

Pricing in Matching Markets*

George J. Mailath[†] Andrew Postlewaite[‡] Larry Samuelson[§]

January 26, 2010

Abstract

Different markets are cleared by different types of prices—a universal price for all buyers and sellers in some markets, seller-specific prices that are uniform across buyers in others, and personalized prices tailored to both the buyer and the seller in yet others. We introduce the notion of premuneration values—the values in the absence of any muneration (payments)—created by the buyer-seller match. We characterize the premuneration values under which uniform-price and personalized-price equilibria agree. In this case, we have efficient allocations, including pre-match investment decisions, without the costs of personalized pricing. We then examine the inefficiencies that arise when the premuneration values preclude the agreement of uniform-price and personalized-price equilibria. We view premuneration values as an important consideration in market design.

Keywords: Directed search, matching, premuneration value, pre-match investments, search.

JEL codes: C78, D40, D41, D50, D83

*We thank Antonio Penta as well as participants at numerous seminars and conferences for helpful comments. We thank the National Science Foundation (grants SES-0350969, SES-0549946, SES-0648780, and SES-0850263) for financial support.

[†]University of Pennsylvania, gmailath@econ.upenn.edu

[‡]University of Pennsylvania, apostlew@econ.upenn.edu

[§]Yale University, Larry.Samuelson@yale.edu

Contents

1	Introduction	1
2	Premuneration Values	6
3	An Example	8
3.1	Efficient Outcome	8
3.2	Personalized Pricing	9
3.3	Uniform Pricing Can Induce Efficient Outcomes	10
3.4	Uniform Pricing Need Not Induce Efficient Outcomes	11
3.5	Premuneration Values and Inefficiency	13
4	The Matching Market	16
5	Equilibrium	19
5.1	Personalized Pricing	19
5.2	Uniform Pricing	23
5.3	Allocating Buyers and Sellers	25
6	Efficiency	27
7	When is Personalization Redundant?	30
8	Existence of Equilibrium	32
8.1	Uniform-Price Equilibrium	32
8.2	Personalized-Price Equilibrium	33
9	Discussion	33
9.1	Premuneration Values	33
9.2	Endogenous Information Acquisition	34
9.2.1	Information Acquisition	34
9.2.2	An Example	35
9.3	Endogenous Information Revelation	38
A	Proof of Lemma 4	39
B	Proof of Proposition 3: Existence of Equilibrium.	40
B.1	Preliminaries	40
B.2	The game Γ^n	40
B.2.1	Strategy spaces	40
B.2.2	Buyer and Price-Setter Payoffs	41
B.2.3	Seller Payoffs	42
B.3	Equilibrium in game Γ^n	45
B.4	The limit $n \rightarrow \infty$	49
B.5	Uniform-Price Equilibria	60
B.6	Nontriviality	61

Pricing in Matching Markets

1 Introduction

Prices. Consider three people and the prices that clear their labor markets:

- Alice works in a daily spot market for casual, unskilled labor. A prototypical job pays her a fixed sum to drive a truck to pick up materials. This market is characterized by a single price, governing the transactions in each buyer/seller pair.
- Carol is a senior executive with an Ivy League degree. She receives job offers at quite different wages from various firms, each of which has made offers to others at wages different than those offered Carol.
- Bob works as a tax preparer, bolstered by a degree from his local junior college. He quotes the same hourly price to all of his clients, though some other tax preparers and accountants charge different prices.

We refer to the prices faced by Alice as *universal* prices. These are the prices typically shown in supply-and-demand diagrams in introductory texts. Carol faces *personalized* prices that depend on both her characteristics and those of her trading partner. Bob faces *uniform* prices that depend on the characteristics of the agent posting the price but not those of the agent on the other side of the transaction. Carol also faced personalized prices when purchasing her education—elite universities offer scholarships to some, while rejecting others who would pay full tuition. Bob faced uniform prices—junior colleges and technical schools typically accept all applicants at their posted prices.

Why do we see universal prices in some markets, uniform prices in others, and personalized prices in yet others? What implications does the type of pricing have for market outcomes? How are these prices linked to market characteristics? This paper addresses these questions, concentrating on uniform and personalized prices.

Premuneration values. An interaction between a buyer and seller entails a cost or benefit to each side, generating a surplus if the sum of the costs and benefits is positive. The surplus can be reallocated via a transfer from one side to the other. The *premuneration values* (from the Latin

munerare, to give or to pay) are the values to the parties *prior* to any transfer. Understanding the nature of the surplus and the premuneration values is key to understanding differences in pricing across markets.

The surplus in Alice's market is reasonably modeled as the sum of two terms, a negative premuneration value for Alice, reflecting her value of foregone leisure, and a premuneration value for her employer reflecting the benefits of having the materials delivered. Because Alice's cost is independent of the employer for whom she works, while the benefit to the employer is independent of the individual who delivers the materials, there is no issue of efficient matching in these markets. It matters that the right people transact, but does not matter who transacts with whom. It is then no surprise that a universal price clears the market.

The surpluses in Carol's markets depend in a complementary fashion on the agents on both sides of the market. Talented executives are likely to be more productive when paired with productive firms than with mediocre firms, and vice versa. Similarly, a good student fares especially well when paired with a good school while the latter is especially effective when working with good students. Clearing such markets requires not only getting the right people to transact, but also making sure that they transact with the right partners. We might then expect to need the adaptability of personalized prices.

Bob's markets also exhibit complementarity. Even below the Fortune 500 and the Ivy League, there are gains from matching skilled professionals with the right firms and good students with good schools. Then why do we see uniform prices in Bob's markets and personalized prices in Carol's?

The prices required to clear the market depend upon the status quo given by the premuneration values. Both Carol and her alma mater own some of the surplus created in the match that gave Carol her education. Carol owns her enhanced earning power, but the university owns the increment to its ranking based on her superb SAT score, the increment to its prestige should she become a Supreme Court Justice, and the increment to its endowment should she become a wealthy donor. In her employment match, Carol's employer owns the revenue her services will generate, but she owns the value of the contacts that she makes and the increase in the value of her human capital stemming from working at this firm before starting her own company. In contrast, Bob's junior college anticipates no benefit from Bob beyond his tuition, while Bob is indifferent over whose taxes he prepares, so long as the client pays.

Investments. To examine the connection between premuneration values and prices, we consider a model in which buyers and sellers invest in attributes prior to entering the market. Alice requires a commercial driver’s license, Bob an associate’s degree in accounting, and Carol an MBA. On the other side of the market, Carol’s employer makes complementary investments in capital and a client base, Bob’s clients develop their financial positions, and Alice’s employer undertakes the construction project for which her delivered materials are to be used.

Once the agents are in the market, they are matched in pairs to create surpluses that depend supermodularly on the attributes chosen by both agents. In the absence of any transfers, the division of this surplus would be determined by the agents’ premuneration values. Prices reallocate this division to generate a final allocation.

Because the agents’ attributes are complements in the creation of surplus, the efficient outcome once the agents have entered the market is straightforward—agents should be matched positively assortatively—and the frictionless matching process we examine will produce such an outcome. It is less clear that the investment incentives created by this market will lead to efficient attribute choices. We find that efficiency now hinges on whether prices are personalized or uniform, and on how the individuals’ premuneration values depend on the party with whom they match.

When prices can be personalized, there exists an equilibrium in which the resulting outcome is efficient, both in matching and ex ante investments. Moreover, the equilibrium division of the surplus in this market is *independent* of the agents’ premuneration values. It makes no difference to Carol, both in terms of her match and her payoff, whether she owns all or none of the surplus from the match. When prices are restricted to being uniform, in contrast, efficient equilibria exist if and only if the premuneration values on the side of the market setting prices are independent of the attributes of the agents on the other side of the market (as is the case in Bob’s but not Carol’s markets). Equivalently, if (and only if) this constant premuneration value condition fails, the attributes of the agents with whom Bob matches, as well as his payoff, depend critically on premuneration values, and there are efficiency gains to be had from personalizing prices.

Our result is not simply that uniform prices suffice when the *surplus* exhibits no complementarities (and hence the efficient outcome exhibits no matching problem). Instead, it is that an efficient outcome, including both investments and matching, can be supported by uniform prices even when the surplus depends supermodularly on attribute choices and hence matching is an issue, as long as the price-setter’s premuneration value does not

depend on the price-setter's match.

Why are prices important? In a world devoid of pricing frictions, personalized prices would be the norm and there would accordingly be little reason to be concerned with remuneration values.¹ But the world is not frictionless. A seller posting personalized prices must ascertain potential buyers' attributes, a process that can be quite costly. For example, estimates from 11 highly selective liberal arts colleges indicate that they spent about \$3,000 on admissions per matriculating student in 2004.² The cost for identifying whether a high school diploma comes from a legitimate high school is \$100.³ There may thus be substantial savings from posting uniform prices and letting buyers sort themselves (as Bob's clients do), if the remuneration values are such that uniform prices can do this sorting. Alternatively, if the remuneration 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.⁴

¹Our emphasis on remuneration values thus reflects no disagreement with Coase's (1960) observation that property rights would be irrelevant in a world without transactions costs.

²Expenditures for the 11 colleges, all but one of which continually appear in the *U. S. News and World Report* top 25 liberal arts colleges, were \$370 per applicant for the 1995-1996 admissions season. Publicly available data on subsequent expenditure growth rates projects an expenditure of \$625 per applicant in the 2004-2005 academic year. The 2002 admission rate for these schools was 34%. Coupling this with an estimated enrollment rate of 60% yields a cost of \$3000 per matriculating student. (Memorandum, Office of Institutional Research and Analysis, University of Pennsylvania, July 2004. We thank Bernie Lentz for his help with these data.)

³"Vetting Those Foreign College Applications," *New York Times*, September 29, 2004, page A21.

⁴For example, Bulow and Levin (2006) note that the National Residency Matching Program matching 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. However, Nicholson (2003) argues that the result is an inefficient allocation of residents to specialties. Medical students who do their residency acquire training that 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 specialties. This is as compared to their stipend, which was \$44,700 in 2007/8 (Association of American Medical College Survey of Household Stipends, Benefits and Funding, Autumn 2007 Report).

Designing markets. The remuneration values in a market can be designed as part of the institutional and legal environment of the market. For example, 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 in principle another institutional arrangement could grant them to the faculty. Indeed, 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. In a similar vein, one could arrange the remuneration 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 owns the future income of students to whom it gives degrees, but 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, 2001) attest to the possibility of such an arrangement.⁵

Our results suggest that appropriately designed remuneration values can be valuable, allowing efficient equilibria to be supported by uniform prices and hence avoiding the costs of personalized pricing or the costs of inefficient uniform pricing. Unfortunately, there are often constraints on the design of remuneration values. If universities owned 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?⁶ Such an arrangement might also require changes in labor laws that preclude involuntary

⁵Basketball star Yao Ming (Houston Rockets) has a contract with the China Basketball Association calling for 30% of his NBA earnings to be paid to the Chinese Basketball Association (in which he played prior to joining the Rockets), while another 20% will go 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. These contracts suggest that the remuneration values could be designed to assign some future earnings to the team.

