

Pricing and investments in matching markets

GEORGE J. MAILATH

Department of Economics, University of Pennsylvania

ANDREW POSTLEWAITE

Department of Economics, University of Pennsylvania

LARRY SAMUELSON

Department of Economics, Yale University

Different markets are cleared by different types of prices: seller-specific prices that are uniform across buyers in some markets and personalized prices tailored to the buyer in others. We examine a setting in which buyers and sellers make investments before matching in a competitive market. We introduce the notion of premuneration values—the values to the transacting agents prior to any transfers—created by a buyer–seller match. Personalized-price equilibrium outcomes are independent of premuneration values and exhibit inefficiencies only in the event of “coordination failures,” while uniform-price equilibria depend on premuneration values and, in general, feature inefficient investments even without coordination failures. There is thus a trade-off between the costs of personalizing prices and the inefficient investments under uniform prices. We characterize the premuneration values under which uniform-price equilibria similarly exhibit inefficiencies only in the event of coordination failures.

KEYWORDS. Directed search, matching, premuneration value, prematch investments, search.

JEL CLASSIFICATION. C78, D40, D41, D50, D83.

1. INTRODUCTION

1.1 *Investment and matching markets*

We analyze a model in which agents match to generate a surplus that they then split. Prior to matching, the agents make investments that affect the size of the surplus.

George J. Mailath: gmailath@econ.upenn.edu

Andrew Postlewaite: apostlew@econ.upenn.edu

Larry Samuelson: Larry.Samuelson@yale.edu

The first working paper version of this paper was titled “Pricing in Matching Markets” (2010). Earlier versions (without the existence proof) were distributed at various presentations beginning in 2003. We thank Philipp Kircher, Ben Lester, Antonio Penta, and participants at numerous seminars and conferences for helpful comments, the co-editor (Gadi Barlevy) and referees for their thoughtful reports, and Zehao Hu for excellent research assistance. We thank the National Science Foundation (Grants SES-0350969, SES-0549946, SES-0648780, and SES-0850263) for financial support.

Copyright © 2013 George J. Mailath, Andrew Postlewaite, and Larry Samuelson. Licensed under the [Creative Commons Attribution-NonCommercial License 3.0](https://creativecommons.org/licenses/by-nc/3.0/). Available at <http://econtheory.org>.

DOI: [10.3982/TE1189](https://doi.org/10.3982/TE1189)

Consider, for example, a continuum of workers and a continuum of firms, each with unit mass. Each worker and firm first makes a costly investment in an *attribute*: firms invest in technology while workers invest in human capital. In the second stage, workers and firms match and generate a surplus. In the absence of any monetary transfers, the firm owns the output produced by the worker, while the worker bears the cost of the effort exerted in the course of production and owns the value of the skills learned in the course of production. We call these costs and benefits the agents' *premuneration values* (from *pre* plus the Latin *munerare*, to give or to pay). Both the surplus and its division between buyer and seller premuneration values depend on the attributes the agents choose. The worker's human capital may enhance the quality of the output owned by the firm, and the firm's technology may enhance the value of on-the-job learning to the worker. The final division of the surplus between the worker and the firm is determined by the premuneration values and a subsequent monetary transfer.

A large literature examines settings in which agents make investments before trading in a market. One extreme, which Williamson (1975) discusses, treats the case of a single buyer and seller. The agents' post-investment market power then gives rise to a "holdup" problem that prompts inefficient investments. At the other extreme, Cole et al. (2001) and Peters and Siow (2002) examine models with competitive post-investment markets, featuring a continuum of heterogeneous buyers and sellers and frictionless trading, showing that equilibria with efficient investments exist.

Our analysis falls between these extremes. Our post-investment markets feature continua of heterogeneous agents, but we introduce a key friction into the trading process, namely that firms (continuing with our example) cannot observe workers' attribute choices.

1.2 *Personalized pricing*

The appropriate equilibrium notion in our setting is not obvious, to a large extent because we must determine the returns to attributes that nobody chooses. Continuing with our example, it is helpful to first consider the case in which firms *can* observe workers' investments. We refer to this as *personalized pricing*, since wages can be conditioned on the chosen attributes of *both* the firm and the worker. In this setting, an equilibrium is a specification of the attribute chosen by each firm and worker—a wage function and a matching of firms and workers such that no agent can increase his utility by changing his decision and such that markets clear.

This equilibrium notion is similar to Walrasian equilibrium, except that the wage function attaches a value only to pairs of firm and worker attributes that are chosen in the investment stage, and not to unchosen attributes. In the language of Walrasian equilibrium, the price vector includes a price for every good present in the market, but not for nonexistent goods. We address the latter with a requirement that no firm (say) can unilaterally deviate to adopt some currently unchosen attribute and then match with a worker at her existing attribute, while splitting the surplus in such a way as to make both better off.

Environments in which people decide which goods to bring to market or which investments to make before entering the market readily give rise to coordination failures.

In the extreme, there is an autarkic equilibrium in which neither firms nor workers invest because no one expects the other side to invest. We could preclude such coordination failures by simply assuming that prices exist for *all* attributes, in and out of the market. On the one hand, we find the existence of such prices counterintuitive. On the other hand and more importantly, like Makowski and Ostroy (1995), we expect coordination failures to be endemic when people must decide what goods to market, and hence we think it important to work with a model that does not preclude them.

Personalized-price equilibria can be shown to exist using a variant of the existence argument in Cole et al. (2001). There exist coordination-failure equilibria with inefficient investments, but there also exist efficient equilibria in which no worker–firm pair, matched or unmatched, could be made jointly better off, even if they could commit to their investments prior to matching. Premuneration values are irrelevant, in the sense that every personalized-price equilibrium outcome remains an equilibrium outcome irrespective of the allocation of premuneration values.

1.3 Uniform pricing

We are interested in the case in which firms cannot observe workers' attribute choices. Wages can then depend only on firms' attributes, and we speak of *uniform pricing* to emphasize that workers who have chosen different attributes must be offered the same wage. Our equilibrium notion is a specification of the attribute chosen by each firm and worker, a wage function, and a choice of firm on the part of each worker, such that no agent can increase his utility by changing his decision and such that markets clear. Analogous to personalized-price equilibrium, the possibility of coordination failures again arises.

We show that a uniform-price equilibrium exists. However, these equilibria are, in general, inefficient, even if they exhibit no coordination failures. There exist efficient uniform-price equilibria if, and essentially only if, firms' premuneration values are independent of workers' attributes. Hence, premuneration values matter for uniform-price equilibria.

An efficient equilibrium always exists under personalized pricing, but personalizing prices may be costly. Expanding beyond our worker–firm example, estimates from 11 highly selective liberal arts colleges indicate that they spent about \$3,000 on admissions, i.e., ascertaining students' attributes in an attempt to personalize prices, per matriculating student in 2004.¹ The cost to identify whether a foreign high school diploma comes from a legitimate high school is \$100.² There may thus be substantial savings from posting uniform prices and letting buyers sort themselves, *if* the premuneration values are such that uniform prices can do this sorting. Alternatively, if the premuneration values are such that uniform prices cannot duplicate the allocation of personalized prices, and if transactions costs or institutional considerations preclude personalized prices, then market outcomes will be inefficient.

¹Memorandum, Office of Institutional Research and Analysis, University of Pennsylvania, July 2004. We thank Bernie Lentz for his help with these data.

As we discuss in Section 6.1.3, the admissions decision itself amounts to a personalization of prices.

²"Vetting Those Foreign College Applications," *New York Times*, September 29, 2004, p. A21.

1.4 *Premuneration values*

The premuneration values of the firms in our motivating example typically depend on their employees' attributes—better skilled and more productive employees enhance the quality and quantity of a firm's output. The business pages are filled with announcements of the good news that a firm has hired a particularly prized employee. Moving beyond this example, students are matched with universities after students incur substantial preparation costs and universities hire faculty. Both sides care about the investments the other side makes. Universities reap benefits well beyond tuition revenues from talented students, and students clamor for spots at elite universities. Similarly, an aspiring faculty member cares about the investments a university makes in facilities and other faculty, while the university cares about the investment in knowledge and research capabilities of the potential recruit.

The central message of this paper is that there is a trade-off between the costs of personalizing pricing and the inefficiency of uniform pricing. One might hope to ameliorate this trade-off by reallocating the premuneration values. In particular, premuneration values are affected by the explicit and implicit property rights to the costs and benefits that flow from a match. For example, one could arrange the premuneration values in a university–student interaction so that the university owns all of the surplus. This would require a somewhat unconventional arrangement in which the university shares in the future income of students to whom it gives degrees. However, income-contingent loans in a number of countries (including Australia, Sweden, and New Zealand) that effectively give the lender a share of students' future income (Johnstone 2004) attest to the possibility of such an arrangement.³

There are often, however, constraints on the design of premuneration values. Moral hazard problems loom especially large. If universities owned a large share of students' enhanced future income streams, why would the students exert the effort required to realize this future income? How are we to measure and collect the increment to income attributable to the university education? Such an arrangement might also require changes in labor laws that preclude involuntary servitude. More generally, laws concerning workplace safety, the (in)ability to surrender legal rights, the division of marital assets, and the custody and sale of children constrain the allocation of premuneration values. Our analysis points to the cost of such constraints or institutional arrangements, in the form of personalization costs or inefficient uniform pricing.

³In the summer of 2010, the United Kingdom debated the possibility of partially funding higher education through a “graduate tax” levied on college graduates' income (<http://www.bbc.co.uk/news/education-10649459>). Luigi Zingales, in a recent column in the *New York Times* (“The College Graduate as Collateral,” *New York Times*, June 14, 2012, p. A35) argued that the U.S. “Internal Revenue Service could perform collection services on the behalf of private lenders.” Basketball star Yao Ming (Houston Rockets) has a contract with the Chinese Basketball Association calling for 30% of his National Basketball Association earnings to be paid to the Chinese Basketball Association (in which he played prior to joining the Rockets), while another 20% goes to the Chinese government. Similar arrangements hold for Wang Zhizhi (Dallas Mavericks) and Menk Bateer (Denver Nuggets and San Antonio Spurs). (See the *Detroit News*, April 26, 2002, <http://www.detnews.com/2002/pistons/0204/27/sports-475199.htm/>.) We can view the initial match between Yao Ming and his Chinese team as producing a surplus that includes the enhanced value of his earnings as a result of developing his basketball skills, and view the contract as setting premuneration values.

1.5 *Related literature*

Our model is related to the literature on competitive search (see [Guerrieri et al. 2010](#) and [Peters 2010](#) for recent contributions and for pointers to the literature). We depart from a standard complete-information competitive search model in three respects. First, we include a first stage at which investments are made, whereas most competitive search models begin with buyers and sellers who have exogenously given attributes. Second, we examine matching outcomes that turn out to be ex post efficient (given equilibrium investments), rather than focussing on the inefficiencies caused by coordination failures in the search process. Third, like [Guerrieri et al. \(2010\)](#), we introduce a key friction into the competitive search model—*asymmetric information*—in the sense that sellers cannot condition prices on buyers' characteristics.

Our analysis differs from that of [Guerrieri et al. \(2010\)](#) most notably in the nature of the prematching investment choice. In their model, only sellers make investments and they consist of paying a fixed cost to participate in the second stage. Sellers who enter the second stage are homogenous, making it more difficult to screen buyers than in our model. Premuneration values play no role in their model and coordination failures cannot arise. The resulting equilibria are inefficient, and the inefficiencies arise not at the investment stage, but out of constraints on the ability to screen workers. In contrast, in our model, the continuum of possible investments available to agents on both sides of the market is the source of inefficiencies, with the existence and nature of inefficiency depending on the nature of the premuneration values.

Variants of competitive search models are used to accommodate sources of friction other than one-sided asymmetric information. The most obvious such friction is to assume that buyers and sellers cannot instantly match. Instead, buyers must engage in costly search, including the prospects of being either temporarily or permanently unable to find a seller (e.g., [Niederle and Yariv 2008](#) and [Peters 2010](#)). We forgo including such considerations so as to focus on one friction at a time, in our case asymmetric information. [Bhaskar and Hopkins \(2011\)](#) and [Gall et al. \(2006, 2009\)](#) examine an alternative class of models in which information is complete and, hence, different prices can be set for different workers, but inefficiencies arise out of exogenous limitations on the ability to reallocate the surplus in a match via transfers, including cases in which no transfers can be made. In contrast to these models, *any* division of the surplus between any matched pair of agents is possible in our setting; the only constraint we place on the reallocation is that the implementing transfer is independent of the buyers' attributes.

Our focus on creating incentives for efficient investments is shared by a number of other papers.⁴ [Acemoglu and Shimer \(1999\)](#) analyze a worker–firm model in which firms (only) make ex ante investments. If wages are determined by post-match bargaining, then the resulting effective market power gives rise to a standard holdup problem that induces firms to underinvest. The holdup problem disappears if workers have no bargaining power, but then there is excess entry on the part of firms. [Acemoglu and Shimer](#)

⁴Early indications that competitive search under complete information might create investment incentives appear in [Hosios \(1990\)](#), [Moen \(1997\)](#), and [Shi \(2001\)](#). [Masters \(2011\)](#) examines a model with two-sided investments.

show that efficient outcomes can be achieved if the bargaining process is replaced by wage posting on the part of firms, followed by competitive search. [de Meza and Lockwood \(2010\)](#) examine an investment and matching model that gives rise to *excess* investment. Their overinvestment possibility rests on a discrete set of investment choices and the presence of bargaining power in a noncompetitive post-investment stage. In contrast, the competitive post-investment markets of [Cole et al. \(2001\)](#) and [Peters and Siow \(2002\)](#) lead to efficient two-sided investments.

Moving from complete-information to incomplete-information matching models typically gives rise to issues of either screening, as considered here, or signaling. See [Cole et al. \(1995\)](#), [Hopkins \(2012\)](#), [Hoppe et al. \(2009\)](#), and [Rege \(2008\)](#) for models that incorporate signaling into matching models with investments.

Finally, in [Mailath et al. \(2012\)](#) we obtain comparative statics results in a simplified version of the model studied here. In particular, we investigate the impact of changes in premuneration values in a parametric model in which only buyers or only sellers can choose attributes.

2. THE MODEL

2.1 *The market*

There is a unit measure of buyers whose types are indexed by β and distributed uniformly on $[0, 1]$, and a unit measure of sellers whose types are indexed by σ and distributed uniformly on $[0, 1]$. For ease of reference, buyers are female and sellers are male.

Buyers and sellers have an outside option (with payoff zero) that precludes participation in the matching process. If they do not take this option, they make choices in two stages. First, simultaneously each buyer chooses an attribute $b \in \mathbb{R}_+$ and each seller chooses an attribute $s \in \mathbb{R}_+$. Second, buyers and sellers match, with each match generating a surplus to be split between the participating agents. Under a strict buyer–seller interpretation, a match is a trade. More generally, a matched pair of agents engages in some productive activity.

Attributes are costly, but enhance the surplus generated in the second stage. To keep the analysis tractable, we assume that agents' types affect the first-stage cost of investment but not the second-stage surplus, which depends only on the attributes chosen by the agents. In particular, the cost of attribute $b \in \mathbb{R}_+$ to buyer β is given by $c_B(b, \beta)$ and the cost of attribute $s \in \mathbb{R}_+$ to seller σ is given by $c_S(s, \sigma)$.

Suppose that a buyer and seller match, but (presumably counterfactually) no transfers are made. Both the buyer and the seller receive some value from the match, which we call their *premuneration value*. A firm that does not pay its employee captures the value of the employee's production. The employee's premuneration value includes the cost of her effort, but may also include the value of her enhanced human capital stemming from her association with the firm. We let $h_B(b, s)$ denote the premuneration value of the buyer and $h_S(b, s)$ denote the premuneration value of the seller from a match between a buyer with attribute b and a seller with attribute s . The surplus of the match is

then

$$v(b, s) \equiv h_B(b, s) + h_S(b, s).$$

In other words, the surplus from the match, in the absence of transfers, is still divided between the buyer and the seller, with the premuneration values describing that division.⁵

The premuneration values depend on the nature of the interaction between the two agents and the legal and institutional environment in which that interaction takes place. For example, the law may stipulate that the employer owns the output produced by an employee and owns any patents that emerge from the employee's work, but that the employee owns the value of any contacts she makes while on the job.

The important point is that a match creates a surplus, independent of transfers. Some of this surplus is owned by the seller and the rest by the buyer, as specified by the premuneration values. Premuneration values are thus the counterparts of endowments in standard general equilibrium models.

Transfers alter the division of the surplus. A match between a buyer and a seller with attribute choices (b, s) at a price p yields a gross (i.e., ignoring investment costs) buyer payoff of

$$h_B(b, s) - p$$

and a gross seller payoff of

$$h_S(b, s) + p.$$

The attribute s chosen by a seller is observable, and prices, in general, reflect this in that sellers with different values of s trade at different prices. However, we assume that trade must take place at *uniform prices*. Two restrictions are embodied in this restriction. First, we think of the agents as being faced with *prices* at which they are free to trade or not. We exclude the possibility that sellers offer contracts or menus of contracts that combine different prices with various probabilities of trade or combine different prices with variation along some other margin. This is potentially restrictive, as such contracts can effectively screen agents whose characteristics cannot be observed. Indeed, we suggest that sellers unable to observe buyers' characteristics distort their attribute choices so as to screen buyers more effectively. We rule out such contracts because they typically require a commitment not to trade under circumstances in which profitable trading opportunities are known to exist. We are on the one hand, skeptical of such commitments, and on the other, would want any such model to be explicit about the subsequent trading possibilities.

Second, prices must be *uniform*, meaning that prices can be conditioned *only* on seller attributes. Any buyer who trades with a given seller does so at the same price,

⁵Notice that premuneration values are *not* the agents' outside options should they reject a match, but quite to the contrary are the values the agents receive if they *do* match, but make no transfers. If transfers are impossible, the premuneration values are the values of the match. Utility is perfectly transferable in our model, but premuneration values are nonetheless well defined, describing values in the absence of transfers.

regardless of the buyer's attribute (though trades involving different sellers may occur at different prices).

Several factors could constrain prices to be uniform. First, it may be prohibitively expensive for sellers to observe buyers' characteristics. For example, firms may be unable to observe whether potential employees have invested in effective work habits. Second, tailoring prices to buyers' attribute choices may entail prohibitive menu costs. A college may prefer to set uniform prices rather than bear the cost of an admissions department to vet applicants carefully. Similarly, it may be costless to use generic contract forms to make a standard offer to every buyer who appears, while tailoring offers to buyers' characteristics requires a costly legal process. Third, legal restrictions may prescribe uniform pricing. For example, employers may be prohibited from discriminating against potential employees whose attributes make them potentially expensive health risks or union contracts may prohibit wage discrimination.

In each case, the constraints that give rise to uniform pricing also determine on which of the two parties' attributes prices can be conditioned. If buyer attributes are unobservable, then the only possibility is to condition prices on seller attributes. It is convenient to consistently call the side of the market on which prices can be conditioned *sellers*. Prices may then be either positive or negative, and the agent we call a seller may in ordinary parlance be called either a buyer or seller.

2.2 Example: Basic structure

We introduce here an example that we carry throughout the analysis. The remuneration values are

$$h_B(b, s) = \theta bs \quad \text{and} \quad h_S(b, s) = (1 - \theta)bs$$

for $\theta \in (0, 1]$ (footnote 6 explains why $\theta = 0$ is excluded), so that the surplus function is given by $v(b, s) = bs$. The cost functions are

$$c_B(b, \beta) = \frac{b^3}{3\beta} \quad \text{and} \quad c_S(s, \sigma) = \frac{s^3}{3\sigma}.$$

Given the supermodularity of the surplus function bs , it is a familiar result that efficiency requires positive assortative matching and, hence, that a buyer of type β is matched with a seller of type $\sigma = \beta$. Efficiency further requires that such a pair choose attributes b and s so as to maximize their net (of investment cost) surplus. Appendix A.1 presents the straightforward calculation that the efficient outcome entails attribute-choice functions

$$\mathbf{b}(\beta) = \beta \quad \text{and} \quad \mathbf{s}(\sigma) = \sigma,$$

so that the pair produces total surplus σ^2 for a total net surplus $\frac{1}{3}\sigma^2$.

3. EQUILIBRIUM

3.1 Assumptions

ASSUMPTION 1 (Supermodularity). *The premuneration values $h_B: \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}$ and $h_S: \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}$ are C^2 , increasing in b and s , and satisfy⁶*

$$\frac{\partial^2 h_B}{\partial b \partial s} > 0 \quad \text{and} \quad \frac{\partial^2 h_S}{\partial b \partial s} \geq 0.$$

There is a simple class of problems for which this assumption holds that includes our example: premuneration values constitute fixed shares of the surplus, or $h_B(b, s) = \theta v(b, s)$ and $h_S(b, s) = (1 - \theta)v(b, s)$ for some $\theta \in (0, 1]$, and the surplus function $v: \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}$ is strictly supermodular ($\partial^2 v / \partial b \partial s > 0$) as well as (twice continuously) differentiable and increasing in b and s .

Our next assumption is a “no free surplus” requirement that matches are not profitable without investments:

ASSUMPTION 2 (Essentiality). *The premuneration values $h_B(b, 0)$ and $h_B(0, s)$ are constant in b and s , respectively, and*

$$h_B(0, 0) + h_S(0, 0) = 0.$$

The following single-crossing condition requires that higher-index buyers and sellers are more productive, in the sense that they have lower investment costs.

