized after production takes place.

[5] This ensures that matching is individually rational for all workers. If their outside option were not zero income, then the only change in the analysis would be to pin down the mass of workers who prefer to remain unmatched.

Let $a : [0,1] \to \{0,1\}$ be a measurable investment function, where $a(\theta) = 0$ if a worker with cost $c(\theta)$ does not invest, and $a(\theta) = 1$ if he does.[6] Given a, the resulting distribution of attribute x is $H(\cdot, a)$, which is a mixture of H_1 and H_0 with weights given by the measure of workers who invest and do not invest, respectively. Formally,

$$H(x,a) = \left(\int_{\{\theta : a(\theta)=1\}} dQ(\theta) \right) H_1(x) + \left(1 - \int_{\{\theta : a(\theta)=1\}} dQ(\theta) \right) H_0(x). \tag{1}$$

Let $\mu(\cdot, a) : [0,1] \to [0,1]$ be a measurable and measure-preserving matching function that assigns each worker with x to a firm with $y = \mu(x,a)$, when the investment function is a. Let $w(\cdot, \cdot, a) : [0,1]^2 \to \mathbb{R}_+$ be a measurable wage function for each a, so that $w(x, \alpha, a)$ is the wage of a worker with x when the realization of the shock is α and the investment function is a.

An *equilibrium* is a triple (a, μ, w) such that, given (w, μ), workers invest optimally, i.e., for all θ, $a(\theta) = 1$ if and only if $U_1(w) - c(\theta) \geq U_0(w)$, where

$$U_i(w) = \int \int u(w(x, \alpha, a)) dH_i(x) dL(\alpha), \quad i = 0, 1; \tag{2}$$

and for any measurable a (which determines $H(\cdot, a)$) and any realization of α, the wage and matching functions $(w(\cdot, \alpha, a), \mu(\cdot, a))$ constitute a Walrasian equilibrium of the labor market.

We have made several simplifying assumptions: α realizes before matching; only one side is risk averse; investment is binary; only one side invests; and the model is closed via a matching market. Section 6 discusses these assumptions in detail and provides generalizations.

3 Equilibrium Analysis

We begin by showing that an equilibrium exists, proceeding essentially by construction: for each given investment function, we obtain the Walrasian equilibrium of the labor market; and given the Walrasian equilibrium that will prevail in the labor market under a conjectured investment choices of all the workers, we obtain the equilibrium pre-match investment decision of each worker. We then shed light on whether equilibrium is unique, and also show that investment is inefficient.

We solve the model backwards analyzing first the matching and then the investment stage. Consider the *matching stage* given investment a (so that $H(\cdot, a)$ is the cdf of x) and shock realization α. Since α is revealed before firms and workers match, matching takes place under certainty, and thus this is an assignment game similar to Becker (1973).

Since f is supermodular in (x, y) for each α, the optimal matching exhibits positive sorting.[7] Hence, $\mu(\cdot, a)$ solves $H(x, a) = G(\mu(x, a))$ and thus $\mu(x, a) = G^{-1}(H(x, a))$, which is strictly

[6] Measurability is always with respect to the appropriate Borel sigma field.

[7] A twice continuously differentiable function $z : \mathbb{R}^2 \to \mathbb{R}$ is supermodular in (x, y) if $z_{xy} \geq 0$.

increasing in x. To support this allocation as a Walrasian equilibrium, let $w(x,\alpha,a)$ be the wage of a worker with skill x given (α,a). Then a firm with productivity y solves

$$\max_x f(x,y,\alpha) - w(x,\alpha,a).$$

The first-order condition of this problem is $f_x(x,y,\alpha) = w_x(x,\alpha,a)$, and hence in *equilibrium* $f_x(x,\mu(x,a),\alpha) = w_x(x,\alpha,a)$ for all x, which yields the following wage function:

$$w(x,\alpha,a) = w(0,\alpha,a) + \int_0^x f_x(s,\mu(s,a),\alpha)ds = \int_0^x f_x\left(s, G^{-1}\left(H(s,a)\right),\alpha\right)ds, \tag{3}$$

where we used that if $x=0$ then $\mu(0,a)=0$ and thus both output and $w(0,\alpha,a)$ are zero. Appendix A.1 shows that matching function $\mu(\cdot,a) = G^{-1}\left(H(\cdot,a)\right)$ and wage function (3) constitute a Walrasian equilibrium given (α,a). The payoff of a matched pair (x,y) is then $u(w(x,\alpha,a))$ for the worker with skill x (minus the sunk investment cost if he invested) and $f(\mu^{-1}(y,a),y,\alpha) - w(\mu^{-1}(y,a),\alpha,a)$ for the firm with productivity y.

Consider now the *investment stage*. Given the wage function w that workers anticipate in the matching stage, a worker with cost $c(\theta)$ invests if and only if $U_1(w) - U_0(w) \geq c(\theta)$. The left side is independent of θ since workers take (μ,w) as given. And because w is strictly increasing in x and H_1 strictly FOSD H_0, we obtain $U_1(w) - U_0(w) > 0$. Thus, there is a unique threshold $\theta^* \in (0,1)$ such that a worker with ability θ invests if and only if $\theta \geq \theta^*$, where θ^* solves $U_1(w) - U_0(w) = c(\theta^*)$ (recall the properties of c). So in any equilibrium we have

$$a(\theta) = \begin{cases} 1 & \text{if } \theta \geq \theta^* \\ 0 & \text{if } \theta < \theta^*. \end{cases}$$

Thus, (1) becomes $H(x,\theta^*) = (1-Q(\theta^*))H_1(x) + Q(\theta^*)H_0(x)$ and, similarly,

$$\mu(x,\theta^*) = G^{-1}\left(H(x,\theta^*)\right) \tag{4}$$

$$w(x,\alpha,\theta^*) = \int_0^x f_x\left(s, G^{-1}\left(H(x,\theta^*)\right),\alpha\right)ds. \tag{5}$$

With a slight abuse of notation, we will replace $U_i(w)$ by $U_i(\theta^*)$, $i=0,1$, where

$$U_i(\theta^*) = \int\left(\int u\left(w(x,\alpha,\theta^*)\right)dH_i(x)\right)dL(\alpha).$$

Equilibrium existence reduces to finding a threshold θ^* that satisfies $U_1(\theta^*) - U_0(\theta^*) = c(\theta^*)$ which, using integration by parts, can be written as

$$\int\int u'(w(x,\alpha,\theta^*))w_x(x,\alpha,\theta^*)\Delta H(x)dxdL(\alpha) = c(\theta^*), \tag{6}$$

where $\Delta H \equiv H_0 - H_1$.[8] Once we solve for θ^*, the induced function a determines $H(\cdot, \theta^*)$, which then yields a Walrasian equilibrium in the matching stage given by $(\mu(\cdot, \theta^*), w(\cdot, \alpha, \theta^*))$ for each realization of α. This pins down $U_1(\theta^*) - U_0(\theta^*)$ in the investment stage, which is higher than $c(\theta)$ for $\theta \geq \theta^*$ and lower otherwise. This rationalizes $H(\cdot, \theta^*)$, completing the equilibrium construction.

It remains to show that there is a solution to $U_1(\theta^*) - U_0(\theta^*) = c(\theta^*)$. The left side is positive and continuous for all $\theta^* \in [0, 1]$, and c diverges to infinity when θ^* goes to zero and it is zero at $\theta^* = 1$. By the Intermediate Value Theorem, there is at least one solution to $U_1(\theta^*) - U_0(\theta^*) = c(\theta^*)$ where $U_1 - U_0$ crosses c *from below*. In fact, *both* the smallest and the largest equilibria have this feature. Moreover, in any solution $0 < \theta^* < 1$.[9]

A natural question is when equilibrium is unique. From the properties of c, it is unique if $U_1 - U_0$ is increasing in θ^* (Figure 1.A). This is also *necessary* if we want uniqueness for *all* c's that satisfy our assumptions. Uniqueness is easier if c is also convex. For if c is convex enough, so that its graph is "close to the axes," then only the first crossing of c by $U_1 - U_0$ survives.

Let $R \equiv -u''/u'$ be the coefficient of absolute risk aversion.

Proposition 1 (Existence and Uniqueness) *An equilibrium exists, and all equilibria exhibit* $\theta^* \in (0, 1)$*. It is unique if agents' absolute risk aversion* R *is (uniformly) sufficiently small, or if the cost function* c *is sufficiently convex.*

The proof is in Appendix A.2. For example, if $f(x, y, \alpha) = \alpha xy$, y is uniformly distributed, and $\mathbb{E}[\alpha] > 0$, then equilibrium is unique if $R(w) \leq \mathbb{E}[\alpha] \int (\Delta H)^2 / \left(\int \Delta H\right)^2$ for all w (see Appendix A.2). This upper bound is pinned down by primitives L, H_0, and H_1. Similarly, if c is given by $c(\theta) = (1/\theta^{\frac{1}{j}}) - 1$, then $\lim_{j \to \infty} (1/\theta^{\frac{1}{j}}) - 1 = 0$ for all $\theta \in (0, 1]$, and $c_\theta(\theta) = -(1/j)(1/\theta^{\frac{1}{j}+1})$ diverges to $-\infty$ as θ goes to 0. Equilibrium is unique for sufficiently large j.

We now turn to the analysis of equilibrium efficiency. Assume that the planner has to respect the timing in which matching takes place after investment and after the uncertainty has been resolved. He can only choose θ^* in the first stage to maximize a weighted (by λ and $1 - \lambda$) sum of (ex-ante) workers and firms expected payoffs. Formally, the planner solves

$$\max_{\theta^* \in [0,1]} \lambda \left(Q(\theta^*) \int \int u(w(x, \alpha, \theta^*)) dL(\alpha) dH_0(x) + (1 - Q(\theta^*)) \int \int u(w(x, \alpha, \theta^*)) dL(\alpha) dH_1(x) \right.$$
$$\left. - \int_{\theta^*}^1 c(\theta) dQ(\theta) \right) + (1 - \lambda) \left(\int \int \left(f(x, \mu(x, \theta^*), \alpha) - w(x, \alpha, \theta^*) \right) dL(\alpha) dH(x, \theta^*) \right). \quad (7)$$

The first term in parenthesis in the objective function represents the expected utilities of the workers as a function of θ^*, while the second is the corresponding expression for the firms.

[8] At several points in the analysis we will pass the derivative through the integral and use integration by parts, as in (6). To avoid technical detours in the text, we justify these operations in Appendix A.10.

[9] If we did not assume $\lim_{\theta \to 0} c(\theta) = \infty$, or more generally that $c(0)$ is sufficiently large, then there could be an equilibrium with $\theta^* = 0$, i.e., everyone invests.

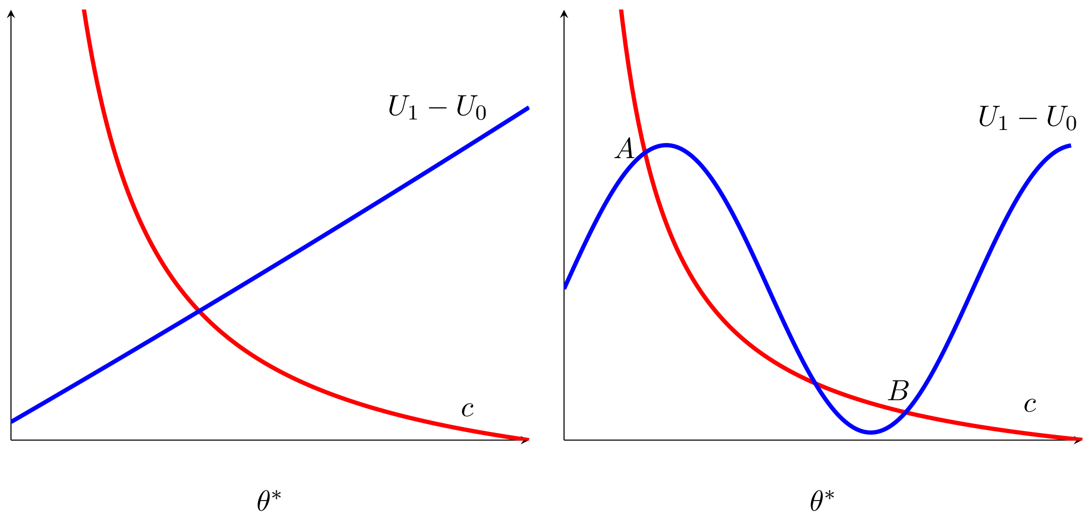

Figure 1: A. Unique equilibrium. B. Multiple equilibria.

Proposition 2 (Inefficiency) *Equilibrium investment is generically inefficient. The planner's optimal threshold θ_p^* in an interior solution satisfies the following equation:*

$$\int \int u' w_x \Delta H dx dL = c + \frac{1}{q} \int \int \left(u' - \left(\frac{1-\lambda}{\lambda} \right) \right) w_{\theta^*} dH dL. \tag{8}$$

Any equilibrium threshold θ^* is interior and satisfies (6); in comparison, equation (8) has an extra term, which reflects the inefficiency. The basic reason is that workers do not take into account that investment has an effect on equilibrium matching and thus on wages (wages decrease the more workers invest). As a result, the equilibrium can exhibit over or under-investment in skills, depending on risk aversion and the weights the planner puts on the utility of workers versus firms. For instance, for any given u, if λ is sufficiently close to one so that workers receive most weight in the planner's objective, then $\int (u' - ((1-\lambda)/\lambda)) w_{\theta^*} dH > 0$, and thus the equilibrium exhibits over-investment: The planner prefers lower investment in order to increase wages. Appendix A.3 proves the proposition and relates it to the literature on pre-match investments.

4 Comparative Statics

We now turn to our main exercise, the equilibrium comparative statics of the model. The equilibrium objects are θ^*, μ, and w, which depend on the distributions L, H_0, H_1, G, and Q. Our focus is the analysis of the impact of changes in these distributions – a FOSD or an IR shift – on θ^*, which determines how the measure of workers who invest changes.[10] These are non-obvious

[10] An IR shift asserts that a random variable $X \in [0,1]$ is riskier than a random variable $Y \in [0,1]$ if $\int_0^t F_X(s)ds \geq \int_0^t F_Y(s)ds$ for all t, with equality at $t = 1$ (so both have the same mean). Alternatively, this holds if and only if

and unexplored comparative statics. Once we pin down the effect on the investment threshold θ^*, we assess the changes in the matching function μ and the wage function w.

It is instructive to distinguish changes in aggregate and idiosyncratic *risk* from changes in *heterogeneity*. Such a distinction is relevant, for instance, for the analysis of life-time earnings, consumption and income inequality and, as a result, for designing policies, issues we discuss in Section 5. Distributions G and Q represent the ex-ante heterogeneity in our model. In turn, L, H_0, and H_1 represent risk, since workers face a *controllable* (through investment) idiosyncratic risk x, determining the ex-post heterogeneity of workers, and a *background* aggregate risk α.

In each case, we provide conditions on primitives under which a stochastically better or riskier aggregate shock, ability, skill, or firm-productivity distribution induce more workers to invest, which in many settings is the intuitive results to conjecture. For instance, one would think that a worker would be more inclined to invest to take advantage of a stochastically better labor market (FOSD shift in L), or to shield himself from a more uncertain labor market (IR shift in L). As mentioned in the Introduction, these comparative statics results hinge on a subtle trade-off between complementarities in production and workers' attitudes towards risk, which in principle can go either way and is not a priori obvious how to discipline.

Our results will apply to all *stable* equilibria, where "$U_1 - U_0$ crosses c from below," see equilibria A and B in Figure 1.B. Proposition 1 shows that there is at least one, and that both the *smallest* and *largest* equilibria are stable. They are stable in a natural (tâtonnement) sense: if θ^* is to the left of the crossing point, then $U_1(\theta^*) - U_0(\theta^*) < c(\theta^*)$ and θ^* will go up since some workers above θ^* will not find it optimal to invest. The opposite happens for θ^* to the right of the crossing point. Equilibria where the crossing occurs from above are unstable (see Figure 2.A).

Our approach is to index the appropriate distribution by a parameter $t \in [0, 1]$, where an increase in t represents a FOSD or an IR shift. We then consider how θ^* changes with t by analyzing the behavior of $U_1(\theta^*, t) - U_0(\theta^*, t) = c(\theta^*)$, where for clarity we added t as an argument of U_i, $i = 0, 1$. Figure 2.B reveals that a *sufficient* condition for θ^* to decrease when t increases, so that more workers invest, is that $U_1 - U_0$ increases in t for *all* values of θ^*. This condition is also *necessary* for the results to hold for all the cost functions c in the class we consider. This is because if it strictly decreased in an interval, then one could find a c in our class that crosses $U_1 - U_0$ from below, and thus θ^* would increase in t for any such crossing in that interval.

4.1 Changes in Aggregate Risk

We first investigate how θ^* changes with labor market risk α, that is when labor market conditions become either stochastically better or more variable. We index L by $t \in [0, 1]$, so $L(\cdot|t)$, with $L(\alpha|\cdot)$ continuously differentiable in t for each α. Let L_t be the derivative of L with respect to t.

$\mathbb{E}[z(X)] \leq \mathbb{E}[z(Y)]$ for all concave functions z. This is called the convex order in the stochastic order literature (see Chapter 1 in Muller and Stoyan (2002)).

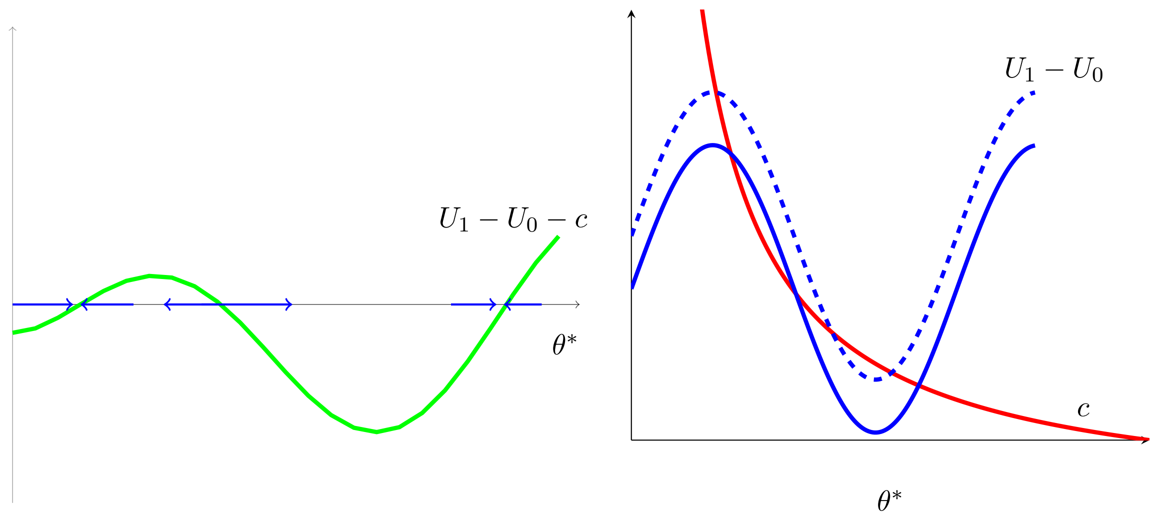

Figure 2: A. Two stable equilibria at the extremes, one unstable equilibrium in between. B. Comparative Statics of $U_1 - U_0$ with respect to changes in risk.

Our aim is to find conditions on primitives (preferences and technology) such that $U_1 - U_0$ increases in t *for all* θ^*. To simplify the notation, let $\phi(x,t) \equiv \int u(w(x,\alpha,\theta^*))dL(\alpha|t)$, where we omitted θ^* as an argument since it will be fixed in what follows. Then from (6) we have

$$U_1(\theta^*,t) - U_0(\theta^*,t) = \int \phi_x(x,t)\Delta H(x)dx,$$

and thus $U_1 - U_0$ increases in t if and only if $\int \phi_{xt}\Delta H dx \geq 0$, for which it suffices that ϕ is supermodular in (x,t) since $\Delta H = H_0 - H_1 \geq 0$.[11]

When is ϕ supermodular? Since $\phi_x = \int u'(w)w_x dL$, a sufficient condition when $L(\alpha|\cdot)$ is ordered by FOSD is that $u'(w)w_x$ is increasing in α. Similarly, when $L(\alpha|\cdot)$ is ordered by IR, it suffices that $u'(w)w_x$ is convex in α.[12] *Hence, a FOSD shift reduces θ^* in any stable equilibrium – more workers invest – if $u'(w)w_x$ is increasing in α, and for an IR shift if it is convex in α.*

For an intuition, consider an IR shift and a worker with skill x. For the sake of argument, assume that $u(w)$ is concave in α. Then an increase in t reduces the worker's expected utility, but less so for higher x if $\phi_{xt} \geq 0$, which holds if $u'(w)w_x$ is convex in α. Since investment yields stochastically higher skills x, it follows that an IR shift in L induces more workers to invest (θ^* decreases). A similar intuition holds for a FOSD shift.

We now derive conditions on primitives that ensure that $u'(w)w_x$ is increasing or convex in α.

[11]With an arbitrary w this condition is necessary for the result to hold for *all* H_i, $i = 0,1$, ordered by strict FOSD. For if $\phi_{xt} < 0$ on an interval of x, then we can choose H_0 and H_1 such that $H_1(x) < H_0(x)$ only for x in that set, and the result fails. By (5), the wage function also depends on H_0 and H_1 through μ, so this argument does not apply. But this suggests that it would be hard to relax our conditions.

[12]These conditions are necessary if we want ϕ supermodular in (x,t) for all cdf's L ordered by FOSD or IR.

FOSD Shift in L. Under FOSD, $L(\alpha|\cdot)$ is decreasing in t for all α, so higher realizations are more likely. Equivalently, the expectation of any increasing function of α is increasing in t.

Proposition 3 (FOSD Shift in L**)** *In any stable equilibrium, more workers invest in skills in response to a FOSD shift in L if either of the following conditions hold:*

(i) Absolute risk aversion R is (uniformly) sufficiently small; or

(ii) For all values of w, $R(w) \leq 1/w$, and f_x is log-supermodular in (x, α) for each y, and in (y, α) for each x.[13]

The basic idea of the proof (see Appendix A.4 for details) is simple. We want $u'(w)w_x$ to be increasing in α: the derivative with respect to α is

$$u''(w)w_\alpha w_x + u'(w)w_{x\alpha}. \tag{9}$$

When workers are risk neutral, (9) is positive if and only if $w_{x\alpha} \geq 0$, and this holds strictly from (5) and the assumptions on f. By continuity, (9) remains positive when risk aversion is small, which yields part (i). More generally, if we divide (9) by $u'(w)$ and by $w_\alpha w_x$, which are positive almost everywhere due to the assumptions on f, then the sign of (9) depends on the sign of

$$\left(\frac{w_{x\alpha}w}{w_\alpha w_x}\right)\frac{1}{w} - R(w),$$

and so (9) is positive if risk aversion is bounded above by a positive expression that depends on the complementarities in w. Part (ii) provides the required conditions on risk aversion and match output complementarities that deliver the comparative statics result.