⁶Measurement and collection both pose difficulties. The University of New Mexico sued a former researcher for rights to patents that he applied for four 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.

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 may 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.

2 Premuneration Values

Our model is adapted from Cole, Mailath, and Postlewaite (2001). 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, the buyer is female and seller 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, buyers and sellers simultaneously choose attributes. We denote the cost of attribute $b \in \mathbb{R}_+$ to buyer β by $c_B(b, \beta)$ and the cost of attribute $s \in \mathbb{R}_+$ to seller σ by $c_S(s, \sigma)$.

Buyers and sellers match in the second stage. Matching and the resulting division of the surplus is mediated through prices. A match between a buyer and seller with attribute choices (b, s) at a price p yields a gross buyer payoff of

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

where $h_B(b, s)$ is the *buyer premuneration value*, and a gross seller payoff of

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

where $h_S(b, s)$ is the *seller premuneration value*. The total surplus from the match is given by


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

In the simplest buyer-seller problem with exogenous attributes, the seller's premuneration value is his cost and the buyer's is her value.

Intuitively, one can think of sellers as posting prices in the second stage, after which each buyer chooses a seller, with market clearing requiring that each seller is matched with one buyer. Prices may be positive or negative. The examples in Section 1 include cases in which prices were posted by those relinquishing a good (educational services) as well as by those receiving a good (labor services), each of whom would be designated the seller in our model.

Explicitly modeling the price formation process gives rise to a range of technical issues that obfuscate the underlying economics. We make our argument more transparent by adopting a Walrasian perspective. We accordingly directly define profitable deviations for buyers and sellers, with an equilibrium being an outcome and pricing function admitting no profitable deviations. The resulting definitions of equilibrium are Walrasian in spirit.

We are interested in two types of pricing. We say that pricing is *uniform* if each seller posts a price at which *any* buyer is free to purchase, regardless of that buyer's attribute choice (though different sellers may set different prices). Prices are *personalized* if sellers can charge different prices to buyers with different attribute choices. Sellers may be forced to charge uniform prices if they cannot observe buyers' attribute choices, or if legal restrictions preclude personalization, or if personalization is prohibitively expensive.⁷ We use the information-based motivation for uniform prices when developing the model, while remembering that other factors may also lead to uniform pricing.

Remark 1 (Productive Types) The surplus generated by a match in our model depends only on the attendant attribute choices. In other cases, the surplus might depend on the agents' types as well as attribute choices. Harvard may care not only about an applicant's accomplishments (attribute choice), but also about the applicant's "cost of acquiring" such accomplishments (type). If a buyer's attribute choice and type can be summarized by a one-dimensional augmented attribute that enters the remuneration values (replacing attribute choice), then we need only treat this augmented attribute as the new "attribute choice" in order for our analysis to apply. Personalized pricing is possible if this summary variable can be observed, while our model of uniform pricing applies if it cannot. If sellers can observe attribute choices but not types, but care about both (or only about types), then attribute choices take on a dual role, directly enhancing the value of a match while also providing signals of types. Cole, Mailath, and Postlewaite (1995), Hopkins (forthcoming), and Hoppe, Moldovanu, and Sela (2009) examine such models. 

Remark 2 (Directed Search) Our paper is related to the literature on

⁷For example, employers may be prohibited from discriminating against potential employees whose weight makes them a potentially expensive health risk. Alternatively, it may be costless to use generic contract forms to offer a standard deal to every buyer who appears, while tailoring offers to buyers' characteristics may require a costly legal process.

directed search, but incorporates prematch investments on both sides of the market and heterogeneity of both firms and workers.⁸

Our personalized-price equilibrium, in which workers' investments (attributes) are observable, can be thought of as each firm posting a wage for each possible worker attribute and each worker directing her search to the firm for which matching is most profitable, resulting in a one-to-one match. Our uniform-price equilibrium treats the case of unobservable (to firms) heterogeneity in workers' attributes. In a uniform-price equilibrium, each firm sets a single price and each worker applies to the (unique) firm for which matching is most profitable, again resulting in a one-to-one match. \blacklozenge

3 An Example

Before introducing the formalities, we illustrate the economic forces at play under personalized and uniform pricing in a simple, informal example.

The remuneration values are such that a fixed share $\theta \in (0, 1]$ of the surplus goes to the buyer, so that

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

where the surplus function is given by

$$v(b, s) = bs,$$

and the cost functions by

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

3.1 Efficient Outcome

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

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

⁸See Guerrieri, Shimer, and Wright (2009) and Peters (2009) for directed search models with one-sided heterogeneity and asymmetric information.

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

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

3.2 Personalized Pricing

Suppose that sellers observe buyers' attribute choices and so can personalize their prices. We first show that if buyers and sellers optimize given the personalized-price function

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

the result is a feasible and efficient outcome. Note 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.

Given the pricing function (1), 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

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

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

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

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) and the resulting outcome is efficient. It is straightforward to verify that this is a personalized-price equilibrium as defined in Section 5.1 below.

Equilibrium payoffs to the seller and buyer are

$$\begin{aligned} \frac{(\mathbf{s}(\sigma))^2}{2} - \frac{(\mathbf{s}(\sigma))^3}{3\sigma} &= \frac{\sigma^2}{2} - \frac{\sigma^3}{3\sigma} = \frac{1}{6}\sigma^2 \\ \text{and} \quad \frac{(\mathbf{b}(\beta))^2}{2} - \frac{(\mathbf{b}(\beta))^3}{3\beta} &= \frac{\beta^2}{2} - \frac{\beta^3}{3\beta} = \frac{1}{6}\beta^2. \end{aligned}$$

Note that the attribute-choice functions \mathbf{b} and \mathbf{s} , with the attendant matching, are part of a personalized-price equilibrium irrespective of the value of θ , with the personalized prices changing as θ varies. Moreover, all such equilibria give the same payoffs. Hence, remuneration values are irrelevant. It does not matter who “owns” the technology that combines buyer and seller attribute choices to create the surplus when there is a competitive market with personalized prices for the attributes.

3.3 Uniform Pricing Can Induce Efficient Outcomes

Suppose now that each seller cannot observe buyers’ chosen attributes, and so must set a uniform price that depends only on his own attribute (which is observable to buyers). Let $p_U(s)$ be the uniform-price function that attaches to each seller attribute choice s the equilibrium price the seller receives in the personalized-price equilibrium constructed above:

$$p_U(s) = p_P(\tilde{b}(s), s) = \frac{s^2}{2} - (1 - \theta)\tilde{b}(s)s.$$

What would happen if buyers and sellers made optimal attribute choices, given these prices? Suppose first that sellers own none of the surplus (i.e., $\theta = 1$, and hence $h_S(b, s) = 0$). In this case, facing prices $p_U(s) = s^2/2$, buyers and sellers choose precisely the attributes they chose in the personalized price case above. The result is an efficient outcome, identical to the personalized-price equilibrium outcome. Consequently, no seller would gain by personalizing his price even if he could. In this case, the ability to personalize prices is irrelevant. Proposition 1 below shows that this is not a coincidence, being an implication of the property $dh_S(b, s)/db = 0$.

On the other hand, when $\theta > 1$, as we show in the next subsection, p_U is not part of a uniform-price equilibrium.

3.4 Uniform Pricing Need Not Induce Efficient Outcomes

For general θ , under uniform pricing, buyer β chooses a buyer attribute b and a seller attribute s to solve

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

Assuming p_U is differentiable, the first-order conditions are

$$\begin{aligned} \theta s - \frac{b^2}{\beta} &= 0 \\ \text{and} \quad \theta b - p'_U(s) &= 0. \end{aligned}$$

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

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

We can then solve for the equilibrium attribute-choice functions $\mathbf{b}(\beta)$ and $\mathbf{s}(\sigma)$, the uniform price function $p_U(s)$, and the matching function $\tilde{b}(s)$ identifying the buyer attribute $b = \tilde{b}(s)$ matched with a seller who chooses attribute s , finding⁹

⁹To solve, we conjecture that the equilibrium attribute-choice functions are given by the linear functions

$$\mathbf{b}(\beta) = A\beta \tag{2}$$

$$\text{and} \quad \mathbf{s}(\sigma) = B\sigma. \tag{3}$$

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, we 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, we rewrite the buyer's first first-order condition as $\theta Bb/A - b^2/\beta = 0$ and solve for b , yielding

$$b = \frac{\theta B}{A}\beta. \tag{4}$$

$$\begin{aligned}
\mathbf{b}(\beta) &= \theta^{\frac{2}{3}}(2-\theta)^{\frac{1}{3}}\beta, \\
\mathbf{s}(\sigma) &= \theta^{\frac{1}{3}}(2-\theta)^{\frac{2}{3}}\sigma, \\
p_U(s) &= \frac{\theta}{2} \left(\frac{\theta}{2-\theta} \right)^{1/3} s^2, \\
\text{and } \tilde{b}(s) &= \left(\frac{\theta}{2-\theta} \right)^{1/3} s.
\end{aligned}$$

The result is a uniform-price equilibrium (defined in Section 5.2 below).

When $\theta = 1$, this uniform-price equilibrium is efficient. In this case, as we noted in Section 3.3, the uniform-price equilibrium choices are the same as in the personalized-price equilibrium. There is then no need to personalize prices. Notice that in this case the seller's premuneration value is zero—the seller owns none of the surplus—and consequently is obviously independent of the attribute of buyer with whom the seller matches. Conversely, when $\theta < 1$, the uniform-price equilibrium is both inefficient and not the outcome of a personalize-price equilibrium.¹⁰

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

$$s = \frac{(2-\theta)A}{B}\sigma. \quad (5)$$

Combining (2) with (4) and (3) with (5), we solve for $A = \theta^{\frac{2}{3}}(2-\theta)^{\frac{1}{3}}$ and $B = \theta^{\frac{1}{3}}(2-\theta)^{\frac{2}{3}}$, giving the result.