ASSUMPTION 3 (Single crossing). *The cost function $c_B: \mathbb{R}_+ \times [0, 1] \rightarrow \mathbb{R}_+$ is C^2 , strictly increasing and convex in b , with $c_B(0, \beta) = 0 = \partial c_B(0, \beta) / \partial b$, and*

$$\frac{\partial^2 c_B}{\partial b \partial \beta} < 0.$$

The cost function c_S satisfies analogous conditions.

Our next assumption ensures that efficient attribute choices exist and are bounded.

ASSUMPTION 4 (Boundedness). *There exists \bar{b} such that for all $b > \bar{b}$, $s \in \mathbb{R}_+$, $\beta \in [0, 1]$, and $\sigma \in [0, 1]$,*

$$v(b, s) - c_B(b, \beta) - c_S(s, \sigma) < 0.$$

A similar statement, with an analogous \bar{s} , applies to sellers.

⁶The asymmetry in this assumption—it requires a strict inequality on the cross partial of h_B , but only a weak inequality on that of h_S —reflects our convention that sellers set prices. If the derivative for buyers is zero, then every buyer will attempt to purchase from the same seller, destroying all hope of sorting buyers. However, Section 4.2 shows that there exist efficient uniform-price equilibrium outcomes if and only if seller premuneration values do not depend on buyer attribute choices, making it important to include the weak inequality for the seller. As becomes clear, this zero second derivative for the seller poses no difficulty. The asymmetry that appears in the first part of Assumption 2 similarly arises out of the convention that sellers set prices, though this part of the assumption is more technical in nature, allowing us to rule out some troublesome boundary cases.

3.2 Feasible outcomes

We next define feasible matchings between buyers and sellers. We denote by $\mathbf{b}: [0, 1] \rightarrow [0, \bar{b}]$ and $\mathbf{s}: [0, 1] \rightarrow [0, \bar{s}]$ the Lebesgue-measurable functions that describe the attributes chosen by buyers and sellers.

The closures of the sets of attributes chosen by buyers and sellers, respectively, are denoted by $\mathcal{B} \equiv \text{cl}(\mathbf{b}([0, 1]))$ and $\mathcal{S} \equiv \text{cl}(\mathbf{s}([0, 1]))$. We refer to \mathcal{B} and \mathcal{S} as the set of *marketed attributes*. Let $\lambda_{\mathcal{B}}$ and $\lambda_{\mathcal{S}}$ be the measures induced on \mathcal{B} and \mathcal{S} by the agents' attribute choices: for Borel sets $\mathcal{B}' \subset \mathcal{B}$ and $\mathcal{S}' \subset \mathcal{S}$,

$$\lambda_{\mathcal{B}}(\mathcal{B}') = \lambda\{\beta \in [0, 1] : \mathbf{b}(\beta) \in \mathcal{B}'\}$$

and

$$\lambda_{\mathcal{S}}(\mathcal{S}') = \lambda\{\sigma \in [0, 1] : \mathbf{s}(\sigma) \in \mathcal{S}'\},$$

where λ is Lebesgue measure.

We simplify the analysis by focussing on equilibria in which attribute-choice functions are strictly increasing when positive (i.e., $\mathbf{b}(\beta) > 0$ and $\beta' > \beta$ imply $\mathbf{b}(\beta') > \mathbf{b}(\beta)$, and similarly for \mathbf{s}), and assign equal masses of buyers and sellers to zero attribute choices. We show that equilibria exist with attribute-choice functions that satisfy these restrictions. More general feasible matchings can be defined, but at the cost of considerable technical complication.

The measures of buyers and of sellers who choose the zero attribute are denoted by $\underline{\beta} \equiv \sup\{\beta : \mathbf{b}(\beta) = 0\}$ and $\underline{\sigma} \equiv \sup\{\sigma : \mathbf{s}(\sigma) = 0\}$.

DEFINITION 3.1. Suppose \mathbf{b} and \mathbf{s} are strictly increasing when positive and that $\underline{\sigma} = \underline{\beta}$. A *feasible matching* is a pair of measure-preserving functions $\tilde{b}: (\mathcal{S}, \lambda_{\mathcal{S}}) \rightarrow (\mathcal{B}, \lambda_{\mathcal{B}})$ and $\tilde{s}: (\mathcal{B}, \lambda_{\mathcal{B}}) \rightarrow (\mathcal{S}, \lambda_{\mathcal{S}})$ that satisfy

$$\tilde{s}(\tilde{b}(s)) = s \quad \text{for all } s \in \mathbf{s}((\underline{\sigma}, 1]) \tag{1}$$

and

$$\tilde{b}(\tilde{s}(b)) = b \quad \text{for all } b \in \mathbf{b}((\underline{\beta}, 1]). \tag{2}$$

Given a feasible matching (\tilde{b}, \tilde{s}) , the buyer attribute $\tilde{b}(s)$ is matched to a seller with attribute s , and the seller attribute $\tilde{s}(b)$ is matched to a buyer with attribute b . Observe that (1) and (2) imply that \tilde{s} is one-to-one on $\mathbf{b}((\underline{\beta}, 1])$ and \tilde{b} is one-to-one on $\mathbf{s}((\underline{\sigma}, 1])$. The measure-preserving requirement on \tilde{b} ensures that the measure of any set of sellers is equal to the measure of the set of buyers with whom they are matched, i.e., $\lambda_{\mathcal{B}}(\tilde{b}(S')) = \lambda_{\mathcal{S}}(S')$ for all Borel $S' \subset \mathcal{S}$ (and similarly for \tilde{s}).

We have simplified the analysis by defining the matching functions \tilde{b} and \tilde{s} on the closures \mathcal{S} and \mathcal{B} of the sets of chosen attributes. In many cases of interest, efficient

attribute-choice functions are discontinuous (see Cole et al. 2001, Section 2 for an example of discontinuous attribute-choice functions with personalized pricing (cf. Section 6.1)). Since the sets \mathcal{B} and \mathcal{S} are the closures of the sets of attribute choices, a seller σ (with attribute choice $s(\sigma)$) may be matched with a buyer attribute choice b that is not chosen by any buyer. We interpret such a seller as matching with a buyer whose attribute choice is arbitrarily close to b , while retaining the convenience of saying that $s(\sigma)$ matches with b . Defining feasible matchings on either the agents directly or on the sets of attributes (rather than their closures) avoids this interpretation, at the cost of requiring the equivalent but more complicated formulation used in Cole et al. (2001).

DEFINITION 3.2. A *feasible outcome* $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ is a pair of attribute-choice functions \mathbf{b} and \mathbf{s} that are strictly increasing when positive and satisfy $\underline{\sigma} = \underline{\beta}$, along with a feasible matching (\tilde{b}, \tilde{s}) .

3.3 Uniform pricing

Sellers post prices that depend on their own attribute choices, but not the attributes of buyers. We describe these prices by a *uniform-price function* $p_U : \mathcal{S} \rightarrow \mathbb{R}$.

Fix a uniform-price function p_U . The payoffs to a buyer β who chooses $b \in \mathcal{B}$ and a seller σ who chooses $s \in \mathcal{S}$ and matching according to the feasible matching (\tilde{b}, \tilde{s}) are

$$\Pi_B(b, \beta) \equiv h_B(b, \tilde{s}(b)) - p_U(\tilde{s}(b)) - c_B(b, \beta)$$

and

$$\Pi_S(s, \sigma) \equiv h_S(\tilde{b}(s), s) + p_U(s) - c_S(s, \sigma).$$

Under uniform pricing, sellers cannot condition on buyer attributes. Consequently, sellers choose only their own attributes. Buyers, alternatively, choose attributes and can choose any marketed seller attribute regardless of their own attribute choice. These choices should maximize payoffs. A buyer β *optimizes (at \mathbf{b}) given p_U* if

$$\Pi_B(\mathbf{b}(\beta), \beta) = \max_{(b,s) \in \mathbb{R}_+ \times \mathcal{S}} h_B(b, s) - p_U(s) - c_B(b, \beta) \geq 0. \tag{3}$$

Similarly, a seller σ *optimizes (at \mathbf{s}) given p_U* if

$$\Pi_S(\mathbf{s}(\sigma), \sigma) = \max_{s \in \mathcal{S}} h_S(\tilde{b}(s), s) + p_U(s) - c_S(s, \sigma) \geq 0. \tag{4}$$

Note that our notion of optimization (at \mathbf{b} or \mathbf{s}) given p_U imposes individual rationality on buyers and sellers, ensuring that we are not inappropriately forcing our agents to participate in the market.

3.4 Equilibrium

The uniform-price function p_U determines the payoff to a buyer for *any* attribute he chooses and any seller he matches with, since prices do not depend on the buyers' attribute choices. It also determines the payoff to any seller who chooses a marketed attribute (i.e., $s \in \mathcal{S}$). It does not do so for nonmarketed seller attributes, since such attributes are not priced by the function p_U . We think of a seller who chooses a nonmarketed attribute as naming the price at which he is willing to trade and then trading with one of the buyers willing to trade at that price, if there are any. However, this attribute and price combination potentially attracts many buyer attributes, all of which are indistinguishable to the seller. The following definition requires that the seller's deviation to (s', p) with $s' \notin \mathcal{S}$ be profitable *irrespective* of what type of buyer is attracted.⁷

DEFINITION 3.3. Given $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_U)$, there is a *profitable seller deviation* if there exists σ for which there is an unmarketed attribute choice $s' \notin \mathcal{S}$, a price $p \in \mathbb{R}$, and at least one buyer $b' \in \mathcal{B}$ such that

$$h_B(b', s') - p > h_B(b', \tilde{s}(b')) - p_U(\tilde{s}(b')),$$

and for any such b' ,

$$h_S(b', s') + p - c_S(s', \sigma) > \Pi_S(\mathbf{s}(\sigma), \sigma).$$

REMARK 3.1 (Profitable deviations). A seller is defined to have a profitable deviation under uniform pricing only if he is better off when matched with *any* buyer who is attracted to the deviation. Why make sellers so pessimistic? One could alternatively think of requiring only that the seller be better off given a random draw from the set of attracted buyers. Though the details of the calculations (and the existence proof) would differ considerably, the qualitative forces behind our results would remain. In particular, the essence of uniform pricing is that the seller cannot stipulate which buyers he is and is not willing to trade with. This inability affects the seller most starkly when we assume the seller draws the worst buyer from the set of willing buyers, but the effects remain as long as the seller cannot select the *best* buyer.⁸

Adopting the pessimistic formulation that seller deviations must be profitable when matched with the worst willing buyer makes seller deviations less attractive and, hence, enlarges the set of uniform-price equilibria. Our key results (Propositions 4.1(ii) and 4.2), establishing conditions under which uniform price equilibria are efficient, are rendered more powerful by such a permissive definition of equilibrium.

⁷We could extend Definition 3.3 to cover deviations to any seller attribute (rather than simply unmarketed seller attributes) as well as deviations to other prices at the seller's current attribute. Appendix B shows that if buyers optimize given p_U and sellers have no profitable deviations in this extended sense, then sellers must also be optimizing given p_U .

⁸Peters (2010) examines the implications of allowing the seller to select the best buyer in a competitive search model.

DEFINITION 3.4. A feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ and a uniform-price function $p_U : S \rightarrow \mathbb{R}$ constitute a *uniform-price equilibrium* if all agents optimize given p_U and there are no profitable seller deviations.

REMARK 3.2. The definition of a uniform-price equilibrium has two aspects. First, our formulation of equilibrium in the post-investment matching market is analogous to the definition of a Walrasian equilibrium: we require that agents optimize and that markets clear. In particular, we require that prices be such that when each buyer chooses a payoff-maximizing seller given the feasible sets defined by those prices, the result is a feasible match. As is the case in Walrasian equilibrium, we do not explain the origin of such prices; neither do we explain what happens when some buyer chooses to match with a seller who is not payoff maximizing or if prices lead to optimal choices that do not clear the market. For example, suppose that given a candidate equilibrium outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$, buyer β could get a higher payoff by deviating and choosing seller attribute s' rather than the prescribed seller attribute $\tilde{s}(b(\beta))$. This presumably results in there being two buyers matched with seller s' . Rather than probing the details of what happens in such a case, we simply conclude that the candidate equilibrium is indeed *not* an equilibrium. If we were to model this as a well defined game, we would have to specify which buyer ends up matched with the seller. One could provide such specificity, but doing so gives rise to a number of arbitrary choices and technical issues that obscure the underlying economics.

Having this specification of equilibrium at the matching stage, our treatment of first-stage investment decisions is analogous to requiring a subgame perfect equilibrium. Investment choices must be optimal, given the correctly anticipated second-stage implications of these choices.

REMARK 3.3 (Complete pricing). By altering Definition 3.4 to require p_U to have domain $[0, \bar{v}]$, thereby setting a price for *every* seller attribute (whether marketed or not), and expanding to $[0, \bar{v}]$ the set of seller attribute choices over which the buyer optimizes, we obtain a *complete* uniform-price equilibrium. Notice, however, that the matching function is still restricted to marketed attributes and, hence, the seller's payoff when choosing an unmarketed attribute is still defined separately as in Definition 3.3.

REMARK 3.4 (Hedonic pricing). In a uniform-price equilibrium, each buyer faces prices over seller attributes, and so it is tempting to interpret the prices as hedonic prices. However, since sellers care about buyer attributes and the prices are not a function of these attributes, all payoff-relevant characteristics are not priced.⁹ Accordingly, a uniform-price equilibrium is not an equilibrium in hedonic prices.

REMARK 3.5. We consider only equilibria in pure strategies. Bulow and Levin (2006) examine a model in which sellers (firms, in their model) can observe buyers' (workers') attribute choices, are legally constrained to offer the same price to all buyers, buyers have

⁹Of course, in equilibrium, each seller may be able to infer the buyer attribute that is matched with each marketed attribute at the equilibrium price.

zero remuneration values (ensuring that, in terms of our model, equilibrium prices are negative), and sellers can condition their choice among multiple buyers on the buyers' attributes. If the agents in their model were competitive, in the sense of there being a continuum of agents on each side of the market, then the result would be a uniform-rationing version of personalized-pricing equilibrium (see [Section E.4](#)), giving rise to an efficient outcome. Because they work with a finite number of buyers and sellers, their equilibria are in mixed strategies.

3.5 Example: A uniform-price equilibrium

Under uniform pricing, buyer β faces a uniform-price schedule p_U and chooses a buyer attribute b and a seller attribute $s \in \mathcal{S}$ to solve

$$\max_{b,s} \theta bs - p_U(s) - \frac{b^3}{3\beta}.$$

When choosing an attribute s , the seller is selected by a buyer with attribute $b = \tilde{b}(s)$ and receives prices p_U . The seller σ thus solves

$$\max_s (1 - \theta)\tilde{b}(s)s + p_U(s) - \frac{s^3}{3\sigma}.$$

The uniform-price equilibrium is given by the collection (the derivation appears in [Appendix A.2](#))

$$\mathbf{b}(\beta) = \theta^{2/3}(2 - \theta)^{1/3}\beta \tag{5}$$

$$\mathbf{s}(\sigma) = \theta^{1/3}(2 - \theta)^{2/3}\sigma \tag{6}$$

$$p_U(s) = \frac{\theta}{2} \left(\frac{\theta}{2 - \theta} \right)^{1/3} s^2 \tag{7}$$

and

$$\tilde{b}(s) = \left(\frac{\theta}{2 - \theta} \right)^{1/3} s. \tag{8}$$

When $\theta = 1$, this uniform-price equilibrium gives the efficient outcome calculated in [Section 2.2](#). In this case, the restriction to uniform pricing imposes no efficiency costs, and giving sellers the ability to condition prices on buyer attributes has no effect on behavior or payoffs. Conversely, when $\theta < 1$, the uniform-price equilibrium is inefficient, in that the generated surplus of almost all matched pairs is not maximized. We discuss this inefficiency further in [Section 4.3](#).

Note that the equilibrium is not unique. In particular, all buyers and sellers choosing the zero attribute is also an equilibrium outcome.

4. EFFICIENCY

When are uniform-price equilibrium outcomes efficient? Efficiency fails (i.e., total surplus is not maximized) when either the wrong agents are matched or the wrong attributes are chosen by matched agents.

4.1 Efficient matching

Efficiency requires that the second-stage matching be positively assortative in attributes. The supermodularity assumptions on premuneration values guarantee this positive assortativity in equilibrium.

LEMMA 4.1. *In any uniform-price equilibrium $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_U)$, \tilde{b} and \tilde{s} are strictly increasing for strictly positive attributes, and so the matching is positively assortative in attributes.*

PROOF. Suppose \tilde{b} is not strictly increasing. Since \tilde{b} is one-to-one on $\mathbf{s}((\underline{\sigma}, 1])$ (see Definition 3.1 and its following comment), there exists $0 < s_1 < s_2$ with $b_1 \equiv \tilde{b}(s_1) > \tilde{b}(s_2) \equiv b_2$. Adding

$$h_B(b_1, s_1) - p_U(s_1) \geq h_B(b_1, s_2) - p_U(s_2)$$

and

$$h_B(b_2, s_2) - p_U(s_2) \geq h_B(b_2, s_1) - p_U(s_1)$$

gives

$$h_B(b_1, s_1) + h_B(b_2, s_2) \geq h_B(b_1, s_2) + h_B(b_2, s_1),$$

contradicting the strict supermodularity of h_B .

Equation (2) then implies that \tilde{s} is strictly increasing. □

4.2 Efficient investments

Efficiency at the investment stage requires that the attribute choice functions (\mathbf{b}, \mathbf{s}) satisfy

$$(\mathbf{b}(\phi), \mathbf{s}(\phi)) \in \arg \max_{b, s \in \mathbb{R}_+} W(b, s, \phi),$$

where

$$W(b, s, \phi) \equiv v(b, s) - c_B(b, \phi) - c_S(s, \phi).$$

This efficiency is not guaranteed. We begin with some intuition, which is appropriate when equilibrium is characterized by first-order conditions. Fix a uniform-price equilibrium. By standard incentive compatibility arguments, the uniform-price function is differentiable. The first-order conditions for the buyer's choice of attribute b and matching attribute choice s in a uniform-price equilibrium are

$$0 = \frac{dh_B(b, s)}{db} - \frac{dc_B(b, \beta)}{db} \tag{9}$$

and

$$0 = \frac{dh_B(b, s)}{ds} - \frac{dp_U(s)}{ds}, \tag{10}$$

while the seller’s first-order condition for choosing s is (assuming \tilde{b} is differentiable)

$$0 = \frac{dh_S(\tilde{b}(s), s)}{db} \frac{d\tilde{b}(s)}{ds} + \frac{dh_S(\tilde{b}(s), s)}{ds} + \frac{dp_U(s)}{ds} - \frac{dc_S(s, \sigma)}{ds}. \tag{11}$$

Using (10) to eliminate $dp_U(s)/ds$ in (11) and then using the identity $v(b, s) = h_B(b, s) + h_S(b, s)$ in (9) and (11), these three first-order conditions can be reduced to

$$0 = \frac{dv(b, s)}{db} - \frac{dh_S(b, s)}{db} - \frac{dc_B(b, \beta)}{db}$$

and

$$0 = \frac{dh_S(b, s)}{db} \frac{d\tilde{b}(s)}{ds} + \frac{dv(b, s)}{ds} - \frac{dc_S(s, \sigma)}{ds}.$$

Efficiency requires that any matched buyer and seller maximize the difference between the surplus they generate and their investment costs, giving rise to the first-order conditions

$$\begin{aligned} 0 &= \frac{dv(b, s)}{db} - \frac{dc_B(b, \beta)}{db} \\ 0 &= \frac{dv(b, s)}{ds} - \frac{dc_S(s, \sigma)}{ds}. \end{aligned} \tag{12}$$

Comparing these, it is immediate that the solution to the first-order conditions for an efficient allocation is a solution for the first-order conditions for the uniform-price equilibrium if $dh_S(b, s)/db = 0$, that is, if each seller’s premuneration value is independent of the attribute choice of the buyer with whom the seller is matched. Moreover, the same argument shows that when seller premuneration values are independent of buyer attributes, every uniform-price equilibrium is constrained efficient, in that no efficiency gains can be achieved without a simultaneous deviation to unmarketed buyer and seller attributes. In other words, inefficiency arises only out of coordination failure.

These arguments are summarized in the following proposition. The proof follows the preceding intuition (though it requires no differentiability assumptions) and so is relegated to [Appendix C](#).

PROPOSITION 4.1. *Suppose the sellers’ premuneration values do not depend on the buyer’s attribute, and the most efficient buyer and seller are capable of producing a surplus, i.e., there exists $(b, s) \in (0, \bar{b}] \times (0, \bar{s}]$ with $h_B(b, s) + h_S(b, s) - c_B(b, 1) - c_S(s, 1) > 0$. Then,*

- (i) *there exist efficient uniform-price equilibria and*
- (ii) *every uniform-price equilibrium outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ is constrained efficient,*

$$\begin{aligned} W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi) &= \max_{\substack{b \in \mathbf{b}([0, 1]), \\ s \in \mathbb{R}_+}} W(b, s, \phi) \\ &= \max_{\substack{b \in \mathbb{R}_+, \\ s \in \mathbf{s}([0, 1])}} W(b, s, \phi), \end{aligned}$$

and every complete uniform-price equilibrium is efficient.