For an economic intuition, interpret the second term in (9) as a *substitution effect*: some agents switch from not investing to investing because of the complementarities in the wage between x and α, with α now stochastically larger. In turn, the first term resembles an *income effect* where the better shock increases all wages – even for those who do not invest – thereby reducing the incentives to invest. If the substitution effect dominates the income effect, our result obtains.

The conditions in Proposition 3 are easily satisfied, as the following example illustrates.

Example. Let $f(x, y, \alpha) = \alpha xy$. It is easy to verify that part (i) holds for the constant absolute risk aversion (CARA) $u(w) = -e^{-Rw}$, $R > 0$, when R is sufficiently small. Regarding part (ii), note that $f_x = \alpha y$, which is log-supermodular in (x, α) and in (y, α). So for the constant relative risk aversion (CRRA) $u(w) = (w^{1-\sigma} - 1)/(1 - \sigma)$ if $\sigma > 0$, $\sigma \neq 1$, and $u(w) = \log w$ if $\sigma = 1$, we have $R(w) = \sigma/w$, and hence $R(w) \leq 1/w$ for $\sigma \in (0, 1]$.[14]

[13] A twice continuously differentiable function $z : \mathbb{R}^2 \to \mathbb{R}$ is log-supermodular (log-submodular) in (x, y) if $\log z$ is supermodular (submodular), or $z_{xy}z - z_x z_y \geq 0$ (≤ 0). This is stronger than supermodularity when $z_x z_y \geq 0$.

[14] CRRA satisfies part (ii) but not (i) since R cannot be made uniformly small, as it diverges to infinity as w goes to zero. But (i) also holds if $u(w) = ((w + b)^{1-\sigma} - 1)/(1 - \sigma)$ for any $b > 0$ (e.g., an initial wealth).

This example also furnishes a counterexample: if $\sigma > 1$, then the 'intuitive' comparative static breaks down, and a FOSD shift in L reduces the measure of workers who invest. To see this, note that under the multiplicatively separable f assumed we have $w_x w_\alpha = w w_{x\alpha}$, and hence the sign of (9) depends on the sign of $(1/w) - R(w) = (1-\sigma)/w < 0$. Indeed, we can generalize this example as follows: if f is such that $w_x w_\alpha = w w_{x\alpha}$, then a FOSD shift in L *increases* the measure of workers who invest if $R(w) \leq 1/w$ for all w, and *decreases* it if $R(w) \geq 1/w$ for all w.

IR Shift in L. Under IR, $\int_0^\alpha L_t(s|t)ds \geq 0$ for all α and t, and $\int L_t(s|t)ds = 0$ (the mean remains constant), so the aggregate shock α becomes riskier. Equivalently, the expectation of any convex function of α is increasing in t.

Let $P \equiv -u'''/u''$ be the coefficient of absolute prudence, and say that f is a *multiplicatively separable class* if it can be written as $f(x,y,\alpha) = \eta(\alpha)z(x,y)$, with η increasing, convex and log-concave, and twice continuously differentiable on $(0,1]$; and z is increasing on $[0,1]^2$, twice continuously differentiable, strictly increasing and strictly supermodular on $(0,1]^2$, with $z(0,0) = 0$. This class of match output functions is commonly used in applications, and exhibit the convenient property asserted above that $w_x w_\alpha = w w_{x\alpha}$.

Proposition 4 (IR Shift in L) *In any stable equilibrium, more workers invest in skills in response to an IR shift in L if either of the following conditions hold:*

(*i*) *Absolute risk aversion is (uniformly) sufficiently small and $f_{x\alpha\alpha} \geq 0$ with strict inequality on a set of (x,y,α) of positive measure; or*

(*ii*) *For all w, $P(w) \geq 3/w$ and f is a multiplicatively separable class.*

To see the idea of the proof (see Appendix A.5 for details), recall that we seek conditions on primitives such that $u'(w)w_x$ is convex in α. Differentiating twice with respect to α gives

$$u'''(w)w_x w_\alpha^2 + 2u''(w)w_\alpha w_{\alpha x} + u''(w)w_x w_{\alpha\alpha} + u'(w)w_{x\alpha\alpha}. \tag{10}$$

If workers are risk neutral (so $u'' = u''' = 0$), then (10) is positive if and only if $w_{x\alpha\alpha} \geq 0$, and this holds if $f_{x\alpha\alpha} \geq 0$. Intuitively, this also holds when risk-aversion is sufficiently small as long as $w_{x\alpha\alpha} > 0$ on a set of positive measure, which yields part (*i*). Regarding part (*ii*), algebraic manipulation of (10) reveals that its sign depends on the sign of

$$R(w)w_x w_\alpha^2 \left(P(w) - \frac{1}{w}\left(\frac{2w_{\alpha x}w}{w_x w_\alpha} + \frac{w_{\alpha\alpha}w}{w_\alpha^2}\right)\right) + w_{x\alpha\alpha}, \tag{11}$$

If f is a multiplicatively separable class, then the term in the inner parentheses involving complementarities and curvature of w is bounded above by 3, since $w_x w_\alpha = w w_{x\alpha}$ and $w_{\alpha\alpha}w \leq w_\alpha^2$ by log-concavity of η. Hence, the first term in (11) is positive if $P(w) \geq 3/w$, which yields part (*ii*).[15]

[15]Under further assumptions, we can replace the class of f in (*ii*) by $f(x,y,\alpha) = \eta(\alpha)z(x,y) + \delta s(x) + t(y)$,

For an economic intuition, first note that this result again relies on two forces, curvature properties of w (and thus f) and workers' risk attitudes. Consider part (i). Under risk neutrality ($R = 0$), a riskier shock induces more agents to invest if w_x is convex in α, which implies that the mean wage increases for all workers and especially so for those with a higher skill x – something that investment yields stochastically. Intuitively, this result extends to small absolute risk aversion. Regarding the more interesting part (ii), where workers are strictly risk averse, the first term of (11) plays a central role. The difference with the FOSD shift is that here *prudence* is the driving force. If workers are sufficiently prudent, then a riskier shock induces more workers to invest in their skill. As is well-known, in many settings prudence triggers precautionary savings to insure against bad realizations of a shock. Here, a riskier shock induces sufficiently prudent workers to engage in *precautionary investments*. How large prudence needs to be depends on curvature and complementarity properties of the wage function. If the wage function is convex in α, a riskier shock generates more upside compared to downside risk. Since prudent agents particularly dislike downside risk, this reduces the incentives to invest and only the very prudent workers would take the precautionary action. Moreover, if the wage function exhibits complementarities in x and α, a riskier shock exposes workers with high x even more to this risk than those with low x (since for high x the difference in payoffs between high and low shock is much larger), further reducing the incentives for precautionary investment. These forces against investment are tamed but still present if w is submodular and log-concave. Only under enough prudence the net response to increased aggregate risk will be an increase in precautionary investment.

Example. Let $f(x, y, \alpha) = \alpha^\beta xy$, $\beta > 1$, which implies that $w_{x\alpha\alpha} \geq 0$, that $w_{x\alpha}w = w_x w_\alpha$ (so $w_{x\alpha}w/w_x w_\alpha = 1$), and that w is log-concave in α. Part (i) then holds for u CARA when R is sufficiently small. The condition in part (ii) is $P(w) \geq (2 + ((\beta - 1)/\beta))/w$. If $\beta = 2$ then the bound is $P(w) \geq 2.5/w$. If u is CRRA, then $P(w) = (\sigma + 1)/w$, and thus $\sigma \geq 2$ suffices.

In some cases with CRRA we can obtain a sharper characterization. The following example justifies our focus on "small" risk aversion (part (i)) and on "large prudence" (part (ii)).

Example. Let $f(x, y, \alpha) = \alpha^\beta z(x, y)$, $\beta > 1$, and let u be CRRA. Then $U_1 - U_0$ is given by $(\int \alpha^{\beta(1-\sigma)} dL) \int (\int_0^x z_x(s, \mu(s, \theta^*))ds) z_x(x, \mu(x, \theta^*)) \Delta H dx$. An IR shift in L increases investment if and only if $\alpha^{\beta(1-\sigma)}$ is convex in α. The sign of second derivative with respect to α equals the sign of $(1 - \sigma)(1 - (1/\beta) - \sigma)$. This is strictly positive if and only if $\sigma \in [0, 1 - (1/\beta)) \cup (1, \infty)$. Hence, investment increases – that is, θ^* decreases – if σ is small (consistent with part (i) since $R = \sigma/w$) or large enough (consistent with part (ii) since $P = (\sigma + 1)/w$). But θ^* increases for intermediate values of σ. Note that the intuition is as explained above: here σ parameterizes the

which subsumes other cases not covered in (ii) (details are in Appendix A.7). In fact, we could have written part (ii) in a more cumbersome way without assuming that f is a multiplicatively separable class: it suffices that $\varrho = \sup_{x,\alpha}[((2w_{x\alpha}w)/(w_x w_\alpha)) + ((w_{\alpha\alpha}w)/w_\alpha^2)] < \infty$ and $P(w) \geq \varrho/w$ for all w. This subsumes (ii) with $\varrho = 3$.

risk attitudes while β parameterizes the complementarities (in (x, α)) and curvature (in α) of the wage function. The trade off between the two determines the sign of (11).

Parts (*ii*) of Propositions (3)–(4) yields the following corollary regarding a decreasing convex order (DCO) shift in L, which both reduces the mean and increases the riskiness of α.[16]

Corollary 1 (DCO Shift in L) *If f is a multiplicatively separable class, $R(w) \geq 1/w$ and $P(w) \geq 3/w$ for all w, then a DCO shift in L increases the measure of workers who invest.*

The proof is in Appendix A.6, but its intuition is clear: under the stated conditions, $u'(w)w_x$ is both decreasing and convex, and thus the decrease in the mean of α and in its riskiness when t increases work in the same direction, providing more incentives to invest.

We highlight this particular corollary because it leads to a testable prediction. Indeed, Bloom, Floetotto, Jaimovic, Saporta-Ekstein, and Terry (2018) uncover the interesting fact that recessions in the US entail both an increase in uncertainty of aggregate shocks to production (a higher variance) and a decrease in their mean. They base their estimates on a Cobb-Douglas f, which clearly satisfies our condition. Moreover, the stated properties of R and P are easy to estimate for, say, CRRA, which is commonly assumed in the empirical literature. Under these conditions our model predicts that, in a recession, one should observe an increase in the mass of individuals who get an education before going on the labor market, which can be empirically tested.

4.2 Changes in Idiosyncratic Risk and Innate Heterogeneity

Equilibrium investment changes when the quality of education or completion risk improves, which can be captured by a FOSD shift in H_1 (which is the idiosyncratic risk in our model). It also changes with the underlying heterogeneity: with a better distribution of firm productivity – for example, due to technological change – or of workers' ability, – for example, due to improved early childhood investment – which can be captured by either a FOSD shift in G or in Q. The natural comparative statics result we seek in these cases is that equilibrium investment increases.

We will index either H_1, G, or Q by $t \in [0, 1]$, that is $H_1(\cdot|t)$, $G(\cdot|t)$, or $Q(\cdot|t)$, and assume continuous differentiability in t. The effect on investment triggered by these distributional changes is harder to sign than in the aggregate shock case, since they impact wages directly. We will here focus on FOSD shifts since they are the most intuitive one in this case, but for completeness we also present results for IR shifts in Appendix A.9. To avoid repetition, we will omit the analysis of a FOSD shift in H_0 since it mirrors that of H_1.

[16] Random variables $X \in [0, 1]$ and $Y \in [0, 1]$ are ordered by DCO if $\int_0^t F_X(s)ds \geq \int_0^t F_Y(s)ds$ for all t (so it holds for $t = 1$ and thus $\mathbb{E}[X] \leq \mathbb{E}[Y]$). Alternatively, this holds if and only if $\mathbb{E}[z(X)] \leq \mathbb{E}[z(Y)]$ for all decreasing and convex functions z (see Shaked and Shanthikumar (2007) p.182). Intuitively, X has more spread and lower mean than Y, so only decision makers with a convex and decreasing utility function prefer X to Y. Note the similarity with second-order stochastic dominance or increasing convex order, except for the reverse monotonicity in z.

Let $\psi(x,t) \equiv \int u(w(x,\alpha,\theta^*,t))dL(\alpha)$, where as before we omit θ^* as an argument, and we note that w depends on t via the matching function (4). With some abuse of notation we will use the same notation for the function ψ whether t enters H_1, G, or Q.

If we consider a shift in productivity distribution G or ability distribution Q, it follows from

$$U_1(\theta^*,t) - U_0(\theta^*,t) = \int \psi_x(x,t)\,(H_0(x) - H_1(x))\,dx,$$

that $U_1 - U_0$ increases in t if $\psi_{xt} \geq 0$, which holds if $u'(w)w_x$ increases in t.

If instead the shift is in the skill distribution H_1, then there is also a direct effect on the integrand. But since $H_1(x|\cdot)$ decreases in t and $\psi_x \geq 0$, the additional effect is positive and $u'(w)w_x$ increasing in t continues to be sufficient. *Hence, a FOSD shift in H_1, G, or Q, reduces θ^* in any stable equilibrium – more workers invest – if $\psi_{xt} \geq 0$ and thus if $u'(w)w_x$ is increasing in t.*

The next result provides conditions on primitives that satisfy this condition.

Proposition 5 (FOSD Shift G, Q, H_1) (i) *In any stable equilibrium, more workers invest in skills in response to a FOSD shift in G if absolute risk aversion is (uniformly) sufficiently small;*

(ii) *In any stable equilibrium, more workers invest in skills in response to a FOSD shift in Q if $wR(w)$ is sufficiently large for all w, $f_x > 0$, $f_{xy} > 0$, $h_0(0) > h_1(0)$, and $Q_t(\cdot|t) < 0$ on $(0,1)$;*

(iii) *In any stable equilibrium, more workers invest in skills in response to a FOSD shift in H_1 if $wR(w)$ is sufficiently large for all w, $f_x > 0$, $f_{xy} > 0$, and $\partial h_1/\partial t < 0$ at $x = 0$.*

The proof is in Appendix A.8. To see the basic idea, differentiate $u'(w)w_x$ with respect to t:

$$u''(w)w_t w_x + u'(w)w_{xt}. \tag{12}$$

If workers are risk neutral (so $u'' = 0$), then (12) is positive if and only if $w_{xt} = f_{xy}\mu_t$ is positive. This actually holds if t shifts G. To see this, recall that μ satisfies $G(\mu(x,t,\theta^*)|t) = Q(\theta^*)H_0(x) + (1-Q(\theta^*))H_1(x)$, and thus $\mu_t = -G_t/g \geq 0$. Intuitively, if G improves, there are more high quality firms in the upper tail, which alleviates the competition among workers for firms with high y. Hence, a worker with any given skill x matches with a better firm. Since the matching improves, so do wages, $w_t \geq 0$, and this effect is stronger for workers with higher x, $w_{xt} = f_{xy}\mu_t \geq 0$, increasing the incentives to invest. By continuity, this also holds for small risk-aversion (part (i)).

In terms of substitution and income effects, the intuition is that small enough risk aversion guarantees that the positive substitution effect of a FOSD shift in G (workers want to invest more to take advantage of better firm matches) dominates the negative income effect (workers earn more even if they do not invest), inducing more investment.

Instead, with a FOSD shift in ability distribution Q, $G(\mu(x,t,\theta^*)) = Q(\theta^*|t)H_0(x) + (1 - Q(\theta^*|t))H_1(x)$ and thus $\mu_t = Q_t\Delta H/g \leq 0$; similarly, with a FOSD shift in skill distribution H_1, $G(\mu(x,t,\theta^*)) = Q(\theta^*)H_0(x) + (1 - Q(\theta^*))H_1(x|t)$ and thus $\mu_t = ((1 - Q(\theta^*))\partial H_1/\partial t)/g \leq 0$. In both cases, there are now more skilled workers competing for the best firms and so workers with any given skill end up with worse firms. Hence, $w_t \leq 0$ and $w_{xt} \leq 0$, so if workers were risk neutral they would be less inclined to invest. Sufficiently large risk aversion helps to offset this reduced incentives to invest, which sheds light on parts (*ii*) and (*iii*). To see this, rewrite (12) as

$$u'(w)w_t w_x \left(\frac{1}{w} \left(\frac{w_{xt} w}{w_t w_x} \right) - R(w) \right). \tag{13}$$

If the positive term, $w_{xt}w/w_t w_x$, is bounded above (ensured by the other conditions in (*ii*) and (iii)), then $w_t \leq 0$ implies that enough risk aversion makes (13) positive, and investment increases.

For further intuition, note that the income and substitution effects have reverse signs compared to the case of G (since w *decreases* in t due to more competition among workers). So for the positive income effect to dominate the negative substitution effect, absolute risk aversion needs to be sufficiently *large* instead of sufficiently small.

These comparative statics results yield testable predictions of how changes in the underlying heterogeneity affect educational investments in a variety of settings. For instance, they provide a rationale for why college attainment has stagnated in light of a decrease in learning ability over time (Castro and Coen-Pirani (2016); captured here by a reverse FOSD shift in Q); or for the increase in vocational training preparation in response to technological change that has raised the relative demand for skilled labor (Cairo (2013); here captured by a FOSD shift in G).

4.3 Effects on Equilibrium Matching and Wages

Having determined the comparative statics for investment θ^*, we can now turn to the impact of changes in L, G, Q, and H_1 on the equilibrium matching function μ and wage function w. In each case we assume that the appropriate sufficient conditions derived in Propositions (3)–(5) hold. The logic underlying all of the following results is that improvements in the workers' skill distribution triggered by more investment lead to a lower matching function: a worker with given x matches with a (weakly) worse firm y due to tougher competition for top firms. A lower matching function then implies a lower wage for each worker with given skill x. These indirect effects through investment must be weighed against the direct effect of a shift of the appropriate distribution. For instance, an improvement in the distribution of firm attributes triggers a direct effect that pushes matching and wages in the opposite direction.

Effects on the Matching Function. Under both FOSD and IR shifts in the distribution of aggregate shock L, the effect on matching function (4) is unambiguous when our comparative statics result holds: and increase in t *reduces* θ^* and thus μ since $\mu_t = (H_0 - H_1)q\theta^*_t/g$ is negative. This is due to the improvement in the (endogenous) skill distribution $H = (1-Q(\theta^*))H_1 + Q(\theta^*)H_0$ as more workers invest, reinforcing the competition for firms with better y. As a result, any worker with skill x is matched to a firm with lower productivity y.

If there is a FOSD shift in firm productivity G, once we take into account the change in θ^*, the derivative of μ with respect to t is $(q\theta_t^*(H_0-H_1)-G_t)/g$, which is ambiguous. This is intuitive since on the one hand the firm quality distribution improves and this increases μ for each x, but on the other hand workers' skill distribution improves with more investment and this worsens the match for each x. If we instead consider the FOSD shift in H_1, then both the change in H_1 (improvement in exogenous skill distribution) and on θ^* (improvement in endogenous skill distribution) reinforce each other and the total derivative of μ with respect to t is $(q\theta_t^*(H_0-H_1)+(1-Q(\theta^*))\partial H_1/\partial t)/g$, which is negative. Hence, more workers investing implies worse matching outcomes for each worker with a given skill x. Similarly with a FOSD shift in Q.

Effects on the Wage Function. Induced changes in the matching function directly translate into changes in the wage function (5). Consider a FOSD or an IR shift in shock distribution L. Then for each (x, α) we have the change in wage $w_t = \int_0^x f_{xy}\mu_{\theta^*}\theta_t^*$, which is negative when our comparative results hold, so that $\theta_t^* \leq 0$. Since more workers invest, the distribution of their attributes improves, deteriorating the matching and thus decreasing the wage for each (x, α).

Under a FOSD shift in G, H_1, or Q, the change in w is given by $w_t = \int_0^x f_{xy}(\mu_t + \mu_{\theta^*}\theta_t^*)$. The term in parenthesis is negative with a FOSD shift in H_1 or Q but ambiguous with a FOSD shift in G. Thus, wages decrease for all (x, α) if Q or H_1 improve but can increase or decrease with an improvement in G. The intuition follows from the effect on the matching function μ.[17]

5 Economic Relevance

This section sheds light on the economic relevance of our theory. Although there are many economic applications of our framework, we will focus on how it can enhance our understanding of the roles of heterogeneity (innate or through choices) versus risk in labor market outcomes.

Distinguishing between cross-sectional heterogeneity and labor market risk is crucial for pinning down the determinants of life-time earnings and consumption dispersion (Keane and Wolpin (1997), Storesletten, Telmer, and Yaron (2004), Huggett, Ventura, and Yaron (2011)), and for analyzing changes in income inequality (Cunha and Heckman (2016)). As a result, it is also important for designing effective redistributive or efficiency-enhancing policies.

We contribute to this research area with a tractable equilibrium framework that features both rich heterogeneity of workers and jobs as well as various sources of risk, in the presence of endogenous investment – a margin important to our conclusions below but often overlooked. We explore three areas where our framework can be useful. First, we use our comparative statics results to gain insight into the rise in household income inequality. Whether risk or innate heterogeneity is the driving force is important for policy design. Second, we show that our model

[17] In turn, the behavior of the *expected wage* is ambiguous, since H_1 or L change as well. For example, a FOSD shift in H_1 improves H, and this increases the expected wage while the point-wise decrease in w decreases it.

naturally lends itself to study the welfare impact of a variety of policies that we see in practice, ranging from social insurance to policies that impact initial heterogeneity or risk directly. Finally, while our analysis shows that the distinction between risk and heterogeneity is important when thinking about policy design, it is usually difficult to separate the two in the data. In this third part, we discuss how our results can be used to identify whether observed changes in outcomes are due to changes in risk versus changes in innate worker and job heterogeneity.

5.1 Inequality

An issue that has received much attention in the recent literature is the rise in U.S. (household) income inequality. The emerging consensus is that this is more due to an increase in the returns to education over time than to an increase in positive assortative matching in marriage (see Eika, Mogstad, and Zafar (2018) and the references therein).

This is an important insight but it is based on descriptive decompositions that do not account for the underlying forces of changes in educational returns or for their equilibrium effects on educational investment – something our framework can do. And although our model does not contain a marriage market, we can still use it to show that accounting for these issues may be important to better understand the increase in inequality (especially since changes in marriage sorting seem to be of secondary importance for the rise in inequality).