¹⁰Consider a buyer and seller with the same attribute choice x . Since this pair is not matched (given $\theta < 1$), if the sum of the buyer and seller equilibrium payoffs is less than x^2 , then the buyer and seller in question could do better matching with each other, and so the outcome is inconsistent with personalized pricing (by Lemma 7). The sum of payoffs is

$$\begin{aligned}
&(1-\theta)\tilde{b}(x)x + p_U(x) + \theta\tilde{s}(x)x - p_U(\tilde{s}(x)) \\
&= (1-\theta) \left(\frac{\theta}{2-\theta} \right)^{\frac{1}{3}} x^2 + \frac{\theta}{2} \left(\frac{\theta}{2-\theta} \right)^{\frac{1}{3}} x^2 \\
&\quad + \theta \left(\frac{2-\theta}{\theta} \right) x^2 - \frac{\theta}{2} \left(\frac{\theta}{2-\theta} \right)^{\frac{1}{3}} \left(\frac{2-\theta}{\theta} \right)^{\frac{2}{3}} x^2 \\
&= \frac{1}{2} \left[(2-\theta)^{\frac{2}{3}}\theta^{\frac{1}{3}} + (2-\theta)^{\frac{1}{3}}\theta^{\frac{2}{3}} \right] x^2.
\end{aligned}$$

The coefficient of x^2 can be bounded as follows:

$$\frac{1}{2}[(2-\theta)^{\frac{2}{3}}\theta^{\frac{1}{3}} + (2-\theta)^{\frac{1}{3}}\theta^{\frac{2}{3}}] = (2-\theta)^{\frac{1}{3}}\theta^{\frac{1}{3}} \frac{1}{2}[(2-\theta)^{\frac{1}{3}} + \theta^{\frac{1}{3}}] < \frac{1}{2}[(2-\theta)^{\frac{1}{3}} + \theta^{\frac{1}{3}}] < 1,$$

where the first inequality follows from $(2-\theta)\theta < 1$ and the second from Hardy, Littlewood,

3.5 Premuneration Values and Inefficiency

Recall that the personalized-price equilibrium outcome (which duplicates the uniform-price equilibrium outcome for $\theta = 1$) is independent of θ .

In the personalized-price equilibrium, 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 lower than in the (efficient) personalized-price equilibrium. The inability to personalize prices prevents sellers from offering buyers lower prices in return for higher buyer attributes. As a result, buyers receive a lower return on their investment under uniform pricing, and so choose lower attributes.

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.

The sellers' attribute choice in the uniform-price equilibrium is 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 $s > b$.

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 $\theta \approx .38$, sellers make precisely the attribute choices under uniform pricing that they would under personalized pricing. 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) case that prices can be personalized.

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 when they can personalize prices) 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 .38.

and Pólya (1952, §2.9, Theorem 16). Hence, the sum of payoffs is less than x^2 , as required.




Figure 1: Uniform-price equilibrium attribute choices as a function of θ , the buyers' premuneration-value share of the surplus. The lower curved line is the coefficient of the (linear) buyer attribute-choice function, while the upper curved line is that of the seller attribute-choice function. Both coefficients are 1 in the personalized-price equilibrium.

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 investment less attractive.

Sellers' incentives to screen buyers lead not only to attribute choices that exceed the efficient investments of personalized pricing, but also to attribute choices that are inefficiently high *given* the buyers' (inefficiently low, compared to personalized pricing) 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 socially optimal attribute choice for seller σ then solves

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

The solution to this problem is

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

smaller than the seller's equilibrium attribute choice ($\sigma\theta^{1/3}(2 - \theta)^{1/3}$).

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 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 don't we see such higher prices under personalized pricing? Suppose that given a uniform-price equilibrium, a single seller had 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.

Similarly, the buyer's payoff is

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

This payoff is always smaller under the uniform than personalized-price equilibrium.

Remark 3 (Who Should Set Prices?) 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 will have to be negative in order to bring buyers into the market, and buyers will choose the seller posting 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. 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 only depend on buyer attributes). While it is more efficient for sellers to be the price setters for large θ , it would be more efficient to have buyers post prices when θ is small. \blacklozenge

4 The Matching Market

We now turn to the formal analysis of the general model, beginning with the assumptions of the model.

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 \mathcal{C}^2 , increasing in b and s , and satisfy*

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

Suppose, for example, the 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 $\theta \in (0, 1]$. Then if the surplus function $v : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}$ is (twice continuously) differentiable, increasing in b and s , and strictly supermodular ($\partial^2 v / \partial b \partial s > 0$), these premuneration values satisfy Assumption 1.

Assumption 2 (Essentiality) *For all s , $h_B(0, s)$ and $h_S(0, s)$ are constant in s , with*

$$h_B(0, s) + h_S(0, s) \leq 0.$$

Assumption 2 requires that a positive buyer attribute is essential to the match. The asymmetric treatments of buyers and sellers in Assumptions 1 and 2 anticipate asymmetries arising out of the fact that sellers post prices. Assumption 2 simplifies the discussion of matching (in its absence, we cannot rule out zero attribute buyers matching with positive attribute sellers, precluding the simple formulation of matching that we use).

Assumption 3 (Single-crossing) *The cost function $c_B : \mathbb{R}_+ \times [0, 1] \rightarrow \mathbb{R}_+$ is \mathcal{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.

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.

Assumption 5 (Nontriviality) *For every $\beta = \sigma \equiv \phi \in (0, 1]$, there exists $(b, s) \in [0, \bar{b}] \times [0, \bar{s}]$ with $v(b, s) - c_B(b, \phi) - c_S(s, \phi) > 0$.*

There is always an equilibrium in which every agent chooses the outside option—it does not pay to be the only one in the market. We consider only equilibria where everyone enters the market. While our notation ignores the outside option, our optimality conditions guarantee that the outside option is not strictly profitable. At the same time, we allow the possibility that many or all types of agents pile up on the zero-cost zero attribute, and present conditions ensuring the existence of equilibria in which this does not happen.

We denote by $\mathbf{b} : [0, 1] \rightarrow [0, \bar{b}]$ and $\mathbf{s} : [0, 1] \rightarrow [0, \bar{s}]$ the Lebesgue-measurable functions describing the attributes chosen by buyers and sellers. We assume the matching between buyer and seller attribute choices depends only on the distribution of such choices in the market.

Our first task is to define feasible matchings between buyers and sellers. We denote by \mathcal{B} and \mathcal{S} the closures of the sets of attributes chosen by buyers and sellers respectively, $\mathcal{B} \equiv \text{cl}(\mathbf{b}([0, 1]))$ and $\mathcal{S} \equiv \text{cl}(\mathbf{s}([0, 1]))$, and refer to \mathcal{B} and \mathcal{S} as the *sets of attributes in the market*. 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 $\mathcal{S}' \subset \mathcal{S}$,

$$\begin{aligned} \lambda_{\mathcal{B}}(\mathcal{B}') &= \lambda\{\beta \in [0, 1] : \mathbf{b}(\beta) \in \mathcal{B}'\} \\ \text{and} \quad \lambda_{\mathcal{S}}(\mathcal{S}') &= \lambda\{\sigma \in [0, 1] : \mathbf{s}(\sigma) \in \mathcal{S}'\}, \end{aligned}$$

where λ is Lebesgue measure.

Definition 1 *Suppose \mathbf{b} and \mathbf{s} 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}). Suppose also that $\underline{\sigma} \equiv \sup\{\sigma : \mathbf{s}(\sigma) = 0\} = \sup\{\beta : \mathbf{b}(\beta) = 0\} \equiv \underline{\beta}$. A feasible matching is a pair of measure-preserving functions $\tilde{b} : (\mathcal{S}, \lambda_{\mathcal{S}}) \rightarrow (\mathcal{B}, \lambda_{\mathcal{B}})$ and $\tilde{s} : (\mathcal{B}, \lambda_{\mathcal{B}}) \rightarrow (\mathcal{S}, \lambda_{\mathcal{S}})$ satisfying*

$$\tilde{s}(\tilde{b}(s)) = s \text{ for all } s \in \mathbf{s}((\underline{\sigma}, 1]), \quad (6)$$

$$\text{and} \quad \tilde{b}(\tilde{s}(b)) = b \text{ for all } b \in \mathbf{b}((\underline{\beta}, 1]). \quad (7)$$

Observe that equations (6) and (7) imply that \tilde{s} is one-to-one on $\mathbf{b}((\underline{\beta}, 1])$ and \tilde{b} is one-to-one on $\mathbf{s}((\underline{\sigma}, 1])$.

Remark 4 (Feasible Matchings) We simplify the analysis by restricting attention to attribute-choice functions that are strictly increasing when positive and that assign equal masses of buyers and sellers to zero attribute choices ($\underline{\beta} = \underline{\sigma}$). We show that equilibria exist having these properties, as do our examples. We could define feasible matchings more generally, but at the cost of considerable technical complication. In particular, our requirements imply that the measure $\lambda_{\mathcal{B}}$ is atomless on $((\underline{\beta}, 1])$ and the measure $\lambda_{\mathcal{S}}$ is atomless on $((\underline{\sigma}, 1])$ and that these sets have equal measure, allowing us to restrict attention to matchings that are one-to-one on these sets.

We further simplify 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, Mailath, and Postlewaite (2001, Section 2) for an example). 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 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, Mailath, and Postlewaite (2001).

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., $\lambda_{\mathcal{B}}(\tilde{b}(\mathcal{S}')) = \lambda_{\mathcal{S}}(\mathcal{S}')$ for all Borel $\mathcal{S}' \subset \mathcal{S}$ (and similarly for \tilde{s}). \blacklozenge

Definition 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}) .

Given a feasible matching (\tilde{b}, \tilde{s}) , $\tilde{b}(s)$ specifies the buyer attribute matched to a seller with attribute s , and $\tilde{s}(b)$ specifies the seller attribute matched to a buyer with attribute b .

5 Equilibrium

5.1 Personalized Pricing

To capture investment incentives, we must model the rewards agents would receive if they chose a “surprise” attribute, i.e., an attribute that is not chosen in equilibrium. In the example in Section 3, we avoided this issue by pricing all matches, in order to focus on the incentives for matching and screening buyers.¹¹ In general, invoking the existence of prices for non-existent attribute choices precludes serious consideration of important issues such as how markets coordinate behavior, and so we now confront this issue directly.

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 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 *chosen* attributes. Informally, a *personalized-price equilibrium* is a feasible outcome and a personalized-price function such that no agent has an incentive to deviate from the behavior specified by the feasible outcome.