The constancy of $h_S(b, s)$ in b is also essentially necessary for personalized-price equilibria to be achieved via uniform pricing. The “essentially” here is that this constancy need not hold for pairs (b, s) that are not matched in equilibrium.¹⁰

PROPOSITION 4.2. *Suppose the efficient outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ can be supported as a uniform-price equilibrium outcome. Then for all $s \in S$,*

$$\frac{dh_S(\tilde{b}(s), s)}{db} = 0.$$

PROOF. It follows from (9) and (12) (again, without any differentiability assumptions beyond those placed on the primitives of the model in Assumptions 1 and 3) that if $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$ is an efficient outcome that can be supported by uniform prices, then

$$\frac{dh_B(\tilde{b}(s), s)}{db} = \frac{dv(\tilde{b}(s), s)}{db},$$

implying $dh_S(\tilde{b}(s), s)/db = 0$. □

4.3 Example: Efficiency

Suppose first that sellers own none of the surplus (i.e., $\theta = 1$ and, hence, $h_S(b, s) = 0$ and $dh_S(b, s)/db = 0$). In this case, the uniform-price equilibrium of Section 3.5 results in an efficient outcome. Consequently, no seller would gain by personalizing his price even if he could and the ability to personalize prices is irrelevant.¹¹

Now consider differing values of θ and, hence, different arrangements of premuneration values, assuming that the total surplus is still given by bs and the corresponding uniform-price equilibria. In the efficient outcome, the buyer’s equilibrium attribute choice is $\mathbf{b}(\beta) = \beta$. Buyer attributes in the uniform-price equilibria are again a linear function of the buyer’s index, with slope $\theta^{2/3}(2 - \theta)^{1/3}$. This slope is below 1 for all $\theta < 1$, that is, buyers’ investments are inefficiently low. The inability to personalize prices prevents sellers from offering buyers lower prices in return for higher buyer attributes. As a result, the return on buyers’ investments under uniform pricing is less than the social return, and buyers choose lower attributes than would be efficient.

The magnitude of the inefficiency decreases as θ increases. The smaller the buyers’ premuneration values, the larger the extent to which their attribute choices fall short of efficient levels.

Sellers’ attribute choices in the uniform-price equilibrium are similarly a linear function of index, with slope $\theta^{1/3}(2 - \theta)^{2/3}$. Since this exceeds the buyer coefficient, buyers choose smaller attributes than sellers, with buyers of attribute choice level b matching with values $s > b$.

¹⁰Analogously, the single-crossing condition is essentially necessary for a separating equilibrium in a signaling model.

¹¹For every $\theta \in (0, 1]$, there is a personalized-price equilibrium (cf. Section 6.1) that gives the same (efficient) investments and payoffs that appear under the uniform-price equilibrium of Section 3.5 when $\theta = 1$.




FIGURE 1. Uniform-price equilibrium attribute choices as a function of θ , the buyers' premuneration-value share of the surplus. Both seller and buyer attribute-choice functions are linear, with the vertical axis measuring the coefficient. The choices are efficient (the coefficients equal 1) when $\theta = 1$.

Perhaps surprisingly, the sellers' investment behavior is not monotonic in θ , as illustrated in Figure 1. For low levels of θ —when the sellers' share of the surplus is near 1—sellers invest very little. This is to be expected since the value of their investment depends on buyers' investment, which is low in this case. The slope of the seller attribute-choice function initially increases in θ , a consequence of the increase in buyers' attribute choices and the increase in the price a seller attribute fetches. When $\theta \approx 0.38$, sellers make precisely the attribute choices under uniform pricing that they would in the efficient outcome. The equilibrium is still inefficient, however, as buyers invest too little. For larger values of θ , uniform pricing leads sellers to invest *more* than they do in the efficient outcome.

To understand this seller behavior, notice that a seller would like to screen the buyers to whom he sells, but the inability to personalize prices precludes doing so directly. The key to screening buyers is that high-attribute buyers have a higher willingness to pay for high-attribute sellers than do low-attribute buyers. Sellers then have an incentive to choose higher attributes (than the efficient level) and charge higher prices. As θ increases, buyer attribute choices increase, making screening all the more valuable to sellers. As a result, seller attribute choices continue to increase above their efficient levels as θ increases above 0.38.

Once θ reaches $2/3$, sellers' attribute choices no longer increase (though seller attribute choices remain above efficient levels). Buyers' attribute choices continue to increase as θ increases, but the decreasing share that sellers receive makes screening less valuable and, hence, makes investment less attractive.

Sellers' incentives to screen buyers lead not only to attribute choices that exceed the efficient investments, but also to attribute choices that are inefficiently high *given* the buyers' (inefficiently low) attribute choices, for all $\theta < 1$. In equilibrium, seller σ is

matched with buyer $\beta = \sigma$, who makes attribute choice $\theta^{2/3}(2 - \theta)^{1/3}\sigma$. The net surplus (ignoring the cost of b) from a match of seller σ with such a buyer is

$$s\theta^{2/3}(2 - \theta)^{1/3}\sigma - \frac{s^3}{3\sigma}.$$

The seller attribute that maximizes this surplus is

$$s(\sigma) = \sigma\theta^{1/3}(2 - \theta)^{1/6},$$

which is smaller than the seller's equilibrium attribute choice of

$$\sigma\theta^{1/3}(2 - \theta)^{1/3}.$$

5. EXISTENCE OF EQUILIBRIUM

Appendix D establishes the existence of uniform-price equilibria by showing the existence of complete uniform-price equilibria (see Remark 3.3).

PROPOSITION 5.1. *If there exists $(b, s) \in (0, \bar{b}] \times (0, \bar{s}]$ with*

$$h_B(b, s) + h_S(0, s) - c_B(b, 1) - c_S(s, 1) > 0, \tag{13}$$

then there exists a complete uniform-price equilibrium in which some buyers and some sellers make strictly positive attribute choices.

Moreover, if for all $\phi \in (0, 1]$, there exists $(b, s) \in (0, \bar{b}] \times (0, \bar{s}]$,

$$h_B(b, s) + h_S(0, s) - c_B(b, \phi) - c_S(s, \phi) > 0, \tag{14}$$

then there exists a complete uniform-price equilibrium with $\mathbf{b}(\beta), \mathbf{s}(\sigma) > 0$ for $\beta, \sigma \in (0, 1]$.

In general, condition (13) is stronger than the requirement that there be a positive surplus for the most efficient match (though (13) is implied by that requirement if $h_S(b, s)$ is independent of b , the condition of Proposition 4.1). Uniform-pricing equilibria are inefficient when $h_S(b, s)$ depends on b , and if this dependence is too extreme, (13) may fail and there may be no investment on either side.

Two significant complications must be confronted in the proof of existence of uniform-price equilibria: Equilibrium attribute-choice functions may be discontinuous and we must preclude profitable deviations to attributes not in the market. These complications preclude the direct application of a fixed point theorem. We proceed indirectly, constructing a simultaneous-move three-player game whose equilibria capture the relevant behavior of uniform-price equilibria. The players include a buyer, whose payoff corresponds to the total buyer payoff in our model, a seller, whose payoff is analogous but who does not set prices, and a price-setter, who is penalized for market imbalance. In constructing this game, we define seller payoffs in a manner that incorporates the pessimism inherent in our definition of uniform-price equilibrium. Glicksberg's

fixed point theorem establishes the existence of Nash equilibria in the three-player game when strategies are constrained to be Lipschitz continuous. We then examine the limit as this constraint is removed, showing that the result corresponds to a uniform-price equilibrium of the underlying economy.

6. DISCUSSION

6.1 Comparison with personalized pricing

6.1.1 Personalized-price equilibrium The obvious point of comparison for a uniform-price equilibrium is a scenario in which prices can be conditioned on both buyer and seller characteristics. In such a scenario, there is a *personalized-price function* $p_P: \mathcal{B} \times \mathcal{S} \rightarrow \mathbb{R}$, where $p_P(b, s)$ is the (possibly negative) price that a seller with attribute choice $s \in \mathcal{S}$ receives when selling to a buyer with attribute choice $b \in \mathcal{B}$. This gives rise to a *personalized-price equilibrium*, analogous to that of a uniform-price equilibrium except that sellers can charge different prices to different buyers, and the possibility of a profitable deviation to an unmarketed attribute is now open to buyers as well as sellers. [Appendix E](#) develops the details, establishing the following results.

- Personalized-price equilibria exist, and are, modulo some technical differences in the specification, equivalent to the *ex post contracting equilibria* of Cole et al. (2001).
- Personalized-price equilibria are constrained efficient, in the sense that there is no alternative, Pareto superior allocation that restricts buyers and sellers to choosing attributes marketed in the equilibrium. Personalized-price equilibria may exhibit “coordination failure” inefficiencies, in which mutual gains could be realized if buyers and sellers both bring currently unmarketed attributes to the market. There exists an efficient personalized-price equilibrium.
- Premuneration values are irrelevant for personalized-price equilibria. For a given specification of premuneration values and attendant personalized-price equilibrium, any other specification of premuneration values admits a personalized-price equilibrium whose outcome, including investments, matching function, and payoffs, duplicates that of the original equilibrium.
- Under the conditions of [Proposition 4.1](#), uniform and personalized-price equilibria coincide. In this case, the ability to personalize prices is irrelevant. Personalization brings sellers no advantage, and even the slightest cost of personalization suffices to ensure that we see uniform pricing.

The essence of our results, culminating in [Propositions 4.1](#) and [4.2](#), is to establish conditions under which personalization is redundant. If these conditions hold, uniform prices also lead to efficient equilibrium outcomes. If not, uniform prices are inextricably linked to inefficient investments. Under uniform pricing, premuneration values matter.

6.1.2 *Example: Personalized pricing* Returning to our example, suppose that sellers observe buyers' attribute choices and so can personalize their prices. If buyers and sellers optimize given the personalized-price function

$$p_P(b, s) = \frac{1}{2}s^2 - (1 - \theta)bs, \tag{15}$$

the result is a feasible and efficient outcome.¹² In particular, given the pricing function (15), buyer β chooses the attribute $b = \mathbf{b}(\beta) = \beta$ and chooses to match with seller attribute $s = \mathbf{s}(\beta)$. The seller chooses attribute $s = \mathbf{s}(\sigma) = \sigma$. The resulting matching of buyers and sellers clears the seller attribute market (in that the distributions of demanded and supplied seller attributes agree), and the resulting outcome is efficient. Appendix A.3 contains the details and confirms that this is a personalized-price equilibrium.

If $\theta < 1$ in our example, then all buyers receive lower payoffs under uniform than under personalized prices.¹³ A natural conjecture is that sellers are necessarily disadvantaged by the inability to personalize prices. The seller's equilibrium payoff in the uniform-price equilibrium is given by

$$(1 - \theta)\tilde{b}(s(\sigma))s(\sigma) + p_U(s(\sigma)) - \frac{(s(\sigma))^3}{3\sigma} = \frac{1}{6}\theta(2 - \theta)^2\sigma^2.$$

When $\theta = 1$, this duplicates the payoff from the personalized-price equilibrium. For θ for which sellers' attributes exceed the personalized-price equilibrium level, every seller actually earns a *higher* payoff under the uniform-price equilibrium. This higher payoff results from the higher prices that buyers are willing to pay for the higher attributes chosen by sellers when they cannot personalize prices.

Why do we not see such higher prices under personalized pricing? Suppose that given a uniform-price equilibrium, a single seller has the ability to personalize prices. Such a seller could profitably reduce his attribute choice and the price at which he trades, using personalization to exclude the undesirable buyers that render such a deviation unprofitable under uniform pricing.

6.1.3 *Which prices are personalized?* Personalizing prices requires a seller to set a price for every buyer attribute in the market. However, personalized-price outcomes can be achieved with much simpler pricing schemes. The apparent absence of complicated pricing schemes thus need not signal the absence of personalized pricing.

The critical feature of personalized pricing is the seller's ability to exclude buyers with attribute choices lower than the seller's equilibrium match. In particular, by charging a sufficiently high price to specific buyer attribute choices, a seller can ensure that buyers with those attributes chose not to buy. We denote this sufficiently high price by P .

¹²Note that for any seller attribute s , the price that a seller would receive in a match with a buyer with attribute b is decreasing in b : higher values of b are more valuable and, hence, sellers are willing to charge less for them.

¹³The buyer's payoff under uniform pricing, $\theta\tilde{s}(b(\beta)) - p_U(\tilde{s}(b(\beta))) - (b(\beta))^3/(3\beta) = \frac{1}{6}\theta^2(2 - \theta)\beta^2$, falls short of the buyer's payoff in the personalized-price equilibrium.

A personalized-price function p_P is a *uniform-rationing price* if it has the form

$$p_P(b, s) = \begin{cases} p_{UR}(s) & \forall b \geq \tilde{b}(s) \\ P & \text{otherwise} \end{cases}$$

for some $p_{UR}: \mathcal{S} \rightarrow \mathbb{R}_+$ and $\tilde{b}: \mathcal{S} \rightarrow \mathcal{B}$. Under uniform-rationing pricing, a seller with attribute choice s sets a uniform price $p(s) = p_{UR}(s)$, but then excludes any buyers with $b < \tilde{b}(s)$.

Appendix E.4 provides the straightforward argument that any personalized-price equilibrium outcome can be supported by a uniform-rationing price. Hence, personalized pricing may be ubiquitous without one observing complete menus of prices. Whenever we observe sellers rejecting some buyers—colleges denying some applicants or firms rejecting some workers as unqualified—we may be observing personalized-pricing outcomes.

6.2 Information

Suppose sellers are constrained to set uniform prices because buyers' attributes are not observable, but that buyers can certify these attributes, perhaps by taking exams or completing internships that demonstrate their skills. One might suspect that if the cost to buyers of certifying their attribute is not too high, the uncertainty might "unravel": high-attribute buyers would reveal themselves, making it optimal for the highest-attribute buyers in the remaining pool to reveal themselves, and so on until all buyers' attributes are known.¹⁴ In addition, it seems that this cascading information revelation must make at least lower-ranked buyers worse off, if not all buyers. Indeed, to avoid such unraveling, Harvard Business School students have successfully lobbied for policies that prohibit students' divulging their grades to potential employers, while the Wharton student government adopted a policy banning the release of grades.¹⁵

In contrast, in our example, *all* buyers may be worse off when information about their attributes is suppressed than when it is known. This result holds no matter what (nonzero) share the buyers own of the surplus, and holds for all buyers. It is the distorted incentives to invest that ensure even the lowest attribute buyers would be made worse off if buyer attribute information were suppressed.

6.3 Who should set prices?

Suppose we could design the informational or legal context so that one side of the market can set prices, but cannot observe the characteristics on the other side of the market. Which should we choose? We return to our example. When $\theta = 0$, so the seller owns all of the surplus, the equilibrium collapses into the trivial equilibrium in which no surplus is generated. In this case, a buyer's payoff is solely the price p_U , which has to be negative so as to bring buyers into the market, and buyers will choose the seller who posts

¹⁴See Grossman (1981), Milgrom (1981), or Okuno-Fujiwara et al. (1990) for analyses of such unraveling.

¹⁵Ostrovsky and Schwarz (2010) investigate the optimal amount of information to disclose from the students' perspective.

the smallest (“largest negative”) price. Because sellers cannot condition prices on buyer attribute choice, every buyer will choose $b = 0$ in equilibrium. Similarly, when θ is positive but small, the equilibrium is markedly inefficient, featuring tiny attribute choices. This is an indication that the “wrong” side of the market is setting prices, that is, the side setting prices owns little of the surplus. Suppose personalization by a price-setter is precluded for some reason other than informational asymmetries (such as legal restrictions or transaction costs), but that an alternative market design would allow buyers to post uniform prices (i.e., prices that depend only on buyer attributes). While it is more efficient for sellers to be the price-setters for $\theta > \frac{1}{2}$, it would be more efficient to have buyers post prices when $\theta < \frac{1}{2}$.

6.4 Overinvestment or underinvestment?

The inefficiencies that arise in holdup problems are well understood. The inefficiencies that arise under uniform pricing are qualitatively different, as is easily seen from the overinvestment by sellers in our example for some values of θ . The inefficiencies in uniform-price equilibria in our model stem from sellers’ use of attribute choice as an instrument to screen buyers, in addition to price and the response of buyers.

Our example shows that these inefficiencies may lead to either over- or underinvestment. To provide insight into the nature of the forces involved, it is useful to analyze why the outcome of a personalized-price equilibrium $(\mathbf{b}^P, \mathbf{s}^P, \tilde{b}^P, \tilde{s}^P, p_P)$ is not a uniform-price equilibrium outcome. Consider the outcome $(\mathbf{b}^P \mathbf{s}^P, \tilde{b}^P, \tilde{s}^P)$ with uniform price $p_U(s) = p_P(\tilde{b}(s), s)$. With this uniform price, buyers who previously matched with low-attribute sellers find sellers with higher attributes more attractive, since the uniform price does not penalize low-attribute buyers. This suggests that in a uniform-price equilibrium, a seller could discourage low-attribute buyers by raising his price, and to avoid losing the high-attribute buyers, also raising his attribute. The supermodularity in premuneration values ensures that it is possible to screen buyers in this way.

However, simply altering the seller attributes and p_U is not sufficient, in general, to obtain a uniform-price equilibrium. There are two distinct issues. First, in a personalized-price equilibrium, from the envelope theorem, the impact on a buyer β of a marginal deviation from $\mathbf{b}^P(\beta)$ is given by

$$\frac{\partial h_B(b, s)}{\partial b} - \frac{\partial p_P(b, s)}{\partial b} - \frac{\partial c_B(b, \beta)}{\partial b},$$

evaluated at $s = \tilde{s}(b)$. In contrast, in a uniform-price equilibrium, the second term $(\partial p^P / \partial b)$ is absent. Second, the seller attributes (and prices) are different than in the personalized-price equilibrium. Either effect may dominate. In our running example, the buyers underinvest, and for θ not too small, the sellers overinvest.

More specific results can be obtained with the help of more structure. [Mailath et al. \(2012\)](#) examine a parameterized family of matching economies, with a focus on the circumstances in which sellers prefer to have larger or smaller premuneration values. An initial expectation is that sellers always prefer to own a larger share of the surplus. However, sellers profit not only from owning a share of the surplus, but also from the prices

paid by buyers. A smaller seller (and hence larger buyer) premuneration value makes matches more valuable to buyers, enhancing the prices they are willing to pay. Mailath et al. (2012) show that if (and only if) buyers are perfectly homogeneous, sellers prefer to own none of the surplus. Homogeneous buyers compete fiercely for desirable sellers, and sellers are only too happy to exacerbate the price-enhancing effects of such competition by ceding ownership to buyers. As buyers become more heterogeneous, they also become less competitive, and sellers would prefer to own a larger and larger share of the surplus (though typically not all of it).

6.5 Premuneration values

Our main result is that under uniform pricing, the decomposition of the total surplus of a match into the buyer and seller premuneration values affects the efficiency of pre-match investments. Appropriately specified premuneration values can allow us to avoid either the cost (or impossibility) of personalizing prices or the inefficiencies of uniform pricing.

Premuneration values can sometimes be rearranged by appropriate legal and institutional innovations. The match of researchers and universities generates a surplus that includes the value of marketable patents from faculty research. Historically, universities have owned these patents, but another institutional arrangement could grant them to the faculty. The feasibility of such ownership is reflected in the decisions of many universities to unilaterally grant professors shares in the revenues from patents stemming from their research. Why are all premuneration values not specified so as to allow efficient uniform-pricing outcomes?

Section 1.4 highlighted moral hazard problems. Monitoring considerations may also play a role. Consider a collection of heterogeneous and risk averse agents who are to be matched with risk neutral principals. One could ensure that the principal's premuneration values are independent of agent characteristics by assigning ownership of the technology to the agents. Uniform pricing per se would then impose no costs, but the agents would inefficiently bear all of the risk associated with the match, leading to inefficient actions and less valuable matches. We could instead let the principal own some or all of the technology, but now the principal's premuneration value will no longer be independent of the characteristics of the agent with whom he is matched.¹⁶ Finally, legal restrictions may be at work.¹⁷

¹⁶Even before the incentive-design stage, simply measuring and contracting on the relevant variables may pose difficulties. The University of New Mexico sued a former researcher for rights to patents that he applied for 4 years after he had left the university, arguing that the patents stemmed from research that he had done before leaving ("Universities Try to Keep Inventions From Going 'Out the Back Door,'" *Chronicle of Higher Education*, May 17, 2002). In principle, the owner of the rights to a song is entitled to a payment each time the song is played on the radio in a bar or health club, but collection is impractical.

¹⁷For example, Bulow and Levin (2006) note that the National Residency Matching Program that matches medical residents and hospitals constrains hospitals to make the same offers to all residents. They argue that the primary effect is not inefficient matching, but a transfer of surplus to the hospitals (with Niederle and Roth 2003, 2005 offering an alternative view). However, Nicholson (2003) argues that the result is an inefficient allocation of residents to specialties. Medical students who do their residency acquire training that

Putting these considerations together, new monitoring and contracting technologies may be valuable, not only because they can create better incentives within a match, but also because they can create more leeway for designing remuneration values and hence better matching.