To this end, we develop a parametric example of our model in Appendix B.1. We compute the ratio of average wages of educated over non-educated workers, $\mathbb{E}[w(x)|a=1]/\mathbb{E}[w(x)|a=0]$, as a measure of returns to education, and the wage variance, $Var[w(x)]$, as a measure of inequality. We first show that a FOSD shift in either idiosyncratic risk H_1, firm heterogeneity G, or aggregate risk L unambiguously increases the returns to education. Thus, the observed rise in these returns could be due to an improvement in schooling quality (FOSD shift in H_1); or in firm productivity (FOSD shift in G) through technological change that turns out to be skill-biased in equilibrium; or in aggregate risk (FOSD shift in L), which due to complementarities in production favors educated workers more. Importantly, while any of those shifts unambiguously increases the returns to education, they have an *ambiguous* effect on the variance of wages and thus on inequality (see Appendix B.1, Figure 3), which can actually *decrease* under some conditions.

What is crucial for these results is the *endogenous* supply response. If instead the investment threshold θ^* was fixed, then a FOSD shift in L or G would have the expected effect of always *increasing* the variance of the distribution of wages. Similarly, with a fixed θ^*, a FOSD shift in H_1 leads to an increase in inequality for most parameter choices. However, with an endogenous increase in investment in response to the distributional shifts, these direct effects are dampened by equilibrium effects: The matching function and thus also the wage function decrease for every worker type x. Educated workers are most affected by this, pushing the wage variance down.

As a result, caution is warranted when making statements about educational returns being responsible for the increase in household wage inequality. By considering only the direct effect of the increase in the returns to education on income inequality, statistical decompositions – that keep *education fixed* when changing returns – likely overstate its impact. And if general equilibrium effects are strong enough, such decompositions can even lead to qualitatively wrong conclusions, picking up a spurious correlation between returns on education and inequality. Moreover, pinning down the underlying cause for the rise in education premium (for example, FOSD shift in G, H_1, or L) is important for designing effective policies. Depending on the main driving force, policy interventions should be directed at providing insurance against initial conditions (worker heterogeneity), technological change (changes in firm heterogeneity) or shocks (risk).

5.2 Efficiency

Our framework is also useful for studying policies affecting initial heterogeneity or risk in the presence of endogenous investment. For instance, we can analyze policies aiming to improve workers' initial conditions (e.g. through public education) or to provide insurance against shocks (e.g. through unemployment insurance). These policies can attenuate the inefficiency in investment and also try to complete the missing insurance markets or improve the environment in which agents interact by changing its riskiness or individuals' initial conditions. An in-depth policy analysis is beyond the scope of this paper. Instead, for some commonly observed policies we highlight interesting trade-offs that are overlooked when taking educational investment as *exogenous*.

Social Insurance Policy. Consider a policy that insures workers with a transfer, s, for low realizations of the aggregate shock, $\alpha < \hat{\alpha}$, financed by risk-neutral firms. This is a stylized way of capturing the expansion of unemployment benefits in recessions.[18] The planner solves[19]

$$\max_s \quad \lambda \left(\int_0^{\hat{\alpha}} \int u(w(x,\alpha,\theta^*) + s) dH(x,\theta^*) dL(\alpha) + \int_{\hat{\alpha}}^1 \int u(w(x,\alpha,\theta^*)) dH(x,\theta^*) dL(\alpha) \; - \int_{\theta^*}^1 c(\theta) dQ(\theta) \right) \\ + (1-\lambda) \left(\int \int \left(f(x,\mu(x,\theta^*),\alpha) - w(x,\alpha,\theta^*) \right) dH(x,\theta^*) dL(\alpha) - sL(\hat{\alpha}) \right).$$

The first term in parentheses represents the expected utilities of workers, while the second term is the corresponding expression for firms, where weighted by λ and $1-\lambda$. After some algebra the first-order condition (FOC) of this problem is (see Appendix B.2.1)

$$0 = \int_0^{\hat{\alpha}} \int \left(u'(w+s) - \frac{1-\lambda}{\lambda} \right) dH dL \\ + \theta_s^* \left[\int_0^{\hat{\alpha}} \int \left(u'(w+s) - \frac{1-\lambda}{\lambda} \right) w_{\theta^*} dH dL + \int_{\hat{\alpha}}^1 \int \left(u'(w) - \frac{1-\lambda}{\lambda} \right) w_{\theta^*} dH dL \right]. \quad (14)$$

[18]If we interpret the shock α as an *idiosyncratic* productivity shock, the natural interpretation of this policy would be that agents receive unemployment benefits when their shock realization is sufficiently low.

[19]We can allow the planner to also choose $\hat{\alpha}$. Since this reveals the same trade-off, we omit this extension.

The planner uses this policy with two goals: to provide insurance to risk-averse workers against bad realizations of the shock (i.e. to complete missing insurance markets); and to address the inefficiency of investment in the decentralized equilibrium stemming from a wage externality that individuals do not internalize (recall Proposition 2). The first term in (14) reflects the insurance motive. She transfers s from risk-neutral firms to risk-averse workers when $\alpha < \hat{\alpha}$ (provided that λ is high enough, in which case an increase in s reduces that term). In turn, the second term in (14) is due to the inefficiency of equilibrium investment and parallels the second term on the right side of (8). If this term is positive, then workers invest too much and the planner uses this policy to also correct this over-investment. A higher transfer *reduces* the incentives to use education as an insurance device (θ^*_s is negative, see Appendix B.2.1), mitigating over-investment. Thus, the crowding out of private insurance by social insurance is welfare-improving. But if there is under-investment in equilibrium, then the planner's motives clash. While a higher transfer fills in for missing insurance markets, which increases welfare, the policy crowds our private insurance which was already inefficiently low. Due to this countervailing force, *less* social insurance is optimal compared to the case when educational investment is exogenous.

Policies Affecting Risk or Initial Conditions. We now allow the planner to directly affect L, H_1, or Q. An example is a regulation that makes the environment less risky (reverse IR shift in L) or stochastically better (FOSD shift in L). Another is an improvement in the quality of public education (FODS in H_1) or early childhood interventions (FOSD in Q). We capture these policies in reduced form by parameterizing the cdf's by $t_j, j \in \{\alpha, x, \theta\}$, which the planner can change at the cost $\kappa(t_j)$, financed by a proportional tax τ on firms' profits.

We here focus on the planner altering L. For concreteness, think of the aggregate shock α as TFP plus noise where this policy raises TFP by facilitating technology adoption, and thus L shifts in a FOSD sense. We derive the remaining policies that influence H_1 and Q in Appendix B.2.

All of these policies have a direct desirable effect, which conforms with the standard intuition about what they should achieve. But there is also a new indirect effect stemming from equilibrium investment and the associated inefficiency. This effect attenuates the welfare-enhancing effects of these policies. We stress that our comparative statics results on L, H_1 and Q are crucial to sign the relevant terms in the planner's FOC and thus to fully understand the impact of those policies.

Assume that $k(0) = 0$, $k' > 0$, and $k'' \geq 0$. The planner solves,

$$\max_{t_\alpha \in [0,\bar{t}]} \lambda \left(Q(\theta^*) \int \int u(w(x,\alpha,\theta^*)) dL(\alpha|t_\alpha) dH_0(x) + (1 - Q(\theta^*)) \int \int u(w(x,\alpha,\theta^*)) dL(\alpha|t_\alpha) dH_1(x) - \int_{\theta^*}^1 c(\theta) dQ(\theta) \right) + (1-\lambda)(1-\tau) \left(\int \int \left(f(x, \mu(x,\theta^*), \alpha) - w(x,\alpha,\theta^*) \right) dL(\alpha|t_\alpha) dH(x,\theta^*) \right)$$

$$s.t. \quad \kappa(t_\alpha) = \tau \int \int (f(x,\mu(x,\theta^*),\alpha) - w(x,\alpha,\theta^*)) dH dL(\alpha|t_\alpha)), \quad 0 \leq \tau \leq 1, \tag{15}$$

where we include the constraints of budget balance and that the proportional tax needs to be between zero and one. After manipulating the constraints, the FOC reads:[20]

$$\int\int\left(u+\frac{1-\lambda}{\lambda}(f-w)\right)dHdL+\theta^*_{t_\alpha}\int\int\left(u'-\frac{1-\lambda}{\lambda}\right)w_{\theta^*}dHdL=\kappa'\frac{2-\lambda}{\lambda}. \tag{16}$$

For concreteness, let t_α shift L in FOSD sense. The first term on the left side of (16) captures the intuitive benefits of this improvement in risk: both profits $f-w$ as well as worker utility $u(w)$ increase in t_α since they increase in the shock α. Without endogenous investment, the planner would simply balance this marginal benefit with the marginal cost, given by the right side of (16).

The second term on the left side of (16), however, arises because investment is endogenous, since the planner also wants to correct its inefficiency. For example, consider the case of over-investment in equilibrium. Then this term is *negative* (the policy induces further over-investment), reducing the marginal benefit from this policy. To see this, note that the integral is positive in this case but – based on our equilibrium comparative statics – θ_{t_α} is negative (better risk leads to *more* investment). Hence, not taking into account the endogenous supply response can overstate the benefits of the policy. An analogous conclusion obtains from a reverse IR shift in L.

Similar effects and trade-offs ensue when the planner's policies directly change H_1 or Q (see Appendix B.2.2 and B.2.3). The common thread of all these policies is that the intuitive beneficial effects of improving the environment can be tamed by the presence of endogenous investment.

5.3 Identifying Changes in Heterogeneity versus Changes in Risk

To design policies such as the ones above it is not only important to understand their effects, but also to distinguish in the data whether outcomes are driven by initial conditions or shocks (see discussion in Heathcote, Storesletten, and Violante (2009)). To be sure, this is a challenging task.

We now illustrate that our model could help disentangle changes in risk from changes in heterogeneity in the data. Our approach does not require to structurally estimate the model but relies on the reduced form wage function of our model, which can be estimated. Endogenous educational investment is crucial to this approach, which we believe is novel.

To obtain a reduced form wage function, we derive the unique equilibrium of the following parametric example in Appendix B.3.1.

Example. Assume risk neutrality, f given by $f(x,y,\alpha)=\alpha^2xy$, c by $c(\theta)=1-\theta$, G is uniform on $[0,t_y]$, H_1 is given by $H_1(x)=x^{t_x}$ and H_0 by $H_0(x)=x$, $t_x>1$, Q is uniform on $[0,t_\theta]$ and L on $[0,t_\alpha]$. The parameters $(t_y,t_x,t_\theta,t_\alpha)$ shift the various distributions in the sense of FOSD.

[20] Combining the two constraints, we can rewrite the second one as $\kappa(t_\alpha)\le\int\int(f(x,\mu(x,\theta^*),\alpha)-w(x,\theta^*,\alpha))dHdL(\alpha|t_\alpha)$. We can always ensure that this is slack by a suitable choice of κ, and so we drop it from our discussion. We simply substitute the budget balance constraint into (15) and derive the FOC.

From the unique equilibrium we obtain the following *reduced-form wage equation*:

$$\tilde{w}(x) = \beta_0 + \beta_1 x^2 + \beta_2 x^{t_x+1} + \epsilon, \tag{17}$$

where we have already substituted in both the equilibrium matching function and investment threshold, and where we added a mean-zero measurement error ϵ to our otherwise deterministic model, which is assumed to be orthogonal to (x^2, x^{t_x+1}). Note that, when taking our model at face value, (possibly unobserved) ability impacts wages only through skill x.

While β_0 is the constant of integration of the equilibrium wage function, β_1 and β_2 depend on the structural parameters $(t_y, t_x, t_\theta, t_\alpha)$, which we would like to identify. In Appendix B.3.1, we use the wage regression (17) to prove that (i) if t_x is known and there is no aggregate risk, so $t_\alpha = 0$, then (t_y, t_θ) are identified; and that (ii) if t_x and t_θ are known, then (t_y, t_α) are identified.

Thus, based on two regression coefficients (the curvature of w), we can typically identify two parameters that shift the risk and heterogeneity distributions. And by estimating (17) at different points in time, the various FOSD shifts can be distinguished by changes in the wage curvature.

We stress that endogenous investment is necessary to disentangle shifts in heterogeneity versus aggregate risk (part (ii)) since t_α only enters the wage through θ^*. Also, (i)–(ii) assume that we can identify the skill distribution parameter t_x, which determines the extent of idiosyncratic risk, from the data and hence that it is known. But even if t_x is unknown, it can be identified using non-linear estimation techniques that pin down the polynomial of unknown degree in (17), and thus by estimating the wage function over time we can also single out changes in idiosyncratic risk.[21]

Similarly, repeated (over time) estimation of the reduced form wage regression allows us to disentangle IR shifts from FOSD in the various distributions. For illustration, we make a slight change to our example and assume instead that L is uniform on $[0.5 - t_\alpha, 0.5 + t_\alpha]$ and Q is uniform on $[.5 - t_\theta, .5 + t_\theta]$. Thus, (t_θ, t_α) now shift the ability and risk distribution in an IR sense.

We again obtain a reduced-form wage equation like (17), but now the coefficients are different functions of the structural parameters. In Appendix B.3.2 we use the wage regression equation to show that (i) if t_x is known and there is no aggregate risk $t_\alpha = 0$, then (t_y, t_θ) are identified; and that (ii) if t_x and t_θ are known, then (t_y, t_α) are identified. That is, we can also use the reduced form wage regression to distinguish FOSD from IR shifts in aggregate risk and heterogeneity.

6 Robustness

This section explores variations of the model where we relax some of the assumptions. In each case we point out how the main insights are affected.

[21] See for instance Cho and Phillips (2016).

6.1 Timing of the Aggregate Shock Realization

We assume that α realizes after investment but before matching. This is the natural timing in many economic applications: when workers invest, they are uncertain about future labor market conditions. Also, we want to capture in a simple way situations where α *cannot* be fully insured. This would not be the case if the shock were realized after the match, since firms would fully insure workers against the risk in α, and changes in its distribution would play no role. An alternative is to assume some friction that precludes full insurance. In Appendix A.11 we explore the case when the division of match output is exogenously fixed meaning utility is non-transferable (which by assumption makes the shock uninsurable), which is a common case analyzed in the literature (e.g., Becker (1973), Smith (2006), Legros and Newman (2010)). Similar results obtain.

A more complex variation left for future research is to incorporate moral hazard or adverse selection at the matching stage, which precludes full insurance. This would call for a nontrivial extension of Legros and Newman (2007) and the use of their sorting condition. So on top of the conditions on risk attitudes and complementarities that we highlight, how workers and firms can transfer utility would play a role. But if the the model has a 'transferable utility' representation (Legros and Newman (2007), Section 6.1), then the analysis would be similar to our case.

A related variation is when both sides are risk averse, as in marriage market applications. Little changes if we maintain the timing of the shock realization, since once uncertainty is realized, payoffs are a strictly monotone transformation of income. The analysis is more involved if the shock realizes after matching, and once again requires the sorting condition of Legros and Newman (2007). If both sides are symmetric, however, then a lot can be said, as we show below.

6.2 Two-Sided Investments

A natural extension is to allow firms to invest as well prior to matching. To this end, let $\omega \in [0, 1]$ be the initial characteristic of each firm, distributed according to R, and let k be the firms' cost of investment, with the same properties as the workers' cost function c. If a firm invests, then its productivity parameter y is drawn from G_1, while if it does not invest it is drawn from G_0, where $G_0 \leq G_1$ with strict inequality on a positive-lengthed interval. The rest is as before.

Intuitively, a firm invests if and only if $\omega \geq \omega^*$. The two equations that now characterize the equilibria of the model are:

$$U_1(\theta^*, \omega^*, t) - U_0(\theta^*, \omega^*, t) = c(\theta^*) \tag{18}$$

$$V_1(\omega^*, \theta^*, t) - V_0(\omega^*, \theta^*, t) = k(\omega^*), \tag{19}$$

where $U_1 - U_0 = \int\int u' w_x \Delta H dx dL$ and, similarly, $V_1 - V_0 = \int\int \pi_y \Delta G dy dL$.

Equation (18) defines a strictly increasing "best response" function r_w given by $\theta^* = r_w(\omega^*)$ for workers, and (19) defines a similar function r_f given by $\omega^* = r_f(\theta^*)$ for firms. Any intersection of r_w and r_f is an equilibrium; it is stable if $\partial r_w^{-1}/\partial\theta^* \geq \partial r_f/\partial\theta^*$ when they cross.

We show in Appendix A.12 that when risk aversion is small a stable equilibrium exists and the results on changes in L hold essentially as before. But those that operate via matching and wages are ambiguous since they have opposite effects on each side of the market.

6.3 Continuous Investments

We assumed that investment is binary. Assume instead that a worker can invest in any amount $a \in [0,1]$, in which case x is drawn from $H(\cdot|a)$, given by $H(x|a) = aH_1(x) + (1-a)H_0(x)$, with $H_1 \leq H_0$ with strict inequality on a positive-lengthed interval. If a worker invests a and his characteristic is θ, then his disutility is $d(a,\theta)$. The function d is continuous in (a,θ), $d(\cdot,\theta)$ is strictly increasing, continuously differentiable, and convex in a for all θ, $d(0,\theta) = 0$ for all θ, and $d(a,\cdot)$ is strictly decreasing in θ for all a. In addition, d is strictly submodular in (a,θ). To simplify some of the arguments and the notation, assume that θ is uniformly distributed on $[0,1]$.

A strategy profile in the first stage with continuous investments is a measurable function $\hat{a} : [0,1] \to [0,1]$, which maps $\theta \in [0,1]$ to actions $a \in [0,1]$.

Let $q(\hat{a}) = \int \hat{a}(\theta)d\theta$ be the average level of investment under $\hat{a}$. Then the distribution of workers in the second stage given investment strategy $\hat{a}$ is $H(x|\hat{a}) = \int H(x|\hat{a}(\theta))d\theta = q(\hat{a})H_1(x) + (1-q(\hat{a}))H_0(x)$. Hence $\mu(x,\hat{a}) = G^{-1}(H(x|\hat{a})) = G^{-1}(q(\hat{a})H_1(x) + (1-q(\hat{a}))H_0(x))$ and thus $w(x,\alpha,\hat{a}) = \int_0^x f_x(s,\mu(s,\hat{a}),\alpha)ds$. Notice that (μ, w) depend on $\hat{a}$ only through q.

An equilibrium reduces to finding a measurable $\hat{a}$ that satisfies the usual best response property. Appendix A.13 shows that an equilibrium exists using an existence theorem for large games in Rath (1992), and that equilibrium investment is inefficient. Regarding the comparative statics results, one can show the results extend for the average investment level in the smallest and largest equilibrium. There is, however, a class in which *all* the results extend, which we explore next.

6.4 The Partnership Model

A variation of our model is the one-population case, where a continuum of agents from a given population match in pairs. This is commonly used in matching models that analyze partnership formation (e.g., Kremer (1993), Kremer and Maskin (1996), and Legros and Newman (2002)).

The model is as before except that there are no firms, only a unit measure of risk averse agents with ability $\theta \in [0,1]$ distributed according to Q and, for each α, the match output function is symmetric in the partners' characteristics, so $f(x,x',\alpha) = f(x',x,\alpha)$ for all $x, x' \in [0,1]$.

Since f is strictly supermodular, there will be positive sorting in the second stage for any investment function a, so the matching function is given by $\mu(x) = x$. Also, partners in equilibrium split match output in half, so $w(x,\alpha) = f(x,x,\alpha)/2$. Note the crucial feature of the second stage equilibrium: matching and wages are *independent* of a.

The following insights emerge in this setting (see Appendix A.14). First, since μ and w are independent of the mass of people who invest, we can easily obtain existence, uniqueness, and unambiguous comparative statics. This applies to both the binary and the continuous investment cases. Second, since the equilibrium of the partnership model can be reinterpreted as a symmetric equilibrium with two identical populations, the fact that both sides are risk averse here does not affect the main results given our timing. Third, this variation is a simple instance where two-sided investments can be easily analyzed. Finally, it provides a tractable setting where we can assume that the aggregate shock is realized after matching and utility is imperfectly transferable.

6.5 Second Stage Matching Market

Finally, we have chosen to pin down the wage function by assuming that the labor market in the second stage is a matching market à la Becker (1973). This gives us a tractable way to understand the equilibrium effects in our set up. But it is not the only plausible way to model the second stage. To keep the door open for other interpretations of the second stage, and to provide some guidance to empirical researchers who want to estimate a partial equilibrium version of the model, we also state the comparative statics results using an arbitrary wage function with certain properties (see Lemmas 1-4 in the Appendix).

7 Concluding Remarks

This paper develops an equilibrium model that features both rich heterogeneity of workers and firms, and idiosyncratic and aggregate risk, in the presence of endogenous educational investment.

Our main contribution is the equilibrium comparative statics analysis of how changes in heterogeneity and in labor market risk affect educational investment, matching and wages. Underlying all the results is a subtle trade off between complementarities in the equilibrium wage function and workers' risk attitudes. We provide intuitive conditions on risk aversion and prudence, and on match output for changes in risk and heterogeneity (either a FOSD or an IR shift) to induce more workers to invest, which we argue is the natural comparative statics result to seek.

We demonstrate the economic relevance of our framework by exploring three areas that our comparative statics can shed light on: the rise in household inequality, the welfare impact of policies we see in practice, and the determination of whether observed changes in outcomes are due to changes in risk versus changes in worker/firm heterogeneity.

Our model overcomes a major shortcoming of standard assignment models, which assume that agents do not adjust their characteristics with changes in the economic environment and, in particular, with changes in risk and heterogeneity. We believe that our framework can serve as a building block for richer environments – e.g., with multi-dimensional heterogeneity or in dynamic settings – that are suitable for structural empirical work.

A Appendix

A.1 Proof of Walrasian Equilibrium

We claim that, given (a, α), $(\mu(\cdot, a), w(\cdot, a, \alpha))$ is a Walrasian equilibrium of the labor market, where $\mu(x, a) = G^{-1}(H(x, a))$ for all x, $w(x, a, \alpha) = \int_0^x f_x(s, \mu(s, a), \alpha)ds$ for all x. That is, the market clears, and agents behave optimally. By construction, $\mu(\cdot, a)$ clears the market.

Consider firm y. It solves

$$\max_{x \in [0,1]} \left(f(x, y, \alpha) - \int_0^x f_x(s, \mu(s, a), \alpha)ds \right),$$

and from the first-order condition (FOC) we obtain $f_x(x, y, \alpha) = f_x(x, \mu(x, a), \alpha)$ and hence $y = \mu(x, a)$ or $x = \mu^{-1}(y, a)$. To show that this is a global optimum, note that for any $x > x'$

$$f(x, y, \alpha) - \int_0^x f_x(s, \mu(s, a), \alpha)ds \geq f(x', y, \alpha) - \int_0^{x'} f_x(s, \mu(s, a), \alpha)ds$$

if and only if

$$f(x, y, \alpha) - f(x', y, \alpha) \geq \int_{x'}^x f_x(s, \mu(s, a), \alpha)ds. \tag{20}$$

Since $y = \mu(x, a)$ and since, by the Fundamental Theorem of Calculus, $f(x, y, \alpha) - f(x', y, \alpha) = \int_{x'}^x f_x(s, y, \alpha)ds$, it follows that (20) is equivalent to $\int_{x'}^x (f_x(s, \mu(x, a), \alpha) - f_x(s, \mu(s, a), \alpha))\, ds \geq 0$, and this holds from the supermodularity of f and from $\mu(x, a) \geq \mu(s, a)$ for all $s \in [x', x]$. Hence, choosing x is preferred than choosing $x' < x$ for y. A similar argument holds for $x' > x$, thus establishing that choosing x is a global optimum for y.