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

$$\begin{aligned} \Pi_B(b, \beta) &\equiv h_B(b, \tilde{s}(b)) - p_P(b, \tilde{s}(b)) - c_B(b, \beta) \\ \text{and} \quad \Pi_S(s, \sigma) &\equiv h_S(\tilde{b}(s), s) + p_P(\tilde{b}(s), s) - c_S(s, \sigma). \end{aligned}$$

We introduce a standard notion of price-taking behavior:

Definition 3 *Given a feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$, buyer β is a price taker under p_P if*

$$(\mathbf{b}(\beta), \tilde{s}(\mathbf{b}(\beta))) \in \arg \max_{(b, s) \in \mathcal{B} \times \mathcal{S}} h_B(b, s) - p_P(b, s) - c_B(b, \beta). \quad (8)$$

Seller σ is a price taker under p_P if

$$(\tilde{b}(\mathbf{s}(\sigma)), \mathbf{s}(\sigma)) \in \arg \max_{(b, s) \in \mathcal{B} \times \mathcal{S}} h_S(b, s) + p_P(b, s) - c_S(s, \sigma). \quad (9)$$

As a special case, seller σ always has the option of selecting the attribute chosen by some seller $\hat{\sigma}$, in the process replacing the buyer with whom σ is

¹¹Such pricing is called *complete* in Definition 10 below.

matched with $\hat{\sigma}$'s buyer $\tilde{b}(\hat{\sigma})$ and altering the price σ receives. Condition (9) requires that this be unprofitable: $\Pi_S(\mathbf{s}(\sigma), \sigma) \geq h_S(\tilde{b}(s), s) + p_P(\tilde{b}(s), s) - c_S(s, \sigma)$ for all $s \in \mathcal{S}$. A similar comment applies, of course, to buyers. More generally, condition (9) requires that no seller σ find it profitable to select the attribute chosen by some seller $\hat{\sigma}$ and then to sell to some buyer b other than $\tilde{b}(\hat{\sigma})$ at the appropriate equilibrium personalized price.

Since the personalized-price function p_P prices only pairs of chosen attributes, it remains to specify what price (if any) the seller will receive, or indeed with whom the seller will be matched, should he choose an attribute not chosen by any other seller. Our notion of equilibrium requires that there exist no chosen buyer attribute for which the surplus the buyer and seller would then generate could be divided between them so as to make both better off:

Definition 4 *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 $\Pi_S(\mathbf{s}(\sigma), \sigma) < 0$ or there exist a seller attribute choice $s \in [0, \bar{s}]$, a (chosen) buyer attribute $b \in \mathcal{B}$, and a price $p \in \mathbb{R}$ such that*

$$h_B(b, \tilde{s}(b)) - p_P(b, \tilde{s}(b)) < h_B(b, s) - p \quad (10)$$

$$\text{and} \quad \Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(b, s) + p - c_S(s, \sigma). \quad (11)$$

If $\Pi_S(\mathbf{s}(\sigma), \sigma) < 0$, the outside option is better for the seller. If $\Pi_S(\mathbf{s}(\sigma), \sigma) \geq 0$, a seller has a profitable deviation when he is able to attract a buyer attribute b (condition (10)) that yields him a higher payoff than had he followed the behavior prescribed by the given outcome (condition (11)). This latter deviation involves targeting a particular buyer attribute. As such, it should be distinguished from the seller deviation under uniform pricing introduced in Definition 8 below (where no such targeting is possible).

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

Definition 5 *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 $\Pi_B(\mathbf{b}(\beta), \beta) < 0$ or there exist a buyer attribute choice $b \in [0, \bar{b}]$, a (chosen) seller attribute $s \in \mathcal{S}$, and a price $p \in \mathbb{R}$ such that*

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

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

We then have the following definition of personalized-price equilibrium:

Definition 6 *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 price takers under p_P and no buyer or seller has a profitable deviation.*

The notions of price taking (Definition 3) and profitable deviation (Definitions 4 and 5) are distinct but related. In particular, Definitions 4 and 5 imply price-taking behavior for the special case mentioned just after Definition 3, in which a seller σ chooses a seller attribute $s \in S$, $s \neq s(\sigma)$, that is, an attribute that already exists in the market, and then trades with the buyer matched to that attribute ($\tilde{b}(s)$) at price $p_P(\tilde{b}(s), s)$. If matching with that buyer attribute at the market price yields a higher net payoff to seller σ , then the seller has a profitable deviation: the same attribute choice s and buyer $\tilde{b}(s)$ at price $p_P(\tilde{b}(s), s) - \varepsilon$, for sufficiently small $\varepsilon > 0$ satisfies (10), while also satisfying (11). Hence, we immediately have the following result, that the absence of profitable deviations ensures that no seller (or buyer) would prefer to mimic that actions of another type of seller (or buyer):

Lemma 1 *Given $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$, if there exists $\sigma \in [0, 1]$ and $s \in \mathcal{S}$ such that $\Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(\tilde{b}(s), s) + p_P(\tilde{b}(s), s) - c_S(s, \sigma)$, then seller σ has a profitable deviation. If there exists $\beta \in [0, 1]$ and $b \in \mathcal{B}$ such that $\Pi_B(\mathbf{b}(\beta), \beta) < h_B(b, \tilde{s}(b)) + p_P(b, \tilde{s}(b)) - c_B(b, \beta)$, then buyer β has a profitable deviation.*

In the following lemma, we show that if there is a feasible outcome and personalized-price function such that there do not exist profitable deviations for any buyer or seller, the price function can be extended to unchosen attributes so that the seller is indifferent over buyer attributes. This facilitates comparisons with uniform pricing (see Definition 7 below). Under the extended price function, one can think of the seller as choosing only his attribute $s \in \mathcal{S}$ to maximize $h_S(\tilde{b}(s), s) + \hat{p}_P(\tilde{b}(s), s) - c_S(s, \sigma)$, while the buyer acts as a price taker in her choice of attribute and seller attribute with which to match (see (8)). The result is symmetric: a price function that makes the buyer indifferent over all seller attributes is consistent with an equilibrium in which the buyer chooses her attribute (with the matching function determining her matched attribute) and the seller chooses pairs of attributes.

Lemma 2 *Suppose $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ is feasible and $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$ admits no seller or buyer profitable deviations. Then $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, \hat{p}_P)$ is a personalized-price equilibrium, where \hat{p}_P is the personalized-price function given by*

$$\hat{p}_P(b, s) = p_P(\tilde{b}(s), s) + h_S(\tilde{b}(s), s) - h_S(b, s), \quad \forall (b, s) \in \mathcal{B} \times \mathcal{S}. \quad (12)$$

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

Proof. Seller indifference is immediate. Lemma 1 then implies that the seller is a price taker under \hat{p}_P .

We then need show only that the buyer is a price taker. Suppose (8) 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 deviation,

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

But this, with (12), yields a contradiction. ■

Remark 5 (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 characterization of a personalized-price equilibrium outcome. The following is a straightforward calculation.

Lemma 3 *Let $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$ be 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 6 (Ex Post Contracting Equilibrium) Cole, Mailath, and Postlewaite (2001) study a continuum 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, Mailath, and Postlewaite (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.¹² Technical differences

¹²Consequently, matching is over buyers and sellers, not attributes as here. This difference results in some technical complications.

between Cole, Mailath, and Postlewaite (2001) and the model here prevent a simple formal statement of the precise relationship between ex post contracting equilibrium and personalized-price equilibrium. However, the set of outcomes and implied payoffs are essentially the same under the two notions. In particular, if all buyers and sellers are price takers under $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$, 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 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. \blacklozenge

5.2 Uniform Pricing

We now consider the case in which sellers cannot set personalized prices. If unable to observe buyers' attributes, for example, a seller posts a price that can depend on his own attribute choice, but not on the buyer's attribute choice. Such a price function is a *uniform-price function* $p_U : \mathcal{S} \rightarrow \mathbb{R}$.

Given a feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ and a uniform price p_U the payoffs to a buyer β who chooses $b \in \mathcal{B}$ and a seller σ who chooses $s \in \mathcal{S}$ are as before:

$$\begin{aligned} \Pi_B(b, \beta) &\equiv h_B(b, \tilde{s}(b)) - p_U(\tilde{s}(b)) - c_B(b, \beta) \\ \text{and} \quad \Pi_S(s, \sigma) &\equiv h_S(\tilde{b}(s), s) + p_U(s) - c_S(s, \sigma). \end{aligned}$$

Under uniform pricing sellers cannot condition on buyer attributes. Consequently, sellers choose only their own attributes, while buyers can choose any seller attribute regardless of their own attribute choice.

Definition 7 *Given a feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$, buyer β is a price taker under p_U if*

$$\Pi_B(\mathbf{b}(\beta), \beta) = \max_{(b,s) \in \mathbb{R}_+ \times \mathcal{S}} h_B(b, s) - p_U(s) - c_B(b, \beta). \quad (13)$$

Seller σ is a price taker under p_U if

$$\Pi_S(\mathbf{s}(\sigma), \sigma) = \max_{s \in \mathcal{S}} \Pi_S(s, \sigma). \quad (14)$$

As with personalized-price equilibria, we need to address seller deviations to attributes not in \mathcal{S} (but need not separately consider corresponding buyer deviations, since sellers cannot condition on buyers' attribute choices).

Because sellers cannot observe buyers' attributes, however, a seller cannot choose an attribute s and a price p targeted at a particular buyer attribute b . Instead, the attribute s and price p may attract a range of buyer attributes. The following definition requires that the seller's deviation to (s, p) be profitable *independently* of the buyer attracted. Alternatives to this definition are discussed in Remark 7 below.

Definition 8 *Given $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_U)$, there is a profitable seller deviation if there exists σ and either $\Pi_S(\mathbf{s}(\sigma), \sigma) < 0$ or there exists s' and a price $p \in \mathbb{R}$ such that there exists $b' \in \mathcal{B}$ with*

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

and for all $b'' \in \mathcal{B}$,

$$\begin{aligned} \text{if } h_B(b'', \tilde{s}(b'')) - p_U(\tilde{s}(b'')) < h_B(b'', s') - p \\ \text{then } \Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(b'', s') + p - c_S(s', \sigma). \end{aligned}$$

If a seller has no profitable deviations (and if buyers are price takers under p_U), then that seller is also a price taker under p_U . This result is stronger than Lemma 1, but the proof is more involved. Under personalized pricing, any seller who envied another seller's transaction could simply offer a slightly lower price to the target buyer, making both the given seller and the target buyer better off, with the ability to personalize prices ensuring that the seller need not be concerned with other buyers. This immediately ensures that a seller without profitable deviations cannot envy the actions of another seller. Under uniform pricing, a seller cannot target a given buyer in this way under. Instead, a seller who mimicked the attribute selection of another seller while undercutting his price attracts not only the buyer matched with the target seller but also buyers with lower attributes, making the mimicking behavior less profitable. Appendix A proves:

Lemma 4 *Suppose the feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ and uniform price p_U satisfy (13). If seller σ has no profitable deviations, then he is a price taker under p_U .*

Definition 9 *A feasible outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ and a uniform-price function p_U constitute a uniform-price equilibrium if all buyers are price takers under p_U and no seller has a profitable deviation.*

Remark 7 (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, or given the seller's most-preferred buyer from this set. Allowing the seller to select his most preferred buyer essentially restores the ability to personalize prices (Lemma 6 below makes this connection precise) so that *some* pessimism is essential if uniform and personalized pricing are to give different outcomes. A more pessimistic formulation makes seller deviations less attractive and hence enlarges the set of uniform-price equilibria. Our key results (Propositions 1 and 2), establishing conditions under which personalized-price and uniform-price equilibria coincide, and under which uniform-price equilibria must be inefficient, are rendered more powerful by such a permissive definition of the latter. \blacklozenge

5.3 Allocating Buyers and Sellers

The single-crossing assumptions on the remuneration values naturally lead to positive assortative matching in equilibrium.