APPENDIX A: EXAMPLE, DETAILED CALCULATIONS

A.1 Efficiency

Efficiency requires that for each matched pair β and σ , attribute choices b and s solve

$$\max_{b,s} bs - \frac{b^3}{3\beta} - \frac{s^3}{3\sigma},$$

giving first-order conditions

$$s - \frac{b^2}{\beta} = 0 \quad \text{and} \quad b - \frac{s^2}{\sigma} = 0.$$

Efficiency also requires positive assortative matching in attribute (and so in index, since the cost functions guarantee that attribute choices is increasing in index). We can accordingly solve by setting $\sigma = \beta$, which in turn implies $s = b$, giving the efficient attribute-choice functions

$$\mathbf{b}(\beta) = \beta \quad \text{and} \quad \mathbf{s}(\sigma) = \sigma.$$

A.2 Derivation of (5)–(8)

The buyer chooses an attribute b and a seller attribute s with whom to match so as to solve

$$\max_{b,s} \theta bs + p_U(s) - \frac{b^3}{3\beta}.$$

Assuming p_U is differentiable, the first-order conditions for the buyer are

$$\theta s - \frac{b^2}{\beta} = 0$$

and

$$\theta b - p'_U(s) = 0.$$

dramatically increases their future earnings. Nicholson argues that this part of the surplus from the match (which is owned by the student) is so large in some specialties (such as dermatology, general surgery, orthopedic surgery, and radiology) that if personalized prices were employed, medical students would pay hospitals handsomely for the opportunity to do their residency in these specialities. This is as compared to their stipend, which was \$44,700 in 2007–2008 (Association of American Medical College Survey of Household Stipends, Benefits and Funding, Autumn 2007 Report).

When choosing an attribute s , the seller is selected by a buyer with attribute $b = \tilde{b}(s)$. The seller σ thus solves

$$\max_s (1 - \theta)\tilde{b}(s)s + p_U(s) - \frac{s^3}{3\sigma},$$

implying (assuming \tilde{b} is differentiable) the first-order condition

$$(1 - \theta)[\tilde{b}'(s)s + \tilde{b}(s)] + p'_U(s) - \frac{s^2}{\sigma} = 0.$$

Begin by conjecturing that the equilibrium attribute-choice functions are given by the linear functions

$$\mathbf{b}(\beta) = A\beta \tag{A.1}$$

and

$$\mathbf{s}(\sigma) = B\sigma. \tag{A.2}$$

Then assuming that in equilibrium, a buyer of type β matches with seller of type $\sigma = \beta$, we have $\tilde{b}(s) = As/B$. Using this, rewrite the buyer's second first-order condition as $\theta As/B - p'_U(s) = 0$ and solve for the price function

$$p_U(s) = \frac{\theta A}{2B}s^2.$$

The requirement that low-index traders be willing to participate in the market implies that the constant of integration equals 0. Similarly, rewrite the buyer's first first-order condition as $\theta Bb/A - b^2/\beta = 0$ and solve for b , yielding

$$b = \frac{\theta B}{A}\beta. \tag{A.3}$$

Turning to the seller, write the first-order condition as $2(1 - \theta)As/B + \theta As/B - s^2/\sigma = 0$ and solve for s :

$$s = \frac{(2 - \theta)A}{B}\sigma. \tag{A.4}$$

Combining (A.1) with (A.3) and (A.2) with (A.4) yields $A = \theta^{2/3}(2 - \theta)^{1/3}$ and $B = \theta^{1/3}(2 - \theta)^{2/3}$. It is straightforward to verify that the second-order conditions are satisfied and so the conjecture is verified.

A.3 Personalized prices

Suppose that sellers observe buyers' attribute choices and so can personalize their prices. Consider the candidate personalized-price function

$$p_P(b, s) = \frac{1}{2}s^2 - (1 - \theta)bs. \tag{A.5}$$

Given the pricing function (A.5), buyer β chooses a buyer attribute b and a seller attribute s (i.e., chooses to match with a seller with that attribute) to solve

$$\max_{b,s} \theta bs - \frac{1}{2}s^2 + (1 - \theta)bs - \frac{b^3}{3\beta} = \max_{(b,s)} bs - \frac{1}{2}s^2 - \frac{b^3}{3\beta}.$$

Hence, buyer β chooses the attribute $b = \mathbf{b}(\beta) = \beta$ and chooses to match with seller attribute $s = \mathbf{b}(\beta)$. The implied distribution of demanded seller attributes is uniform on $[0, 1]$.

When choosing an attribute s , the seller is selected by a buyer with attribute $b = \tilde{b}(s) = s$. The seller σ thus solves

$$\max_s (1 - \theta)\tilde{b}(s)s + \frac{1}{2}s^2 - (1 - \theta)\tilde{b}(s)s - \frac{s^3}{3\sigma} = \max_s \frac{1}{2}s^2 - \frac{s^3}{3\sigma},$$

yielding the attribute choice $s = \mathbf{s}(\sigma) = \sigma$. The implied distribution of supplied seller attributes is uniform on $[0, 1]$.

The resulting matching of buyers and sellers clears the seller attribute market (in that the distributions of demanded and supplied seller attributes agree). We thus have a personalized-price equilibrium.

Equilibrium payoffs to the seller and buyer are

$$\frac{1}{2}(\mathbf{s}(\sigma))^2 - \frac{(\mathbf{s}(\sigma))^3}{3\sigma} = \frac{1}{2}\sigma^2 - \frac{\sigma^3}{3\sigma} = \frac{1}{6}\sigma^2$$

and

$$\frac{1}{2}(\mathbf{b}(\beta))^2 - \frac{(\mathbf{b}(\beta))^3}{3\beta} = \frac{1}{2}\beta^2 - \frac{\beta^3}{3\beta} = \frac{1}{6}\beta^2.$$

APPENDIX B: THE ABSENCE OF PROFITABLE DEVIATIONS AND OPTIMIZATION GIVEN p_U

Say that a seller has an extended profitable deviation if either he has a profitable seller deviation in the sense of Definition 3.3 or there exists an attribute $s \in \mathcal{S}$ for which the two inequalities of Definition 3.3 hold. Note that this includes the possibility of charging a different price for $\mathbf{s}(\sigma)$.

LEMMA B.1. *Fix a feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ and a uniform price p_U . Suppose all buyers optimize at \mathbf{b} given p_U (i.e., (3) holds). If seller σ has no extended profitable deviation, then he is optimizing at \mathbf{s} given p_U .*

PROOF. Suppose there exists a seller σ and attribute choice $s' \in \mathcal{S}$ such that

$$\Pi_S(\mathbf{s}(\sigma), \sigma) < \Pi_S(s', \sigma) = h_S(\tilde{b}(s'), s') + p_U(s') - c_S(s', \sigma).$$

Let $\varepsilon = [\Pi_S(s', \sigma) - \Pi_S(\mathbf{s}(\sigma), \sigma)]/4 > 0$. Then there exists $\delta > 0$ such that for all $b \geq \tilde{b}(s') - \delta$,

$$h_S(b, s') + p_U(s') - c_S(s', \sigma) > \Pi_S(\mathbf{s}(\sigma), \sigma) + 3\varepsilon. \tag{B.1}$$

Denote by p'' the price for an attribute s'' that makes the buyer with attribute $\tilde{b}(s')$ indifferent between s' (her equilibrium match) and s'' , i.e.,

$$h_B(\tilde{b}(s'), s'') - p'' = h_B(\tilde{b}(s'), s') - p_U(s').$$

Choose $s'' > s'$ sufficiently close to s' so that

$$|h_S(b, s'') - c_S(s'', \sigma) - h_S(b, s') + c_S(s', \sigma)| < \varepsilon \quad \forall b \in \mathcal{B} \tag{B.2}$$

holds and $|p'' - p_U(s')| < \varepsilon/2$. From **Assumption 1** (Supermodularity),

$$h_B(\tilde{b}(s') - \delta, s'') - p'' < h_B(\tilde{b}(s') - \delta, s') - p_U(s').$$

For $\hat{p} < p''$ sufficiently close to p'' , we have $p'' - \hat{p} > \varepsilon/2$ and

$$h_B(\tilde{b}(s') - \delta, s'') - \hat{p} < h_B(\tilde{b}(s') - \delta, s') - p_U(s').$$

Moreover, the buyer with attribute $\tilde{b}(s')$ receives strictly higher payoff from (s'', \hat{p}) than from $(s', p_U(s'))$. Another application of **Assumption 1** shows that for all $b \leq \tilde{b}(s') - \delta$,

$$h_B(b, s'') - \hat{p} < h_B(b, s') - p_U(s').$$

From (3), for all $b \in \mathcal{B}$,

$$h_B(b, s') - p_U(s') \leq h_B(b, \tilde{s}(b)) - p_U(\tilde{s}(b)),$$

and so no buyer with attribute $b \leq \tilde{b}(s') - \delta$ finds (s'', \hat{p}) attractive. Thus, the pair (s'', \hat{p}) is a profitable deviation for seller σ , since

$$\begin{aligned} h_S(\tilde{b}(s') - \delta, s'') + \hat{p} - c_S(s'', \sigma) &> h_S(\tilde{b}(s') - \delta, s') + \hat{p} - c_S(s', \sigma) - \varepsilon \\ &\geq \Pi_S(\mathbf{s}(\sigma), \sigma) + 3\varepsilon + (\hat{p} - p'') + (p'' - p_U(s')) - \varepsilon \\ &= \Pi_S(\mathbf{s}(\sigma), \sigma) + \varepsilon, \end{aligned}$$

where the first inequality follows from (B.2) and the second follows from (B.1). □

APPENDIX C: PROOF OF PROPOSITION 4.1: EFFICIENT UNIFORM PRICING

Let $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$ be an efficient personalized-price equilibrium (see **Appendix E**; existence is established in **Proposition E.1**). We first show that the price function can be altered so that the seller is indifferent over buyer attributes. In particular, $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, \hat{p}_P)$ is a personalized-price equilibrium, where \hat{p}_P is the personalized-price function given by

$$\hat{p}_P(b, s) = p_P(\tilde{b}(s), s) + h_S(\tilde{b}(s), s) - h_S(b, s) \quad \forall (b, s) \in \mathcal{B} \times \mathcal{S}. \tag{C.1}$$

Moreover, under \hat{p}_P , the seller is indifferent over all marketed buyer attributes.

To verify this, note that seller indifference is immediate, and it is then immediate that the seller is optimizing given \hat{p}_P . We then need show only that the buyer is optimizing

given p_U . Suppose (E.1) fails at some β . Then, for some $(b, s) \in \mathcal{B} \times \mathcal{S}$ and for sufficiently small $\varepsilon > 0$,

$$h_B(b, s) - (\hat{p}_P(b, s) + \varepsilon) - c_B(s, \beta) > \Pi_B(\mathbf{b}(\beta), \beta).$$

Since no buyer has a profitable out-of-market deviation, then

$$h_S(\tilde{b}(s), s) + \hat{p}_P(\tilde{b}(s), s) \geq h_S(b, s) + \hat{p}_P(b, s) + \varepsilon.$$

But this, with (C.1), yields a contradiction.

We now notice that if $h_S(b, s)$ does not depend on b , then neither does \hat{p}_P , implying that $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_U)$ for $p_U(s) = \hat{p}_P(\cdot, s)$ is a uniform-price equilibrium.

The constrained efficiency of uniform-price equilibria and the efficiency of complete uniform-price equilibria, when the seller remuneration values do not depend on buyer attributes, follow from the observation that such equilibria are also personalized-price equilibria, and from Lemmas E.2 and E.3.

APPENDIX D: PROOF OF PROPOSITION 5.1: EXISTENCE OF EQUILIBRIUM

The existence proof is involved and indirect. We want to construct a game Γ whose equilibria induce uniform-price equilibria. However, it is not obvious what this game should be, since the notion of uniform-price equilibrium does not specify who sets prices and does not identify the implications of out-of-equilibrium actions. In effect, constructing such a game requires us to provide the extensive-form details that the notion of uniform-price equilibrium is designed to avoid (cf. Remark 3.2). In addition, were we to formulate such a game, say Γ , the double continuum of players would make its analysis cumbersome.

Our response is to formulate a sequence of approximating games Γ^n . Each game Γ^n has a single buyer and a single seller, each of whom we think of as having types randomly drawn from the unit interval, as well as a price-setter. The price-setter endeavors simply to balance supply and demand, and the existence of this agent allows us to avoid strategic considerations that arise out of price setting. However, complications in defining payoffs when the functions mapping buyer or seller types into attribute choices have jumps (and hence the sets of attributes in the matching market have gaps). We avoid this by imposing the constraint that in game Γ^n , these functions must be continuous, with slopes bounded by $1/n$ and n . This gives us a game that is relatively easy to work with. We then take the limit as $n \rightarrow \infty$ and argue that the limiting strategy profile induces a uniform-price equilibrium. Intuitively, taking this limit allows us to accommodate (in the limit) the possibility of jumps in the attribute-choice functions (precluded in game Γ^n).

Characterizing the limiting strategies is tedious, to some extent because of the potential that jumps appear in the limit (cf. Section D.4). An alternative model would allow agents' realized attributes to be randomly drawn, according to a distribution that shifts in response to their investment (as in Gall et al. 2009). Assuming that the distributions from which realizations are drawn have full support, there would then be no gaps

in the sets of realized attributes. However, in our view, the gaps are an important feature of the model. Such gaps arise in some efficient attribute-choice functions (see Cole et al. 2001 for an example), and we think of such gaps as capturing potentially important economic behavior. In addition, the alternative model is sufficiently far from the personalized-pricing model of Cole et al. (2001) that it is difficult to isolate the effects of uniform pricing.

D.1 Preliminaries

Let $P = \max\{h_B(\bar{b}, \bar{s}), h_S(\bar{b}, \bar{s})\}$. Then P is sufficiently large that no buyer is willing to purchase any seller attribute choice $s \in [0, \bar{s}]$ at a price exceeding P ; neither is any seller willing to sell to a buyer $b \in [0, \bar{b}]$ at price less than $-P$. We can thus limit prices to the interval $[-P, P]$.

Since buyer remuneration values are \mathcal{C}^2 , there is a Lipschitz constant Δ such that for all $\varepsilon > 0, s \in [0, \bar{s} - \varepsilon]$, and $b \in [0, \bar{b}]$, we have $h_B(b, s + \varepsilon) - h_B(b, s) < \Delta\varepsilon$. As a result, given a choice between seller s and seller $s + \varepsilon$ at a price higher by $\Delta\varepsilon$, buyers always choose the former. Equilibrium prices thus never need to increase at a rate faster than Δ .

D.2 The game Γ^n

Each game Γ^n has three players, consisting of a buyer, a seller, and a price-setter.

D.2.1 Strategy spaces We begin by defining the strategy spaces for Γ^n .

The buyer chooses a pair of functions, $(\mathbf{b}, \mathbf{s}_B)$, where $\mathbf{b}: [0, 1] \rightarrow [0, \bar{b}]$ specifies a buyer attribute choice and $\mathbf{s}_B: [0, 1] \rightarrow [0, \bar{s}]$ specifies a seller attribute with which to match, each as a function of the buyer's type. We denote the set of pairs of increasing functions $(\mathbf{b}, \mathbf{s}_B)$ normed by the sum of the L^1 norms on the component functions by Y_B . In Γ^n , the buyer is restricted to the subset of Y_B , denoted by Y_B^n , of functions satisfying

$$(\beta' - \beta)/n \leq \mathbf{b}(\beta') - \mathbf{b}(\beta) \leq n(\beta' - \beta) \quad \forall \beta < \beta' \in [0, 1] \tag{D.1}$$

and

$$(\beta' - \beta)/n \leq \mathbf{s}_B(\beta') - \mathbf{s}_B(\beta) \leq n(\beta' - \beta) \quad \forall \beta < \beta' \in [0, 1]. \tag{D.2}$$

The seller chooses an increasing function \mathbf{s} , where $\mathbf{s}: [0, 1] \rightarrow [0, \bar{s}]$ specifies a seller attribute choice as a function of seller's type. We denote the set of increasing functions \mathbf{s} endowed with the L^1 norm by Y_S . In Γ^n , the seller is restricted to the subset of Y_S , denoted by Y_S^n , of functions satisfying

$$(\sigma' - \sigma)/n \leq \mathbf{s}(\sigma') - \mathbf{s}(\sigma) \leq n(\sigma' - \sigma) \quad \forall \sigma < \sigma' \in [0, 1].$$

The price-setter chooses an increasing function $p_U: [0, \bar{s}] \rightarrow [-P, P]$ satisfying

$$p_U(s') - p_U(s) < 2\Delta(s' - s) \tag{D.3}$$

for all $s < s' \in [0, \bar{s}]$. Denote the set of increasing functions p_U satisfying (D.3), endowed with the sup norm, by Y_P (note that Y_P is not indexed by n). Every function in Y_P is continuous; indeed the collection Y_P is equicontinuous.

The set $Y \equiv Y_B \times Y_S \times Y_P$, when normed by the sum of the three constituent norms, is a compact metric space.¹⁸ It is immediate that $Y^n \equiv Y_B^n \times Y_S^n \times Y_P$ is a closed subset of Y , and so also compact.

D.2.2 Buyer and price-setter payoffs

The buyer The buyer's payoff from $(\mathbf{b}, \mathbf{s}_B) \in Y_B^n$, when the price-setter chooses $p_U \in Y_P$ is

$$\int (h_B(\mathbf{b}(\beta), \mathbf{s}_B(\beta)) - p_U(\mathbf{s}_B(\beta)) - c_B(\mathbf{b}(\beta), \beta)) d\beta. \tag{D.4}$$

Note that the buyer's payoff is independent of seller behavior.

For any \mathbf{s}_B and \mathbf{s} , define

$$F_B(s) \equiv \lambda\{\beta : \mathbf{s}_B(\beta) \leq s\}$$

and

$$F_S(s) \equiv \lambda\{\sigma : \mathbf{s}(\sigma) \leq s\}.$$

The price-setter The price-setter's payoff from $p_U \in Y_P$, when the buyer and seller choose $(\mathbf{b}, \mathbf{s}_B, \mathbf{s}) \in Y_B^n \times Y_S^n$ is given by

$$\int_0^{\bar{s}} p_U(s)(F_B(s) - F_S(s)) ds. \tag{D.5}$$

Hence, the price-setter has an incentive to raise the price of seller attribute choices in excess demand and lower the price of seller attribute choices in excess supply.

D.2.3 Seller payoffs The specification of the seller's payoff is complicated by the need to incorporate incentives that arise from the possibility of profitable seller deviations to attribute choices outside of S . Given an attribute choice s , price p , and price function p_U , set

$$B(s, p, p_U) \equiv \left\{ b \in [0, \bar{b}] : h_B(b, s) - p \geq \max_{s' \in [0, \bar{s}]} \{h_B(b, s') - p_U(s')\} \right\}. \tag{D.6}$$

Hence, $B(s, p, p_U)$ is the set of buyer attributes that find attribute s at price p (weakly) more attractive than any attribute $s' \in [0, \bar{s}]$ at price $p_U(s')$. Note that since the buyer is constrained in Γ^n to choose seller attributes so that (D.2) is satisfied, a maximizing buyer's payoff from an attribute b (ignoring costs) need not be given by $\max_{s' \in [0, \bar{s}]} \{h_B(b, s') - p_U(s')\}$. Note also that for all s and $p_U \in Y_P$, since there is no a priori restriction on p , $B(s, p, p_U)$ is nonempty for low p (possibly requiring $p < -P$, e.g., if $p_U \equiv -P$), and it is empty if $p > p_U(s)$. Indeed, for sufficiently low p , $B(s, p, p_U) = [0, \bar{b}]$.

¹⁸It suffices for this conclusion to show that Y is sequentially compact, since sequential compactness is equivalent to compactness for metric spaces (Dunford and Schwartz 1988, p. 20). An argument analogous to that of Helly's theorem (Billingsley 1986, Theorem 25.9) shows that Y is sequentially compact.

LEMMA D.1. (i) If $B(s, p, p_U) \neq \emptyset$, then $B(s, p, p_U) = [b_1, b_2]$ with $b_1 \leq b_2$.

(ii) For fixed s and p_U , let $\bar{p}(s, p_U) \equiv \sup\{p : B(s, p, p_U) \neq \emptyset\}$ and write $[b_1(p), b_2(p)]$ for $B(s, p, p_U)$ when $p \leq \bar{p}(s, p_U)$. Denote the set of discontinuity points in the domain of $b_j(p)$ by $\mathcal{D}_j(s, p_U)$ ($j = 1, 2$). The set $\{s : \mathcal{D}_j(s, p_U) \neq \emptyset\}$ has zero Lebesgue measure.

(iii) Suppose $\{(s^\ell, p^\ell, p_U^\ell)\}_\ell$ is a sequence converging to (s, p, p_U) with $\emptyset \neq B(s^\ell, p^\ell, p_U^\ell) \equiv [b_1^\ell, b_2^\ell]$. Then $B(s, p, p_U) \neq \emptyset$ and so $B(s, p, p_U) = [b_1, b_2]$, where

$$b_1 \leq \liminf_{\ell} b_1^\ell \leq \limsup_{\ell} b_2^\ell \leq b_2. \tag{D.7}$$

(iv) Moreover, if $p \notin \mathcal{D}_j(s, p_U) \cup \{\bar{p}(s, p_U)\}$, then $b_j = \lim_{\ell} b_j^\ell$.

PROOF. (i) Suppose $b_1, b_2 \in B(s, p, p_U)$ with $b_1 < b_2$, and $\hat{b} \notin B(s, p, p_U)$ for some $\hat{b} \in (b_1, b_2)$. Then there exists $\hat{s} \in [0, \bar{s}]$ such that

$$h_B(\hat{b}, s) - p < h_B(\hat{b}, \hat{s}) - p_U(\hat{s}).$$

If $\hat{s} > s$, then Assumption 1 implies

$$\begin{aligned} h_B(b_2, \hat{s}) - h_B(b_2, s) &\geq h_B(\hat{b}, \hat{s}) - h_B(\hat{b}, s) \\ &> p_U(\hat{s}) - p, \end{aligned}$$

contradicting $b_2 \in B(s, p, p_U)$. Similarly, $\hat{s} < s$ contradicts $b_1 \in B(s, p, p_U)$ and so $\hat{s} = s$. But $b_2 \in B(s, p, p_U)$ then implies $p_U(s) \geq p$, while $\hat{b} \notin B(s, p, p_U)$ implies $p_U(s) < p$, the final contradiction, and so $\hat{b} \in B(s, p, p_U)$. It is immediate that $B(s, p, p_U)$ is closed.