In turn, a worker with attribute x obtains $u(w(x, a, \alpha)) \geq u(0)$. Let $\pi(y, a, \alpha)$ be the profit of a firm with attribute y. If y is matched with x, we have $\pi(y, a, \alpha) = f(\mu^{-1}(y, a), y, \alpha) - w(\mu^{-1}(y, a), a, \alpha)$. By the Envelope Theorem, we can rewrite it as $\pi(y, a, \alpha) = \int_0^y f_y(\mu^{-1}(s, a), s, \alpha)ds$, where we omitted $\pi(0, a, \alpha)$ since it is zero (recall that $f(0, 0, \alpha) = 0$).

It is now easy to check that workers behave optimally too (it is implied by the above analysis), in the sense that each x maximizes her payoff at $y = \mu(x, a)$. We can think of each x as solving

$$\max_{y \in [0,1]} u \left(f(x, y, \alpha) - \int_0^y f_y(\mu^{-1}(s, a), s, \alpha)ds \right).$$

Since the objective function is a strictly increasing transformation of $f(x, y, \alpha) - \pi(y, a, \alpha)$, the FOC holds if and only if $f_y(x, y, \alpha) = f_y(\mu^{-1}(y, a), y, \alpha)$, so $x = \mu^{-1}(y, a)$ or $y = \mu(x, a)$. The proof of global optimality of this choice follows as in the firms' case and is omitted. □

A.2 Proof of Proposition 1

Existence of an equilibrium was proven in the text. Regarding uniqueness, it suffices to show that $U_1 - U_0 = \int\int u'w_x(H_0 - H_1)dxdL$ is increasing in θ^* when $R(w) = -u''(w)/u'(w)$ is sufficiently small (uniformly in w). To this end, let $r \in [0,1]$ be a parameter indexing u, so that $u(w,0) = w$ (risk neutral), and increases in r leads to a strictly increasing and concave transformation of u (so if $r > r'$, then $R(w,r) \geq R(w,r')$ for all w). Since R cannot be uniformly small in w if it is unbounded above, we will assume that R is uniformly bounded above for all r, and also that u, u', and u'' are continuous in r for each w.

Differentiating $\int\int u'w_x\Delta HdxdL$ with respect to θ^* yields

$$\begin{aligned}\int\int\left(u''w_xw_{\theta^*} + u'w_{x\theta^*}\right)\Delta HdxdL &= \int\int u'(w,r)\left(w_{x\theta^*} - R(w,r)w_xw_{\theta^*}\right)\Delta HdxdL \\ &\geq \int\int u'(w,r)\left(w_{x\theta^*} - R(w,r)\max_{x,\theta^*,\alpha}(w_xw_{\theta^*})\right)\Delta HdxdL \\ &= \int\int u'(w,r)\left(w_{x\theta^*} - R(w,r)\gamma\right)\Delta HdxdL,\end{aligned}$$

where we have set $\gamma \equiv \max_{x,\theta^*,\alpha}(w_xw_{\theta^*}) > 0$, which is positive and finite since w_x and w_{θ^*} are positive almost everywhere and continuous on $[0,1]^3$.

At $r = 0$, $R(w,0) = 0$ and $u'(w,0) = 1$ for all w, so the last expression above becomes $\int\int w_{x\theta^*}\Delta HdxdL$, which is positive since $w_{x\theta^*}$ is positive almost everywhere. Hence, equilibrium is unique in the risk neutral case. From the continuity of $u'(w,\cdot)$ and $R(w,\cdot)$ for all w, it follows that $\int\int\left(u''w_xw_{\theta^*} + u'w_{x\theta^*}\right)\Delta HdxdL$ is also positive for r sufficiently small, that is, when the coefficient of absolute risk aversion is uniformly small in w, and equilibrium uniqueness follows.

Finally, to see that equilibrium is also unique if c is sufficiently convex, index c by $j \in \mathbb{R}_+$, and assume that $c(\cdot,j)$ is convex, strictly decreasing, and continuously differentiable on $(0,1)$, with $c(1,j) = 0$ and $\lim_{\theta\to 0} c(\theta,j) = +\infty$ for all j. Also, assume that $c(\theta,\cdot)$ is strictly decreasing in j and converges to 0 for all $\theta \in (0,1)$ as $j \to \infty$.

We will focus on the first (lowest) crossing between $U_1 - U_0$ and c, and show that for j sufficiently large it becomes the unique crossing and thus the unique equilibrium. To this end, note that since $U_1 - U_0 > 0$ for all $\theta^* \in [0,1]$ and continuous in θ^*, $\beta = \min_{\theta^*\in[0,1]}(U_1(\theta^*) - U_0(\theta^*)) > 0$.

We first show that the first crossing converges to 0 as j goes to infinity. Let $\theta^*_\ell(j) = \inf\{\theta^* | U_1(\theta^*) - U_0(\theta^*) = c(\theta^*,j)\}$ be the lowest equilibrium threshold given j. We know that $\theta^*_\ell(j) > 0$ for all j (since c is unbounded near 0 and $U_1 - U_0$ is finite). We claim that as j increases θ^*_ℓ strictly decreases and converges to 0. To see this, note that c strictly decreases in j for each $\theta \in (0,1)$ and so does the first crossing, and this implies that θ^*_ℓ strictly dereases in j. Now, given $\varepsilon > 0$ there exists N_1 such that $0 < c(\theta^*,j) < \beta$ for all $j \geq N_1$ and $\theta^* \geq \varepsilon$. Hence, any equilibrium has $\theta^* \in (0,\varepsilon)$ for all $j \geq N_1$. Since $\varepsilon > 0$ was arbitrary, $\lim_{j\to\infty}\theta^*_\ell(j) = 0$.

Now, $U_1 - U_0$ continuously differentiable in θ^* on $[0,1]$ implies that $|\partial(U_1 - U_0)/\partial\theta^*| < M < \infty$. In turn, the derivative of c is unbounded near $\theta^* = 0$ for all j. To see this, fix j and let $\theta_0^* < \theta_1^*$; by the Mean Value Theorem, there exists $\theta_2^* \in (\theta_0^*, \theta_1^*)$ such that $c(\theta_1^*, j) - c(\theta_0^*, j) = c_\theta(\theta_2^*, j)(\theta_1^* - \theta_0^*)$. As θ_0^* goes to 0, the left side goes to $-\infty$ and hence so does $c_\theta(\theta_2^*, j)$.

We now show that there exists an N such that for all $j \geq N$ equilibrium is unique. To see this, note that there exists an N_2 such that for all $j \geq N_2$, $c_\theta(\theta_l^*(j), j) < -M < \partial(U_1(\theta_\ell^*(j)) - U_0(\theta_\ell^*(j))/\partial\theta^*$, and the same holds in a neighborhood of $\theta_\ell^*(j)$ for all $j \geq N_2$. Take any $j_0 \geq N_2$ and consider $[0, \theta_\ell^*(j_0)]$. There exists an $N > N_2$ such that $c(\theta^*, j) < \beta$ for all $\theta^* \geq \theta_\ell^*(j_0)$ and $j \geq N$, and such that there is a unique crossing below $\theta_\ell^*(j_0)$. The last point follows because $c_\theta(\theta^*, j) < -M < \partial(U_1(\theta^*(j)) - U_0(\theta^*(j))/\partial\theta^*$ for all $j \geq N$ and $\theta^* \in (0, \theta_\ell^*(j_0)]$, so there cannot be more than one crossing in that interval, and there is at least one. $\square$

Regarding the example mentioned in the text right after Proposition 1, let $f = \alpha xy$, so $f_x = \alpha y$ and $f_{xy} = \alpha$. Also, assume that $g = 1$ and $\mathbb{E}[\alpha] > 0$. Finally, set $\overline{R} = \sup_w R(w, 1)$, which is an upper bound for risk averion for all (w, r). By replacing above $R(w, r)$ by $\overline{R}$ (which preserves the inequality), we obtain that $\int\int (u'' w_x w_{\theta^*} + u' w_{x\theta^*}) \Delta H dx dL \geq 0$ if

$$\overline{R} \leq \frac{\int\int u'(w,r) w_{x\theta^*} \Delta H dx dL}{\gamma \int\int u'(w,r) \Delta H dx dL},$$

where we recall that $\gamma = \max_{x,\theta^*,\alpha}(w_x w_{\theta^*}) > 0$. Since $u'(\cdot, r)$ is decreasing in w, it suffices to show that

$$\overline{R} \leq \frac{u'(\overline{w}, r)}{u'(0, r)} \frac{\int\int w_{x\theta^*} \Delta H dx dL}{\gamma \int\int \Delta H dx dL}.$$

Since $w_x = \alpha\mu(x, \theta^*) \leq 1$, with equality at $x = \alpha = 1$, and $w_{\theta^*} = \alpha \int_0^x \mu_{\theta^*}(s, \theta^*) ds = \alpha \int_0^x \Delta H ds \leq \int \Delta H dx$, with equality at $x = \alpha = 1$, we obtain that $\gamma = \int \Delta H dx$. Also, $w_{x\theta^*} = \alpha\mu_{\theta^*}(x, \theta^*) = \alpha \Delta H$, and thus $\int\int w_{x\theta^*} \Delta H dx dL = \mathbb{E}[\alpha] \int (\Delta H)^2 dx$. Finally, from the unique solution on $[0, \overline{w}]$ (for each r) to the differential equation $-u''(w, r)/u'(w, r) = R(w, r)$, we obtain $u'(\overline{w}, r)/u'(0, r) = e^{-\int_0^{\overline{w}} R(s,r) ds} \geq e^{-\overline{R}\overline{w}}$. Thus, $\int\int (u'' w_x w_{\theta^*} + u' w_{x\theta^*}) \Delta H dx dL \geq 0$ if

$$e^{\overline{R}\overline{w}} \, \overline{R} \leq \frac{\mathbb{E}[\alpha] \int (\Delta H)^2 dx}{(\int \Delta H dx)^2},$$

which holds for $\overline{R}$ small, which makes the coefficient of risk aversion uniformly small. The resulting bound is determined by $\overline{w}$, L, H_0, and H_1.

A.3 Proof of Proposition 2

The equilibrium of the second stage is efficient for each (θ^*, α). Indeed, $u(w(x, \theta^*, \alpha))$ for all x and $f(\mu^{-1}(y, \theta^*), y, \alpha) - w(\mu^{-1}(y, \theta^*), \theta^*, \alpha)$ for all y can be interpreted as the core utilities for workers and firms in the second-stage assignment game given (α, θ^*).

Let $\Delta h \equiv h_0 - h_1$. The first order condition of the planner's problem (7) is:

$$\lambda \left(q \left(\int \int u dH_0 dL - \int \int u dH_1 dL + c \right) + \int \int u' w_{\theta^*} dH dL \right)$$
$$+ (1-\lambda) \left(\int \int (f_y \mu_{\theta^*} - w_{\theta^*}) dH dL + q \int \int (f - w) \Delta h dx dL \right) = 0. \quad (21)$$

Integration by parts yields

$$\int u dH_0 - \int u dH_1 = - \int u' w_x \Delta H dx$$
$$\int (f - w) \Delta h dx = - \int (f_x + f_y \mu_x - w_x) \Delta H dx$$
$$= - \int f_x \Delta H - \int f_y \frac{h}{g} \Delta H dx + \int f_x \Delta H,$$

where the last equality uses the derivatives of μ and w and $h \equiv Q(\theta^*) h_0 + (1 - Q(\theta^*)) h_1$. Inserting these expressions into (21) and noting that $\int f_y \mu_{\theta^*} dH = q \int f_y (h/g) \Delta H dx$ yields

$$\lambda \left(q \left(- \int \int u' w_x \Delta H dx dL + c \right) + \int \int u' w_{\theta^*} dH dL \right) - (1-\lambda) \left(\int \int w_{\theta^*} dH dL \right) = 0, \quad (22)$$

which rearranges to (8). It follows that equilibrium is generically inefficient, since any equilibrium threshold θ^* is interior and solves $\int \int u' w_x \Delta H dx = c$, while the planner's optimal threshold θ_p^* is either at a corner or is interior and satisfies (22), which contains an extra term compared to the equilibrium condition, which vanishes if and only if $u' = 1$ and $\lambda = 1/2$. □

The intuition follows easily from (22), which shows that the planner wants to balance the marginal costs and benefits from increasing θ^*. The costs of increasing θ^* (i.e. of lower investment) is that workers forgo the additional utility from investment, the term $-\lambda q \int \int u' w_x \Delta H dx dL$, and firms enjoy lower profits, the term $-(1-\lambda) \left(\int \int w_{\theta^*} dH dL \right)$, since lower investment drives up the wage. In turn, the benefits from higher θ^* are the saved cost, the term $\lambda q c$, and the increased worker utility, the term $\lambda \int \int u' w_{\theta^*} dH dL$. In equilibrium, however, agents do not take into account the effect on wages when they invest, while the planner does. This leads to the two extra terms involving w_{θ^*}, compared to the equilibrium condition $\int \int u' w_x \Delta H dx dL = c$. When $u' = 1$ and the planner puts equal weights on both sides, then the change in wages when θ^* changes is simply a transfer between workers and firms that has no effect on efficiency. Otherwise it does, as risk aversion and/or $\lambda \neq 1/2$ leaves the planner no longer indifferent as to how match output is split.

The equilibrium may entail over or under-investment, since the second term in (8) can be positive or negative: for $w_{\theta^*} = q \int_0^x f_{xy} (\Delta H / g) ds$ increases in x while the first term in the integrand single crosses zero from above, so the overall sign of the integral is unclear. For example, if $u(w) = -(e^{-Rw} - 1)/R$ and $\lambda = 1/2$, then $u'(w) = e^{-Rw} \leq 1$, and thus $\int (u' - 1)\, w_{\theta^*} dH < 0$ and the equilibrium exhibits underinvestment. And for any given u, if λ is sufficiently close to one (so the term representing the effect of wages on workers expected utility receives most of the weight), then $\int (u' - ((1-\lambda)/\lambda))\, w_{\theta^*} dH > 0$, and thus the equilibrium exhibits over-investment.

Using the terminology from Nöldeke and Samuelson (2015), what we have analyzed is the efficiency of our 'ex post' equilibrium, where matching takes place after investments are sunk. Since workers' investment has stochastic returns and results in a match-relevant skill x, our ex post analysis ensures that matching is stable in the second stage. But it precludes risk sharing opportunities that would be available if firms and workers could match ex ante based on θ and y (so firms could insure workers against the risks x and α), which enhances efficiency but leads to incentives to re-match ex post once x is revealed. The inefficiency in our model, however, does not stem from the lack of risk-sharing (our planner, by assumption of our timing, does not take this source of inefficiency into account) but from agents not internalizing the effect of their investment on wages. This is why, contrary to Nöldeke and Samuelson (2015), our equilibrium with one-sided investment in generically inefficient. In their work, investment does not alter the distribution of workers' characteristics as in our case, and hence it does not impact directly the matching (and hence wage) function, thus precluding our source of inefficiency. A more complete analysis of efficiency with pre-match risky investments is an interesting topic for future research.[22]

A.4 Proof of Proposition 3

In the text we explain why supermodularity of ϕ in (x, t) was sufficient for our comparative statics results regarding shifts in L. Implicit in that argument was a monotone comparative static result that we now make explicit, which applies to *all* the changes in distributions that we examine.

Fix any $\theta^* \in [0, 1]$, and let $H(\cdot|a) \equiv aH_1(\cdot) + (1-a)H_0(\cdot)$. Consider first the case where t shifts L. Then we can write the problem a worker faces when choosing $a \in \{0, 1\}$ as follows:

$$\max_{a\in\{0,1\}} \int\int u(w(x, \alpha, \theta^*))dH(x|a)dL(\alpha|t) - ac(\theta), \tag{23}$$

or simply as $\max_{a\in\{0,1\}} J(a, t) - ac(\theta)$. Theorem 2.6.6 in Topkis (1998) implies that the optimal a is increasing in t if and only if J is supermodular in (a, t).[23]

By inverting the order of integration, we can write J as

$$J(a, t) = \int \phi(x, t)dH(x|a).$$

Theorem 3.10.1 in Topkis (1998) implies that J is supermodular in (a, t) if ϕ is supermodular in

[22] An earlier working paper version of Nöldeke and Samuelson (2015) contains a remark related to this issue.

[23] Necessity follows since the range of c is $\mathbb{R}_+$.

(x,t) and H is ordered by FOSD in a (which is indeed the case since $H_1 \leq H_0$). This in turn implies that $U_1 - U_0 = \int \phi_x \Delta H dx$ increases in t for any θ^*, which was our argument in the text.

We first provide conditions on absolute risk aversion R and the wage function w under which ϕ is supermodular in (x,t) and then pin down the result from primitives.

Lemma 1 (FOSD Shift in L) *In any stable equilibrium, more workers invest in skills in response to a FOSD shift in L if either of the following conditions hold:*

(i) Absolute risk aversion is (uniformly) sufficiently small and $w_{x\alpha} > 0$; or

(ii) For all values of w, $R(w) \leq 1/w$, and w is increasing and log-supermodular in (x,α).

Proof. Using integration by parts and the definition of ϕ we obtain

$$\int \phi_{xt} \Delta H dx = \int \int u'(w) \left[w_{x\alpha} - R(w) w_x w_\alpha \right] (-L_t) \, \Delta H d\alpha dx,$$

where $L_t \leq 0$ for all α. If it is 0 for all α, then there is nothing to prove (there is no FOSD shift), so we will assume that it is strictly negative on a set of α's of positive measure.

Regarding part (i), we prove a slightly more general results that only requires $w_{x\alpha} > 0$ on $(0,1]^2$. The proof mimics the uniqueness part of the proof of Proposition 1, with w_α and $w_{x\alpha}$ instead of w_{θ^*} and $w_{x\theta^*}$. It also subsumes the case where workers are risk neutral, so that $u'(w) = 1$ and $R(w) = 0$ for all values of w, and so w supermodular in (x,α) implies that $\int \phi_{xt} \Delta H dx \geq 0$.

Finally, to prove part (ii) it suffices to show that the term in the square brackets in the integrand is positive. But since $R(w) \leq 1/w$ for all values of w, we obtain $w_{x\alpha} - R(w) w_x w_\alpha \geq w_{x\alpha} - \frac{w_x w_\alpha}{w} \geq 0$, where the inequality follows from the log-supermodularity of w in (x,α). □

Proof of Proposition 3. Part (i) is immediate since the assumptions on f imply that the w is strictly increasing and $w_{x\alpha} > 0$ for almost all values of (x,α) and θ^*. Hence, the result follows from Lemma 1 (i) if R sufficiently small uniformly in w.

Consider part (ii). Rewrite the wage function (5) as follows:

$$w(x,\alpha,\theta^*) = \int \mathbb{I}_{[0,x]}(s) f_x(s, \mu(s,\theta^*), \alpha) ds. \tag{24}$$

It is well-known (Karlin and Rinott (1980)) that w is log-spm in (x,α) if the integrand is log-supermodular in (x,s,α) for each θ^*, and this holds if $\mathbb{I}$ is log-supermodular in (x,s) and f_x is log-supermodular in (s,α).

It is easy to verify that f_x in the integrand is log-supermodular in (s,α) if and only if

$$(f_{xx\alpha} f_x - f_{x\alpha} f_{xx}) + (f_{xy\alpha} f_x - f_{x\alpha} f_{xy}) \, \mu_x \geq 0,$$

which holds if the parentheses are positive, i.e., if f_x is log-supermodular in (x,α) and (y,α).

Similarly, if we take any two pairs (s, x), (s', x'), it is easy to check that

$$\mathbb{I}_{[0,x\vee x']}(s\vee s')\mathbb{I}_{[0,x\wedge x']}(s\wedge s')\geq \mathbb{I}_{[0,x]}(s)\mathbb{I}_{[0,x']}(s'),$$

and thus $\mathbb{I}$ is log-supermodular in (x, s).

Hence, w is log-spm in (x, α) for each θ^* and thus $w_{x\alpha}w/w_x w_\alpha \geq 1$. It follows that $R(w) \leq (1/w)(w_{x\alpha}w/w_x w_\alpha)$ if $R(w) \leq 1/w$. Therefore, the conditions in part (ii) of Lemma 1 hold and a FOSD shift in L increases $U_1 - U_0$, thereby decreasing the equilibrium value of θ^*. □

A.5 Proof of Proposition 4

In the following lemma, we first provide conditions on absolute risk aversion and prudence, R and P, and on the wage function w under which ϕ is supermodular in (x, t).

Lemma 2 (IR Shift in L**)** *In any stable equilibrium, more workers invest in skills in response to an IR shift in L if either of the following conditions hold:*

(i) Absolute risk aversion is (uniformly) sufficiently small and $w_{x\alpha\alpha} > 0$*; or*

(ii) For all w*,* $P(w) \geq 3/w$*, with* $w_{x\alpha\alpha} \geq 0$*, and* w *is increasing in* (x, α)*, log-concave in* α *for all* x *and log-submodular in* (x, α)*.*[24]

Proof. Using integration by parts twice on $\phi_x = \int u'(w)w_x dL$ we obtain

$$\begin{aligned}\int \phi_{xt}\Delta H dx &= \int\int \left(\frac{\partial}{\partial\alpha}(u''(w)w_\alpha w_x + u'(w)w_{x\alpha})\right)\left(\int_0^\alpha L_t ds\right)\Delta H d\alpha dx\\ &= \int\int u'(w)\left[R(w)\left(P(w)w_x w_\alpha^2 - 2w_{\alpha x}w_\alpha - w_x w_{\alpha\alpha}\right) + w_{x\alpha\alpha}\right]\left(\int_0^\alpha L_t ds\right)\Delta H d\alpha dx,\end{aligned}$$

where $\int_0^\alpha L_t ds \geq 0$ for all α. If it is 0 for all α, then there is nothing to prove (there is no IR shift), so we will assume that it is strictly positive on a set of α's of positive measure.