Lemma 5

5.1 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.

5.2 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.

Proof. (1) 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 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 (8) for the buyer choosing b_1 and from (9) for the seller choosing s_2 , we have

$$\begin{aligned} h_B(b_1, s_1) - p_P(b_1, s_1) &\geq h_B(b_1, s_2) - p_P(b_1, s_2) \\ \text{and} \quad h_S(b_2, s_2) + p_P(b_2, s_2) &\geq h_S(b_1, s_2) + p_P(b_1, s_2), \end{aligned}$$

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 (8) for the buyer choosing b_2 and from (9) 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.

(2) The argument is similar, though simpler since under uniform pricing it follows from the supermodularity of h_B alone and (13). Using the same notation, 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

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

a contradiction. ■

Remark 8 (Exclusion) Personalized prices allow a seller to accept some buyers while excluding others who would be willing to pay the same price. What is important to the seller is the ability to exclude buyers with lower attribute choices than the seller's equilibrium match. This exclusion can be implemented via appropriate pricing. In particular, by charging a sufficiently high price to specific buyer attribute choices, a seller can ensure that buyers with those attributes will chose not to buy. We denote this sufficiently high price by P . A personalized-price function p_P is a *uniform-rationing price* if it has the form

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

for some $p_{UR} : \mathcal{S} \rightarrow \mathbb{R}_+$ and $b^\dagger : \mathcal{S} \rightarrow \mathcal{B}$. 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 $b < b^\dagger(s)$. Given a personalized-price specification $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$, its associated uniform-rationing price is given by $p_{UR}(s) = p_P(\tilde{b}(s), s)$ and $b^\dagger(s) = \tilde{b}(s)$ for all $s \in \mathcal{S}$. A personalized-price equilibrium outcome that can be supported by a uniform-rationing price is a *uniform-rationing equilibrium outcome*.

Lemma 6 *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. ■



If a uniform-price equilibrium outcome is not also a personalized-price equilibrium outcome, then it must be that some seller would like to change his price in order to attract better buyers, but is deterred from doing so by the specter of less desirable buyers. This seller would then welcome the ability to exclude some buyers by personalizing prices. Section 3.3 provides an example in which personalized-price equilibrium outcomes and uniform-price equilibrium outcomes coincide. In this case, the personalized-price power to exclude buyers is unnecessary—buyers sort themselves among sellers just as sellers would have them do.

6 Efficiency

From Lemma 5, in both personalized-price and uniform-price equilibria, matching 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 or unmatched pair of agents can increase its net surplus without both agents deviating to attribute choices outside the sets \mathcal{B} and \mathcal{S} :¹³

¹³This is essentially Cole, Mailath, and Postlewaite (2001, Proposition 4), which describes a constrained efficiency property of ex post contracting equilibria (see Remark 6). The current formulation allows a more transparent statement and proof.

Lemma 7 *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' ,*

$$\begin{aligned} W(b, s', \phi) &\leq W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi) \\ \text{and} \quad W(b', s, \phi) &\leq W(\mathbf{b}(\phi), \mathbf{s}(\phi), \phi). \end{aligned}$$

Proof. En route to a contradiction, suppose there exists $\phi \in [0, 1]$, $b \in \mathcal{B}$ 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 (8). \blacksquare

In contrast, uniform-price equilibria in general do not satisfy constrained efficiency.

Lemma 7 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.

Definition 10 *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

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

As the names suggest, every complete personalized-price equilibrium outcome is indeed a personalized-price equilibrium outcome.

Lemma 8

(8.1) *Every complete personalized-price equilibrium outcome is a personalized-price equilibrium outcome.*

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

Proof. The efficiency of complete personalized-price equilibria is a straightforward calculation.

Fix a complete personalized-price equilibrium $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$. We discuss seller deviations; the buyer case is analogous. We need only verify that there are no profitable deviations (in the sense of Definition 4) involving an attribute choice $s' \notin \mathcal{S}$. 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

$$\Pi_S(\mathbf{s}(\sigma), \sigma) < h_S(b', s') + p - c_S(s', \sigma) \quad (15)$$

and

$$h_B(b', \tilde{s}(b')) - p_P(b', \tilde{s}(b')) < h_B(b', s') - p. \quad (16)$$

Since $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$ is a complete personalized-price equilibrium, (15) 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 (16) 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. ■

Remark 9 We could similarly define a complete uniform-price equilibrium by requiring a price for all seller attributes in $[0, \bar{s}]$, while expanding to $[0, \bar{s}]$ the set of seller attribute choices over which the buyer optimizes. It is immediate from the definition that a complete uniform-price equilibrium is a uniform-price equilibrium, and apparent from Section 3.4 that a complete uniform-price equilibrium need not be efficient. ◆

7 When is Personalization Redundant?

When can a personalized-price equilibrium outcome be supported by uniform prices? Or, alternatively, under what conditions can matching when sellers set prices without knowing buyers' attributes achieve outcomes attainable as equilibrium outcomes when they are informed? In the example of Section 3, attribute choices were efficient when sellers own none of the surplus. Our sufficient condition generalizes this property.

We begin with some intuition, appropriate when equilibrium is characterized by first-order conditions. Fix a uniform-price equilibrium, including the uniform-price function p_U . We first note that, by standard incentive compatibility arguments, the uniform-price function is differentiable. The first-order conditions implied for the buyer's choice of attribute b and matching attribute choice s in a uniform-price equilibrium are

$$0 = \frac{dh_B(b, s)}{db} - \frac{dc_B(b, \beta)}{db} \quad (17)$$

$$\text{and} \quad 0 = \frac{dh_B(b, s)}{ds} - \frac{dp_U(s)}{ds}, \quad (18)$$

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

$$0 = \frac{dh_S(\tilde{b}(s), s)}{db} \frac{d\tilde{b}(s)}{ds} + \frac{dh_S(\tilde{b}(s), s)}{ds} + \frac{dp_U(s)}{ds} - \frac{dc_S(s, \sigma)}{ds}. \quad (19)$$

Using (18) to eliminate $dp_U(s)/ds$ in (19) and then using the identity $v(b, s) = h_B(b, s) + h_S(b, s)$ in (17) and (19), these three first-order conditions can be reduced to

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

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

From Lemma 7, establishing the constrained efficiency of a personalized-price equilibrium outcomes, we know that a personalized-price equilibrium must be characterized by the first-order conditions:

$$0 = \frac{dv(b, s)}{db} - \frac{dc_B(b, \beta)}{db} \quad (20)$$

$$0 = \frac{dv(b, s)}{ds} - \frac{dc_S(s, \sigma)}{ds}.$$

Comparing these, it is immediate that the solution to the first-order conditions for the personalized-price equilibrium will be a solution for the first-order conditions for the uniform-price equilibrium if $\frac{dh_S(b,s)}{db} = 0$, that is, if each seller's remuneration value is independent of the attribute choice of the buyer with whom the seller is matched. This argument is summarized in the following proposition (which requires no differentiability assumptions).

Proposition 1 *A personalized-price equilibrium outcome can be achieved in a uniform-price equilibrium if the sellers' remuneration values do not depend on the buyer's attribute.*

Proof. Let $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$ be a personalized-price equilibrium. Applying Lemma 2, $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, \hat{p}_P)$ is also a personalized-price equilibrium, where

$$\hat{p}_P(b, s) = p_P(\tilde{b}(s), s) + h_S(\tilde{b}(s), s) - h_S(b, s)$$

for all $(b, s) \in \mathcal{B} \times \mathcal{S}$.

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 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.¹⁴

Proposition 2 *Suppose the outcome $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s})$ of a personalized-price equilibrium $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$ can be supported as a uniform-price equilibrium outcome with price $p_U(s) = p_P(\tilde{b}(s), s)$. Then for all $s \in \mathcal{S}$,*

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

Proof. It follows from (17) and (20) (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 a personalized-price equilibrium that can be supported by uniform prices, then

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

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

¹⁴Analogously, the single-crossing condition is essentially necessary for a separating equilibrium in a signaling model (Mailath, 1987, Theorem 3).

8 Existence of Equilibrium

8.1 Uniform-Price Equilibrium

A fixed-point argument (in Appendix B) allows us to establish existence of uniform-price equilibria, by showing the existence of complete uniform-price equilibria (cf. Remark 9).

Proposition 3 *If there exists $(b, s) \in (0, \bar{b}] \times (0, \bar{s}]$ with*

$$h_B(b, s) + h_S(0, s) - c_B(b, 1) - c_S(s, 1) > 0, \quad (21)$$

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 $(b, s) \in (0, \bar{b}] \times (0, \bar{s}]$

$$h_B(b, s) + h_S(0, s) - c_B(b, \phi) - c_S(s, \phi) > 0, \quad (22)$$

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

Assumption 5 implies inequality (22) if $h_S(b, s)$ is independent of b (in which case personalized and uniform pricing correspond). Condition (21) and (22) can fail, even in the presence of Assumption 5, if $dh_S(b, s)/ds$ is large (e.g., when θ is small in Section 3.4). In such cases, buyers are the appropriate side of the market to be setting prices (cf. Remark 3). Uniform-pricing equilibria are inefficient when $h_S(b, s)$ depends on b . If this dependence is too extreme, (21) may fail and there may be no investment on either side.

Two significant complications must to 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 incorporating 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.

8.2 Personalized-Price Equilibrium

One route to existence is to note that a personalized-price equilibrium is essentially equivalent to Cole, Mailath, and Postlewaite (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 4 *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 3 ensures that there exists a complete uniform-price equilibrium. Proposition 1 ensures that there is a corresponding complete personalized-price equilibrium $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_P)$, which Lemma 8 ensures is efficient. Then setting

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

gives a complete (and hence efficient) personalized-price equilibrium for the market in question. ■

9 Discussion

9.1 Premuneration Values

Our main result is that a necessary and sufficient condition to avoid either the costs of personalization or the inefficiencies of uniform pricing is that sellers' premuneration values should be independent of the buyer to whom they are matched. Why aren't markets and institutions arranged so that premuneration values have this property?

Moral hazard is a key obstacle to such an arrangement. In Section 1, we touched on the moral hazard problems associated with assigning all of the surplus, including the student's future earnings, to a university. For a second example, 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 remuneration 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 remuneration value will no longer be independent of the characteristics of the agent with whom he is matched. An inefficiency then arises either from uniform pricing or from the costs of personalization.