(ii) Since $B(s, p', p_U) \supset B(s, p, p_U)$ for $p' < p$, $b_1(p)$ and $b_2(p)$ are monotonic functions of p , and so are continuous except at a countable number of points. Moreover, we can apply the maximum theorem (since each of the functions in the maximum is continuous) to conclude that the right side of the inequality in (D.6) is continuous in b , and so b_1 and b_2 are left-continuous functions of p (as (D.6) features a weak inequality bounding p from above).

Suppose $p \in \mathcal{D}_1(s, p_U)$ and let $b_1^+ \equiv \lim_{p' \searrow p} b_1(p')$. Since b_1 is left-continuous, $b_1(p) < b_1^+$. Then for all $b \in [b_1(p), b_1^+]$,

$$h_B(b, s) - p = \max_{s' \in [0, \bar{s}]} h_B(b, s') - p_U(s').$$

From the envelope theorem (Milgrom and Segal 2002, Theorem 2), this implies, for all $b \in (b_1(p), b_1^+)$, that

$$\frac{\partial h_B(b, s)}{\partial b} = \frac{\partial h_B(b, s'(b))}{\partial b},$$

where $s'(b) \in \arg \max_{s' \in [0, \bar{s}]} h_B(b, s') - p_U(s')$. Assumption 1 then implies $s = s'(b)$ for all $b \in (b_1(p), b_1^+)$ and so $p = p_U(s)$.

Since $b_1^+ \in B(s, p_U(s), p_U)$, for all $s'' > s$, then

$$h_B(b_1^+, s'') - h_B(b_1^+, s) \leq p_U(s'') - p_U(s)$$

so that

$$\frac{\partial h_B(b_1^+, s)}{\partial s} \leq \liminf_{s'' > s} \frac{p_U(s'') - p_U(s)}{s'' - s}.$$

Alternatively, for all $s' < s$,

$$p_U(s) - p_U(s') \leq h_B(b_1(p), s) - h_B(b_1(p), s'),$$

so that

$$\limsup_{s' < s} \frac{p_U(s) - p_U(s')}{s - s'} \leq \frac{\partial h_B(b_1(p), s)}{\partial s}.$$

Consequently, since

$$\frac{\partial h_B(b_1(p), s)}{\partial s} < \frac{\partial h_B(b_1^+, s)}{\partial s},$$

the price function p_U cannot be differentiable at s . Finally, since p_U is a monotonic function, it is differentiable almost everywhere (Billingsley 1986, Theorem 31.2) and, hence, $\{s : \mathcal{D}_1(s, p_U) \neq \emptyset\}$ has zero Lebesgue measure. A similar argument shows that $\{s : \mathcal{D}_2(s, p_U) \neq \emptyset\}$ has zero Lebesgue measure.

(iii) Suppose $\{(s^\ell, p^\ell, p_U^\ell)\}_\ell$ is a sequence converging to (s, p, p_U) and let $\{b^\ell\}$ be a sequence of attributes with $b^\ell \in B(s^\ell, p^\ell, p_U^\ell)$ for all ℓ . By taking a subsequence if necessary, we can assume $\{b^\ell\}$ is a convergent sequence with limit b . Since

$$h_B(b^\ell, s^\ell) - p^\ell \geq \max_{s' \in [0, \bar{s}]} \{h_B(b^\ell, s') - p_U^\ell(s')\} \quad \forall \ell,$$

taking limits gives

$$h_B(b, s) - p \geq \max_{s' \in [0, \bar{s}]} \{h_B(b, s') - p_U(s')\}$$

and so $b \in B(s, p, p_U)$. Hence, from part (ii) of the lemma, $B(s, p, p_U) = [b_1, b_2]$ and (D.7) then follows from taking sequences $\{b^\ell\}$ with limits $\liminf_\ell b_1^\ell$ and $\limsup_\ell b_2^\ell$.

(iv) Consider b_2 and suppose $p \notin \mathcal{D}_2(s, p_U) \cup \{\bar{p}(s, p_U)\}$ (and so $p < \bar{p}(s, p_U)$). Hence, $b_2 = b_2^+ \equiv \lim_{p' \searrow p} b_2(p')$. Consider $b \in (b_1^+, b_2)$. For $p' > p$ sufficiently close to p , we have $b \in B(s, p', p_U)$ and so

$$h_B(b, s) - p > \max_{s' \in [0, \bar{s}]} \{h_B(b, s') - p_U(s')\}.$$

Consequently, for ℓ sufficiently large,

$$h_B(b, s^\ell) - p^\ell > \max_{s' \in [0, \bar{s}]} \{h_B(b, s') - p_U(s')\},$$

i.e., $b \in B(s^\ell, p^\ell, p_U^\ell)$. This implies that $b_2^\ell(p^\ell) \geq b$ and, hence, $\liminf b_2^\ell(p^\ell) \geq b$. Since this holds for all $b \in (b_1^+, b_2)$ and $\limsup_\ell b_2^\ell \leq b_2$, we have $\lim_\ell b_2^\ell = b_2$. The argument for b_1 is an obvious modification of this argument. \square

Fix (s, p, p_U) and suppose $\lambda(\{\beta : \mathbf{b}(\beta) \in B(s, p, p_U)\}) > 0$. Since \mathbf{b} is strictly increasing and continuous, it then follows from Lemma D.1 that $\mathbf{b}([0, 1]) \cap B(s, p, p_U) = [b'_1, b'_2]$ for some $0 \leq b'_1 \leq b'_2 \leq \bar{b}$. The payoff to the seller from (s, p, \mathbf{b}, p_U) is given by

$$H(s, p, \mathbf{b}, p_U) \equiv h_S(b'_1, s) + p.$$

This function depends on p_U and \mathbf{b} through the dependence of b'_1 on $B(s, p, p_U)$ and \mathbf{b} . For later reference, note that for fixed s, \mathbf{b} , and p_U , the function $H(s, p, \mathbf{b}, p_U)$ is continuous from the left in p (since \mathbf{b} satisfies (D.1) and both $b_1(p)$ and $b_2(p)$, defined just before Lemma D.1, are left-continuous).

We set

$$\tilde{P}(s, \mathbf{b}, p_U) \equiv \{p : \lambda(\{\beta : \mathbf{b}(\beta) \in B(s, p, p_U)\}) > 0\},$$

and noting that this set is nonempty, we define

$$\bar{H}(s, \mathbf{b}, p_U) \equiv \max \left\{ \sup_{p \in \tilde{P}(s, \mathbf{b}, p_U)} H(s, p, \mathbf{b}, p_U), h_S(0, s) + p_U(s) \right\}. \tag{D.8}$$

Notice that if $p \in \tilde{P}(s, \mathbf{b}, p_U)$ for all $p < p_U(s)$, then the first term in (D.8) will be the maximum.¹⁹

The seller's payoff from $\mathbf{s} \in Y_S^n$ when the buyer and price-setter have chosen $(\mathbf{b}, \mathbf{s}_B, p_U) \in Y_B^n \times Y_P$ is then

$$\int (\bar{H}(\mathbf{s}(\sigma), \mathbf{b}, p_U) - c_S(\mathbf{s}(\sigma), \sigma)) d\sigma. \tag{D.9}$$

Taking the maximum over $\sup_{p \in \tilde{P}(s, \mathbf{b}, p_U)} H(s, p, \mathbf{b}, p_U)$ and $h_S(0, s) + p_U(s)$ effectively assumes that the seller can always sell attribute choice s at the posted price $p_U(s)$, though perhaps only attracting buyer attribute choice 0.

Note that the seller, when considering the payoff implications of altering the attribute-choice function over an interval of seller types, can ignore the seller types outside the interval, since feasibility of buyer responses is irrelevant (the comparison in B for buyer attributes is always to her payoffs, which is independent of seller behavior).

D.3 Equilibrium in game Γ^n

Our next task is to show that each game Γ^n has a Nash equilibrium and that the price-setter plays a pure strategy in any such equilibrium. To do this, we first note that the price-setter's payoff is concave in p_U (note that the buyer's and sellers's payoffs need

¹⁹It need not be true that for $s \in \mathbf{s}_B([0, 1])$, $p \in \tilde{P}(s, \mathbf{b}, p_U)$ for all $p < p_U(s)$. Moreover, we may have $\bar{H}(s, \mathbf{b}, p_U) \neq h_S(\mathbf{b}(\mathbf{s}_B^{-1}(s)), s) + p_U(s)$ (see the discussion just before Lemma D.1).

not be quasiconcave). If the payoff functions in game Γ^n are continuous, then Glicksberg’s fixed point theorem, applied to the game where we allow the buyer and seller to randomize, yields a Nash equilibrium in which the buyer and seller may randomize, but the price-setter does not.

LEMMA D.2. *The buyer, price-setter, and seller payoff functions given by (D.4), (D.5), and (D.9) are continuous functions of $(\mathbf{b}, \mathbf{s}_B, \mathbf{s}, p_U)$ on Y^n .*

PROOF. We first note that for increasing, bounded functions on a compact set, L^1 convergence implies convergence almost everywhere.²⁰

Consider first the buyer. The functions \mathbf{b} , \mathbf{s}_B , and p_U are bounded functions on compact sets and, hence, the absolute value of each of these functions is dominated by an integrable function. The continuity of the buyer’s payoff then follows immediately from Lebesgue’s dominated convergence theorem if we can show that the convergence of \mathbf{b} , p_U , and \mathbf{s}_B in the L^1 norm (and hence almost everywhere) implies the convergence almost everywhere of $h_B(\mathbf{b}, \mathbf{s}_B)$, $p_U(\mathbf{s}_B)$, and $c_B(\mathbf{b}(\cdot), \cdot)$ (note that we are talking about sequences of functions within a given game Γ^n). The first and the third of these follow from the continuity of h_B and c_B (from Assumptions 1 and 3), while for the remaining case, it suffices to note that the collection Y_P is equicontinuous.

Consider now the price-setter. Suppose \mathbf{s}^ℓ converges in L^1 and so, almost everywhere, to \mathbf{s} . Then F_S^ℓ converges weakly to F_S (and so a.e.).²¹ Similarly, if \mathbf{s}_B^ℓ converges in L^1 to \mathbf{s}_B , then F_B^ℓ converges a.e. to F_B . Continuity for the price-setter’s payoff then follows from arguments analogous to those applied to the buyer, since we have convergence almost everywhere of $p_U[F_B - F_S]$.

Finally, we turn to the seller, where the proof of continuity is more involved. It suffices to argue that $\overline{H}(s, \mathbf{b}, p_U)$ is continuous in (s, \mathbf{b}, p_U) for almost all s (since \mathbf{s}_B is irrelevant in the determination of the seller’s payoff and the continuity with respect to \mathbf{s} is then obvious, at which point another appeal to Lebesgue’s dominated convergence theorem completes the argument).

Fix a sequence $(s^\ell, \mathbf{b}^\ell, p_U^\ell)$ converging to some point $(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$. Since we need continuity only for almost all $s \in [0, \bar{s}]$, we can assume $\mathcal{D}_1(\hat{s}, \hat{p}_U) \cup \mathcal{D}_2(\hat{s}, \hat{p}_U) = \emptyset$ (or, equivalently, that \hat{p}_U is differentiable at \hat{s} , see the proof of Lemma D.1). We thus need only to prove the following claim.

CLAIM 1. *We have $\lim_{\ell \rightarrow \infty} \overline{H}(s^\ell, \mathbf{b}^\ell, p_U^\ell) = \overline{H}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$.*

²⁰Suppose $\{f_n\}_n$, with each f_n increasing, converges in L^1 norm to an increasing function f without converging almost everywhere. Then since f is discontinuous on a set of measure zero, there exists (for example) a continuity point x of f with $\limsup f_n(x) > f(x)$ (with the case $\liminf f_n(x) < f(x)$ analogous). The continuity of f at x then ensures that for some point $y > x$, some $\varepsilon > 0$, all $z \in [x, y]$, and for infinitely many n , we have $f_n(z) \geq f_n(x) \geq f(y) + \varepsilon \geq f(z) + \varepsilon$. This in turn ensures that $\int |f_n(z) - f(z)| dz > (y - x)\varepsilon$ infinitely often, precluding the L^1 convergence of $\{f_n\}_{n=1}^\infty$ to f .

²¹Fix $\varepsilon > 0$. By Egoroff’s theorem (Royden 1988, p. 73), \mathbf{s}^ℓ converges uniformly to \mathbf{s} on a set E of measure at least $1 - \varepsilon$. Suppose s is a continuity point of F_S . There then exists $\delta > 0$ such that $|F_S(s) - F_S(s')| < \varepsilon$ for all $|s - s'| \leq \delta$. There exists ℓ' such that for all $\sigma \in E$ and for all $\ell > \ell'$, $|\mathbf{s}^\ell(\sigma) - \mathbf{s}(\sigma)| < \delta$. Consequently, $F_S^\ell(s) = \lambda\{\sigma : \mathbf{s}^\ell(\sigma) \leq s\} \leq \lambda\{\sigma : \mathbf{s}(\sigma) - \delta \leq s\} + \varepsilon = F_S(s + \delta) + \varepsilon$ and $F_S(s - \delta) - \varepsilon \leq F_S^\ell(s)$, and so $|F_S^\ell(s) - F_S(s)| < 2\varepsilon$. Hence, F_S^ℓ converges weakly to F_S .

PROOF. Since $\overline{H}^k(s, \mathbf{b}, p_U)$ is the maximum of two terms, it suffices to show that

$$\lim_{\ell \rightarrow \infty} \sup_{p \in \tilde{P}(s^\ell, \mathbf{b}^\ell, p_U^\ell)} H(s^\ell, p, \mathbf{b}^\ell, p_U^\ell) = \sup_{p \in \tilde{P}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)} H(\hat{s}, p, \hat{\mathbf{b}}, \hat{p}_U)$$

and

$$\lim_{\ell \rightarrow \infty} h_S(0, s^\ell) + p_U^\ell(s^\ell) = h_S(0, \hat{s}) + \hat{p}_U(\hat{s}).$$

The second equality is immediate from the continuity of h_S and \hat{p}_U at \hat{s} .

We accordingly turn to the first term. To conserve on notation, we define $\sup_{p \in \tilde{P}(s, \mathbf{b}, p_U)} H(s, p, \mathbf{b}, p_U) \equiv \overline{\overline{H}}(s, \mathbf{b}, p_U)$.

We first show that

$$\liminf_{\ell \rightarrow \infty} \overline{\overline{H}}(s^\ell, \mathbf{b}^\ell, p_U^\ell) \geq \overline{\overline{H}}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U). \tag{D.10}$$

For all $\varepsilon > 0$, there exists $\hat{p} \in \tilde{P}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$ such that

$$H(\hat{s}, \hat{p}, \hat{\mathbf{b}}, \hat{p}_U) + \varepsilon/2 \geq \overline{\overline{H}}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U).$$

Since $H(\hat{s}, p, \hat{\mathbf{b}}, \hat{p}_U)$ is continuous from the left in p , there exists $\hat{p}' \notin \mathcal{D}_1(\hat{s}, \hat{p}_U) \cup \mathcal{D}_2(\hat{s}, \hat{p}_U) \cup \{\overline{p}(\hat{s}, \hat{p}_U)\}$ with $\hat{p}' \leq \hat{p}$ satisfying

$$|H(\hat{s}, \hat{p}, \hat{\mathbf{b}}, \hat{p}_U) - H(\hat{s}, \hat{p}', \hat{\mathbf{b}}, \hat{p}_U)| < \varepsilon/2$$

and so

$$H(\hat{s}, \hat{p}', \hat{\mathbf{b}}, \hat{p}_U) + \varepsilon \geq \overline{\overline{H}}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U).$$

Since $\hat{\mathbf{b}}$ satisfies (D.1), for sufficiently large ℓ , $\hat{p}' \in \tilde{P}(s^\ell, \mathbf{b}^\ell, p_U^\ell)$ and so (applying Lemma D.1)

$$\lim_{\ell \rightarrow \infty} H(s^\ell, \hat{p}', \mathbf{b}^\ell, p_U^\ell) = H(\hat{s}, \hat{p}', \hat{\mathbf{b}}, \hat{p}_U).$$

Hence,

$$\liminf_{\ell \rightarrow \infty} \overline{\overline{H}}(s^\ell, \mathbf{b}^\ell, p_U^\ell) + \varepsilon \geq \overline{\overline{H}}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U) \quad \forall \varepsilon > 0,$$

yielding (D.10).

We now argue that

$$\overline{\overline{H}}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U) \geq \limsup_{\ell \rightarrow \infty} \overline{\overline{H}}(s^\ell, \mathbf{b}^\ell, p_U^\ell), \tag{D.11}$$

which with (D.10) gives continuity.

Fix $\varepsilon > 0$. For each ℓ , there exists $p^\ell \in \tilde{P}(s^\ell, \mathbf{b}^\ell, p_U^\ell)$ such that

$$H(s^\ell, p^\ell, \mathbf{b}^\ell, p_U^\ell) + \varepsilon \geq \overline{\overline{H}}(s^\ell, \mathbf{b}^\ell, p_U^\ell). \tag{D.12}$$

Without loss of generality, we can assume $\{p^\ell\}_\ell$ is a convergent sequence that has limit \hat{p} . Suppose first that $\hat{p} \in \tilde{P}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$. If $\hat{p} \notin \{\bar{p}(\hat{s}, \hat{p}_U)\}$, it is immediate that

$$H(\hat{s}, \hat{p}, \hat{\mathbf{b}}, \hat{p}_U) + \varepsilon \geq \limsup_{\ell \rightarrow \infty} \bar{H}(s^\ell, \mathbf{b}^\ell, p_U^\ell),$$

which (since it holds for all ε) implies (D.11).

Suppose now that $\hat{p} \notin \tilde{P}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$ or $p = \bar{p}(\hat{s}, \hat{p}_U)$. Since \hat{p}_U is differentiable at \hat{s} , there cannot be a nondegenerate interval of buyer attributes that are indifferent between (\hat{s}, \hat{p}) and the unconstrained optimal seller attribute under \hat{p}_U . This implies $\hat{\mathbf{b}} \cap ([0, 1] \cap B(\hat{s}, \hat{p}, \hat{p}_U)) = \{\hat{b}\}$ for some \hat{b} and so

$$\bar{H}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U) \geq h_S(\hat{b}, \hat{s}) + \hat{p}.$$

From Lemma D.1,

$$\lim_{\ell \rightarrow \infty} H(s^\ell, p^\ell, \mathbf{b}^\ell, p_U^\ell) + \varepsilon = h_S(\hat{b}, \hat{s}) + \hat{p} + \varepsilon$$

and so (taking the lim sup of both sides of (D.12))

$$\bar{H}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U) + \varepsilon \geq \limsup_{\ell \rightarrow \infty} \bar{H}(s^\ell, \mathbf{b}^\ell, p_U^\ell),$$

which (since it holds for all $\varepsilon > 0$) implies (D.11) completing the proof of the claim and lemma. □

Allowing the buyer and seller to choose mixed strategies then gives us a game whose best responses consist of closed, convex sets. As a result, we can apply Glicksberg (1952) to conclude that we have a Nash equilibrium in which the price-setter plays a pure strategy, while the buyer and seller may mix.

LEMMA D.3. *The game Γ^n has a Nash equilibrium, $(\xi_B^n, \xi_S^n, p_U^n) \in \Delta(Y_B) \times \Delta(Y_S) \times Y_P$.*

D.4 The limit $n \rightarrow \infty$

We now examine the limit as $n \rightarrow \infty$. In particular, let $\{(\xi_B^n, \xi_S^n, p_U^n)\}_n \subset \Delta(Y_B) \times \Delta(Y_S) \times Y_P$ be a sequence of Nash equilibria of the games Γ^n . Without loss of generality (since the relevant spaces are sequentially compact), we may assume both that the sequence of equilibria converges to some limit $(\xi_B^*, \xi_S^*, p_U^*)$ and that players' payoffs also converge.

We examine the limit $(\xi_B^*, \xi_S^*, p_U^*)$. Intuitively, we would like to think of this profile as the equilibrium of a "limit game." However, the definition of this limit game is not straightforward, since the definition of the seller's payoffs in the game Γ^n relies on the strategies \mathbf{b} , \mathbf{s}_B , and \mathbf{s} having properties (such as strict monotonicity and continuity) that need not carry over to their limits. In establishing properties of $(\xi_B^*, \xi_S^*, p_U^*)$, we accordingly typically begin our argument in the limit and then pass back to the approximating equilibrium profile $(\xi_B^n, \xi_S^n, p_U^n)$ to obtain a contradiction. The latter step of the

argument is notationally cumbersome and we do not always make the approximation explicit.

Note that while the seller is best responding to ξ_B^n in choosing \mathbf{s} , the choice of p implicit in (D.8) is made after $(\mathbf{b}, \mathbf{s}_B)$ is realized.

While the L^1 topology does not distinguish between functions that agree almost everywhere, it is important for some of the later developments that we make the selection indicated in the next lemma from the equivalence classes of functions that agree almost everywhere.