We prove a more general version of part (i) that only requires $w_{x\alpha\alpha} \geq 0$, with strict inequality on a set of (x, α) of positive measure. Assume that workers are risk neutral, so $u'(w) = 1$ and $R(w) = P(w) = 0$ for all w. Then $w_{x\alpha\alpha} \geq 0$ implies that $\int \phi_{xt}\Delta H dx \geq 0$. Using the argument from Proposition 1, the same holds for absolute risk aversion that is uniformly small in w. This is because the condition on $w_{x\alpha\alpha}$ guarantees that $\int w_{x\alpha\alpha}\left(\int_0^\alpha L_t ds\right)\Delta H d\alpha dx > 0$, and thus the inequality remains strict if R is sufficiently small.

To prove part (ii), it suffices to show that the first term inside the square brackets in the

[24] A function z is log-concave in x if $\log z$ is concave in x.

integrand is positive under the premises. Note that

$$\begin{aligned}P(w)w_xw_\alpha^2-2w_{\alpha x}w_\alpha-w_xw_{\alpha\alpha} &\geq 3\frac{w_xw_\alpha^2}{w}-2w_{\alpha x}w_\alpha-w_xw_{\alpha\alpha}\\ &=2\left(\frac{w_xw_\alpha}{w}-w_{\alpha x}\right)w_\alpha+\left(\frac{w_\alpha^2}{w}-w_{\alpha\alpha}\right)w_x\\ &\geq 0,\end{aligned}$$

where the first inequality uses the assumption $P(w)\geq 3/w$ for all w, and the second uses w submodular (x,α) and w log-concave in α for each x, which imply that both terms in parentheses are positive. Hence, since the integrand is positive it follows that $\int\phi_{xt}\Delta H dx\geq 0$. □

Proof of Proposition 4. The proof of part (i) follows from Lemma 2 (i), since $w_{x\alpha\alpha}=f_{x\alpha\alpha}\geq 0$, strictly so on a set of (x,y,α) of positive measure. Thus, the last term in (11) is positive.

Consider part (ii). Under the conditions on η and z, the last term of (11) is positive (it is given by $w_{x\alpha\alpha}=\eta''z_x\geq 0$), and $w_{x\alpha}w/w_xw_\alpha=(z_x\eta'\eta\int_0^x z_x)/(\eta z_x\eta'\int_0^x z_x)=1$. Thus, the expression in parentheses in the first term of (11) is positive if and only if

$$P(w)\geq\frac{1}{w}\left(2+\frac{\eta''(\alpha)\eta(\alpha)}{(\eta'(\alpha))^2}\right).$$

If η is log-concave in α, so $\eta''\eta/\eta^2\leq 1$, then it suffices that $P(w)\geq 3/w$ for all values of w. □

A.6 Proof of Corollary 1

As in the proof of Lemma 2, twice integrating by parts $\phi_x=\int u'(w)w_xdL$ yields

$$\begin{aligned}\int\phi_{xt}\Delta H dx &= -\int\left((u''(w)w_\alpha w_x+u'(w)w_{x\alpha})\int_0^\alpha L_tds\right)\Big|_{\alpha=0}^{\alpha=1}\Delta H dx\\ &\quad+\int\int\left(\frac{\partial}{\partial\alpha}(u''(w)w_\alpha w_x+u'(w)w_{x\alpha})\right)\left(\int_0^\alpha L_tds\right)\Delta H d\alpha dx\\ &= -\int(u''(w)w_\alpha w_x+u'(w)w_{x\alpha})\big|_{\alpha=1}\left(\int L_tds\right)\Delta H dx\\ &\quad+\int\int u'(w)\left[R(w)\left(P(w)w_xw_\alpha^2-2w_{\alpha x}w_\alpha-w_xw_{\alpha\alpha}\right)+w_{x\alpha\alpha}\right]\left(\int_0^\alpha L_tds\right)\Delta H d\alpha dx\end{aligned}$$

where the first term reflects the change in the mean via $\int L_tds=-\partial(\int\alpha dL)/\partial t\geq 0$ since $\int\alpha dL$ decreases in t, and the second term the change in riskiness through $\int_0^\alpha L_tds\geq 0$ for all α.

As in the proof of Proposition 4, the term in square brackets in the last term is positive since f is a multiplicatively separable class and $P(w)\geq 3/w$ for all w.

Regarding the first term, note $\int L_tds\geq 0$, and since $w_\alpha w_x=ww_{x\alpha}$ and $R(w)\geq 1/w$, $u''(w)w_xw_\alpha+u'(w)w_{x\alpha}\leq 0$. Along with the minus sign we obtain that the first term is positive.

Hence, a DCO shift in L increases the measure of workers who invest (θ^* decreases). □

A.7 Separable Match Output f and IR Shift in L

We asserted in footnote 15 that we can state the analogue of Proposition 4 (*ii*) for the output function f given by $f(x, y, \alpha) = \eta(\alpha)z(x, y) + \delta s(x) + t(y)$. We now prove this assertion. To this end, we assume that instead of supermodular z is log-supermodular, it is convex in x for all y, and $z_x(0, 0) > 0$. We also keep the other assumptions on η and z. In addition, we let s be increasing in x with $s(0) = 0$ and t increasing in y with $t(0) = 0$, both continuously differentiable, and assume that $-\min_{x,y,\alpha}(\eta(\alpha)z_x(x, y))/s'(x)) \leq \delta \leq 0$ to ensure that $f_x \geq 0$ for all (x, y, α). Finally, we impose the following 'likelihood ratio condition' on s: $s'(x)/s(x) \geq (z_x(1, 1)/(z_x(0, 0))(1/x)$ for all x.[25] The wage function is given by

$$w(x, \theta^*, \alpha) = \eta(\alpha) \int_0^x z_x(s, \mu(s, \theta^*)ds + \delta \int_0^x s'(\tau)d\tau.$$

Differentiating it with respect to α yields

$$\frac{w_{\alpha\alpha}w}{w_\alpha^2} = \frac{\eta''\eta}{\eta'^2} + \frac{\eta''\delta \int_0^x s'}{\eta'^2 \int_0^x z_x} \leq 1,$$

where the inequality follows from the log-concavity and convexity of η plus $\delta \leq 0$. Also,

$$\frac{w_{x\alpha}w}{w_x w_\alpha} = \frac{z_x \left(\eta \int_0^x z_x + \delta \int_0^x s'\right)}{(\eta z_x + \delta s') \int_0^x z_x}.$$

The right side is one at $\delta = 0$ and decreasing in δ if and only if $z_x / \int_0^x z_x \leq s' / \int_0^x s'$. Notice that $\int_0^x z_x(s, \mu(s))ds \geq z_x(0, 0)x$, and $z_x(x, \mu(x)) \leq z_x(1, 1)$ for all x. Therefore, $z_x / \int_0^x z_x \leq s' / \int_0^x s'$ holds if $s'(x)/s(x) \geq z_x(1, 1)/(z_x(0, 0)x)$ for all x, which is one of our premises. As a result, $w_{x\alpha}w/(w_x w_\alpha) \leq 1$, so if $P(w) \geq 3/w$, then more workers invest with an IR shift in L. □

A.8 Proof of Proposition 5

Let us modify the monotone comparative statics argument derived in Appendix A.4. The analogue of a worker's optimization problem (23) when either G or Q is indexed by t is

$$\max_{a \in \{0,1\}} \int \int u(w(x, \alpha, \theta^*, t))dH(x|a)dL(\alpha) - ac(\theta), \tag{25}$$

where t enters the wage through the matching function. The first term can be written as $\int \psi(x, t)dH(x|a)$, and once again a is increasing in t if ψ is supermodular in (x, t).

[25] For example, if $s(x) = x^\beta$, then this holds if $\beta \geq z_x(1, 1)/(z_x(0, 0)$ since $s'/s = \beta/x$. If in addition $\beta \geq 1$ and $\eta(0) > 0$, then the lower bound for δ is $-\eta(0)z_x(0, 0)/\beta < 0$.

If instead t shifts H_1, then $\max_{a\in\{0,1\}} \int \psi(x,t)dH(x|a,t)dL(\alpha) - ac(\theta)$, where $H(x|a,t) = aH_1(x|t) + (1-a)H_0(x)$. If the integral is supermodular in (a,t), then a increases in t, and again it suffices that ψ is supermodular in (x,t). To see this, note that the derivative with respect to t is $\int \psi_x(x,t)dH(x|a,t) + \int \psi(x,t)dH_t(x|a,t)$. The first term is increasing in a if $\psi_{xt} \geq 0$, and the second is increasing in a since by FOSD $\partial H_1(x|t)/\partial t \leq 0$.

We again first provide conditions on absolute risk aversion R and the wage function w under which ψ is supermodular in (x,t), and then prove it from primitives.

Lemma 3 (FOSD Shift in G, Q, H_1) *In any stable equilibrium, more workers invest in skills in response to a FOSD shift in G, Q, or H_1 if either of the following conditions hold:*

(i) Absolute risk aversion is (uniformly) sufficiently small and $w_{xt} > 0$; or

(ii) There is a $b > 0$ such that, for all w, $R(w) \geq b/w$, and for almost all (x,α), $b \geq (w_{xt}w)/(w_x w_t)$, with $w_x > 0$ and $w_t < 0$.

Proof. For part (i) we prove a slightly more general version that only requires that $w_{xt} > 0$ for each t and almost all (x,α). From the objective function in the problem above, we must show that $\int \psi dH_1 - \int \psi dH_0 = \int \psi_x \Delta H dx = \int\int u'(w)w_x \Delta H dx dL$ is increasing in t. Differentiating with respect to t and parametrizing R by r as in the proof of Proposition 1, we have

$$\begin{aligned}
\int\int \left(u'' w_x w_t + u' w_{xt}\right) \Delta H dx dL &= \int\int u'(w,r)\left(w_{xt} - R(w,r)w_x w_t\right)\Delta H dx dL \\
&\geq \int\int u'(w,r)\left(w_{xt} - R(w,r)\sup_{x,t,\alpha}(w_x w_t)\right)\Delta H dx dL \\
&= \int\int u'(w,r)\left(w_{xt} - R(w,r)\gamma\right)\Delta H dx dL,
\end{aligned}$$

where with some abuse we have set $\gamma = \sup_{x,t,\alpha}(w_x w_t)$. The right side of the above expression is positive for $r = 0$ and by continuity for r small (so absolute risk aversion is sufficiently small uniformly in w). If the FOSD shift is in G or Q, then more workers invest under the assumption on w_{xt}. If the FOSD shift is in H_1, then part (i) holds since $\partial(U_1(\theta^*,t) - U_0(\theta^*,t))/\partial t = \int \psi_{xt}\Delta H_1 dx + \int \psi_x(-\partial H_1/\partial t)dx$, and both terms are positive.

Finally, consider part (ii) and assume $R(w) \geq b/w$ for all w, where $b > 0$. Then

$$\begin{aligned}
\int\int \left(u'' w_x w_t + u' w_{xt}\right) \Delta H dx dL &= \int\int u'(w)\left(R(w)w_x(-w_t) - (-w_{xt})\right)\Delta H dx dL \\
&\geq \int\int \frac{u'(w)}{w}\left(b w_x(-w_t) - w(-w_{xt})\right)\Delta H dx dL \\
&\geq 0,
\end{aligned}$$

where the last inequality follows from the premise that b is large enough such that the integrand is positive almost everywhere. This completes the proof of part (*ii*). □

Proof of Proposition 5. The proof of part (*i*) directly follows from Lemma 3 (*i*) and the fact that in this case $w_{xt} = f_{xy}(-G_t/g) > 0$, and hence it is omitted.

Consider part (*ii*). The proof, which involves risk aversion to be large enough for the comparative statics result to hold, is more delicate. In particular, we need to show that if $wR(w)$ is large enough for all w, say larger than a threshold number $b > 0$, then $R(w) \geq b/w > 0$ and $b \geq ww_{xt}/w_x w_t$ (or, equivalently, $bw_x(-w_t) - w(-w_{xt}) \geq 0$) are sufficient even at small x, i.e. around $x = 0$, so that $w_x(-w_t)$ does not approach zero faster than $w(-w_{xt})$.

Differentiating wage function (5), we obtain $w_x = f_x$, $w_t = \int_0^x f_{xy}\mu_t$, and $w_{xt} = f_{xy}\mu_t$, where $\mu_t = Q_t \Delta H/g \leq 0$, with equality at $x = 0$. Then

$$\mu_{xt} = \frac{Q_t}{g^2}\left(\Delta h g - \Delta H g' \mu_x\right),$$

where $\mu_x = (Qh_0 + (1-Q)h_1)/g > 0$. Note that at $x = 0$, we obtain $\mu_{xt}(0,t) = \frac{Q_t}{g^2}\Delta h g < 0$ since $\Delta h(0) > 0$ and $Q_t < 0$ for $\theta^* \in (0,1)$ by assumption.

Fix $\alpha \in [0,1]$. We want to show that there exists a nonempty right-neighborhood of $x = 0$ such that $bw_x(-w_t) - w(-w_{xt}) \geq 0$ for all values of b above a finite threshold $\hat{b}_1(\alpha)$. From the derivatives of the wage function (5) we can write $\kappa(x,\alpha,t) \equiv bw_x(-w_t) - w(-w_{xt}) = bf_x \int_0^x f_{xy}(-\mu_t) - (\int_0^x f_x) f_{xy}(-\mu_t)$. It suffices to show that $\kappa(x,\alpha,t) > 0$ near $x = 0$. Since $\kappa(0,\alpha,t) = 0$, it would be enough that $\kappa_x(0,\alpha,t) > 0$ for b above a threshold, for then the asserted neighborhood would exist by continuity. Algebra reveals though that $\kappa_x(0,\alpha,t) = 0$.[26] The second derivative $\kappa_{xx}(x,\alpha,t)$, however, when evaluated at $x = 0$ reduces to

$$\kappa_{xx}(0,\alpha,t) = (b-2) f_x(0,\mu(0,t),\alpha) f_{xy}(0,\mu(0,t),\alpha)\left(-\mu_{xt}(0,t)\right) > 0,$$

for all $b > b_1(\alpha) = 2$ since the other terms are positive. Hence, for each $\alpha \in [0,1]$ there is a $\hat{x}(\alpha) > 0$ such that if $b > 2$ we have $\kappa(x,\alpha,t) > 0$ on $(0,\hat{x}(\alpha))$ and thus $\int\int_0^{\hat{x}(\alpha)} (u'(w)/w)\kappa(x,\alpha,t)\Delta H dx dL > 0$. In turn, on $[\hat{x}(\alpha),1]$, the expression $w(-w_{xt}) = f_x(\int_0^x f_{xy}(-\mu_t))$ is positive for all (x,α,t) with $x \in [\hat{x}(\alpha),1]$, and is also continuous. Thus $\iota(\alpha) = \min_{x \geq \hat{x}(\alpha),\alpha}(f_x(\int_0^x f_{xy}(-\mu_t))) > 0$. In addition, $\varphi = \max_{x,\alpha}(\int_0^x f_x \int_0^x f_{xy}(-\mu_t)) < \infty$. Hence, if $b \geq b_2(\alpha) = \varphi/\iota(\alpha)$, then $\kappa(x,\alpha,t) \geq 0$ for all $x \geq \hat{x}(\alpha)$. And since $\hat{x}(\alpha) > 0$ for all $\alpha \in [0,1]$, it follows that $\hat{b} = \sup_\alpha\{2, b_2(\alpha)\} < \infty$. Therefore, if $b \geq \hat{b}$, then $\int\int \psi_{xt}\Delta H dx dL \geq 0$, and this completes the proof of a FOSD shift in H_1.

Finally, consider part (*iii*). In this case $\mu_t = (1-Q)(\partial H_1/\partial t)/g \leq 0$, with equality at $x = 0$

[26] The reason is that every term in the derivative contains either an integral $\int_0^x$ or μ_t, and both vanish at $x = 0$.

since $H_1(0|t) = 0$ for all t, and also

$$\mu_{xt} = \frac{1-Q}{g^2}\left(\frac{\partial h_1}{\partial t} g - \frac{\partial H_1}{\partial t} g' \mu_x\right),$$

where $\mu_x = (Qh_0 + (1-Q)h_1)/g > 0$. At $x = 0$ we obtain $\mu_{xt}(0,t) = (1-Q)(\partial h_1(0|t)/\partial t)/g < 0$ since $\partial h_1(0|t)/\partial t < 0$ and $g(0) > 0$ (by assumption), and $0 < \theta^* < 1$ in any equilibrium. Now follow the same steps as in part (*ii*) to complete the proof. □

A.9 IR Shift in G, Q, H_1

In some applications one may be interested in, say, the effects of an increase in firm productivity dispersion.[27] Or in the effects of an increase in the spread of the distribution of abilities, for example, if parental wealth inequality affects pre-school investments and thus inequality in abilities. Or in the effects of a change in the variance of college completion risk across educational institutions. Questions of this kind call for IR shifts in G, Q, or H_1, which are more difficult to sign than the FOSD ones due to the changes in sign of the relevant distribution when t changes.

Consider first an IR shift in G or H_1. The derivatives w_t and w_{xt} change signs as a function of x, precluding supermodularity of ψ in (x,t), so we cannot use the same argument as in the FOSD shifts. Instead, we derive a sufficient condition in terms of the integral $\int u'(w) w_x \Delta H dx$: if it is increasing in t for all (α, t), then an IR shift in G increases investment in any stable equilibrium. This condition still presents nontrivial challenges, since the known results for signing integrals do not apply to $\int (u''(w) w_x w_t + u'(w) w_{xt}) \Delta H dx$ (e.g., ΔH is neither monotone nor does it exhibit a useful single crossing property to couple with $u''(w) w_x w_t + u'(w) w_{xt}$). And if the shift is in H_1, we also need to sign $\int \psi_x (-\partial H_1/\partial t) dx$, and $\partial H_1/\partial t$ changes signs as well. A similar difficulty is present when there is an IR shift in Q, but in this case we derive a different sufficient condition.

We first provide conditions on absolute risk aversion R and the wage function w under which an IR shift in G or H_1 leads to more investment, and then derive the conditions from primitives.

Lemma 4 (IR Shift in G, H_1) (*i*) *In any stable equilibrium, more workers invest in skills in response to an IR shift in G if $\int w_{xt} dx > 0$ for all (α, t), $H_0 - H_1$ is sufficiently close to one, and R is (uniformly) sufficiently small;*

(*ii*) *In any stable equilibrium, more workers invest in skills in response to an IR shift in Q if $w_{xt} > 0$ for almost all (x, α, t), and R is (uniformly) sufficiently small;*

(*iii*) *In any stable equilibrium, more workers invest in skills in response to an IR shift in H_1 if w is convex in x for all (α, t) and strictly so on an interval for each (α, t), $H_0 - H_1$ is sufficiently close to zero, and R is (uniformly) sufficiently small.*

[27] This can be a relevant factor in the recent debate about increasing between-firm wage inequality Song, Price, Guvenen, Bloom, and von Wachter (2018).

Proof. (*i*) We must show that, for each (α, t)

$$\int u'(w)(w_{xt} - R(w)w_x w_t)\Delta H dx \geq 0.$$

If workers are risk neutral, then $\int w_{xt}dx > 0$ and ΔH sufficiently close to one (since H_0 is close to a unit mass point at $x = 0$ and H_1 at $x = 1$) imply that $\int w_{xt}\Delta H dx > 0$. If workers are strictly risk averse, then as in Proposition 1, index u and thus R by a parameter $r \in [0, 1]$ where $R(w, 0) = 0$ for all w. Then an increase in r is a uniform increase in R, and since the integral is positive at $r = 0$, by continuity the same holds for sufficiently small r.

(*ii*) This follows from (*i*), since now for any ΔH the integral is positive for R small enough.

(*iii*) An IR shift in H_1 has two effects, not only an indirect one via the wage function but also a direct one via the impact of t on H_1. Consider the direct effect; integration by parts yields

$$\begin{aligned}\int\int u'(w)w_x\left(-\frac{\partial H_1}{\partial t}\right)dxdL &= \int\int (u''(w)w_x^2 + u'(w)w_{xx})\left(\int_0^x \frac{\partial H_1}{\partial t}ds\right)dxdL \\ &= \int\int u'(w)(w_{xx} - R(w)w_x^2)\left(\int_0^x \frac{\partial H_1}{\partial t}ds\right)dxdL.\end{aligned}$$

By IR, $\int_0^x (\partial H_1/\partial t)ds$ is positive, and to avoid a triviality it is strictly positive on a positive-lengthed interval of values of x. Then if workers are risk neutral, $w_{xx} \geq 0$ and strictly so on an interval for each (α, t) yields the result. If workers are strictly risk averse, the argument in part (*i*) shows that the same holds if r is sufficiently small. Finally, the indirect effect $\int u'(w)(w_{xt} - R(w)w_x w_t)\Delta H dx$ can be made arbitrarily small by making $H_0 - H_1$ sufficiently close to zero. □

Consider the wage function (5). Then $w_t = \int_0^x f_{xy}\mu_t ds$ and $w_{xt} = f_{xy}\mu_t$. With an IR shift in G, $\mu_t = -G_t/g$, which alternates signs starting with negative near $y = \mu(0, \theta^*, t) = 0$. If the IR shift is in H_1, $\mu_t = ((1 - \theta^*)\partial H_1/\partial t)/g$, which also alternates sign with positive sign near $x = 0$. We will use these properties in the next result to derive sufficient conditions from primitives.

Proposition 6 (IR Shift in G, Q, H_1) (*i*) *In any stable equilibrium, more workers invest in skills in response to an IR shift in G if G_t strictly single crosses zero from above, f_{xy} is increasing in (x, y) for each α and strictly so on an interval, g is decreasing in y, H_0 concentrates most of its mass near $x = 0$ and H_1 near $x = 1$, and R is (uniformly) sufficiently small;*

(*ii*) *More workers invest in skills in response to an IR shift in Q if Q_t strictly single crosses zero from above, R is (uniformly) sufficiently small, and c is sufficiently convex to ensure a unique equilibrium with θ^* close to zero;*

(*iii*) *In any stable equilibrium, more workers invest in skills in response to an IR shift in H_1 if f_{xx} is convex in x and strictly so on an interval for each (y, α), $H_0 - H_1$ is sufficiently close to zero, and R is (uniformly) sufficiently small.*

Proof. (*i*) Based on Lemma 4 (*i*), we only need to show that $\int w_{xt} = \int f_{xy}(-G_t/g)dx > 0$ under the stated assumptions. By definition of IR, $\int G_t dx = 0$ (constant mean), and by assumption, G_t has a strict single crossing property from positive to negative (an increase in t leads to a mean preserving spread in G), so $-G_t$ strictly single crosses zero from negative to positive. Since $f_{xy}(\cdot, \mu(\cdot, \theta^*, t), \alpha)/g(\mu(\cdot, \theta^*, t))$ increases in x under the stated assumptions, it follows by a standard integral inequality that $\int (f_{xy}/g)(-G_t)dx > 0$.