We thus regard moral hazard as imposing fundamental constraints on the design of remuneration values. This in turn can make new monitoring and contracting technologies 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.

9.2 Endogenous Information Acquisition

Suppose that lack of information about buyers' attribute choices poses the primary obstacle to personalizing prices. Can sellers take actions to ameliorate this informational asymmetry?

9.2.1 Information Acquisition

Suppose that before attributes are chosen, each seller σ can pay a cost $K(\sigma)$, in which case the seller can observe buyers' attribute choices. It may be that $K(\sigma)$ is constant in σ , but there are many other plausible configurations. For example, the same characteristics that make attributes less costly for larger values of σ may also make informativeness less costly.

We refer to the resulting model as the "endogenously-informed-seller" model. Natural specifications of such a model would lead to the following types of results:

- Let $h_S(b, s)$ be independent of b . Then it is an equilibrium of the endogenously-informed-seller model for each seller to choose not to obtain the monitoring technology, coupled with a uniform-price equilibrium.
- Moreover, this equilibrium is robust: If the seller remuneration value function is almost independent of b and K is bounded away from 0, then it is still an equilibrium for each seller to choose not to obtain the monitoring technology.

The first statement reiterates our basic conclusion—that markets can be (efficiently) cleared by uniform prices when sellers’ remuneration values are independent of match. The next statement notes that this is not a “razor-edge” result. The benefit to a seller of becoming informed is to be able to discriminate among the potential buyers with whom he might transact. But if the differences in the size of $h_S(b, s)$ across potential buyers is sufficiently small, the benefits from acquiring the technology will be less than paying the cost to obtain the technology, and consequently it will be an equilibrium in the endogenously-informed-seller model for no seller to purchase the monitoring technology and for each seller to set a uniform price.

The reverse of this is also true. If the monitoring cost is sufficiently small and $h_S(b, s)$ is not independent of b , then we will not have completely uniform pricing. In particular, suppose the outcome of a uniform-price equilibrium $(\mathbf{b}, \mathbf{s}, \tilde{b}, \tilde{s}, p_U)$ cannot be supported in a personalized-price equilibrium. Then there must exist a seller who can use personalized prices to construct a profitable deviation, and who would do so were the cost of personalization sufficiently small.

One might conjecture that if $h_S(b, s)$ is not independent of b , then we will have a personalized-price equilibrium if the cost of personalization is sufficiently small, without further assumptions on how it is small. However, this is not the case. Let $K(\sigma) = K > 0$. Then the endogenously-informed-seller model does not have an equilibrium in which all sellers acquire the monitoring technology. The lowest type of seller attribute choice necessarily matches with the lowest buyer investment, and hence has no buyers to exclude. It then cannot be in this seller’s best interests to acquire the monitoring technology. For values $K > 0$, this applies to an interval of lowest-type sellers, which precludes the existence of a personalized-price equilibrium.

9.2.2 An Example

We expand the example of Section 3 to illustrate an equilibrium of the endogenously-informed-seller model with a mixture of uniform and personalized prices. We assume that K is decreasing in σ , with $K(1) = 0$, and consider the class of cost functions $\alpha K(\sigma) \equiv k(\alpha, \sigma)$ for $\alpha > 0$. We find that if the cost of the monitoring technology decreases sufficiently quickly in seller index (i.e., α is sufficiently large), then there is an equilibrium in which low-index sellers set uniform prices, while high-index sellers personalize. However, we also find that, even though high-index sellers seemingly should benefit more than low-index sellers from the monitoring technology, if the cost of the technology is independent of seller index, then there is no

such equilibrium. We interpret this nonexistence as an indication that a richer model of endogenous information acquisition is necessary, an important area for further research.

Let the sum of the payoffs to a buyer and seller of types $\beta = \sigma = \phi$ in the personalized-price and uniform-price equilibria be denoted $v_P(\phi, \phi)$ and $v_U(\phi, \phi)$. Then in our example,

$$v_P(\phi, \phi) = \frac{1}{3}\phi^2,$$

and

$$v_U(\phi, \phi) = \frac{1}{6}[\theta(2 - \theta)^2 + \theta^2(2 - \theta)]\phi^2 = \frac{1}{3}\theta(2 - \theta)\phi^2.$$

Let ψ satisfy

$$v^P(\psi, \psi) - v^U(\psi, \psi) = k(\alpha, \psi).$$

A match between two agents of type ψ is then the switch-point at which the efficiency of the uniform-price equilibrium just suffices to warrant paying the cost k of the technology. Agents with types below ψ will not purchase the monitoring technology and will behave as in the uniform-price equilibrium. Agents above ψ will purchase the monitoring technology, and will behave as in the uniform-price equilibrium, with the exception that the price will now be given by

$$p_P(b, s) = \frac{s^2}{2} - (1 - \theta)bs + \Delta.$$

The constant Δ affects none of the incentives in the uniform-price equilibrium. It is chosen to equalize the payoffs of the marginal seller $\sigma = \psi$ in the two equilibria. This is the required condition for this seller to be indifferent between buying and not buying the monitoring technology. We have

$$\begin{aligned} \Delta &= \frac{1}{3}\psi^2 - \frac{1}{3}\theta(2 - \theta)\psi^2 \\ &= \frac{1}{3}(1 - \theta)^2\psi^2 > 0. \end{aligned}$$

Hence, the division of the surplus is pushed in the seller's favor, compared to the uniform-price equilibrium, in response to seller $\sigma = \psi$'s outside option of saving the cost of the monitoring technology by entering the uniform-pricing segment of the market.

The seller's attribute choice drops as σ increases past ψ while the buyer's

jumps up. The price jumps down:

$$\begin{aligned} p_P(\mathbf{b}(\psi), \mathbf{s}(\psi)) &= (\theta - \frac{1}{2})\psi^2 + \Delta \\ &< \frac{\theta}{2} \left(\frac{\theta}{2 - \theta} \right)^{\frac{1}{3}} \theta^{\frac{2}{3}} (2 - \theta)^{\frac{4}{3}} \psi^2 \\ &= p_U(\mathbf{b}(\psi), \mathbf{s}(\psi)). \end{aligned}$$

The inequality is equivalent to

$$(\theta - \frac{1}{2}) + \frac{1}{3}(1 - \theta)^2 < \frac{\theta^2}{2}(2 - \theta),$$

which is satisfied for $\theta \in [0, 1]$. At the switch point ψ , the marginal buyer thus trades off a high-attribute seller and a high price (just below ψ) against a relatively low-attribute low-price seller (above ψ). Sellers above ψ are able to pay less for higher-attribute buyers, but pay for the monitoring technology. Notice that some buyers below ψ would like to buy from sellers above ψ , at the observed prices, without increasing their investments, but the personalized prices of the latter preclude the buyers doing so.

The only optimality condition that is not obvious in this formulation concerns the information-acquisition behavior of sellers near the critical type ψ . Seller ψ is indifferent between acquiring and not acquiring the monitoring technology, which may initially appear to suffice for optimality. However, we have noted that the seller attribute choice falls at type ψ . If K is independent of σ , then sellers' types enter their payoffs only through the cost function c_S . Given the single-crossing property satisfied by c_S , the equilibrium seller attribute choice must be increasing in type, ensuring that the proposed strategies are not an equilibrium. Seller ψ can be indifferent between a large attribute choice coupled with uniform pricing and a small attribute choice coupled with personalized pricing, without seller $\psi - \varepsilon$ for small ε strictly preferring the latter (disrupting the equilibrium) only if seller ψ has a cost advantage in purchasing the monitoring technology, i.e., only if $k(\alpha, \sigma)$ declines sufficiently rapidly in σ , i.e., if α is sufficiently large. This will be the case, and we will have an equilibrium, for all α sufficiently large.¹⁵

¹⁵Let $\bar{c} = \lim_{\sigma \uparrow \psi} c_S(\mathbf{s}(\sigma), \sigma)$ and $\underline{c} = \lim_{\sigma \downarrow \psi} c_S(\mathbf{s}(\sigma), \sigma)$. Then we need

$$\frac{d(\bar{c} - \underline{c})}{d\psi} > \frac{dk(\alpha, \psi)}{d\psi} = \alpha \frac{dK(\psi)}{d\psi}.$$

9.3 Endogenous Information Revelation

The information required to personalize prices may come to light not through the information-acquisition efforts of sellers but because buyers reveal it. Applicants to universities typically take SAT or ACT exams, often beyond the minimum required, that at least partially reveal the attribute of interest. Students of high ability in particular find it in their interest to take them in an attempt to certify their attribute. Thus a more general model would include a richer set of technologies by which either buyers or sellers could make attributes known to all participants.

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.¹⁶ 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.¹⁷ In contrast, in the example of Section 3, 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.

¹⁶See Grossman (1981), Milgrom (1981), or Okuno-Fujiwara, Postlewaite, and Suzumura (1990) for analyses of this.

¹⁷Ostrovsky and Schwarz (forthcoming) investigate the optimal amount of information to disclose from the students’ perspective.

A Proof of Lemma 4

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 $b \geq \tilde{b}(s') - \delta$,

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

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

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

holds and $|p'' - p_U(s')| < \varepsilon/2$.

From single crossing,

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

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 single crossing shows that for all $b \leq \tilde{b}(s') - \delta$,

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

From (13), 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 (A.2) and the second from (A.1).

B Proof of Proposition 3: Existence of Equilibrium.

The existence proof is involved and indirect. We would like to construct a game Γ whose equilibria induce uniform-price equilibria. However, the obvious such game Γ is itself difficult to handle, so we work with an approximating sequence of games Γ^n . We verify that each Γ^n has an equilibrium, take limits, and show that the limiting strategy profile induces a uniform-price equilibrium. Loosely, the n index allows us to accommodate (in the limit) the possibility of jumps in the attribute-choice functions (precluded in game Γ^n).

B.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 would be willing to purchase any seller attribute choice $s \in [0, \bar{s}]$ at a price exceeding P , nor would any seller be 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 would always choose the former. Equilibrium prices will thus never need to increase at a rate faster than Δ .