LEMMA D.4. *The limit profile $(\xi_B^*, \xi_S^*, p_U^*)$ is pure, which we denote by $(\mathbf{b}^*, \mathbf{s}_B^*, \mathbf{s}^*, p_U^*)$. The limit functions can be (and subsequently are) taken to be increasing, and the functions \mathbf{b}^* , \mathbf{s}_B^* , and \mathbf{s}^* can be (and subsequently are) taken to be continuous from the left.*

PROOF. Consider the buyer (the case of the seller is analogous). Toward a contradiction, suppose the buyer’s strategy (ξ_B^*, ξ_S^*) is not pure. Let $\xi_{B,b}^*$ and $\xi_{B,s}^*$ denote the marginal distributions induced on buyer and seller attributes chosen by the buyer. Then define a pair of increasing functions $\mathbf{b}' : [0, 1] \rightarrow [0, \bar{b}]$ and $\mathbf{s}'_B : [0, 1] \rightarrow [0, \bar{s}]$ by

$$\mathbf{b}'(\beta) = \inf\{b : \xi_{B,b}^*(b) \geq \beta\}$$

and

$$\mathbf{s}'_B(\beta) = \inf\{s : \xi_{B,s}^*(s) \geq \beta\}.$$

These functions constitute pure strategies for the buyer yielding the same distribution of buyer and seller attributes as chosen by the buyer under $\xi_{B,s}^*$. (For example, for any set $[\underline{b}, \bar{b}]$ of buyer attributes with $\underline{b} > 0$, $\xi_{B,b}^*[\underline{b}, \bar{b}] = \beta'' - \beta'$, where $\mathbf{b}'(\beta'') = \bar{b}$ and $\mathbf{b}'(\beta') = \underline{b}$.) However, the \mathbf{b}' and \mathbf{s}'_B feature positive assortativity between the buyer’s types and attribute choice, and between the buyer’s and the seller’s attributes with which the buyer matches, while $\xi_{B,b}^*$ and $\xi_{B,s}^*$, being mixed, do not. From Assumptions 1 and 3, this positive assortativity increases the buyer’s payoff and so the constructed pure strategy strictly increases the buyer’s payoff. It then follows from straightforward continuity arguments that for sufficiently large n , i.e., for a game in which the slope requirements on the buyer’s strategy are sufficiently weak and the equilibrium profile $(\xi_B^n, \xi_S^n, p_U^n)$ is sufficiently close to $(\xi_B^*, \xi_S^*, p_U^*)$, there is a pure strategy sufficiently close to \mathbf{b}' and \mathbf{s}'_B that gives the buyer a payoff higher than his supposed equilibrium payoff in Γ^n —a contradiction. Hence, the buyer cannot mix.

The conclusion that each function is increasing is an implication of the observation that if a sequence of increasing functions $\{f_n\}$ converges in L^1 to a function f , then that function is increasing. □

It is helpful to keep in mind the nature of convergence in $\Delta(Y_B) \times \Delta(Y_S) \times Y_P$. Recalling that Y_B and Y_S are each endowed with the L^1 norm, and Y_P is endowed with the sup norm, and recalling the definition of the Prohorov metric (which metrizes weak

convergence), $(\xi_B^n, \xi_S^n, p_U^n)$ converges to the pure profile $(\mathbf{b}^*, \mathbf{s}_B^*, \mathbf{s}^*, p_U^*)$ if, and only if, the following statement holds: For all $\varepsilon > 0$, there exists n' such that for all $n \geq n'$,

$$\xi_B^n \left(\left\{ (\mathbf{b}, \mathbf{s}_B) \in Y_B^n : \int |\mathbf{b}(\beta) - \mathbf{b}^*(\beta)| d\beta < \varepsilon, \int |\mathbf{s}_B(\beta) - \mathbf{s}_B^*(\beta)| d\beta < \varepsilon \right\} \right) \geq 1 - \varepsilon,$$

$$\xi_S^n \left(\left\{ \mathbf{s} \in Y_S^n : \int |\mathbf{s}(\sigma) - \mathbf{s}^*(\sigma)| d\sigma < \varepsilon \right\} \right) \geq 1 - \varepsilon,$$

and

$$\sup |p_U^n(s) - p_U^*(s)| < \varepsilon.$$

We next restate the nature of convergence in a more useful form.

LEMMA D.5. *For all $\varepsilon > 0$, there exists a set $E^\varepsilon \subset [0, 1]$ with $\lambda(E^\varepsilon) \geq 1 - \varepsilon$ and n_ε such that for all $n \geq n_\varepsilon$,*

$$\xi_B^n \left(\left\{ (\mathbf{b}, \mathbf{s}_B) \in Y_B^n : |\mathbf{b}(\beta) - \mathbf{b}^*(\beta)| < \varepsilon, |\mathbf{s}_B(\beta) - \mathbf{s}_B^*(\beta)| < \varepsilon, \forall \beta \in E^\varepsilon \right\} \right) \geq 1 - \varepsilon,$$

$$\xi_S^n \left(\left\{ \mathbf{s} \in Y_S^n : |\mathbf{s}(\sigma) - \mathbf{s}^*(\sigma)| < \varepsilon, \forall \sigma \in E^\varepsilon \right\} \right) \geq 1 - \varepsilon,$$

and

$$|p_U^n(s) - p_U^*(s)| < \varepsilon \quad \forall s.$$

Moreover, the sets E^ε are nested: $E^{\varepsilon'} \subset E^\varepsilon$ if $\varepsilon < \varepsilon'$.

PROOF. Fix $\varepsilon > 0$. We prove that there is a set E_S^ε with $\lambda(E_S^\varepsilon) > 1 - \varepsilon/3$ and an integer n'_S such that

$$\xi_S^n \left(\left\{ \mathbf{s} \in Y_S^n : |\mathbf{s}(\sigma) - \mathbf{s}^*(\sigma)| < \varepsilon, \forall \sigma \in E_S^\varepsilon \right\} \right) \geq 1 - \varepsilon \tag{D.13}$$

for all $n > n'_S$. The same argument implies a set E_B^ε and integer n'_B for the function \mathbf{b}^* , and a \hat{E}_B^ε and n''_B for the function \mathbf{s}_B^* .²² The desired set is $E^\varepsilon = E_S^\varepsilon \cap E_B^\varepsilon \cap \hat{E}_B^\varepsilon$ and integer is $n_\varepsilon = \max\{n'_S, n'_B, n''_B\}$.

Let $\{\sigma^k\}$ be an enumeration of the discontinuities of \mathbf{s}^* . Since \mathbf{s}^* is bounded, there exists K such that the total size of the discontinuities over $\{\sigma^k\}_{k>K}$ is less than $\varepsilon/6$.

²²More precisely, the sets can be chosen so that for $n > n'_B$,

$$\xi_B^n \left(\left\{ (\mathbf{b}, \mathbf{s}_B) \in Y_B^n : |\mathbf{b}(\beta) - \mathbf{b}^*(\beta)| < \varepsilon, \forall \beta \in E_B^\varepsilon \right\} \right) \geq 1 - \varepsilon/2$$

and for $n > n''_B$,

$$\xi_B^n \left(\left\{ (\mathbf{b}, \mathbf{s}_B) \in Y_B^n : |\mathbf{s}_B(\beta) - \mathbf{s}_B^*(\beta)| < \varepsilon, \forall \beta \in \hat{E}_B^\varepsilon \right\} \right) \geq 1 - \varepsilon/2,$$

so that for $n > \max\{n'_B, n''_B\}$,

$$\xi_B^n \left(\left\{ (\mathbf{b}, \mathbf{s}_B) \in Y_B^n : |\mathbf{b}(\beta) - \mathbf{b}^*(\beta)| < \varepsilon, |\mathbf{s}_B(\beta) - \mathbf{s}_B^*(\beta)| < \varepsilon, \forall \beta \in E_B^\varepsilon \cap \hat{E}_B^\varepsilon \right\} \right) \geq 1 - \varepsilon.$$

Fix $L > 2$ such that $\{(\sigma^k - 2^{-kL}, \sigma^k + 2^{-kL})\}_{k=1}^K$ is pairwise disjoint and $2^{1-L} < \varepsilon/6$. Defining

$$E_S^\varepsilon = [0, 1] \setminus \bigcup_k (\sigma^k - 2^{-kL}, \sigma^k + 2^{-kL})$$

yields a set of measure at least $1 - \varepsilon/3$.

Let E_S^K be the set given by $[0, 1] \setminus \bigcup_{k=1}^K (\sigma^k - 2^{-kL}, \sigma^k + 2^{-kL})$; clearly $E_S^\varepsilon \subset E_S^K$. The set E_S^K can be written as the disjoint union of closed intervals I_k , $k = 0, 1, \dots, K$. There exists an $\eta > 0$ such that for all k and for all $\sigma, \sigma' \in I_k$, if $|\sigma - \sigma'| < \eta$, then $|\mathbf{s}^*(\sigma) - \mathbf{s}^*(\sigma')| < \varepsilon/3$.

Let $\{x_\ell\} \subset I_k$ be an η -grid of I_k , i.e., $x_{\ell+1} - \eta < x_\ell < x_{\ell+1}$ for all ℓ .

Consider an increasing function \mathbf{s} that satisfies $\int |\mathbf{s} - \mathbf{s}^*| < \varepsilon\eta/3$. We claim that for all $\sigma \in E_S^K$ (and so for all $\sigma \in E_S^\varepsilon$), $|\mathbf{s} - \mathbf{s}^*| < \varepsilon$. Observe that (D.13) then follows, since n'_S can be chosen so that $\xi_S^n(\{\mathbf{s} \in Y_S^n : \int |\mathbf{s} - \mathbf{s}^*| < \varepsilon\eta/3\}) \geq 1 - \varepsilon$ holds for all $n > n'_S$.

The claim follows from two observations.

1. We have $|\mathbf{s}(x_\ell) - \mathbf{s}^*(x_\ell)| < 2\varepsilon/3$: Suppose $\mathbf{s}(x_\ell) \geq \mathbf{s}^*(x_\ell) + 2\varepsilon/3$ (the other possibility is handled mutatis mutandis). Then for all $\sigma \in (x_\ell, x_{\ell+1})$,

$$\mathbf{s}(\sigma) \geq \mathbf{s}(x_\ell) \geq \mathbf{s}^*(x_\ell) + 2\varepsilon/3 > \mathbf{s}^*(\sigma) + \varepsilon/3.$$

But this is impossible, since it implies $\int |\mathbf{s} - \mathbf{s}^*| > \varepsilon\eta/3$.

2. For all ℓ and all $\sigma \in (x_\ell, x_{\ell+1})$, $|\mathbf{s}(\sigma) - \mathbf{s}^*(\sigma)| < \varepsilon$: Suppose $\mathbf{s}(\sigma) \geq \mathbf{s}^*(\sigma) + \varepsilon$ (the other possibility is handled mutatis mutandis). Then

$$\mathbf{s}(x_{\ell+1}) \geq \mathbf{s}(\sigma) \geq \mathbf{s}^*(\sigma) + \varepsilon \geq \mathbf{s}^*(x_{\ell+1}) + 2\varepsilon/3,$$

contradicting the previous observation.

The last assertion of Lemma D.5 is immediate from the definition of E_S^ε . □

LEMMA D.6. *The profile $(\mathbf{b}^*, \mathbf{s}_B^*, \mathbf{s}^*, p_U^*)$ balances the market, i.e., $F_B^*(s) = F_S^*(s)$ for all s . Hence, $\mathbf{s}_B^*(x) = \mathbf{s}^*(x)$ for almost all $x \in [0, 1]$.*

PROOF. Since F_B^* and F_S^* are continuous from the right, it suffices to show that they agree almost everywhere. We first argue that $F_B^*(s) - F_S^*(s) \leq 0$ almost everywhere. Suppose this is not the case, so there exists $\hat{s} < \bar{s}$ with $F_B^*(\hat{s}) - F_S^*(\hat{s}) = \varepsilon > 0$ and with \hat{s} a continuity point of $F_B^* - F_S^*$. Then there exists s_1 and s_2 with $\hat{s} \in (s_1, s_2)$, $F_B^*(s) - F_S^*(s) \geq \varepsilon/2$ on $[s_1, s_2]$, and either $s_1 = 0$ or, for every $\eta > 0$, there is a value $s_\eta \in [s_1 - \eta, s_1)$ with $F_B^*(s_\eta) - F_S^*(s_\eta) < \varepsilon/2$ (note that $F_B^*(s_\eta) - F_S^*(s_\eta)$ may be negative and so is bounded below by -1). We consider the case in which $s_1 > 0$ and $p_U^*(s_1) < p_U^*(s_2)$; the remaining cases are a straightforward simplification.

Since $F_B^*(s) - F_S^*(s) > 0$ on $[s_1, s_2]$, for fixed $p_U^*(s_1)$ and $p_U^*(s_2)$, the price-setter must be setting prices as large as possible on this interval. If not, there is a price function $\hat{p}_U \in Y_P$ with $\hat{p}_U(s) \geq p_U^*(s)$ for all s and $\hat{p}_U(s) > p_U^*(s)$ for some s yielding strictly higher

payoffs to the price-setter than p_U^* in Γ^n for sufficiently large n , when the buyer and seller choose (ξ_B^n, ξ_S^n) . But this contradicts the equilibrium property of $(\xi_B^n, \xi_S^n, p_U^n)$.

Hence, there exists $s' \in [s_1, s_2]$ such that $dp_U^*(s)/ds = 2\Delta$ on (s_1, s') and $p_U^*(s) = p_U^*(s_2)$ for $s \in [s', s_2]$. That is, prices increase at the maximum rate possible until hitting $p_U^*(s_2)$ (with $s' = s_2$ possible, but since $p_U^*(s_1) < p_U^*(s_2)$, we have $s_1 < s'$). Consequently, $\mathbf{s}_B([0, 1]) \cap [s_1, s_2] \subset \{s_1, s_2\}$, i.e., buyers demand only seller attribute choices s_1 and s_2 from this interval. (Since all seller attribute choices in $[s', s_2]$ command the same price, buyers demand only attribute choice s_2 from this set, while the price of a seller attribute choice increases sufficiently quickly on $[s_1, s']$ that from this set buyers demand only s_1 .)

Since for every $\eta > 0$, there exists $s_\eta \in [s_1 - \eta, s_1]$ with $F_B^*(s_\eta) - F_S^*(s_\eta) < \varepsilon/2$ and yet $F_B^*(s_1) - F_S^*(s_1) \geq \varepsilon$, the buyer must choose attributes arbitrarily close to s_1 for some buyer types. This implies that there is a range of seller attributes just below s_1 with prices that are not too low, that is, there exists $\eta' > 0$ such that

$$p_U^*(s) > p_U^*(s_1) - \Delta(s_1 - s)$$

for all $s \in [s_1 - \eta', s_1]$. Consider now the price function $p_U^\eta \in Y_P$ given by

$$p_U^\eta(s) \equiv \begin{cases} p_U^*(s) & \text{if } s \geq s' \\ \min\{p_U^*(s_1 - \eta) + 2\Delta(s - s_1 + \eta), p_U^*(s')\} & \text{if } s \in (s_1 - \eta, s') \\ p_U^*(s) & \text{if } s \leq s_1 - \eta \end{cases}$$

and note that $p_U^0 = p_U^*$. Since $p_U^\eta \geq p_U^*$, the price-setter's payoff from choosing $p_U^\eta \in Y_P$ less the payoff from p_U^* is bounded below by

$$-\int_{s_1 - \eta}^{s_1} (p_U^\eta(s) - p_U^*(s)) ds + \int_{s_1}^{s'} (p_U^\eta(s) - p_U^*(s)) \varepsilon/2 ds. \tag{D.14}$$

For $\eta < \eta'$ and $s \in (s_1 - \eta, s_1)$,

$$\begin{aligned} p_U^\eta(s) - p_U^*(s) &\leq p_U^*(s_1 - \eta) + 2\Delta(s - s_1 + \eta) - p_U^*(s_1) + \Delta(s_1 - s) \\ &= p_U^*(s_1 - \eta) - p_U^*(s_1) - \Delta(s_1 - s) + 2\Delta\eta \\ &\leq 2\Delta\eta. \end{aligned}$$

Moreover, for $s \in (s_1, s_1 + (s' - s_1)/2)$, if η is sufficiently close to 0, we have $p_U^\eta(s) < p_U^*(s')$ and so

$$\begin{aligned} p_U^\eta(s) - p_U^*(s) &= p_U^*(s_1 - \eta) + 2\Delta(s - s_1 + \eta) - p_U^*(s) \\ &\geq p_U^*(s_1) - \Delta\eta + 2\Delta(s - s_1 + \eta) - p_U^*(s) \\ &= p_U^*(s_1) - \Delta\eta + 2\Delta(s - s_1 + \eta) - p_U^*(s_1) - 2\Delta(s - s_1) \\ &= \Delta\eta. \end{aligned}$$

Since $p_U^\eta(s) \geq p_U^*(s)$ for all s , the expression in (D.14) is bounded below by

$$-\int_{s_1-\eta}^{s_1} 2\Delta\eta ds + \int_{s_1}^{s_1+(s'-s_1)/2} \Delta\eta\varepsilon/2 ds,$$

which is clearly positive for sufficiently small η . Since the lower bound is strictly positive, the price-setter has a profitable deviation (in Γ^n for large n), a contradiction.

We conclude that $F_B^*(s) - F_S^*(s) \leq 0$ for almost all s . It remains to argue that it is not negative on a set of positive measure. Suppose it is. Then there must exist a seller characteristic $\hat{s} > 0$ such that $p_U(s) = -P$ for $s < \hat{s}$, $F_B^*(s) - F_S^*(s) < 0$ for a positive-measure subset of $[0, \hat{s}]$, and $F_B^*(s) - F_S^*(s) = 0$ for almost all $s > \hat{s}$. But then no seller would choose attributes in $[0, \hat{s})$, a contradiction. \square

We now seek a characterization of the seller's payoffs. Intuitively, we would like to use Lemma D.6 and the monotonicity of \mathbf{b}^* and \mathbf{s}_B^* to conclude that there is positive assortative matching, and, indeed, that a seller of type σ matches with a buyer of type $\beta = \sigma$. However, these properties may not hold if \mathbf{b}^* and \mathbf{s}_B^* are not strictly increasing. Moreover, even if we had such a matching, the specification of the seller's payoffs given by (D.9) leaves open the possibility that the (gross) payoff to a seller of type σ who chooses attribute s may not be given by $h_S(\tilde{b}(s), s) + p_U(s)$. Hence, the buyers who sellers are implicitly choosing in their payoff calculations may not duplicate those whose seller choices balance the market.

Our first step in addressing these issues is to show that the buyer's limiting attribute-choice function is indeed strictly increasing. Intuitively, if a positive measure of buyer types choose the same attribute, by having some higher types in the pool choose a slightly higher attribute and having some lower types choose a slightly lower attribute, we can keep the average attribute unchanged while saving costs (from Assumption 3).

LEMMA D.7. *The function \mathbf{b}^* is strictly increasing when it is nonzero.*

PROOF. By construction, \mathbf{b}^* is weakly increasing. We show that $\beta'' > \beta'$ and $\mathbf{b}^*(\beta') > 0$ imply $\mathbf{b}^*(\beta'') > \mathbf{b}^*(\beta')$. Suppose, to the contrary, that $b = \mathbf{b}^*(\beta) > 0$ for two distinct values of β .

Define $\beta_1 \equiv \inf\{\beta : \mathbf{b}^*(\beta) = b\}$, $\beta_2 \equiv \sup\{\beta : \mathbf{b}^*(\beta) = b\}$, and $\bar{\beta} = (\beta_1 + \beta_2)/2$. We assume $0 < \beta_1$ and $\beta_2 < 1$ (if equality holds in either case, then the argument is modified in the obvious manner). We now define a new attribute-choice function (as a function of a parameter $\eta > 0$) that is strictly increasing on a neighborhood of $[\beta_1, \beta_2]$ and agrees with \mathbf{b}^* outside that neighborhood. First, define

$$\beta_1^\eta = \inf\{\beta \leq \beta_1 : \mathbf{b}^*(\beta) \geq b + \eta(\beta - \bar{\beta})\}$$

and

$$\beta_2^\eta = \sup\{\beta \geq \beta_2 : \mathbf{b}^*(\beta) \leq b + \eta(\beta - \bar{\beta})\}.$$

Note that as $\eta \rightarrow 0$, $\beta_j^\eta \rightarrow \beta_j$ for $j = 1, 2$. Finally, define

$$\mathbf{b}^\eta(\beta) \equiv \begin{cases} \mathbf{b}^*(\beta) & \text{if } \beta > \beta_2^\eta \\ b + \eta(\beta - \bar{\beta}) & \text{if } \beta \in [\beta_1^\eta, \beta_2^\eta] \\ \mathbf{b}^*(\beta) & \text{if } \beta < \beta_1^\eta. \end{cases}$$

The difference in payoffs to the buyer under \mathbf{b}^η and under \mathbf{b}^* is given by

$$\int_{\beta_1^\eta}^{\beta_2^\eta} h_B(\mathbf{b}^\eta(\beta), \mathbf{s}_B^*(\beta)) - h_B(\mathbf{b}^*(\beta), \mathbf{s}_B^*(\beta)) - [c_B(\mathbf{b}^\eta(\beta), \beta) - c_B(\mathbf{b}^*(\beta), \beta)] d\beta. \quad (\text{D.15})$$

Now

$$\begin{aligned} & \int_{\beta_1}^{\beta_2} [c_B(\mathbf{b}^\eta(\beta), \beta) - c_B(\mathbf{b}^*(\beta), \beta)] d\beta \\ &= \int_{\beta_1}^{\beta_2} \left[\frac{\partial c_B(b, \beta)}{\partial b} \eta(\beta - \bar{\beta}) + o(\eta) \right] d\beta \\ &= \eta \int_0^{(\beta_2 - \beta_1)/2} \left[\frac{\partial c_B(b, \bar{\beta} + x)}{\partial b} - \frac{\partial c_B(b, \bar{\beta} - x)}{\partial b} \right] x dx + o(\eta). \end{aligned}$$

From Assumption 3, the integrand is strictly negative, and so the integral is strictly negative and independent of η . Since \mathbf{s}_B^* is increasing, a similar argument applied to the difference in the premuneration values shows that

$$\begin{aligned} & \int_{\beta_1}^{\beta_2} h_B(\mathbf{b}^\eta(\beta), \mathbf{s}_B^*(\beta)) - h_B(\mathbf{b}^*(\beta), \mathbf{s}_B^*(\beta)) - [c_B(\mathbf{b}^\eta(\beta), \beta) - c_B(\mathbf{b}^*(\beta), \beta)] d\beta \\ & \geq \eta \int_0^{(\beta_2 - \beta_1)/2} \left[\frac{\partial c_B(b, \bar{\beta} - x)}{\partial b} - \frac{\partial c_B(b, \bar{\beta} + x)}{\partial b} \right] x d\beta + o(\eta). \end{aligned}$$

It remains to argue that the contribution to (D.15) from the intervals $[\beta_1^\eta, \beta_1)$ and $(\beta_2, \beta_2^\eta]$ is of $o(\eta)$. But this is immediate, since $|\mathbf{b}^\eta(\beta) - \mathbf{b}^*(\beta)| \leq \eta$ and $\beta_j^\eta \rightarrow \beta_j$ as $\eta \rightarrow 0$ (for $j = 1, 2$). Hence, for $\eta > 0$ sufficiently small, \mathbf{b}^η gives the buyer a strictly higher payoff under (D.4) than \mathbf{b}^* . But then by a now familiar argument, the buyer has a profitable deviation in Γ^n for sufficiently large n , a contradiction. So \mathbf{b}^* is strictly increasing when it is nonzero. \square

We next show that the seller's payoffs converge to the payoff one would expect the seller to receive by matching with his corresponding buyer type.