(*ii*) Since $U_1 - U_0 = \int \psi dH_1 - \int \psi dH_0 = \int \psi_x \Delta H$, it is increasing in t if ψ_x is increasing in t, so that $\psi_{xt} = \int (u''(w)w_t w_x + u'(w)w_{xt})dL \geq 0$. Since Q enters w through μ and depends only on θ^*, $w_t = Q_t \int_0^x f_{xy}\Delta H/g ds$ and $w_{xt} = Q_t f_{xy}\Delta H/g$. Hence

$$\psi_{xt} = Q_t(\theta^*|t) \int u'(w) \left(f_{xy}\Delta H/g - R(w) f_x \left(\int_0^x f_{xy}\Delta H/g ds \right) \right) dL. \tag{26}$$

Note that if one can ensure that θ^* is such that $Q_t(\theta^*|t) > 0$, then $w_{xt} > 0$ for all (x, α, t) and we can apply Lemma 4 (*ii*). If R is sufficiently small, then the integral in (26) is positive, but Q_t alternate signs. Since Q_t strictly crosses zero once and from above, it must be strictly positive for all θ in a right-neighborhood of zero. If the equilibrium θ^* is sufficiently small, then such an IR shift increases investment. From the uniqueness part of the proof of Proposition 1, it follows that if c is sufficiently convex then θ^* is close to zero, and more workers invest with an IR shift in Q.

(*iii*) Notice that $w_{xx} = f_{xx} + f_{xy}\mu_x$. Since f is supermodular in (x, y) and convex in x (with strict inequality on an interval), it follows that the premise of Lemma 4 (*iii*) holds. □

The intuition of part (*i*) is that, by not investing, a worker draws his skill from a distribution whose mean is very small, while if he invests then the average skill is close to one. This is not enough for more workers to invest with an IR shift in G, since the difference in the expected utilities from investing and not investing may go down due to the behavior of w as a function of (x, t). But if the expected change in w_x increases in t when H_0 concentrates all its mass at $x = 0$ and H_1 at $x = 1$, then the same holds for less extreme H_i, $i = 0, 1$, and for small risk aversion.

A simple example that satisfies the condition on ΔH in part (*iii*) is when, if a worker invests, then x is drawn from $pH_0 + (1-p)H_1$, $0 < p < 1$. If p is large the indirect effect of an IR shift in H_1 on $U_1 - U_0$ via the change in μ is negligible compared to the direct effect on H_1. Convexity of f in x implies that the wage function is strictly convex in x, and thus the direct effect of an IR shift in H_1 has a positive effect on the incentives to invest when the agent is not too risk averse.

Regarding the effects of these IR shifts on μ (and thus on w), note that with an IR shift in G, the change in the matching function is given by $\mu_{\theta^*}\theta_t^* + \mu_t$ (where $\mu_t = -G_t/g$), and G_t changes sign as a function of x (through μ), and hence only in those intervals where $G_t \geq 0$, we can assert that μ decreases (since the effect on G reinforces that on θ^*). Similarly for an IR shift in H_1, where matching function μ decreases for those x's such that $\partial H_1/\partial t \leq 0$. Finally, with an IR shift in Q with θ^* sufficiently small, we have $\mu_t \geq 0$ but $\theta_t^* \leq 0$, so $\mu_{\theta^*}\theta_t^* + \mu_t$ is ambiguous.

A.10 Remarks on Integration and Differentiation

At several points we have interchanged differentiation and integration, and used integration by parts. We now discuss the assumptions on primitives that justify these operations.

Note that if u were three times continuously differentiable for *all* values of w, then a continuous differentiability condition on H_1 with respect to t on $[0,1]^2$ and the same for L with respect to t on $[0,1]^2$ would suffice to justify all the instances where we use integration by parts and differentiate through the integral (this follows from straightforward applications of Theorems 16.8 (ii) and 18.4 in Billingsley (1995)). But if the derivatives of u diverge at $w(0,\alpha) = 0$, which we allow for so as to cover cases like CRRA, then some discussion is in order.

For each $(\alpha, \theta^*) \in [0,1]^2$, $u(w(\cdot, \alpha, \theta^*))$ is absolutely continuous on $[0,1]$. This is because of the assumptions on u and the fact that the wage function w given by $w(x, \alpha, \theta^*) = \int_0^x f_x(s, \mu(s, \theta^*), \alpha) ds$, is continuously differentiable in x. Hence, we can apply integration by parts (Billingsley (1995) Theorem 18.4) to obtain $\int u(w) dH_1 - \int u(w) dH_0 = \int u'(w) w_x \Delta H dx$, which we used repeatedly in the text. And if L has a density l (so L is absolutely continuous on $[0,1]$), we can also integrate by parts $\int u(w) dL$ to obtain $u(w)|_{\alpha=1} - \int u'(w) w_\alpha dL$. Then it is straightforward to justify the steps in Lemmas 1–3, Propositions 1-5, and Corollary 1. And the arguments that use a uniformly small risk aversion bound assume that R is bounded and thus integrable.

Also, when we pass the derivative through the integral and the integral contains the derivatives of u and w, we can justify it with Theorem 16.8 (ii) in Billingsley (1995), whose conditions only need to hold outside a set of measure zero. Since we allow u', u'', and u''' to diverge to infinity at $w(0,\alpha)$, this occurs on a set of values of (x, α) of measure zero (given our assumptions on H_i, $i = 0, 1$, and L), and outside that set the derivatives with respect to α or θ^* or t (depending on the case under consideration) exist, and their integrability with respect to the probability measure over (x, α) can be fulfilled. In turn, the derivatives of the wage function are all integrable under our assumptions. For instance, $w_{\theta^*} = \int_0^x f_{xy} (\Delta H / g) ds$ is finite as f_{xy} is continuous and thus bounded, $\Delta H \leq 1$, and g is positive on $[0,1]$ and thus bounded too. Similarly for w_α and the cross partials $w_{x\theta^*}$ and $w_{x\alpha}$. In cases where $\partial H_1 / \partial t$ appears, as in w_t and w_{xt} in Proposition 5, it suffices that H_1 be continuously differentiable in t on $[0,1]^2$ for it to be integrable. Similarly for L_t, G_t, and Q_t, which appear in a couple of instances in the comparative statics analysis.

A.11 Nontransferable Utility Case

We have assumed that utility is transferable. This need not always be the case in applications. An extreme case is when there is no flexibility at all to transfer utility (Becker (1973)). One interpretation is that there are no transfers and each party enjoys an attribute-dependent utility from matching with different partners. Another one is to assume that matched pairs divide match output in an exogenous way, say equally. For definiteness we consider this case, and hence given

θ^* and μ, the wage of a worker with x is $w(x, \theta^*, \alpha) = f(x, \mu(x, \theta^*), \alpha)/2$. Since f is increasing in a partner's attribute, there is positive sorting in equilibrium and thus μ is given by (4).

We now adapt the main results to the nontransferable utility case. Existence of a stable equilibrium follows exactly as before. Regarding Proposition 1, equilibrium is unique if c is convex enough; the proof is the same as before. Also, if $w_{x\theta^*} > 0$ almost everywhere, then uniqueness follows if absolute risk aversion is sufficiently small uniformly in w. But $w_{x\theta^*} = 0.5(f_{xy}\mu_{\theta^*} + f_{yy}\mu_x\mu_{\theta^*} + f_y\mu_{x\theta^*})$ need not be positive since $\mu_{x\theta^*}$ can be positive or negative. It is, however, bounded below by a finite number. Thus, if f_{yy}/f_y (the degree of convexity of f) is large enough uniformly in (x, θ^*, α), then $w_{x\theta^*} > 0$ almost everywhere and uniqueness follows.

Lemmas 1–3 and 4 hold since they are stated for an arbitrary w, which includes $w = f/2$.

We now show how Propositions 3-6 need to be adjusted.

Regarding Proposition 3, part (i) holds without modification, and (ii) holds if instead of f_x log-supermodular in (x, α) and (y, α) we have f log-supermodular in (x, α) and (y, α).

Proposition 4 holds if in addition to $f_{x\alpha\alpha} \geq 0$ we assume that $f_{y\alpha\alpha} \geq 0$. In particular, part (ii) holds beyond the multiplicatively separable case: it suffices that f is log-submodular in (x, α) (for each y) and (y, α) (for each x), and convex and log-concave in α (for each (x, y)). To see this, notice that w is log-submodular (and hence $w_{x\alpha}w/w_x w_\alpha \leq 1$) if and only if

$$w_{x\alpha}w - w_x w_\alpha = (f_{x\alpha}f - f_x f_\alpha) + (f_{y\alpha}f - f_y f_\alpha)\mu_x \leq 0,$$

and this holds if f is log-submodular in (x, α) and in (y, α). Similarly, w is convex and log-concave in α if and only if $w_{\alpha\alpha} = f_{\alpha\alpha}/2 \geq 0$ and $w_{\alpha\alpha}w/w_\alpha^2 = f_{\alpha\alpha}f/f_\alpha^2 \leq 1$.

Regarding Proposition 5 (i), note that $w_{xt} = 0.5(f_{xy}\mu_t + f_{yy}\mu_t\mu_x + f_y\mu_{xt})$. With a FOSD shift in G, w_{xt} is positive if f is convex in y for all (x, α) and $\mu_{xt} = \partial(-G_t/g)/\partial x \geq 0$, which holds if $\partial(-G_t/g)/\partial y \geq 0$ (e.g., this holds if G is exponential with parameter $1/t$, or truncated on $[0, \overline{y}]$).

If instead there is a FOSD in Q, we must show that $w_{xt} - R(w)w_t w_x \geq 0$ if $wR(w) \geq b$ for b large enough. In this case, $w = 0.5f$, $w_x = 0.5(f_x + f_y((Qh_0 + (1-Q)h_1)/g))$, $w_t = 0.5 f_y Q_t \Delta H/g$, and $w_{xt} = 0.5(Q_t f_y/g)(\Delta H(f_{xy} + f_{yy}((Qh_0 + (1-Q)h_1)/g) - ((Qh_0 + (1-Q)h_1)g'/g^2)) + \Delta h)$. Using these expressions and $R(w) \geq b/w$, we have $w_{xt} - R(w)w_t w_x \geq 0$ for almost all (x, α, t) if

$$b\Delta H\left(\frac{f_x}{f} + \frac{f_y}{f}\frac{Qh_0 + (1-Q)h_1}{g}\right) - \Delta H\left(f_{xy} + f_{yy}\frac{Qh_0 + (1-Q)h_1}{g} - \frac{(Qh_0 + (1-Q)h_1)g'}{g}\right) - \Delta h \geq 0.$$

Note that the second and third terms are bounded and independent of b. So if the term involving b, which consists of positive elements, were bounded away from zero, it would follow that for large enough b the inequality would hold for all (x, α, t). The only problem is that at $x = 0$ the first term may vanish, while $\Delta h \geq 0$. Fix α and t and consider the limit as x goes to zero. If either $\lim_{x\to 0} \Delta H f_x/f > 0$ or $\lim_{x\to 0} \Delta H f_y/f > 0$ for all (α, t), which depend solely on primitives, then it is clear that for b large enough $w_{xt} - R(w)w_t w_x \geq 0$, and hence a FOSD shift in Q leads to an increase in the measure of workers who invest.

Finally, consider a FOSD in H_1. The proof proceeds exactly as in the proof of Proposition 5 (*iii*), but with $\kappa(x, \alpha, t) = (1/4)(b(f_x + f_y\mu_x)f_y(-\mu_t) - f(f_{xy}(-\mu_t) + f_{yy}\mu_x(-\mu_t) + f_y(-\mu_{tx}))$. Now, $\kappa(0, \alpha, t) = 0$ (since μ_t and f vanish at $x = 0$) but

$$\kappa_x(0, \alpha, t) = \frac{(b-1)}{4}(f_x(0, \mu(0, \theta^*, t), \alpha)+f_y(0, \mu(0, \theta^*, t), \alpha)\mu_x(0, \theta^*, t))f_y(0, \mu(0, \theta^*, t), \alpha)(-\mu_{xt}(0, \theta^*, t)),$$

which is positive if $b > 1$, $f_x > 0$, $f_y > 0$, and $\partial h_1(0|t)/\partial t < 0$ (so that $\mu_{xt}(0, \theta^*, t) = (\partial h_1(0|t)/\partial t)/g > 0$). Hence, there is a $\hat{x}(\alpha) > 0$ such that $bw_x(-w_t) - w(-w_{xt}) > 0$ for $x \in (0, \hat{x}(\alpha))$. The rest is as in the proof of Proposition 5 (*iii*).

Turning to Proposition 6 (*i*) (IR shift in G) note that now

$$\int w_{xt}dx = 0.5\left(\int\left(f_{xy}\left(-\frac{G_t}{g}\right) + f_{yy}\frac{\theta^* h_0 + (1-\theta^*)h_1}{g}\left(-\frac{G_t}{g}\right) + f_y\mu_{xt}\right)dx\right),$$

where $\mu_{xt} = -(1/g^2)(\theta^* h_0 + (1-\theta^*)h_1)(g'(-G_t/g) + g_t)$. Assume that $f_{yy} = 0$ so we can omit the middle integral. Let g be such that $g_t(0|t) = 0$ and $g_{xt}(0|t) < 0$; then one can show that $\mu_{xt} > 0$ for x near zero. Then for each α there is an $\hat{x}(\alpha) > 0$ such that $\int_0^{\hat{x}(\alpha)} w_{xt}dx \geq 0$. And since $-G_t$ single-crosses zero from below and $g' \leq 0$, $\int_{\hat{x}(\alpha)}^1 w_{xt}dx > 0$ if f_{xy}/f_y is increasing in (x, y) for each α and sufficiently large. By continuity, we can replace $f_{yy} = 0$ by f_{yy} positive but small enough.

Regarding Proposition 6 (*ii*) (IR shift in Q), some algebra reveals that $\psi_{xt} = \int u'(w)(w_{xt} - R(w)w_t w_x)dL$ can be written as follows:

$$\psi_{xt} = \frac{Q_t}{2}\left(\int\frac{u'(w)f_y}{g}\left(\left(\frac{f_{xy}}{f_y} + \frac{f_{yy}}{f_y}\frac{Qh_0 + (1-Q)h_1}{g}\right)\Delta H - \Delta h - (Qh_0 + (1-Q)h_1)\frac{g'}{g}\Delta H\right)dL\right.$$
$$\left. + \int\frac{u'(w)f_y}{2g}R(w)\frac{f_y\Delta H}{g}\left(f_x + f_y\frac{Qh_0 + (1-Q)h_1}{g}\right)dL\right).$$

If c is convex enough, then θ^* will be small and, assuming Q_t strictly single-crosses zero from above, $Q_t(\theta^*) > 0$. Assume that $R = 0$, so the last integral vanishes. Then if $g' \leq 0$ and either f_{xy}/f_y or f_{yy}/f_y are sufficiently large, then the first integral is positive and $\psi_{xt} > 0$, which implies that θ^* decreases with an IR shift in Q. By continuity, the result holds for uniformly small R.

Finally, to adapt Proposition 6 (*iii*) (IR shift in H_1), note that

$$w_{xx} = 0.5\left(f_{xx} + 2f_{xy}\frac{\theta h_0 + (1-\theta)h_1}{g} + f_{yy}\mu_x^2 + f_y\mu_{xx}\right),$$

where $\mu_{xx} = (1/g^2)[(\theta^* h_0' + (1-\theta^*)h_1')g - (\theta^* h_0 + (1-\theta^*)h_1)^2(g'/g)]$. Now, μ_{xx} is positive near $x = 0$ if $h_i'(0) \geq 0$, $i = 0, 1$, with one of them strict, and $g'(0) \leq 0$. Moreover, μ_{xx} is bounded. Let

f be convex in x for each (y, α) and in y for each (x, α). Then for each α there exists a $\hat{x}(\alpha) > 0$ such that $w_{xx} \geq 0$ for $x \in [0, \hat{x}(\alpha))$, and positive on $[\hat{x}(\alpha), 1]$ if either f_{xx}/f_y, f_{yy}/f_y or f_{xy}/f_y are sufficiently large. The rest of Proposition 6 (ii) remains the same.

A.12 Two-Sided Investments

We will prove the assertions stated in Section 6.2. We first show the results for the risk neutral case and then invoke the usual continuity argument to extend them to small risk aversion.

We first derive the "best response functions" for each side of the market. From

$$\pi(y, \theta^*, \omega^*, \alpha) = f(\mu^{-1}(y, \theta^*, \omega^*), y, \alpha) - w(\mu^{-1}(y, \theta^*, \omega^*), \theta^*, \omega^*, \alpha),$$

it follows that $\pi_y(y, \theta^*, \omega^*, \alpha) = f_y(\mu^{-1}(y, \theta^*, \omega^*), y, \alpha)$, since $f_x \mu_y^{-1}$ and $w_x \mu_y^{-1}$ cancel.

Consider the workers' side. An increase in ω^* increases θ^*. To see this, differentiate $U_1 - U_0 = c$ with respect to ω^* to obtain

$$\frac{\partial \theta^*}{\partial \omega^*} = \frac{-\int\int w_{x\omega^*} \Delta H dx dL}{\int\int w_{x\theta^*} \Delta H dx dL - c'} > 0,$$

which follows from $w_{x\theta^*} > 0$ and $w_{x\omega^*} = f_{xy}\mu_{\omega^*} < 0$, since $\mu_{\omega^*} = -\Delta G/(\omega^* g_0 + (1-\omega^*)g_1) < 0$ on an interval of values of x. Similarly, the best response function of firms is increasing in θ^*, since differentiation of $V_1 - V_0 = k$ yields

$$\frac{\partial \omega^*}{\partial \theta^*} = \frac{-\int\int \pi_{y\theta^*} \Delta G dy dL}{\int\int \pi_{y\omega^*} \Delta G dy dL - k'} > 0,$$

which follows since μ^{-1} is decreasing in θ^* and increasing in ω^* (both strictly so on an interval).

To show that a stable equilibrium exists, note that if no firm invests, $U_1 - U_0 = c$ has a positive solution $\theta^* \in (0, 1)$. Thus, r_w^{-1} starts a positive value $\theta^* > 0$. In turn, if no worker invests, $V_1 - V_0 = k$ has a positive solution $\omega^* \in (0, 1)$ and hence r_f has a positive intercept. Graphically, r_f is above r_w^{-1} at $\theta^* = 0$. At the other end, if every firm invests, then $\theta^* < 1$ given our assumptions on c. Similarly, if every worker invests then $\omega^* < 1$. Graphically, r_w^{-1} is above r_f at $\theta^* = 1$. Since both best responses are continuous functions, there is an equilibrium where r_w^{-1} crosses r_f from below, and hence a stable equilibrium exists.

Regarding comparative statics, let us differentiate (18)–(19) with respect to t to obtain

$$\theta_t^* = \frac{-\frac{\partial \int\int w_x \Delta H dx dL}{\partial t}\left(\int\int \pi_{y\omega^*} \Delta G dy dL - k'\right) + \int\int w_{x\omega^*} \Delta H dx dL \frac{\partial \int\int \pi_y \Delta G dy dL}{\partial t}}{D} \tag{27}$$

$$\omega_t^* = \frac{-\frac{\partial \int\int \pi_y \Delta G dy dL}{\partial t}\left(\int\int w_{x\theta^*} \Delta H dx dL - c'\right) + \int\int \pi_{y\theta^*} \Delta G dy dL \frac{\partial \int\int w_x \Delta H dx dL}{\partial t}}{D}, \tag{28}$$

where D is given by

$$D = \left(\int\int w_{x\theta^*}\Delta H dx dL - c'\right)\left(\int\int \pi_{y\omega^*}\Delta G dy dL - k'\right) - \int\int w_{x\omega^*}\Delta H dx dL \int\int \pi_{y\theta^*}\Delta G dy dL.$$

At a stable equilibrium $D > 0$ since it is equivalent to the stability condition $\partial r_w^{-1}/\partial\theta^* \geq \partial r_f/\partial\theta^*$. Thus, the signs of θ_t^* and ω_t^* depend on the sign of the numerators of (27)–(28). Note that in the numerator of (27) $\int\int \pi_{y\omega^*}\Delta G dy dL - k' > 0$ and $\int\int w_{x\omega^*}\Delta H dx dL < 0$, while in the numerator of (28) $\int\int w_{x\theta^*}\Delta H dx dL - c' > 0$ and $\int\int \pi_{y\theta^*}\Delta G dy dL < 0$.

Consider a FOSD in L. Since $w_{x\alpha} = f_{x\alpha} \geq 0$ and $\pi_{y\alpha} = f_{y\alpha} \geq 0$ (both inequalities are strict on a set of positive measure), the numerator of (27) is strictly negative and thus $\theta_t^* < 0$, so more workers invest upon a FOSD shift in L. Similarly, the numerator of (28) is strictly negative and thus $\omega_t^* < 0$, so more firms invest with a FOSD shift in L. The same holds for an IR shift in L if $f_{x\alpha\alpha} \geq 0$ and $f_{y\alpha\alpha} \geq 0$, with strict inequality on a set of positive measure.

In turn, a FOSD in G_1, Q, or H_1 have ambiguous effects. Regarding G_1, the effect of an increase in t on $V_1 - V_0$ is ambiguous. A FOSD shift in Q is ambiguous since $\pi_{yt} \geq 0$ and $w_{xt} \leq 0$. Finally, the ambiguous effect of a shift in H_1 is due to the same reason as in the one-sided case.

As in previous results, all these results hold with sufficiently (uniformly) small risk aversion.

A.13 Continuous Investments

Recall that $q(\hat{a}) = \int \hat{a}(\theta)d\theta$ is the average level of investment under $\hat{a}$. Let us denote the matching and wage functions by $\tilde{\mu}(s, q(\hat{a}))$ and $\tilde{w}(x, q(\hat{a}), \alpha)$ to emphasize that they depend on $\hat{a}$ through q. Given an investment function $\hat{a}$, a worker with θ solves $\max_{a\in[0,1]} \mathcal{U}(a, \theta, q(\hat{a}))$, where

$$\begin{aligned}\mathcal{U}(\theta, a, q(\hat{a})) = a &\int\int u(\tilde{w}(x, \alpha, q(\hat{a})))dL(\alpha)dH_1(x) \\ &+ (1-a)\int\int u(\tilde{w}(x, \alpha, q(\hat{a})))dL(\alpha)dH_0(x) - d(a, \theta).\end{aligned} \tag{29}$$

Since d is strictly submodular in (a, θ), the optimal choice for a worker increases in θ.

An equilibrium reduces to finding a measurable $\hat{a}$ that satisfies the usual best response property. That is, a measurable function $\hat{a}$ is an equilibrium if

$$\mathcal{U}(\theta, \hat{a}(\theta), q(\hat{a})) \geq \mathcal{U}(\theta, a, q(\hat{a})) \quad \forall a \in [0,1], \quad \text{for almost all} \quad \theta \in [0,1].$$

We will show that this large game satisfies the hypothesis of the existence theorem in Rath (1992), and hence a pure strategy Nash equilibrium exists, which implies existence of an equilibrium in our continuous investment setting.