B.2 The game Γ^n

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

B.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 $\mathbf{s}_B : [0, 1] \rightarrow [0, \bar{s}]$ 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 Υ_B . In Γ^n , the buyer is restricted to the subset of Υ_B , denoted by Υ_B^n , of functions satisfying (B.3) and (B.4):

$$(\beta' - \beta)/n \leq \mathbf{b}(\beta') - \mathbf{b}(\beta) \leq n(\beta' - \beta), \quad \forall \beta < \beta' \in [0, 1], \quad (\text{B.3})$$

and

$$(\beta' - \beta)/n \leq \mathbf{s}_B(\beta') - \mathbf{s}_B(\beta) \leq n(\beta' - \beta), \quad \forall \beta < \beta' \in [0, 1]. \quad (\text{B.4})$$

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 Υ_S . In Γ^n , the seller is restricted to the subset of Υ_S , denoted by Υ_S^n , of functions satisfying (B.5),

$$(\sigma' - \sigma)/n \leq \mathbf{s}(\sigma') - \mathbf{s}(\sigma) \leq n(\sigma' - \sigma), \quad \forall \sigma < \sigma' \in [0, 1]. \quad (\text{B.5})$$

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

$$p_U(s') - p_U(s) < 2\Delta(s' - s) \quad (\text{B.6})$$

for all $s < s' \in [0, \bar{s}]$. Denote the set of increasing functions p_U satisfying (B.6), endowed with the sup norm, by Υ_P (note that Υ_P is not indexed by n). Every function in Υ_P is continuous; indeed the collection Υ_P is equicontinuous.

The set $\Upsilon \equiv \Upsilon_B \times \Upsilon_S \times \Upsilon_P$, when normed by the sum of the three constituent norms, is a compact metric space.¹⁸ It is immediate that $\Upsilon^n \equiv \Upsilon_B^n \times \Upsilon_S^n \times \Upsilon_P$ is a closed subset of Υ , and so also compact.

B.2.2 Buyer and Price-Setter Payoffs

The buyer. The buyer's payoff from $(\mathbf{b}, \mathbf{s}_B) \in \Upsilon_B^n$, when the price-setter has chosen $p_U \in \Upsilon_P$ is

$$\int (h_B(\mathbf{b}(\beta), \mathbf{s}_B(\beta)) - p_U(\mathbf{s}_B(\beta)) - c_B(\mathbf{b}(\beta), \beta)) d\beta. \quad (\text{B.7})$$

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

¹⁸It suffices for this conclusion to show that Υ 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, 1995, Theorem 25.9) shows Υ is sequentially compact. In particular, given a sequence $\{(\mathbf{b}^m, \mathbf{s}_B^m, \mathbf{s}^m, p_U^m)\}$, we can choose a subsequence along which each function converges at every rational value in its domain to a limit $\{(\mathbf{b}^\infty, \mathbf{s}_B^\infty, \mathbf{s}^\infty, p_U^\infty)\}$. Because each function in the sequence $\{(\mathbf{b}^m, \mathbf{s}_B^m, \mathbf{s}^m, p_U^m)\}$ is increasing, so must be each limiting function $\{(\mathbf{b}^\infty, \mathbf{s}_B^\infty, \mathbf{s}^\infty, p_U^\infty)\}$. This ensures convergence at every continuity point of the limit functions, and hence almost everywhere for the functions \mathbf{b}^m , \mathbf{s}_B^m and \mathbf{s}^m and everywhere for the functions p_U^m , sufficing (for bounded functions) for L^1 convergence in the former three cases and convergence in the sup norm in the latter.

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 $p_U \in \Upsilon_P$, when the buyer and seller have chosen $(\mathbf{b}, \mathbf{s}_B, \mathbf{s}) \in \Upsilon_B^n \times \Upsilon_S^n$ is given by

$$\int_0^{\bar{s}} p_U(s) (F_B(s) - F_S(s)) ds. \quad (\text{B.8})$$

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.

B.2.3 Seller Payoffs

The specification of the seller's payoff is complicated by the need to incorporate incentives arising from the possibility of profitable seller deviations. Given an attribute choice s , price p , and price function p_U , set

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

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 (B.4) 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 $p_U \in \Upsilon_P$, 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 , $B(s, p, p_U) = [0, \bar{b}]$.

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

(2) 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 $\mathcal{D}_j(s, p_U)$. The set $\{s : \mathcal{D}_j(s, p_U) \neq \emptyset\}$ has zero Lebesgue measure.

(3) 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

$$b_1 \leq \liminf_\ell b_1^\ell \leq \limsup_\ell b_2^\ell \leq b_2.$$

(4) 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. (1) 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.

(2) 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, from the maximum theorem, both b_1 and b_2 are left-continuous.

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 $b \in [b_1(p), b_1^+]$,

$$h_B(b, s) - p = \max_{s' \in [0, \bar{s}]} h_B(b, s') - p_U(s'). \quad (\text{B.9})$$

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

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

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 $s'' > s$,

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

so that

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

On the other hand, 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, 1995, 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.

(3) 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 . Without loss of generality, we 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 $p \leq \bar{p}(s, p_U)$.

(4) 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. ■

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 A that

$\mathbf{b}([0, 1]) \cap B(s, p, p_U) = [b'_1, b'_2]$ for some $0 \leq b'_1 < b'_2 \leq \bar{b}$. The payoff to the seller from (s, p, \mathbf{b}, p_U) is given by

$$H(s, p, \mathbf{b}, p_U) \equiv h_S(b'_1, s) + p. \quad (\text{B.10})$$

This function depends upon 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 (B.3) and both $b_1(p)$ and $b_2(p)$, defined just before Lemma A, 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, define

$$\bar{H}(s, \mathbf{b}, p_U) \equiv \max \left\{ \sup_{p \in \tilde{P}(s, \mathbf{b}, p_U)} H(s, p, \mathbf{b}, p_U), h_S(0, s) + p_U(s) \right\}. \quad (\text{B.11})$$

Notice that if $p \in \tilde{P}(s, \mathbf{b}, p_U)$ for all $p < p_U(s)$, then the first term in (B.11) will be the maximum.¹⁹

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

$$\int (\bar{H}(\mathbf{s}(\sigma), \mathbf{b}, p_U) - c_S(\mathbf{s}(\sigma), \sigma)) d\sigma. \quad (\text{B.12})$$

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).

B.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,

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 A).

we first note that the price-setter's payoff is concave in p_U (note that the buyer's and sellers's payoffs need not be even 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 B *The buyer, price-setter and seller payoff functions given by (B.7), (B.8) and (B.12), are continuous functions of $(\mathbf{b}, \mathbf{s}_B, \mathbf{s}, p_U)$ on Υ^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 follows 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 Υ_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 $p_U[F_B - F_S]$.

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 $\sigma \in E$, 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 .

Finally, we turn to the seller, where the proof of continuity is more involved. It suffices to argue that $\bar{H}(s, \mathbf{b}, p_U)$ 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 for only 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 A.2). We thus need only prove the following claim.

Claim 1 $\lim_{\ell \rightarrow \infty} \bar{H}(s^\ell, \mathbf{b}^\ell, p_U^\ell) = \bar{H}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$.

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

$$\begin{aligned} \limsup_{\ell \rightarrow \infty} \sup_{p \in \tilde{P}(s^\ell, \mathbf{b}^\ell, p_U^\ell)} H(s^\ell, p, \mathbf{b}^\ell, p_U^\ell) &= \sup_{p \in \tilde{P}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)} H(\hat{s}, p, \hat{\mathbf{b}}, \hat{p}_U) \\ \text{and} \quad \lim_{\ell \rightarrow \infty} h_S(0, s^\ell) + p_U^\ell(s^\ell) &= h_S(0, \hat{s}) + \hat{p}_U(\hat{s}). \end{aligned}$$

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

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

We first show that

$$\liminf_{\ell \rightarrow \infty} \bar{\bar{H}}(s^\ell, \mathbf{b}^\ell, p_U^\ell) \geq \bar{\bar{H}}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U). \quad (\text{B.13})$$

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

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

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

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

Since $\hat{\mathbf{b}}$ satisfies (B.3), for sufficiently large ℓ , $\hat{p}' \in \tilde{P}(s^\ell, \mathbf{b}^\ell, p_U^\ell)$, and so (applying Lemma A.3)

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

Hence,

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

yielding (B.13).

We now argue that

$$\bar{H}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U) \geq \limsup_{\ell \rightarrow \infty} \bar{H}(s^\ell, \mathbf{b}^\ell, p_U^\ell), \quad (\text{B.14})$$

which with (B.13) 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

$$H(s^\ell, p^\ell, \mathbf{b}^\ell, p_U^\ell) + \varepsilon \geq \bar{H}(s^\ell, \mathbf{b}^\ell, p_U^\ell). \quad (\text{B.15})$$

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

$$H(\hat{s}, \hat{p}, \hat{\mathbf{b}}, \hat{p}_U) + \varepsilon \geq \limsup_{\ell \rightarrow \infty} \bar{H}(s^\ell, \mathbf{b}^\ell, p_U^\ell), \quad (\text{B.16})$$

which (since it holds for all ε) implies (B.14).

Suppose now that $\hat{p} \notin \tilde{P}(\hat{s}, \hat{\mathbf{b}}, \hat{p}_U)$ or $\hat{p} = \bar{p}(\hat{s}, \hat{p}_U)$. Since \hat{p}_U is differentiable at \hat{s} , there cannot be a nondegenerate interval of buyer attributes indifferent between (\hat{s}, \hat{p}) and the unconstrained optimal seller attribute under \hat{p}_U . This implies $\hat{\mathbf{b}}([0, 1]) \cap B(\hat{s}, \hat{p}, \hat{p}_U) = \{\hat{b}\}$ 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 A.3,

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

and so (taking the lim sup of both sides of (B.15))

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

which (since it holds for all $\varepsilon > 0$) implies (B.14). ■ ■

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 C *The game Γ^n has a Nash Equilibrium, $(\xi_B^n, \xi_S^n, p_U^n) \in \Delta(\Upsilon_B) \times \Delta(\Upsilon_S) \times \Upsilon_P$.*

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

We now examine the limit as $n \rightarrow \infty$. In particular, let $\{(\xi_B^n, \xi_S^n, p_U^n)\}_n \subset \Delta(\Upsilon_B) \times \Delta(\Upsilon_S) \times \Upsilon_P$ 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 that both the sequence of equilibria converges to some limit $(\xi_B^*, \xi_S^*, p_U^*)$, and that players' payoffs also converge.

We now 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 (B.11) 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 will be 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 *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). Let $\xi_{B,b}^*$ and $\xi_{B,s}^*$ denote the marginal distributions induced on buyer and seller attributes chosen by the buyer.

Suppose the buyer's strategy is not pure. 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 give the same distribution of buyer and seller attributes chosen by the buyer, but feature positive assortativity between the buyer's types and attribute choice, and between the buyer's and the seller's attribute

with which the buyer matches, both of which strictly increase the buyer's payoff. Hence, this 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 giving 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 $\Delta(\Upsilon_B) \times \Delta(\Upsilon_S) \times \Upsilon_P$. Recalling that Υ_B, Υ_S , are each endowed with the L^1 norm and Υ_P with the sup norm, and 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 holds: For all $\varepsilon > 0$ there exists n' such that for all $n \geq n'$,

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

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

and

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

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

Lemma E *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 $n \geq n_\varepsilon$,*

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

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

and

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

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

$$\xi_S^n(\{\mathbf{s} \in \Upsilon_S^n : |\mathbf{s}(\sigma) - \mathbf{s}^*(\sigma)| < \varepsilon, \forall \sigma \in E_S^\varepsilon\}) \geq 1 - \varepsilon \quad (\text{B.17})$$

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 $\varepsilon/6$.