LEMMA D.8. For almost all σ satisfying $\mathbf{b}^*(\sigma) > 0$,

$$\begin{aligned} \lim_n \int \bar{H}(\mathbf{s}(\sigma), \mathbf{b}, p_U^n) - c_S(\mathbf{s}(\sigma), \sigma) d\xi^n \\ = h_S(\mathbf{b}^*(\sigma), \mathbf{s}^*(\sigma)) + p_U^*(\mathbf{s}^*(\sigma)) - c_S(\mathbf{s}^*(\sigma), \sigma). \end{aligned}$$

The functions \mathbf{s} and \mathbf{b} on the left side of this expression are strategies in the game Γ^n , and are the objects over which the equilibrium ξ^n mixes.

PROOF OF LEMMA D.8. Suppose the claim is false. Then, since the limit exists, there exists n'' and $\eta > 0$ such that for all σ in a set G of sellers of measure at least η whose “matched” buyers choose positive attributes (i.e., $\mathbf{b}^*(\sigma) > 0$), for all $n > n''$,

$$\int \overline{H}(\mathbf{s}(\sigma), \mathbf{b}, p_U^n) - c_S(\mathbf{s}(\sigma), \sigma) d\xi^n$$

is at least η distant from

$$h_S(\mathbf{b}^*(\sigma), \mathbf{s}^*(\sigma)) + p_U^*(\mathbf{s}^*(\sigma)) - c_S(\mathbf{s}^*(\sigma), \sigma).$$

Since G has positive measure, we can assume that every index in G is a continuity point of the limit functions $(\mathbf{b}^*, \mathbf{s}_B^*, \mathbf{s}^*)$.

For any $\varepsilon > 0$, let $E^\varepsilon \subset [0, 1]$ be the set from Lemma D.5 that satisfies $\lambda(E^\varepsilon) \geq 1 - \varepsilon$.

Fix an index $\sigma' \in G \cap E^{\varepsilon'}$ for some $\varepsilon' > 0$ (since E^ε is monotonic in ε , $\sigma' \in G \cap E^\varepsilon$ for all smaller ε). Since \mathbf{b}^* is strictly increasing, without loss of generality, we can assume that, for all $\zeta > 0$, there is a positive measure set of buyers with $\mathbf{b}^*(\beta) \in (\mathbf{b}^*(\sigma') - \zeta, \mathbf{b}^*(\sigma'))$. Indeed, a positive measure set of buyers in E^ε does so for all ε sufficiently small. Formally,

$$\forall \zeta > 0 \exists \varepsilon'' \forall \varepsilon < \varepsilon'', \quad \lambda\{\beta \in E^\varepsilon : \mathbf{b}^*(\beta) \in (\mathbf{b}^*(\sigma') - \zeta, \mathbf{b}^*(\sigma'))\} > 0. \tag{D.16}$$

Consider some $\varepsilon < \varepsilon'$ and suppose $n > \max\{n_\varepsilon, n''\}$, where n_ε is from Lemma D.5. Let $(\mathbf{b}, \mathbf{s}_B, \mathbf{s}) \in Y_B^n \times Y_S^n$ be a triple of functions with the property that $|\mathbf{b}(\beta) - \mathbf{b}^*(\beta)| < \varepsilon$ and $|\mathbf{s}_B(\beta) - \mathbf{s}_B^*(\beta)| < \varepsilon$ for all $\beta \in E^\varepsilon$, and $|\mathbf{s}(\sigma) - \mathbf{s}^*(\sigma)| < \varepsilon$ for all $\sigma \in E^\varepsilon$. (Recall from Lemma D.5 that ξ^n assigns high probability to such functions for large n .)

By Lemma D.6, \mathbf{s}^* and \mathbf{s}_B^* are equal almost surely, so without loss of generality, we may assume that $\mathbf{s}^*(x) = \mathbf{s}_B^*(x)$ for all $x \in E$.

Observe first that if the max in (D.8) is achieved by $h_S(0, \mathbf{s}(\sigma)) + p_U^n(\mathbf{s}(\sigma))$, then

$$\begin{aligned} \overline{H}(\mathbf{s}(\sigma), \mathbf{b}, p_U^n) - c_S(\mathbf{s}(\sigma), \sigma) &= h_S(0, \mathbf{s}(\sigma)) + p_U^n(\mathbf{s}(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma) \\ &\leq h_S(\mathbf{b}(\sigma), \mathbf{s}(\sigma)) + p_U^n(\mathbf{s}(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma). \end{aligned}$$

We claim that for sufficiently small $\varepsilon > 0$, the set $\tilde{P}(\mathbf{s}(\sigma), \mathbf{b}, p_U^n)$ contains all $p < p_U^n(\mathbf{s}(\sigma))$. This follows from (D.16) and the observation that buyers in E^ε receive a payoff (ignoring costs) that is arbitrarily close to $h_B(\mathbf{b}^*(\beta), \mathbf{s}^*(\beta)) - p_U^*(\mathbf{s}^*(\beta))$.

Consequently, for p sufficiently close to $p_U^n(\mathbf{s}(\sigma))$, single crossing (Assumption 1) implies that a buyer β with an attribute that satisfies $\mathbf{b}^*(\beta) < \mathbf{b}^*(\sigma)$ will not be attracted (for sufficiently large n). This implies that

$$\sup_{p \in \tilde{P}(\mathbf{s}(\sigma), \mathbf{b}, p_U^n)} H(\mathbf{s}(\sigma), p, \mathbf{b}, p_U^n) = h_S(\mathbf{b}(\sigma), \mathbf{s}(\sigma)) + p_U^n(\mathbf{s}(\sigma)).$$

By choosing ε small (or, equivalently, n large), the right side can be made arbitrarily close to

$$h_S(b_1(s'), s') + p_U^*(s') = h_S(b_1(\mathbf{s}^*(\sigma)), \mathbf{s}^*(\sigma)) + p_U^*(\mathbf{s}^*(\sigma)).$$

Hence, the max in (D.8) is achieved by the first term and we have a contradiction. \square

With this payoff characterization in hand, we can show that seller attribute choices are strictly increasing in types (when positive), as are the types of sellers with whom buyers attempt to match.

LEMMA D.9. *The functions \mathbf{s}_B^* and \mathbf{s}^* are strictly increasing on $\{\beta : \mathbf{b}^*(\beta) > 0\}$.*

PROOF. From Lemma D.6, $\mathbf{s}_B^*(x) = \mathbf{s}^*(x)$ for almost all $x \in [0, 1]$, and so it suffices to prove the result for \mathbf{s}^* . Suppose, to the contrary, there is a strictly positive constant \hat{s} and an associated maximal nondegenerate interval (σ_1, σ_2) with $\mathbf{s}^*(\sigma) = \hat{s}$ and $\mathbf{b}^*(\sigma) > 0$ for all $\sigma \in (\sigma_1, \sigma_2)$. From Lemma D.6, we also have $\mathbf{s}_B^*(\beta) = \hat{s}$ for all $\beta \in (\sigma_1, \sigma_2)$.

Define $b_1 \equiv \lim_{\beta \downarrow \sigma_1} \mathbf{b}^*(\beta)$ and $b_2 \equiv \lim_{\beta \uparrow \sigma_2} \mathbf{b}^*(\beta)$.

Define $\sigma(\eta) \equiv \inf\{\sigma : \mathbf{s}^*(\sigma) \geq \hat{s} + \eta\}$ and notice that $\lim_{\eta \rightarrow 0} \sigma(\eta) = \sigma_2$. The seller attribute-choice function \mathbf{s}' given by

$$\mathbf{s}'(s) = \begin{cases} \mathbf{s}^*(\sigma) & \text{if } \sigma \notin (\sigma_1, \sigma(\eta)) \\ \hat{s} + \eta & \text{if } \sigma \in (\sigma_1, \sigma(\eta)) \end{cases}$$

is weakly increasing. Consider the price $\hat{p} > p_U^*(\hat{s})$ for attribute $\hat{s} + \eta$ that satisfies

$$\hat{p} = \sup\{p : B(\hat{s} + \eta, p, p_U^*) \neq \emptyset\}.$$

(This is $\bar{p}(\hat{s} + \eta, p, p_U^*)$ from Lemma D.1(i).) The price \hat{p} is at least as high as the value p' that satisfies $h_B(b_2, \hat{s}) - p_U(\hat{s}) = h_B(b_2, \hat{s} + \eta) - p'$. At the price \hat{p} for attribute choice $\hat{s} + \eta$, the seller ensures that attribute choice $\hat{s} + \eta$ is chosen by a buyer at least as high as b_2 (the single-crossing condition on buyer premuneration values ensures that no lower-attribute buyers choose $\hat{s} + \eta$). From Lemma D.8, we then have that the switch to attribute-choice function \mathbf{s}' increases the seller's payoff by at least

$$\begin{aligned} & \int_{\sigma_1}^{\sigma_2} (h_S(b_2, \hat{s} + \eta) + \hat{p}) d\sigma - \int_{\sigma_1}^{\sigma_2} (h_S(b_1, \hat{s}) + p_U^*(\hat{s})) d\sigma \\ & \quad - \int_{\sigma_1}^{\sigma(\eta)} (c_S(\hat{s} + \eta, \sigma) - c_S(\mathbf{s}^*(\sigma), \sigma)) d\sigma \\ & > (\sigma_2 - \sigma_1)[h_S(b_2, \hat{s} + \eta) - h_S(b_1, \hat{s})] \\ & \quad - (\sigma(\eta) - \sigma_1)[c_S(\hat{s} + \eta, \sigma_1) - c_S(\hat{s}, \sigma_1)]. \end{aligned}$$

The first term after the inequality is bounded away from zero as η approaches zero, while the second term approaches zero as does η , ensuring that there is some $\eta > 0$ for which

the payoff difference is positive. Intuitively, each seller in the interval (σ_1, σ_2) experiences a discontinuous increase in expected buyer (at a higher price) when increasing her attribute choice, while sellers in the interval $(\sigma_2, \sigma(\eta))$ experience a continuous increase in cost. The attribute-choice function \mathbf{s}' increases the seller's payoff for sufficiently small η , yielding the result. \square

The limiting masses of buyers and seller who choose zero attributes are equal.

LEMMA D.10.

$$\lambda(\{\sigma : \mathbf{s}^*(\sigma) = 0\}) = \lambda(\{\beta : \mathbf{b}^*(\beta) = 0\}).$$

PROOF. First, suppose $\lambda(\{\beta : \mathbf{b}^*(\beta) = 0\}) > \lambda(\{\sigma : \mathbf{s}^*(\sigma) = 0\})$. Then because $\mathbf{s}_B^* = \mathbf{s}^*$ almost everywhere, there exists a positive mass of buyers for whom $\mathbf{b}^*(\beta) = 0$ and $\mathbf{s}_B^*(\beta) > 0$. By Assumption 2, $h_B(0, s)$ is independent of s , and so since p_U^* is strictly increasing, the buyers choosing $b = 0$ can increase their payoff by choosing $s = 0$. The buyer's equilibrium strategy must then be suboptimal in the game Γ^n for sufficiently large n , a contradiction.

Now, suppose $\lambda(\{\beta : \mathbf{b}^*(\beta) = 0\}) < \lambda(\{\sigma : \mathbf{s}^*(\sigma) = 0\})$. Then there exists a positive mass of buyers for whom $\mathbf{b}^*(\beta) > 0$ and $\mathbf{s}_B^*(\beta) = 0$. By Lemma D.6, there is then a positive mass of buyers who choose a zero seller attribute and positive buyer attribute. Since $h_B(b, 0)$ is independent of b (Assumption 2) and $c_B(b, \beta)$ is strictly increasing in b , such buyers can increase their payoff by choosing $b = 0$. The buyer's equilibrium strategy must then be suboptimal in the game Γ^n for sufficiently large n , a contradiction. \square

We now turn to feasible matchings. For $b \in [0, \mathbf{b}^*(1)]$ and $s \in [0, \mathbf{s}^*(1)]$, we define

$$\tilde{b}^*(s) \equiv \begin{cases} \mathbf{b}^*((\mathbf{s}^*)^{-1}(s)) & s \in \mathbf{s}^*([0, 1]), s > 0 \\ \max\{0, \sup_{b \in \mathcal{B}}\{b < \mathbf{b}^*(\inf\{\sigma : \mathbf{s}^*(\sigma) > s\})\}\} & \text{otherwise} \end{cases}$$

and

$$\tilde{s}^*(b) \equiv \begin{cases} \mathbf{s}^*((\mathbf{b}^*)^{-1}(b)) & b \in \mathbf{b}^*([0, 1]), b > 0 \\ \max\{0, \sup_{s \in \mathcal{S}}\{s < \mathbf{s}^*(\inf\{\beta : \mathbf{b}^*(\beta) > b\})\}\} & \text{otherwise.} \end{cases}$$

The maximum in the specification of \tilde{b}^* (with \tilde{s}^* similar) ensures that \tilde{b}^* is well defined when \mathbf{s}^* is continuous at $\inf\{\sigma | \mathbf{s}^*(\sigma) > 0\}$ (in which case, the supremum is taken over the empty set and so has value $-\infty$).

LEMMA D.11. *The pair $(\tilde{b}^*, \tilde{s}^*)$ is a feasible matching. In addition, for all values $b > 0$ and $s > 0$, we have*

$$\tilde{s}^*(b) = \mathbf{s}^*((\mathbf{b}^*)^{-1}(b)),$$

where

$$(\mathbf{b}^*)^{-1}(b) = \begin{cases} \inf\{\beta : \mathbf{b}^*(\beta) > b\} & \text{for } b \leq \mathbf{b}^*(1) \\ 1 & \text{for } b > \mathbf{b}^*(1) \end{cases}$$

and have

$$\tilde{b}^*(s) = \mathbf{b}^*((\mathbf{s}^*)^{-1}(s)),$$

where

$$(\mathbf{s}^*)^{-1}(s) = \begin{cases} \inf\{\sigma : \mathbf{s}^*(\sigma) > s\} & \text{for } s \leq \mathbf{s}^*(1) \\ 1 & \text{for } s > \mathbf{s}^*(1). \end{cases}$$

PROOF. From Lemma D.10, we can assume that \mathbf{b}^* and \mathbf{s}^* share a common set $[0, x]$ on which they are zero. It is then immediate that $(\tilde{b}^*, \tilde{s}^*)$ is a feasible matching.

The final two statements follow immediately from the left continuity of the attribute-choice functions (see Lemma D.4) and the definitions of \tilde{s}^* and \tilde{b}^* . \square

Finally, we show that the seller’s payoff satisfies an optimality condition.

LEMMA D.12. For almost all σ ,

$$\begin{aligned} \lim_n \int \bar{H}(\mathbf{s}(\sigma), \mathbf{b}, p_U^n) - c_S(\mathbf{s}(\sigma), \sigma) d\xi^n & \\ &= h_S(\mathbf{b}^*(\sigma), \mathbf{s}^*(\sigma)) + p_U^*(\mathbf{s}^*(\sigma)) - c_S(\mathbf{s}^*(\sigma), \sigma) \\ &= \max_{s \in S} h_S(\tilde{b}^*(s), s) + p_U^*(s) - c_S(s, \sigma). \end{aligned}$$

PROOF. The first equality duplicates Lemma D.8.

Single crossing (Assumption 3) implies that the attribute choices that maximize $h_S(\tilde{b}^*(s), s) + p_U^*(s) - c_S(s, \sigma)$ are increasing in σ , and so if the second equality fails, in games Γ^n for sufficiently large n , the seller has a profitable deviation. \square

D.5 Uniform-price equilibria

We finally argue that the profile $(\mathbf{b}^*, \mathbf{s}^*, \tilde{b}^*, \tilde{s}^*, p_U^*)$ induces a uniform-price equilibrium of the matching market with identical attribute choices and matching function (but perhaps a vertical shift in the price function).

The first task is to show that equilibrium payoffs are nonnegative, so that agents would not prefer to be out of the market. Suppose $\{\xi_B^n, \xi_S^n, p_U^n\}_n$ is the sequence whose limit induces $(\mathbf{b}^*, \mathbf{s}_B^*, \tilde{b}^*, \tilde{s}^*, p_U^*)$. We have

$$\begin{aligned} h_B(0, 0) - p_U^*(0) &= h_B(0, 0) - p_U^*(0) - c_B(0, \beta) \\ &\leq h_B(\mathbf{b}(\beta), \tilde{s}^*(\mathbf{b}^*(\beta))) - p_U^*(\tilde{s}^*(\mathbf{b}^*(\beta))) - c_B(\mathbf{b}(\beta), \beta) \end{aligned} \tag{D.17}$$

and

$$\begin{aligned} h_S(0, 0) + p_U^*(0) &= h_S(0, 0) + p_U^*(0) - c_S(0, \sigma) \\ &\leq h_S(\tilde{b}^*(\mathbf{s}(\sigma)), \mathbf{s}(\sigma)) + p_U^*(\mathbf{s}(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma). \end{aligned} \tag{D.18}$$

Let

$$\kappa^* \equiv h_B(0, 0) - p_U^*(0) \geq -h_S(0, 0) - p_U^*(0)$$

(where the inequality follows from [Assumption 2](#)) and replace the price function p_U^* with $p_U^* + \kappa^*$. Both ξ_B^n and ξ_S^n remain best responses given price $p_U^n + \kappa^*$ and markets still clear in the limit of $n \rightarrow \infty$. Moreover, replacing p_U^* with $p_U^* + \kappa^*$ in [\(D.17\)](#) and [\(D.18\)](#) gives nonnegative payoffs.

It is immediate from the formulation of the buyer’s payoffs in the game and from [Lemma D.12](#) that almost all buyers and sellers are optimizing given p_U^* .

It remains to consider deviations by a seller of type σ to a value s not chosen by any seller under \mathbf{s}^* . If there is a profitable such deviation for seller σ , then there is a price p such that $B(s, p, p_U^*)$ is nonempty and for all $b \in B(s, p, p_U^*)$,

$$\Pi_S(\mathbf{s}^*(\sigma), \sigma) < h_S(b, s) + p - c_S(s, \sigma).$$

But then for all sufficiently large n , $B(s, p', p_U^n)$ is again nonempty for p' less than but close to p , contradicting the fact that $\Pi_S(\mathbf{s}^*(\sigma), \sigma)$ is close to $\int \bar{H}(\mathbf{s}(\sigma), \mathbf{b}, p_U^n) - c_S(\mathbf{s}(\sigma), \sigma) d\xi^n$.

D.6 Nontriviality

Partial nontriviality We now show that under [\(13\)](#), the profile $(\mathbf{b}^*, \mathbf{s}^*, \tilde{b}^*, \tilde{s}^*, p_U^*)$ is non-trivial. If the equilibrium is trivial, \mathbf{b}^* and \mathbf{s}^* are identically zero, so there is no agent for whom it is profitable to trade at price p_U ; hence, for all $(b, s) \in (0, \bar{b}] \times (0, \bar{s}]$,

$$h_S(0, s) + p_U(s) - c_S(s, 1) \leq 0$$

and

$$h_B(b, s) - p_U(s) - c_B(b, 1) \leq 0,$$

where we focus on agents $\beta = 1 = \sigma$ since they are the most likely to want to trade. Notice that here we are using the maximum that appears in the building block [\(D.8\)](#) for the specification of the seller’s payoff, and this effectively allows the seller to sell any attribute choice $s \in [0, \bar{s}]$ at price $p_U(s)$, assuming in the process that he can attract at least a zero-attribute buyer. For these two inequalities to hold, it must be that

$$h_B(b, s) + h_S(0, s) \leq c_B(b, 1) + c_S(s, 1),$$

contradicting [\(13\)](#).