Before proving existence, note that there are two important cases subsumed by this formulation. One is the *linear disutility case* $d(a, \theta) = ac(\theta)$, with $a \in [0, 1]$. Since for any $\hat{a}$ the function $\mathcal{U}(\theta, \cdot, q(\hat{a}))$ is linear in a, the optimal choice for all θ except for one is either $a = 1$ or $a = 0$. The only exception is θ^*, who is indifferent. The second one is the *strictly convex disutility case*, with $d(\cdot, \theta)$ twice continuously differentiable in a and $d_{aa} > 0$ for all $a \in (0, 1)$. In this case – and assuming interior solutions for all types – each agent has a unique best response for *any* investment strategy profile (pure or mixed). The following existence result subsumes both cases.

Proposition 7 *An equilibrium exists. In any equilibrium, investment is generically inefficient.*

Proof. Since for each a a Walrasian equilibrium exists in the second stage, it suffices to show that an equilibrium in the investment game among the workers exists. Let $C([0, 1]^2)$ be the set of continuous functions endowed with the sup norm, and let $\mathcal{B}(C([0, 1]^2))$ be the Borel sigma field on that metric space. We will verify the conditions for the application of Theorem 2 in Rath (1992).

First, we have a continuum of players uniformly distributed in $[0, 1]$ (indexed by θ).

Second, the set of actions available to each agent is compact, given by $[0, 1]$.

Third, the payoff function of each player θ, $\mathcal{U}(\theta, \cdot, \cdot) : [0, 1]^2 \to \mathbb{R}$, is continuous in (a, q). To see this, fix θ and notice from (29) that we can write $\mathcal{U}(\theta, a, q) = \gamma(a, q) - d(a, \theta)$, where γ is the first term on the right side of (29). Since $d(\cdot, \theta)$ is continuous, it suffices to show that γ is continuous in (a, q). But γ is the sum of products of continuous functions (each term is a linear function in a times an integral that — by the Lebesgue Dominated Convergence Theorem — is continuous in q). Hence, γ is continuous.

Finally, we will show that the function $z : [0, 1] \to C([0, 1]^2)$ that defines our game, given by $z(\theta) = \mathcal{U}(\theta, \cdot, \cdot)$, is measurable. Indeed, we will prove that it is continuous, from which measurability follows. Given θ and θ' in $[0, 1]$, we have

$$|\mathcal{U}(\theta', a, q) - \mathcal{U}(\theta, a, q)| = |d(a, \theta') - d(a, \theta)| \leq \max_{a \in [0,1]} |d(a, \theta') - d(a, \theta)| = \nu(\theta', \theta).$$

It follows from the Theorem of the Maximum that ν is continuous and converges to zero as $\theta' \to \theta$. Hence, $||\mathcal{U}(\theta', a, q) - \mathcal{U}(\theta, a, q)|| \to 0$ as $\theta' \to \theta$, proving continuity of z.[28]

We have verified all the conditions for the application of Theorem 2 in Rath (1992) hold. It follows that there exists an equilibrium in our model.

Regarding the inefficiency of investment, the planner's problem is to choose a measurable

[28] For this argument to work we cannot allow d to diverge to infinity as θ goes to zero. But this is not important, since all that matters in the binary case is that workers with θ close to zero do not invest, and this can be ensured with a large but finite bound for d as θ approaches 0.

investment function $a : [0,1] \to [0,1]$ that maximizes

$$\max_a \lambda \int \left(\int \int u(w(x, q(a), \alpha)) dL(\alpha) dH(x|a(\theta)) - d(a(\theta), \theta) \right) d\theta$$
$$+ (1-\lambda) \int \left(\int \int \left(f(x, \mu(x, q(a)), \alpha) - w(x, q(a), \alpha) \right) dL(\alpha) dH(x|a(\theta)) \right) d\theta.$$

Since $H(x|a(\theta)) = a(\theta) H_1(x) + (1 - a(\theta)) H_0(x)$, this can be rewritten as

$$\max_a \lambda q(a) \int \int u(w(x, q(a), \alpha)) dL(\alpha) dH_1(x) + \lambda(1 - q(a)) \int \int u(w(x, q(a), \alpha)) dL(\alpha) dH_0(x)$$
$$- \lambda \int d(a(\theta), \theta) d\theta + (1-\lambda) q(a) \int \int \left(f(x, \mu(x, q(a)), \alpha) - w(x, q(a), \alpha) \right) dL(\alpha) dH_1(x)$$
$$+ (1-\lambda)(1 - q(a)) \int \int \left(f(x, \mu(x, q(a)), \alpha) - w(x, q(a), \alpha) \right) dL(\alpha) dH_0(x).$$

Assume that a_p solves the planner's problem, and consider a variation $a(\cdot, \varepsilon)$ given by $a(\theta, \varepsilon) = a_p(\theta) + \varepsilon \nu(\theta)$, where ν is an arbitrary continuous function. By optimality, the derivative of the planner's problem with respect to ε evaluated at $\varepsilon = 0$ must be equal to zero, since otherwise the planner could improve his payoff by modifying a_p, contradicting that a_p is optimal.

The value of the objective function under $a(\cdot, \varepsilon)$ is

$$\lambda q(a(\cdot, \varepsilon)) \int \int u(w(x, q(a(\cdot, \varepsilon)), \alpha)) dL(\alpha) dH_1(x) + \lambda(1 - q(a(\cdot, \varepsilon))) \int \int u(w(x, q(a(\cdot, \varepsilon)), \alpha)) dL(\alpha) dH_0(x)$$
$$- \lambda \int d(a(\theta, \varepsilon), \theta) d\theta + (1-\lambda) q(a(\cdot, \varepsilon)) \int \int \left(f(x, \mu(x, q(a(\cdot, \varepsilon))), \alpha) - w(x, q(a(\cdot, \varepsilon)), \alpha) \right) dL(\alpha) dH_1(x)$$
$$+ (1-\lambda)(1 - q(a(\cdot, \varepsilon))) \int \int \left(f(x, \mu(x, q(a(\cdot, \varepsilon))), \alpha) - w(x, q(a(\cdot, \varepsilon)), \alpha) \right) dL(\alpha) dH_0(x).$$

Differentiating this expression with respect to ε, evaluating it at $\varepsilon = 0$, integrating $\int \int (f - w) \Delta h dx dL$ as in the proof of Proposition 2, simplifying, and rearranging yields

$$\int \Bigg[\lambda \left(U_1(a_p) - U_0(a_p) - d_a(a_p(\theta), \theta) + \int \int u'(w(x, \mu(x, q(a_p)), \alpha)) w_q(x, \mu(x, q(a_p)), \alpha) dL(\alpha) dH(x|a_p) \right)$$
$$-(1-\lambda) \int \int w_q(x, \mu(x, q(a_p)), \alpha) dL(\alpha) dH(x|a_p) \Bigg] \nu(\theta) d\theta, \tag{30}$$

where with some abuse of notation we have set $U_i(a_p) = \int \int u(w) dL dH_i$, $i = 0, 1$, and where $H(x|a_p) = q(a_p) H_1(x) + (1 - q(a_p)) H_0(x)$. Since ν is an arbitrary continuous function, it follows that the only way to make (30) equal to zero is to set the term in square brackets equal to zero for almost every θ. That is, a necessary condition for optimality in the planner's problem is that, for almost all θ, the following equation holds:

$$U_1(a_p) - U_0(a_p) = d_a(a_p(\theta), \theta) + \int\int \left(u'(w(x, \mu(x, q(a_p)), \alpha)) - \frac{1-\lambda}{\lambda} \right) w_q(x, \mu(x, q(a_p)), \alpha) dL(\alpha) dH(x|a_p),$$

resembling the planner's FOC (8) in Proposition 2. In equilibrium, if $d(\cdot, \theta)$ is linear in a then the solution is at a corner for all θ except for θ^*, while in any other case the first-order condition of each worker with θ, given that everybody uses a, is $U_1(a) - U_0(a) = d_a(a(\theta), \theta)$. In either case equilibrium investment is generically inefficient. □

A.14 The Partnership Model

We will prove all the claims made in Section 6.4. As usual, in the first stage an agent with characteristic θ invests if and only if $U_1 - U_0 \geq c(\theta)$. Following the same steps as in the baseline model and using the expression for w we obtain the following equilibrium condition:

$$\int\int u' \left(\frac{f(x, x, \alpha)}{2} \right) \frac{f_x(x, x, \alpha)}{2} \Delta H dx dL(\alpha) = c(\theta^*).$$

Since the left side does not depend on θ^*, it trivially follows that there is a unique equilibrium. Regarding Proposition 5, note that u' and f_x do not depend on Q, H_1, or θ^*, which simplifies the comparative statics analysis, leading to unambiguous results. First, since the left side is independent of Q, a FOSD shift in Q does not affect θ^*: all it does is to increase the measure of agents with $\theta \geq \theta^*$ who invest (since now Q puts more mass on higher θ's). Second, since matching is independent of H_1, a FOSD shift in H_1 has only a direct effect on the left side, which increases as a result and hence θ^* decreases, so more agents invest. Furthermore, and similar to Proposition 3, a FOSD shift in L increases the measure of agents who invest if either risk aversion is sufficiently small or if $R(w) \leq 1/w$ for all w and f is log-supermodular in its first and third coordinate for each value of the second one (which implies that w is log-supermodular in (x, α)). Finally, an IR shift in L increases the measure of agents who invest under exactly the same conditions as in Proposition 4.

Note that all these results carry over to the continuous investment case. To see this, replace $c(\theta^*)$ above by $d_a(\hat{a}(\theta), \theta)$. Since the left side is independent of $\hat{a}$, existence and uniqueness are trivial, and the comparative statics results follow exactly as in the binary case.

It is clear that the analysis of this one-population version of the model can be reinterpreted as a symmetric equilibrium in a setting with two identical populations with risk averse agents and two-sided investments (for example, a marriage market application). This illustrates something intuitive: since uncertainty resolves before matching, whether one or both populations are risk averse does not matter since payoffs are simply monotone transformations of income.

Less obvious is that the partnership version can also accommodate a different timing, where

α is realized after matching. Since now transfers are contingent and agents maximize expected utility, the model becomes one with imperfectly transferable utility (as in Legros and Newman (2007)). In this case, however, one can show that if risk aversion is sufficiently small then positive sorting ensues, and splitting output in half is the equilibrium "wage function." The comparative statics with respect to Q and H_1 go through as above, and the same holds for the comparative statics with respect to L for small risk aversion. The tractability of the partnership model does not carry over to the general imperfectly transferable case with two risk averse populations with possibly different levels of risk aversion, whose analysis is beyond the scope of this paper.

B Online Appendix: Omitted Algebra of Section 5

B.1 Wage Inequality and FOSD Shifts in H_1, G, L

We solve in closed form the following version of our model. Agents are risk neutral, match output f is given by $f(x, y, \alpha) = \alpha^2 xy$, the cost function c by $c(\theta) = 1 - \theta$, the productivity distribution G is uniform on $[0, t_y]$, the skill distributions H_1 for educated and H_0 non-educated workers are $H_1(x) = x^{t_x}$ and $H_0(x) = x$, where $t_x > 1$, ability distribution Q is uniform on $[0, t_\theta]$, and shock distribution L is uniform on $[0, t_\alpha]$. Under these assumptions the equilibrium (a, μ, w) is

$$\mu(x, \theta^*, t_y, t_x, t_\theta) = t_y \left(\left(1 - \frac{\theta^*}{t_\theta}\right) x^{t_x} + \frac{\theta^*}{t_\theta} x \right), \tag{31}$$

$$w(x, \theta^*, t_y, t_x, t_\theta) = \alpha^2 \int_0^x \mu(s, \theta^*, t_y, t_x, t_\theta) ds, \tag{32}$$

$$\theta^*(t_y, t_x, t_\theta, t_\alpha) = \frac{3t_\theta(2 + t_x)(1 + 2t_x) - 3\left(\frac{t_\alpha^2}{3}\right)(-1 + t_x)t_\theta t_y}{3t_\theta(2 + t_x)(1 + 2t_x) + 2\left(\frac{t_\alpha^2}{3}\right)(-1 + t_x)^2 t_y}. \tag{33}$$

FOSD IN H_1. Set $t_y = 1, t_\theta = 1, t_\alpha = 0$ and recall that $t_x > 1$ for H_1 to FOSD H_0. The equilibrium investment threshold is given by:

$$\theta^* = \frac{3 - 3t_x + 3(2 + t_x)(1 + 2t_x)}{2(-1 + t_x)^2 + 3(2 + t_x)(1 + 2t_x)}.$$

Let the return to education be the ratio of average wage of educated over non educated workers:

$$
\begin{aligned}
\frac{\mathbb{E}[w(x)|a=1]}{\mathbb{E}[w(x)|a=0]} &= \frac{(1-\theta^*)\int w(x)h_1(x)dx}{\theta^*\int w(x)h_0(x)dx} \\
&= \frac{(1-\theta^*)\int \left(\int_0^x(1-\theta^*)s^{t_x}+\theta^* s\right) t_x x^{t_x-1}dx}{\theta^*\int\left(\int_0^x(1-\theta^*)s^{t_x}+\theta^* s\right)dx} \\
&= \frac{(-1+t_x)t_x(1+2t_x)(5+t_x(23+20t_x+6t_x^2))}{3(3+2t_x(2+t_x))(4+t_x(1+t_x)(15+2t_x(4+t_x)))}. \qquad (34)
\end{aligned}
$$

The education premium is increasing in t_x since

$$
\begin{aligned}
\frac{\partial \frac{\mathbb{E}[w(x)|a=1]}{\mathbb{E}[w(x)|a=0]}}{\partial t_x} &= \frac{-2(-1+t_x)^6 t_x^2+3(1+t_x)^2(2+t_x)^3(1+2t_x)^4(-2+5t_x)}{3(1-t_x+(2+t_x)(1+2t_x))^2((-1+t_x)^2t_x-(1+t_x)(2+t_x)^2(1+2t_x))^2} \\
&+ \frac{(-1+t_x)^4(2+t_x)(1+2t_x)(2+3t_x)(1-2t_x+4t_x^2)}{3(1-t_x+(2+t_x)(1+2t_x))^2((-1+t_x)^2t_x-(1+t_x)(2+t_x)^2(1+2t_x))^2} \\
&- \frac{2(-1+t_x)^2(2+t_x)^2(1+2t_x)^2(2+t_x(23+5t_x(8+t_x(4+t_x))))}{3(1-t_x+(2+t_x)(1+2t_x))^2((-1+t_x)^2t_x-(1+t_x)(2+t_x)^2(1+2t_x))^2},
\end{aligned}
$$

where the common denominator is positive. One can show that this expression is positive for all $t_x > 1$. Thus, the skill premium is increasing in t_x.

The variance of the wage distribution is given by:

$$
\begin{aligned}
Var[w(x)] &= \int(w(x))^2h(x)dx - \left(\int w(x)h(x)dx\right)^2 \\
&= \frac{9(2+t_x)^6(4+t_x)(3+2t_x)(2+3t_x)(+2t_x)^4+18(-1+t_x)^2(2+t_x)^4(3+2t_x)(2+3t_x)(+2t_x)^3(8+t_x(3+2t_x))}{5(2+t_x)^2(4+t_x)(3+2t_x)(2+3t_x)(2(-1+t_x)^2+3(2+t_x)(1+2t_x))^4} \\
&- \frac{12^2(-1+t_x)^4(2+t_x)^2(1+2t_x)^2(2+3t_x)(-60+t_x(-63+8(-1+t_x)t_x))}{5(2+t_x)^2(4+t_x)(3+2t_x)(2+3t_x)(2(-1+t_x)^2+3(2+t_x)(1+2t_x))^4} \\
&+ \frac{4(-1+t_x)^8(24+t_x(54+t_x(31+10t_x)))+24(-1+t_x)^6(2+t_x)(1+2t_x)(28+t_x(76+t_x(65+22t_x)))}{5(2+t_x)^2(4+t_x)(3+2t_x)(2+3t_x)(2(-1+t_x)^2+3(2+t_x)(1+2t_x))^4},
\end{aligned}
$$

which can be shown to be increasing in t_x only if t_x is large enough.

FOSD IN G. Set $t_\theta = 1, t_\alpha = 0$; recall that $t_x > 1$ for H_1 to FOSD H_0. Here we only impose $t_y \geq 1$. Moreover, for convenience, we introduce one more parameter in the cost function that will help us ensure that θ^* is positive in this example: $c(\theta) = k(1-\theta)$, $k > 0$. We are interested in comparative statics with respect to t_y.

The equilibrium investment threshold is given by:

$$
\theta^* = \frac{3(k(2+t_x)(1+2t_x)+t_y-t_xt_y)}{3k(2+t_x)(1+2t_x)+2(-1+t_x)^2t_y)}.
$$

The return to education is given by

$$\frac{\mathbb{E}[w(x)|a=1]}{\mathbb{E}[w(x)|a=0]} = \frac{(1-\theta^*)\int w(x)h_1(x)dx}{\theta^*\int w(x)h_0(x)dx}$$
$$= \frac{(-1+t_x)t_x(1+2t_x)t_y(3k(1+t_x)(2+t_x)(1+2t_x)-(-1+t_x)^2t_y)}{3(1+t_x)(2+t_x)^3(k+2kt_x)^2-6k(-1+t_x)(2+t_x)(1+2t_x)(1+t_x+t_x^2)t_y+3(-1+t_x)^3t_xt_y^2}.$$

The derivative of the education premium with respect to t_y is

$$\frac{\partial\frac{\mathbb{E}[w(x)|a=1]}{\mathbb{E}[w(x)|a=0]}}{\partial t_y} = \frac{k(-1+t_x)t_x(2+t_x)^2(1+2t_x)^2(3k^2(1+t_x)^2(2+t_x)^2(1+2t_x)^2}{3((1+t_x)(2+t_x)^3(k+2kt_x)^2-2k(-1+t_x)(2+t_x)(1+2t_x)(1+t_x+t_x^2)t_y+(-1+t_x)^3t_xt_y^2)^2}$$
$$+\frac{-2k(-1+t_x)^2(1+t_x)(2+t_x)(1+2t_x)t_y-(-1+t_x)^4t_y^2)}{3((1+t_x)(2+t_x)^3(k+2kt_x)^2-2k(-1+t_x)(2+t_x)(1+2t_x)(1+t_x+t_x^2)t_y+(-1+t_x)^3t_xt_y^2)^2},$$

which is positive if k is large enough, that is, if $k > (-t_y+t_xt_y)/(2+5t_x+2t_x^2)$, which is identical to the condition on k that guarantees that θ^* is positive and thus we will assume it throughout.

The variance of the wage distribution is given by:

$$Var[w(x)] = t_y^2\times$$
$$\left\{\frac{(9(2+t_x)^6(4+t_x)(3+2t_x)(2+3t_x)(k+2kt_x)^4+18(-1+t_x)^2(2+t_x)^4(3+2t_x)(2+3t_x)(k+2kt_x)^3(8+t_x(3+2t_x))t_y}{5(2+t_x)^2(4+t_x)(3+2t_x)(2+3t_x)(3k(2+t_x)(1+2t_x)+2(-1+t_x)^2t_y)^4}\right.$$
$$-\frac{(12k^2(-1+t_x)^4(2+t_x)^2(1+2t_x)^2(2+3t_x)(-60+t_x(-63+8(-1+t_x)t_x))t_y^2)}{5(2+t_x)^2(4+t_x)(3+2t_x)(2+3t_x)(3k(2+t_x)(1+2t_x)+2(-1+t_x)^2t_y)^4}$$
$$\left.+\frac{(24k(-1+t_x)^6(2+t_x)(1+2t_x)(28+t_x(76+t_x(65+22t_x)))t_y^3+4(-1+t_x)^8(24+t_x(54+t_x(31+10t_x)))t_y^4)}{5(2+t_x)^2(4+t_x)(3+2t_x)(2+3t_x)(3k(2+t_x)(1+2t_x)+2(-1+t_x)^2t_y)^4}\right\}.$$

The direct effect of t_y on μ and w is given by the first term t_y^2 on the right side, which pushes towards more wage inequality. Taking into account the effect of t_y on the wage variance through θ^*, however, results in an expression that is non-monotone in t_y despite the increase in the return to education through a rise in t_y.

FOSD in L. Set $t_y = 1, t_\theta = 1$, and recall that $t_x > 1$ for H_1 to FOSD H_0. Here we only impose $t_\alpha \geq 1$. Once again to help ensure that θ^* is positive in this example we assume $c(\theta) = k(1-\theta)$, $k > 0$. We consider a FOSD shift in L, by increasing t_α.

The equilibrium investment threshold is given by

$$\theta^* = \frac{9k(2+t_x)(1+2t_x)-3(-1+t_x)t_\alpha^2}{9k(2+t_x)(1+2t_x)+2(-1+t_x)^2t_\alpha^2},$$

where $k > (-t_\alpha^2+t_xt_\alpha^2)/(6+15t_x+6t_x^2)$ guarantees that θ^* is positive, which we impose throughout.

As before, the return to education is

$$\frac{\mathbb{E}[w(x,\alpha)|a=1]}{\mathbb{E}[w(x,\alpha)|a=0]} = \frac{(1-\theta^*)\int\int w(x,\alpha)h_1(x)dxdL}{\theta^*\int\int w(x,\alpha)h_0(x)dxdL}$$
$$= \frac{(-1+t_x)t_x(1+2t_x)t_\alpha^2(9k(1+t_x)(2+t_x)(1+2t_x)-(-1+t_x)^2t_\alpha^2)}{3(9(1+t_x)(2+t_x)^3(k+2kt_x)^2-6k(-1+t_x)(2+t_x)(1+2t_x)(1+t_x+t_x^2)t_\alpha^2+(-1+t_x)^3t_xt_\alpha^4)},$$

where we note that wages now also depend on the realization of α.

The derivative of the education premium with respect to t_α is

$$\frac{\partial\frac{\mathbb{E}[w(x,\alpha)|a=1]}{\mathbb{E}[w(x,\alpha)|a=0]}}{\partial t_\alpha} = 2k(-1+t_x)t_x(2+t_x)^2(1+2t_x)^2t_\alpha$$
$$\times\frac{(27k^2(1+t_x)^2(2+t_x)^2(1+2t_x)^2-6k(-1+t_x)^2(1+t_x)(2+t_x)(1+2t_x)t_\alpha^2-(-1+t_x)^4t_\alpha^4)}{(9(1+t_x)(2+t_x)^3(k+2kt_x)^2-6k(-1+t_x)(2+t_x)(1+2t_x)(1+t_x+t_x^2)t_\alpha^2+(-1+t_x)^3t_xt_\alpha^4)^2},$$

which is positive if $k > (t_\alpha^2 - 2t_xt_\alpha^2 + t_x^2t_\alpha^2)/(6+21t_x+21t_x^2+6t_x^3), t_x > 1, t_\alpha > 1$, which in turn holds under the imposed parametric restrictions and the condition $k > (-t_\alpha^2 + t_xt_\alpha^2)/(6+15t_x+6t_x^2)$ that also ensures that $\theta^* > 0$. Thus, the education premium increases with a FOSD shift in L.