Full nontriviality We now assume [\(14\)](#) holds. Suppose that there is an interval of seller types $[0, \sigma']$ with $\sigma' > 0$ who choose zero attributes. By [Lemma D.9](#), we then have $\mathbf{b}^*(\beta) = 0$ for all $\beta \in [0, \sigma']$. If neither agent of type $\phi \in (0, \beta')$ chooses a strictly positive attribute, it must be that

$$h_S(0, s) + p_U(s) - c_S(s, \phi) \leq 0$$

and

$$h_B(b, s) - p_U(s) - c_B(b, \phi) \leq 0,$$

where (b, s) are a pair of attributes that satisfy (14). But summing these two inequalities yields an inequality that contradicts (14).

APPENDIX E: PERSONALIZED PRICING

E.1 Prices

A *personalized-price function* is a function $p_P : \mathcal{B} \times \mathcal{S} \rightarrow \mathbb{R}$, where $p_P(b, s)$ is the (possibly negative) price that a seller with attribute choice $s \in \mathcal{S}$ receives when selling to a buyer with attribute choice $b \in \mathcal{B}$. We emphasize that a personalized-price function prices only matches between *marketed* attributes.

E.2 Equilibrium

Given a feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ and a personalized price p_P , the payoffs to a buyer β who chooses $b \in \mathcal{B}$ and to a seller σ who chooses $s \in \mathcal{S}$ are given by

$$\Pi_B(b, \beta) \equiv h_B(b, \tilde{s}(b)) - p_P(b, \tilde{s}(b)) - c_B(b, \beta)$$

and

$$\Pi_S(s, \sigma) \equiv h_S(\tilde{b}(s), s) + p_P(\tilde{b}(s), s) - c_S(s, \sigma).$$

The specification of equilibrium begins with appropriate modifications of the notions of buyer and seller optimization given p_U (i.e., (3) and (4)).

DEFINITION E.1. Given a feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$, buyer β is *optimizing at \mathbf{b} given p_P* if

$$(\mathbf{b}(\beta), \tilde{s}(\mathbf{b}(\beta))) \in \arg \max_{(b,s) \in \mathcal{B} \times \mathcal{S}} h_B(b, s) - p_P(b, s) - c_B(b, \beta). \tag{E.1}$$

Seller σ is *optimizing at \mathbf{s} given p_P* if

$$(\tilde{b}(\mathbf{s}(\sigma)), \mathbf{s}(\sigma)) \in \arg \max_{(b,s) \in \mathcal{B} \times \mathcal{S}} h_S(b, s) + p_P(b, s) - c_S(s, \sigma). \tag{E.2}$$

Since the personalized-price function p_P prices only pairs of marketed attributes, the stipulation that a seller must optimize given p_P says nothing about what might happen if this seller chooses an attribute $s \notin \mathcal{S}$. We require that no seller can choose an attribute $s \notin \mathcal{S}$, find a target buyer attribute $b \in \mathcal{B}$ with whom to match, and find a way to split the resulting surplus so that the seller and the target buyer are both better off than in equilibrium.²³

²³Similar to footnote 7, Definitions E.2 and E.3 can be extended to cover deviations to any seller attribute as well as to prices that differ from the personalized-price function. See Mailath et al. (2010) for details.

DEFINITION E.2. Given $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$, there is a *profitable seller deviation* if there exists a seller σ such that either (i) $\Pi_S(\mathbf{s}(\sigma), \sigma) < 0$ or (ii) there exist an unmarketed seller attribute choice $s \notin \mathcal{S}$, a marketed buyer attribute $b \in \mathcal{B}$, and a price $p \in \mathbb{R}$ such that

$$h_B(b, \tilde{s}(b)) - p_P(b, \tilde{s}(b)) < h_B(b, s) - p$$

and

$$\Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(b, s) + p - c_S(s, \sigma).$$

The definition of a buyer's profitable deviation is similar.

DEFINITION E.3. Given $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$, there is a *profitable buyer deviation* if there exists a buyer β such that either (i) $\Pi_B(\mathbf{b}(\beta), \beta) < 0$ or (ii) there exist an unmarketed buyer attribute choice $b \notin \mathcal{B}$, a marketed seller attribute $s \in \mathcal{S}$, and a price $p \in \mathbb{R}$ such that

$$\Pi_B(\mathbf{b}(\beta), \beta) < h_B(b, s) - p - c_B(b, \beta)$$

and

$$h_S(\tilde{b}(s), s) + p_P(\tilde{b}(s), s) < h_S(b, s) + p.$$

DEFINITION E.4. A feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ and a personalized-price function p_P constitute a *personalized-price equilibrium* if all buyers and sellers are optimizing given p_P and no buyer or seller has a profitable deviation.

REMARK E.1 (Premuneration values). Since personalized prices can compensate for any alterations of the division of $v(b, s)$, the decomposition of the surplus $v(b, s)$ between the buyer's and seller's premuneration values plays no role in the efficiency of a personalized-price equilibrium outcome. In particular, it is a straightforward calculation that if $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$ is a personalized-price equilibrium with premuneration values $h_B(b, s)$ and $h_S(b, s)$, then $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p'_P)$ is a personalized-price equilibrium with premuneration values $h'_B(b, s)$ and $h'_S(b, s)$, where

$$p'_P(b, s) = p_P(b, s) + h'_B(b, s) - h_B(b, s) = p_P(b, s) + h_S(b, s) - h'_S(b, s).$$

REMARK E.2 (Ex post contracting equilibrium). Cole et al. (2001) study continua of buyers and sellers who first simultaneously choose attributes (as here), and then match and bargain to divide the resulting surplus $v(b, s)$, with the matching–bargaining stage being modeled as a cooperative game (more specifically, an assignment game). An *ex post contracting equilibrium* in Cole et al. (2001) is a Nash equilibrium of the noncooperative attribute-choice game, where the payoffs from the attribute choices are determined by stable (equivalently, core) allocations in the induced assignment game.²⁴ The set of outcomes and implied payoffs is essentially the same under the two notions (modulo

²⁴Consequently, matching is over buyers and sellers, not attributes as here. This difference results in some technical complications.

some technical differences). In particular, if all buyers and sellers are optimizing in the market given p_P , then no buyer–seller pair with attributes $(b, s) \in \mathcal{B} \times \mathcal{S}$ can block the equilibrium. Moreover, a seller σ has a profitable out-of-market deviation if and only if there is a blocking pair consisting of that seller (with some attribute s) and some buyer with an attribute $b \in \mathcal{B}$. An analogous comment applies to buyers.

E.3 Efficiency

LEMMA E.1. *In any personalized-price equilibrium $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$, \tilde{b} and \tilde{s} are strictly increasing for strictly positive attributes.*

PROOF. We consider only \tilde{b} (since \tilde{s} is almost identical). Suppose \tilde{b} is not strictly increasing. Since \tilde{b} is one-to-one on $\mathbf{s}((\underline{\sigma}, 1])$ (see Definition 3.1 and its following comment), there exists $0 < s_1 < s_2$ with $b_1 \equiv \tilde{b}(s_1) > \tilde{b}(s_2) \equiv b_2$. From (E.1) for the buyer choosing b_1 and from (E.2) for the seller choosing s_2 , we have

$$h_B(b_1, s_1) - p_P(b_1, s_1) \geq h_B(b_1, s_2) - p_P(b_1, s_2)$$

and

$$h_S(b_2, s_2) + p_P(b_2, s_2) \geq h_S(b_1, s_2) + p_P(b_1, s_2),$$

and so

$$h_B(b_1, s_1) + h_S(b_2, s_2) - p_P(b_1, s_1) + p_P(b_2, s_2) \geq v(b_1, s_2).$$

Adding this to the analogous inequality obtained from (E.1) for the buyer choosing b_2 and from (E.2) for the seller choosing s_1 , we obtain

$$v(b_1, s_1) + v(b_2, s_2) \geq v(b_1, s_2) + v(b_2, s_1).$$

But Assumption 1 requires the reverse (strict) inequality, a contradiction. □

From Lemma E.1, matching in a personalized-price equilibrium is positively assortative in attributes. Since the attribute-choice functions are strictly increasing in index when positive, we can accordingly define the ex ante surplus for buyer and seller types $\beta = \sigma = \phi \in [0, 1]$ as

$$\begin{aligned} W(b, s, \phi) &\equiv h_B(b, s) + h_S(b, s) - c_B(b, \phi) - c_S(s, \phi) \\ &= v(b, s) - c_B(b, \phi) - c_S(s, \phi). \end{aligned}$$

An efficient choice of attributes maximizes $W(b, s, \phi)$ for (almost) all ϕ .

Personalized-price equilibrium outcomes are constrained efficient in the sense that no matched pair of agents can increase its net surplus without both agents deviating to attribute choices outside the sets of marketed attributes \mathcal{B} and \mathcal{S} .²⁵

²⁵This is essentially Cole et al. (2001), Lemma 2, which describes a constrained efficiency property of ex post contracting equilibria (see Remark E.2). The current formulation allows a more transparent statement and proof.

LEMMA E.2. *Suppose $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$ is a personalized-price equilibrium. Then for all $\phi \in [0, 1]$, $b \in \mathcal{B}$, $s \in \mathcal{S}$, and all b' and s' ,*

$$W(b, s', \phi) \leq W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi)$$

and

$$W(b', s, \phi) \leq W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi).$$

PROOF. En route to a contradiction, suppose there exist $\phi \in [0, 1]$, $b \in \mathcal{B}$, and $s' \in [0, \bar{s}]$ such that $W(b, s', \phi) > W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi)$. The other possibility is handled analogously.

Let $\varepsilon = [W(b, s', \phi) - W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi)]/3 > 0$ and set $p = h_B(b, s') - h_B(b, \tilde{s}(b)) + p_P(b, \tilde{s}(b)) - \varepsilon$. The seller of type $\sigma = \phi$ can induce a buyer with attribute choice b to buy from him by choosing s' and offering a price p . Moreover, this deviation is strictly preferred by the seller ϕ ,

$$\begin{aligned} & h_S(b, s') + p - c_S(s', \phi) \\ &= h_S(b, s') + h_B(b, s') - h_B(b, \tilde{s}(b)) + p_P(b, \tilde{s}(b)) - \varepsilon - c_S(s', \phi) \\ &> \Pi_S(\mathbf{s}(\phi), \phi) + [h_B(\mathbf{b}(\phi), \mathbf{s}(\phi)) - p_P(\mathbf{b}(\phi), \mathbf{s}(\phi)) - c_B(\mathbf{b}(\phi), \phi)] \\ &\quad - [h_B(b, \tilde{s}(b)) - p_P(b, \tilde{s}(b)) - c_B(b, \phi)] + \varepsilon \\ &\geq \Pi_S(\mathbf{s}(\phi), \phi) + \varepsilon, \end{aligned}$$

where the equality uses the definition of p , the strict inequality follows from $W(b, s', \phi) > W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi) + 2\varepsilon$, and the last inequality is an implication of (E.1). \square

Lemma E.2 does not ensure that a personalized-price equilibrium outcome is efficient. The possibility remains that $W(b, s, \phi)$ may be maximized by a pair of values $b \notin \mathcal{B}$ and $s \notin \mathcal{S}$. In this sense, the inefficiency is the result of a coordination failure. For example, for the premuneration values $h_B(b, s) = \theta bs$ and $h_S(b, s) = (1 - \theta)bs$, it is an equilibrium for all agents to choose attribute 0, giving a constrained-efficient outcome that is, in fact, quite inefficient. The possible inefficiency of a uniform-price equilibrium can be viewed as reflecting incomplete markets.

We could ensure efficiency by ensuring that a price exists for every attribute combination, whether marketed or not.

DEFINITION E.5. The feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ and personalized price p_P is a *complete personalized-price equilibrium* if there is an extension of p_P to $[0, \bar{b}] \times [0, \bar{s}]$ (also denoted by p_P) such that for all β and all σ ,

$$0 \leq \Pi_B(\mathbf{b}(\beta), \beta) = \sup_{(b,s) \in [0, \bar{b}] \times [0, \bar{s}]} h_B(b, s) - p_P(b, s) - c_B(b, \beta)$$

and

$$0 \leq \Pi_S(\mathbf{s}(\sigma), \sigma) = \sup_{(b,s) \in [0, \bar{b}] \times [0, \bar{s}]} h_S(b, s) + p_P(b, s) - c_S(s, \sigma).$$

Though the names suggest that every complete personalized-price equilibrium outcome is indeed a personalized-price equilibrium outcome, this is not immediate, as we have replaced the prohibition on profitable deviations with the requirement that agents be optimizing given p_P with respect to all attribute choices. Nonetheless, we have the following result.

LEMMA E.3. (i) *Every complete personalized-price equilibrium outcome is a personalized-price equilibrium outcome.*

(ii) *A complete personalized-price equilibrium outcome is efficient.*

PROOF. Fix a complete personalized-price equilibrium $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$. To show that this is a personalized-price equilibrium, we must show there are no profitable deviations. We discuss seller deviations; the buyer case is analogous. Suppose the seller has a profitable deviation, so there exists a type σ and an attribute choice $s' \notin \mathcal{S}$, a price $p \in \mathbb{R}$, and $b' \in \mathcal{B}$ with

$$\Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(b', s') + p - c_S(s', \sigma) \tag{E.3}$$

and

$$h_B(b', \tilde{s}(b')) - p_P(b', \tilde{s}(b')) < h_B(b', s') - p. \tag{E.4}$$

Since $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$ is a complete personalized-price equilibrium, (E.3) implies $p > p_P(b', s')$.

There exists some $\beta \in [0, 1]$ for which $b' = \mathbf{b}(\beta)$, and so subtracting $c_B(b', \beta)$ from both sides of (E.4) and again using the assumption that $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$ is a complete personalized-price equilibrium gives $p < p_P(b', s')$, a contradiction.

Since every pair of attributes is priced, the efficiency of complete personalized-price equilibria is a straightforward calculation. □

One route to existence is to note that a personalized-price equilibrium is essentially equivalent to Cole et al.'s (2001) ex post contracting equilibrium, and then to refer to that paper for conditions for the existence of an ex post contracting equilibria. We take an alternative route here, building on the relationship between personalized-price and uniform-price equilibria.

PROPOSITION E.1. *There exists an efficient personalized-price equilibrium.*

PROOF. Suppose first that $h_S(b, s) = 0$ and, hence, $h_B(b, s) = v(b, s)$ for all pairs (b, s) . Proposition 5.1 ensures that there exists a complete uniform-price equilibrium. Since $h_S(b, s) = 0$, this is also a complete personalized-price equilibrium. Then if $h_S(b, s) \neq 0$, by setting

$$p'_P(b, s) = p_P(b, s) - h_S(b, s) = p_P(b, s) + h_B(b, s) - v(b, s),$$

we again have a complete (and, hence, efficient) personalized-price equilibrium. □

E.4 *Uniform-rationing equilibria*

LEMMA E.4. *Any personalized-price equilibrium outcome is a uniform-rationing equilibrium outcome.*

PROOF. Let $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ be a personalized-price equilibrium outcome and consider its associated uniform-rationing price. The conditions for the latter to be a personalized-price equilibrium are implied by the former, with the exception that there may now be profitable deviations by a buyer β with attribute choice $\mathbf{b}(\beta)$ to match with a seller with $s < \tilde{s}(\mathbf{b}(\beta))$ (and, hence, $\tilde{b}(s) < \mathbf{b}(\beta)$). But since $h_S(b, s)$ is increasing in b , the seller in question would welcome such a match. Hence, if this match is a profitable deviation in the uniform-rationing equilibrium, it is a profitable deviation in the personalized-price equilibrium, a contradiction. \square

REFERENCES

- Acemoglu, Daron and Robert Shimer (1999), "Holdups and efficiency with search frictions." *International Economic Review*, 40, 827–849. [539]
- Bhaskar, V. and Ed Hopkins (2011), "Marriage as a rat race: Noisy pre-marital investments with assortative matching." Discussion Paper 2011-65, Scottish Institute for Research in Economics. [539]
- Billingsley, Patrick (1986), *Probability and Measure*, second edition. Wiley, New York. [565, 567]
- Bulow, Jeremy and Jonathan Levin (2006), "Matching and price competition." *American Economic Review*, 96, 652–668. [547, 558]
- Cole, Harold L., George J. Mailath, and Andrew Postlewaite (1995), "Incorporating concern for relative wealth into economic models." *Quarterly Review (Federal Reserve Bank of Minneapolis)*, 19, 12–21. [540]
- Cole, Harold L., George J. Mailath, and Andrew Postlewaite (2001), "Efficient non-contractible investments in large economies." *Journal of Economic Theory*, 101, 333–373. [536, 537, 540, 545, 554, 564, 584, 585, 587]
- de Meza, David and Ben Lockwood (2010), "Too much investment? A problem of endogenous outside options." *Games and Economic Behavior*, 69, 503–511. [540]
- Dunford, Nelson and Jacob T. Schwartz (1988), *Linear Operators Part I: General Theory*. Wiley, New York. [565]
- Gall, Thomas, Patrick Legros, and Andrew F. Newman (2006), "The Timing of Education," *Journal of the European Economic Association*, 4, 427–435. [539]
- Gall, Thomas, Patrick Legros, and Andrew F. Newman (2009), "Mismatch, rematch, and investment." Unpublished paper. [539, 563]
- Glicksberg, Irving L. (1952), "A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points." *Proceedings of the American Mathematical Society*, 3, 170–174. [571]

- Grossman, Sanford J. (1981), "The informational role of warranties and private disclosure about product quality." *Journal of Law and Economics*, 24, 461–483. [556]
- Guerrieri, Veronica, Robert Shimer, and Randall Wright (2010), "Adverse selection in competitive search equilibrium." *Econometrica*, 78, 1823–1862. [539]
- Hopkins, Ed (2012), "Job market signalling of relative position, or Becker married to Spence." *Journal of the European Economic Association*, 10, 290–322. [540]
- Hoppe, Heidrun C., Benny Moldovanu, and Aner Sela (2009), "The theory of assortative matching based on costly signals." *Review of Economic Studies*, 76, 253–281. [540]
- Hosios, Arthur J. (1990), "On the efficiency of matching and related models of search and unemployment." *Review of Economic Studies*, 57, 279–298. [539]
- Johnstone, D. Bruce (2004), "Cost-sharing and equity in higher education: Implications of income contingent loans." In *Markets in Higher Education* (Pedro Teixeira, Ben Jongbloed, David Dill, and Alberto Amaral, eds.), 37–59, Kluwer, Dordrecht. [538]
- Mailath, George J., Andrew Postlewaite, and Larry Samuelson (2010), "Pricing in matching markets." Working Paper 10-003, Penn Institute for Economic Research. [583]
- Mailath, George J., Andrew Postlewaite, and Larry Samuelson (2012), "Premuneration values and investments in matching markets." Working Paper 12-008, Penn Institute for Economic Research. [540, 557, 558]
- Makowski, Louis and Joseph M. Ostroy (1995), "Appropriation and efficiency: A revision of the first theorem of welfare economics." *American Economic Review*, 85, 808–827. [537]
- Masters, Adrian (2011), "Commitment, advertising and efficiency of two-sided investment in competitive search equilibrium." *Journal of Economic Dynamics and Control*, 35, 1017–1031. [539]
- Milgrom, Paul R. (1981), "Good news and bad news: Representation theorems and applications." *Bell Journal of Economics*, 12, 380–391. [556]
- Milgrom, Paul R. and Ilya Segal (2002), "Envelope theorems for arbitrary choice sets." *Econometrica*, 70, 583–601. [566]
- Moen, Espen R. (1997), "Competitive search equilibrium." *Journal of Political Economy*, 105, 385–411. [539]
- Nicholson, Sean (2003), "Barriers to entering medical specialties." Working Paper 9649, National Bureau of Economic Research. [558]
- Niederle, Muriel and Alvin E. Roth (2003), "Relationship between wages and presence of a match in medical fellowships." *Journal of the American Medical Association*, 290, 1153–1154. [558]
- Niederle, Muriel and Alvin E. Roth (2005), "The gastroenterology fellowship market: Should there be a match?" *American Economic Review*, 95, 372–375. [558]

Niederle, Muriel and Leeat Yariv (2008), "Matching through decentralized markets." Unpublished paper. [539]

Okuno-Fujiwara, Masahiro, Andrew Postlewaite, and Kotaro Suzumura (1990), "Strategic information revelation." *Review of Economic Studies*, 57, 25–47. [556]

Ostrovsky, Michael and Michael Schwarz (2010), "Information disclosure and unraveling in matching markets." *American Economic Journal: Microeconomics*, 2 (2) 34–63. [556]

Peters, Michael (2010), "Noncontractible heterogeneity in directed search." *Econometrica*, 78, 1173–1200. [539, 546]

Peters, Michael and Aloysius Siow (2002), "Competing premarital investments." *Journal of Political Economy*, 110, 592–608. [536, 540]

Rege, Mari (2008), "Why do people care about social status?" *Journal of Economic Behavior and Organization*, 66, 233–242. [540]

Royden, Halsey L. (1988), *Real Analysis*, third edition. Prentice-Hall, Engelwood Cliffs, New Jersey. [569]

Shi, Shouyong (2001), "Frictional assignment. I. Efficiency." *Journal of Economic Theory*, 98, 232–260. [539]

Williamson, Oliver E. (1975), *Markets and Hierarchies*. Free Press, New York. [536]

Submitted 2012-3-12. Final version accepted 2012-8-2. Available online 2012-8-3.