The variance of the wage distribution is given by

$$Var[w(x,\alpha)] = t_\alpha^4 \times 1/(45(2+t_x)^2(4+t_x)(3+2t_x)(2+3t_x)(9k(2+t_x)(1+2t_x)+2(-1+t_x)^2t_\alpha^2)^4)$$
$$\Big\{729(2+t_x)^6(4+t_x)(3+2t_x)(2+3t_x)(k+2kt_x)^4+486(-1+t_x)^2(2+t_x)^4(3+2t_x)(2+3t_x)(k+2kt_x)^3(8+t_x(3+2t_x))t_\alpha^2$$
$$+(-108k^2(-1+t_x)^4(2+t_x)^2(1+2t_x)^2(2+3t_x)(-60+t_x(-63+8(-1+t_x)t_x))t_\alpha^4)$$
$$+(72k(-1+t_x)^6(2+t_x)(1+2t_x)(28+t_x(76+t_x(65+22t_x)))t_\alpha^6)$$
$$+(4(-1+t_x)^8(24+t_x(54+t_x(31+10t_x)))t_\alpha^8))\Big\},$$

which is non-monotone in t_α but would be increasing if there was only the direct effect of t_α on the wage (captured by the first term t_α^4 multiplying the expression in braces). It is the indirect equilibrium effect through θ^* that leads to the nonmonotonicity of the variance in t_α.

Figure 3 illustrates the non-monotonicity of the wage variance as a function of (t_x, t_y, t_α). In all these cases, the return to education increases in the corresponding shift.

B.2 Effects of Alternative Policies

B.2.1 Social Insurance Policy

To obtain the FOC (14), differentiate the planner's objective function with respect to s, taking into account both the direct effect and the indirect one via θ^*. From the resulting expression, first integrate by parts $\int\int(f-w)\Delta hdLdx$ to obtain $\int\int(f_y\mu_{\theta^*} - w_{\theta^*})dLdH + \int\int(f-w)\Delta hdxdL = -\int\int w_{\theta^*}dLdH$; then note that $L(\hat{\alpha}) = \int_0^{\hat{\alpha}}\int dLdH$; finally, use that in equilibrium $U_1 - U_0 = c$

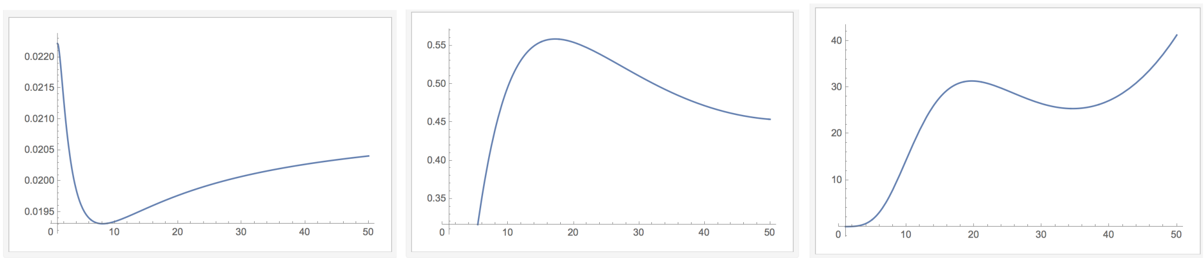

Figure 3: A. Wage variance as a function of t_x with $k = 1$ (left). B. Wage variance as a function of t_y with $k = 2, t_x = 40$ (middle) C. Wage variance as a function of t_α with $k = 15, t_x = 40$ (right). All figures show that the wage variance is not an increasing function of the shifters.

and that $U_1 - U_0$ can be written as $\int_0^{\hat{\alpha}} \int u(w+s)\Delta h dL dx + \int_{\hat{\alpha}}^1 \int u(w)\Delta h dL dx$ to cancel $U_1 - U_0$ and c. Rearrange the remaining terms in the FOC to obtain (14).

The sign of θ^*_s is determined by the effect of s on $U_1 - U_0$, given by

$$\frac{\partial(U_1 - U_0)}{\partial s} = \int_0^{\hat{\alpha}} \left(\int u' dH_1 - \int u' dH_0 \right) dL,$$

which is strictly negative for all $\hat{\alpha} > 0$ since u' is a decreasing function of x and H_1 FOSD H_0.

B.2.2 Policy Shifting H_1

Assume the planner chooses t_x, which leads to a FOSD shift H_1, at cost $\kappa(t_x)$, with $\kappa' > 0, \kappa'' > 0$ and $\kappa(0) = 0$. We also assume budget balance. The planner solves,

$$\max_{t_x \in [0,\bar{t}]} \lambda \left(Q(\theta^*) \int \int u(w(x,\alpha,t_x,\theta^*)) dL(\alpha) dH_0(x) + (1 - Q(\theta^*)) \int \int u(w(x,\alpha,t_x,\theta^*)) dL(\alpha) dH_1(x|t_x) \right.$$
$$\left. - \int_{\theta^*}^1 c(\theta) dQ(\theta) \right) + (1-\lambda)(1-\tau) \left(\int \int \left(f(x,\mu(x,t_x,\theta^*),\alpha) - w(x,\alpha,t_x,\theta^*) \right) dL(\alpha) dH(x,\theta^*,t_x) \right) - \kappa(t_x)$$
$$s.t. \quad \kappa(t_x) = \tau \int \int (f(x,\mu(x,t_x,\theta^*),\alpha) - w(x,\alpha,t_x,\theta^*)) dH(x,\theta^*,t_x) dL(\alpha)), \quad 0 \le \tau \le 1$$

Proceeding as outlined in the social insurance case, we obtain the following FOC:

$$(1 - Q(\theta^*)) \int \int (u + \frac{1-\lambda}{\lambda}(f - w)) dH_{t_x} dL$$
$$+ \int \int \left(u' w_t + \frac{1-\lambda}{\lambda}(f_y \mu_t - w_t) \right) dH dL + \theta^*_{t_x} \int \int \left(u' - \frac{1-\lambda}{\lambda} \right) w_{\theta^*} dH dL = \kappa' \frac{2-\lambda}{\lambda}$$

where the first effect is the direct beneficial effect of a FOSD shift in H_1 (a productivity improvement benefits both workers and firms). But there are also two general equilibrium effects that can attenuate this marginal benefit. The first effect in the second line reflects that a shift in the *distribution* H_1 impacts wages and profits and the second term in the second line reflects that the planner also wants to mitigate the inefficiency of investment.

B.2.3 Policy Shifting Q

Finally, let the planner choose t_θ, which leads to a FOSD shift in Q, at cost $\kappa(t_\theta)$, with $\kappa' > 0, \kappa'' > 0$ and $\kappa(0) = 0$. We also assume budget balance. The planner solves,

$$\max_{t_\theta \in [0,\bar{t}]} \lambda \left(Q(\theta^*|t_\theta) \int\int u(w(x,\alpha,t_\theta,\theta^*))dL(\alpha)dH_0(x) + (1 - Q(\theta^*|t_\theta)) \int\int u(w(x,\alpha,t_\theta,\theta^*))dL(\alpha)dH_1(x) \right.$$
$$\left. - \int_{\theta^*}^{1} c(\theta)dQ(\theta|t_\theta) \right) + (1-\lambda)(1-\tau) \left(\int\int \left(f(x,\mu(x,t_\theta,\theta^*),\alpha) - w(x,\alpha,t_\theta,\theta^*)\right) dL(\alpha)dH(x,\theta^*,t_\theta) \right) - \kappa(t_\theta)$$
$$s.t. \quad \kappa(t_\theta) = \tau \int\int (f(x,\mu(x,t_\theta,\theta^*),\alpha) - w(x,\alpha,t_\theta,\theta^*))dL(\alpha)dH(x,\theta^*,t_\theta), \quad 0 \le \tau \le 1$$

After some algebra, the FOC can be expressed as:

$$- Q_{t_\theta} \left(\int\int udLdH_1 - \int\int udLdH_0 \right) - \int_{\theta^*}^{1} c(\theta)dQ_{t_\theta}$$
$$+ \int\int \left(u'w_t + \frac{1-\lambda}{\lambda}(f_y\mu_t - w_t) \right) dHdL + \theta^*_{t_\theta} \int\int \left(u' - \frac{1-\lambda}{\lambda} \right) w_{\theta^*} dHdL = \kappa' \frac{2-\lambda}{\lambda}$$

where the first line is the direct beneficial effect of of a FOSD shift in Q (stemming from having more individuals invest in education for a given θ^* and from a lower cost for those who invest). But there are also two general equilibrium effects that can attenuate this marginal benefit. The first effect in the second line reflects that a shift in the *distribution* Q impacts wages and profits and the second term reflects that the planner also wants to mitigate the inefficiency of investment.

B.3 Identification

B.3.1 Example (FOSD Shifts in Distribution)

Under the parametric assumptions in the example in the text, the unique equilibrium is given by:

$$\mu(x, t_y, t_x, t_\theta, \theta^*,) = t_y \left(\left(1 - \frac{\theta^*}{t_\theta} \right) x^{t_x} + \frac{\theta^*}{t_\theta} x \right) \tag{35}$$

$$w(x, t_y, t_x, t_\theta, \theta^*) = \alpha^2 \int_0^x \mu(s, t_y, t_x, t_\theta, \theta^*)ds \tag{36}$$

$$\theta^*(t_y, t_x, t_\theta, t_\alpha) = \frac{3t_\theta(2+t_x)(1+2t_x) - 3(t_\alpha^2/3)(-1+t_x)t_\theta t_y}{3t_\theta(2+t_x)(1+2t_x) + 2(t_\alpha^2/3)(-1+t_x)^2 t_y} \tag{37}$$

Substituting (35) and (37) into (36), yields the reduced form wage function:

$$
\begin{aligned}
&\tilde{w}(x, t_y, t_x, t_\theta, t_\alpha) \\
&= C + \frac{3\alpha^2 t_y(3(2+t_x)(1+2t_x) - t_\alpha^2(-1+t_x)t_y)}{18t_\theta(2+t_x)(1+2t_x) + 4t_\alpha^2(-1+t_x)^2 t_y} x^2 + \frac{\alpha^2(1+2t_x)t_y(9(-1+t_\theta)(2+t_x) + t_\alpha^2(-1+t_x)t_y)}{(1+t_x)(9t_\theta(2+t_x)(1+2t_x) + 2t_\alpha^2(-1+t_x)^2 t_y)} x^{t_x+1} \\
&:= \beta_0 + \beta_1 x^2 + \beta_2 x^{t_x+1}
\end{aligned}
$$

where C denotes the constant of integration.

We now prove each the two assertions in the text:

(i) If t_x is known and there is no aggregate risk $t_\alpha = 0, \alpha = 1$, then solve the regression coefficients β_1 and β_2 for (t_y, t_θ) to obtain:

$$
\begin{aligned}
t_y &= \beta_2 + 2\beta_1 + \beta_2 t_x \\
t_\theta &= \frac{(\beta_2 + 2\beta_1 + \beta_2 t_x)(9(2+t_x)(1+2t_x) - (-1+t_x)(2\beta_1 + 4\beta_1 t_x + 3\beta_2(1+t_x)))}{18\beta_1(2+t_x)(1+2t_x)}
\end{aligned}
$$

(ii) If t_x and t_θ are known, then solve regression coefficients β_1 and β_2 for (t_y, t_α) to obtain:

$$
\begin{aligned}
t_y &= \frac{\beta_2 + 2\beta_1 + \beta_2 t_x}{\alpha^2} \\
t_\alpha &= \frac{3\alpha\sqrt{\beta_2 + 2\beta_1 - 2\beta_1 t_\theta + \beta_2 t_x}\sqrt{2 + 5t_x + 2t_x^2}}{\sqrt{-3\beta_2^2 - 8\beta_2\beta_1 - 4\beta_1^2 - 3\beta_2^2 t_x - 4\beta_2\beta_1 t_x - 4\beta_1^2 t_x + 3\beta_2^2 t_x^2 + 8\beta_2\beta_1 t_x^2 + 8\beta_1^2 t_x^2 + 3\beta_2^2 t_x^3 + 4\beta_2\beta_1 t_x^3}}
\end{aligned}
$$

These solutions are unique, given that we restrict the shifters of the distributions to be positive (this is a natural restriction since it says that the distributions have positive support) and given that we know the realization of α in (ii) (which can, for instance, be observed TFP), as we assume in the model. Otherwise, it would not be possible to compute the wage function.

B.3.2 Example (FOSD and IR Shifts in Distributions)

Under the modified example with L uniform on $[0.5-t_\alpha, 0.5+t_\alpha]$ and Q uniform on $[0.5-t_\theta, 0.5+t_\theta]$, the unique equilibrium is given by:

$$
\mu(x, t_y, t_x, t_\theta, \theta^*) = t_y\left(\left(1 - \frac{\theta^* - \frac{1}{2} + t_\theta}{2t_\theta}\right)x^{t_x} + \frac{\theta^* - \frac{1}{2} + t_\theta}{2t_\theta}x\right) \tag{38}
$$

$$
w(x, t_y, t_x, t_\theta, \theta^*) = \alpha^2 \int_0^x \mu(s, t_y, t_x, t_\theta, \theta^*) ds \tag{39}
$$

$$
\theta^*(t_y, t_x, t_\theta, t_\alpha) = \frac{72t_\theta(2+t_x)(1+2t_x) - (3+4t_\alpha^2)(-1+t_x)(1-t_x+2t_\theta(2+t_x))t_y}{72t_\theta(2+t_x)(1+2t_x) + 2(3+4t_\alpha^2)(-1+t_x)^2 t_y} \tag{40}
$$

Substituting (38) and (40) into (39), yields the reduced form wage function (with some abuse we also call it $\tilde{w}$ and name the coefficients β_i, $i = 0, 1, 2$):

$$\begin{aligned}
&\tilde{w}(x, t_y, t_x, t_\theta, t_\alpha) \\
&\quad = C + \frac{3\alpha^2 t_y(6(1+2t_\theta)(2+t_x)(1+2t_x) - (3+4t_\alpha^2)(-1+t_x)t_y)}{144t_\theta(2+t_x)(1+2t_x) + 4(3+4t_\alpha^2)(-1+t_x)^2 t_y} x^2 \\
&\qquad\quad + \frac{\alpha^2(1+2t_x)t_y(18(-1+2t_\theta)(2+t_x) + (3+4t_\alpha^2)(-1+t_x)t_y)}{2(1+t_x)(36t_\theta(2+t_x)(1+2t_x) + (3+4t_\alpha^2)(-1+t_x)^2 t_y)} x^{t_x+1} \\
&\quad := \beta_0 + \beta_1 x^2 + \beta_2 x^{t_x+1}
\end{aligned}$$

where C denotes the constant of integration.

We now prove each of the two assertions made in the text:

(i) If t_x is known and there is no aggregate risk $t_\alpha = 0, \alpha = 1$, then solve regression coefficients β_1 and β_2 for (t_y, t_θ) to obtain:

$$\begin{aligned}
t_y &= \beta_2 + 2\beta_1 + \beta_2 t_x \\
t_\theta &= -\frac{(\beta_2 + 2\beta_1 + \beta_2 t_x)(-3\beta_2 3(-1+t_x^2) - 2(1+2t_x)(\beta_1 3(-1+t_x) - 9(2+t_x)))}{36(\beta_2 - 2\beta_1 + \beta_2 t_x)(2 + 5t_x + 2t_x^2))}
\end{aligned}$$

(ii) If t_x and t_θ are known, then solve regression coefficients β_1 and β_2 for (t_y, t_α) to obtain:

$$\begin{aligned}
t_y &= \frac{\beta_2 + 2\beta_1 + \beta_2 t_x}{\alpha^2} \\
t_\alpha &= \frac{\sqrt{(-9\beta_2^2(-1+t_x)(1+t_x)^2 - 12\beta_1(1+2t_x)(\beta_1(-1+t_x) + 3\alpha^2(-1+2t_\theta)(2+t_x)) - 6\beta_2(1+t_x)(2+t_x)(2\beta_1(-1+t_x) - 3\alpha^2(1+2t_\theta)(1+2t_x))}}{2\sqrt{(-1+t_x)(\beta_2 + 2\beta_1 + \beta_2 t_x)(2\beta_1 + 4\beta_1 t_x + 3\beta_2(1+t_x))}}
\end{aligned}$$

Again, these solutions are unique, given that we restrict the shifters to be positive (this is a natural restriction since it says that the distributions have positive support) and maintaining the assumption that we know the realization of α in (ii).

References

Acemoglu, D., and D. Autor (2011): *Skills, Tasks and Technologies: Implications for Employment and Earnings*vol. 4 of *Handbook of Labor Economics*, chap. 12, pp. 1043–1171. Elsevier.

Becker, G. S. (1973): "A Theory of Marriage: Part I," *Journal of Political Economy*, 81(4), 813–46.

Bhaskar, V., and E. Hopkins (2016): "Marriage as a Rat Race: Noisy Premarital Investments with Assortative Matching," *Journal of Political Economy*, 124, 992–1045.

Billingsley, P. (1995): *Probability and Measure*, Wiley Series in Probability and Statistics. Wiley.

Bloom, N., M. Floetotto, N. Jaimovic, I. Saporta-Ekstein, and S. Terry (2018): "Really Uncertain Business Cycles," *Econometrica*, 86, 1031–1065.

Cairo, I. (2013): "The Slowdown in Business Employment Dynamics: The Role of Changing Skill Demands," *Working Paper*.

Castro, R., and D. Coen-Pirani (2016): "Explaining the Evolution of Educational Attainment in the United States," *American Economic Journal: Macroeconomics*, 8(3), 77–112.

Chade, H., J. Eeckhout, and L. Smith (2017): "Sorting through Search and Matching Models in Economics," *Journal of Economic Literature*, 55(2), 493–544.

Chetty, R., J. N. Friedman, and J. E. Rockoff (2014): "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates," *The American Economic Review*, 104(9), 2593–2632.

Chiappori, P.-A., M. C. Dias, and C. Meghir (2018): "The Marriage Market, Labor Supply, and Education Choice," *Journal of Political Economy*, 126(S1), S26–S72.

Chiappori, P.-A., M. Iyigun, and Y. Weiss (2009): "Investment in Schooling and the Marriage Market," *American Economic Review*, 99(5), 1689–1713.

Cole, H. L., G. J. Mailath, and A. Postlewaite (2001): "Efficient Non-Contractible Investments in Large Economies," *Journal of Economic Theory*, 101(2), 333–373.

Cunha, F., and J. Heckman (2016): "Decomposing Trends in Inequality in Earnings into Forecastable and Uncertain Components," *Journal of Labor Economics*, 34(S2), S31–S65.

Del Boca, D., C. Flinn, and M. Wiswall (2014): "Household Choices and Child Development," *The Review of Economic Studies*, 81(1), 137–185.

Dunne, T., L. Foster, J. Haltiwanger, and K. Troske (2004): "Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment," *Journal of Labor Economics*, 22(2), 397–430.

Eika, L., M. Mogstad, and B. Zafar (2018): "Educational assortative mating and household income inequality," *Journal of Political Economy*, forthcoming.

Gale, D., and L. Shapley (1962): "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, 69, 9–15.

Gollier, C. (2004): *The Economics of Risk and Time*. MIT Press.

Heathcote, J., K. Storesletten, and G. L. Violante (2009): "Quantitative Macroeconomics with Heterogeneous Households," *Annual Review of Economics*, 1(1), 319–354.

Huggett, M., G. Ventura, and A. Yaron (2011): "Sources of Lifetime Inequality," *American Economic Review*, 101(7), 2923–54.

Karlin, S., and Y. Rinott (1980): "Classes of orderings of measures and related correlation inequalities. I. Multivariate totally positive distributions," *Journal of Multivariate Analysis*, 10, 467–498.

Keane, M., and K. I. Wolpin (1997): "The Career Decisions of Young Men," *Journal of Political Economy*, 105(3), 473–522.

Kremer, M. (1993): "The O-ring theory of economic development," *Quarterly Journal of Economics*, 3, 551–575.

Kremer, M., and E. Maskin (1996): "Wage inequality and segregation by skill," NBER.

Legros, P., and A. Newman (2010): "Co-ranking mates: Assortative matching in marriage markets," *Economics Letters*, 106, 177–179.

Legros, P., and A. F. Newman (2002): "Monotone matching in perfect and imperfect worlds," *The Review of Economic Studies*, 69, 925–942.

Legros, P., and A. F. Newman (2007): "Beauty Is a Beast, Frog Is a Prince: Assortative Matching with Nontransferabilities," *Econometrica*, 75(4), 1073–1102.

Mailath, G. J., L. Samuelson, and A. Postlewaite (2013): "Pricing and investments in matching markets," *Theoretical Economics*, 8(2).

Muller, A., and D. Stoyan (2002): *Comparison Methods for Stochastic Models and Risks*. Wiley and Sons.

Nöldeke, G., and L. Samuelson (2015): "Investment and Competitive Matching," *Econometrica*, 83(3), 835–896.

Peters, M., and A. Siow (2002): "Competing Premarital Investments," *Journal of Political Economy*, 110(3), 592–608.

Rath, K. P. (1992): "A direct proof of the existence of pure strategy equilibria in games with a continuum of players," *Economic Theory*, 2, 427–433.

Rothschild, M., and J. Stiglitz (1970): "Increasing Risk: I. A Definition," *Journal of Economic Theory*, 2, 225–43.

Shaked, M., and J. G. Shanthikumar (2007): *Stochastic Orders*. Springer.

Shapley, L., and M. Shubik (1972): "The Assignment Game I: The Core," *International Journal of Game Theory*, pp. 111–130.

Smith, L. (2006): "The Marriage Model with Search Frictions," *Journal of Political Economy*, 114, 1124–1144.

Song, J., D. Price, F. Guvenen, N. Bloom, and T. von Wachter (2018): "Firming Up Inequality," Working Papers 750, Federal Reserve Bank of Minneapolis.

Storesletten, K., C. Telmer, and A. Yaron (2004): "Consumption and risk sharing over the life cycle," *Journal of Monetary Economics*, 51(3), 609–633.

Topkis, D. M. (1998): *Supermodularity and complementarity*. Princeton university press.

Zhang, H. (2015): "Pre-Matching Gambles," Michigan State University Working Paper.