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 $[0, 1] \setminus \cup_{k=1}^K (\sigma^k - 2^{-kL}, \sigma^k + 2^{-kL})$; 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} satisfying $\int |\mathbf{s} - \mathbf{s}^*| < \varepsilon\eta/3$. 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 (B.17) then follows, since n'_S can be chosen so that $\xi_S^n(\{\mathbf{s} \in \Upsilon_S^n : \int |\mathbf{s} - \mathbf{s}^*| < \varepsilon\eta/3\}) \geq 1 - \varepsilon$ holds for all $n > n'_S$.

The claim follows from two observations:

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

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

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

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

and, for $n > n''_B$,

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

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

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

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

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 E is immediate from the definition of E_S^ε . \blacksquare

Lemma F *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 $p_U^*(s_1) < p_U^*(s_2)$, with the remaining cases 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 $\hat{p}_U \in \Upsilon_P$ 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 $p_U^\eta \in \Upsilon_P$ 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 $p_U^\eta \in \Upsilon_P$ less the payoff from p_U^* is bounded below by

$$- \int_{s_1 - \eta}^{s_1} (p_U^\eta(s) - p_U^*(s)) ds + \int_{s_1}^{s'} (p_U^\eta(s) - p_U^*(s)) \varepsilon/2 ds. \quad (\text{B.18})$$

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 (B.18) 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. ■

We now seek a characterization of the seller's payoffs. Intuitively, we would like to use Lemma F 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 (B.12) leaves open the possibility that the (gross) payoff to a seller of type σ choosing attribute s may not be given by $h_S(\bar{b}(s), s) + p_U(s)$. Hence, the buyers that 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 some lower types choose a slightly lower attribute, we can keep the average attribute unchanged, while saving costs (from Assumption 3).

Lemma G *The function \mathbf{b}^* is strictly increasing when 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

$$\begin{aligned} \beta_1^\eta &= \inf\{\beta \leq \beta_1 : \mathbf{b}^*(\beta) \geq b + \eta(\beta - \bar{\beta})\} \\ \text{and} \quad \beta_2^\eta &= \sup\{\beta \geq \beta_2 : \mathbf{b}^*(\beta) \leq b + \eta(\beta - \bar{\beta})\}. \end{aligned}$$

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

$$\int_{\beta_1^\eta}^{\beta_2^\eta} h_B(\mathbf{b}^\eta(\beta), \mathbf{s}_B^*(\beta)) - h_B(\mathbf{b}^*(\beta), \mathbf{s}_B^*(\beta)) - [c_B(\mathbf{b}^\eta(\beta), \beta) - c_B(\mathbf{b}^*(\beta), \beta)] d\beta. \quad (\text{B.19})$$

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 (B.19) from the intervals $[\beta_1^\eta, \beta_1]$ and $(\beta_2, \beta_2^\eta]$ is of order $o(\eta)$. 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 (B.7) 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 nonzero. \blacksquare

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 H *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. 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 $n > n''$,

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

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 may 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 E satisfying $\lambda(E^\varepsilon) \geq 1 - \varepsilon$.

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 may 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,

$$\forall \zeta > 0 \exists \varepsilon'' \forall \varepsilon < \varepsilon'', \quad \lambda\{\beta \in E^\varepsilon : \mathbf{b}^*(\beta) \in (\mathbf{b}^*(\sigma') - \zeta, \mathbf{b}^*(\sigma'))\} > 0. \quad (\text{B.20})$$

Consider some $\varepsilon < \varepsilon'$ and suppose $n > \max\{n_\varepsilon, n''\}$, where n_ε is from Lemma E. Let $(\mathbf{b}, \mathbf{s}_B, \mathbf{s}) \in \Upsilon_B^n \times \Upsilon_S^n$ 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 that, from Lemma E, ξ^n assigns high probability to such functions for large n .)

By Lemma F, \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 (B.11) is achieved by $h_S(0, \mathbf{s}(\sigma)) + p_U^n(\mathbf{s}(\sigma))$, then

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

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 (B.20) and the observation that buyers

in E^ε receive a payoff (ignoring costs) 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 attribute satisfying $\mathbf{b}^*(\beta) < \mathbf{b}^*(\sigma)$ will not be attracted (for sufficiently large n). This implies that

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

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 (B.11) is achieved by the first term, and we have a contradiction. \blacksquare

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

Proof. From Lemma F, $\mathbf{s}_B^*(x) = \mathbf{s}^*(x)$ 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 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 F, 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 $\hat{s} + \eta$ satisfying

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

(This is $\bar{p}(\hat{s} + \eta, p, p_U^*)$ from Lemma A(1).) The price \hat{p} is at least as high as the value p' satisfying $h_B(b_2, \hat{s}) - p_U(\hat{s}) = h_B(b_2, \hat{s} + \eta) - p'$. 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 remuneration values ensures that no lower attribute buyers will choose $\hat{s} + \eta$). From Lemma H, we have 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 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. ■

The limiting mass of buyers and seller choosing zero attributes are equal:

Lemma J

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

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, 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 $\lambda(\{\beta : \mathbf{b}^*(\beta) = 0\}) < \lambda(\{\sigma : \mathbf{s}^*(\sigma) = 0\})$. Let $[\sigma', \sigma'']$ be the interval of indices for which $\mathbf{s}^*(\sigma) = 0$ and $\mathbf{b}^*(\sigma) > 0$. Since an interval of low-index buyers are willing to match with the zero seller attribute, and the least-desirable willing buyer enters the seller's payoff calculation (from (B.11)), sellers' payoffs in Γ^n converge to $h_S(\mathbf{b}^*(0), \mathbf{s}^*(\sigma)) + p_U^*(\mathbf{s}^*(\sigma)) - c_S(\mathbf{s}^*(\sigma), \sigma)$, for $\sigma \in [\sigma', \sigma'']$. But, by Lemma H, these sellers have payoffs

converging to $h_S(\mathbf{b}^*(\sigma), \mathbf{s}^*(\sigma)) + p_U^*(\mathbf{s}^*(\sigma)) - c_S(\mathbf{s}^*(\sigma), \sigma)$, a contradiction. ■

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 K *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

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

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 J, 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) and the definitions of \tilde{s}^* and \tilde{b}^* . ■

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

Lemma L For almost all σ ,

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

Proof. The first inequality duplicates Lemma H.

Single-crossing (Assumption 3) implies that the attribute choices maximizing $h_S(\tilde{b}^*(s), s) + p_U^*(s) - c_S(s, \sigma)$ 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. ■

B.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

$$\begin{aligned} h_B(0, 0) - p_U^*(0) &= h_B(0, 0) - p_U^*(0) - c_B(0, \beta) \\ &\leq h_B(\mathbf{b}(\beta), \tilde{s}^*(\mathbf{b}^*(\beta))) - p_U^*(\tilde{s}^*(\mathbf{b}^*(\beta))) - c_B(\mathbf{b}(\beta), \beta) \end{aligned} \quad (\text{B.21})$$

and

$$\begin{aligned} h_S(0, 0) + p_U^*(0) &= h_S(0, 0) + p_U^*(0) - c_S(0, \sigma) \\ &\leq h_S(\tilde{b}^*(\mathbf{s}(\sigma)), \mathbf{s}(\sigma)) + p_U^*(\mathbf{s}(\sigma)) - c_S(\mathbf{s}(\sigma), \sigma). \end{aligned} \quad (\text{B.22})$$

Let

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

(where the inequality follows from Assumption 5) 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 (B.21)–(B.22) gives nonnegative payoffs.

It is immediate from the formulation of the buyer's payoffs in the game and from Lemma L that almost all buyers and sellers are price takers under 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 $\int \bar{H}(\mathbf{s}(\sigma), \mathbf{b}, p_U^n) - c_S(\mathbf{s}(\sigma), \sigma) d\xi^n$.

B.6 Nontriviality

Partial nontriviality. We now show that under (21), the profile $(\mathbf{b}^*, \mathbf{s}^*, \tilde{b}^*, \tilde{s}^*, p_U^*)$ is nontrivial. If the equilibrium is trivial, \mathbf{b}^* and \mathbf{s}^* are identically zero, so that there is no agent for whom it is profitable to trade at price p_U , and hence for all $(b, s) \in (0, \bar{b}] \times (0, \bar{s}]$,

$$\begin{aligned} & h_S(0, s) + p_U(s) - c_S(s, 1) \leq 0 \\ \text{and} \quad & h_B(b, s) - p_U(s) - c_B(b, 1) \leq 0, \end{aligned}$$

where we focus on agents $\beta = 1 = \sigma$ since they are the most likely to want to trade. Notice that we are using here the maximum that appears in the building block (B.11) for the specification of the seller's payoff, and which 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 (21).

Full nontriviality. We now assume (22) holds. Suppose that there is an interval of seller types $[0, \sigma']$ with $\sigma' > 0$ who choose zero attributes. By Lemma I, 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

$$\begin{aligned} & h_S(0, s) + p_U(s) - c_S(s, \phi) \leq 0 \\ \text{and} \quad & h_B(b, s) - p_U(s) - c_B(b, \phi) \leq 0, \end{aligned}$$

where (b, s) are a pair of attributes satisfying (22). But summing these two inequalities yields an inequality contradicting (22).

References

- BILLINGSLEY, P. (1995): *Probability and Measure*. John Wiley and Sons, New York, 3rd edn.
- BULOW, J., AND J. LEVIN (2006): “Matching and Price Competition,” *American Economic Review*, 96(3), 652–668.
- COASE, R. H. (1960): “The Problem of Social Cost,” *Journal of Law and Economics*, 2(1), 1–40.
- COLE, H. L., G. J. MAILATH, AND A. POSTLEWAITE (1995): “Incorporating Concern for Relative Wealth into Economic Models,” *Quarterly Review*, 19(3), 333–373.
- (2001): “Efficient Non-Contractible Investments in Large Economies,” *Journal of Economic Theory*, 101(2), 333–373.
- DUNFORD, N., AND J. T. SCHWARTZ (1988): *Linear Operators Part I: General Theory*. John Wiley and Sons, New York, Wiley Classics Library Edition.
- GLICKSBERG, I. L. (1952): “A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points,” *Proceedings of the American Mathematical Society*, 3(1), 170–174.
- GROSSMAN, S. J. (1981): “The Informational Role of Warranties and Private Disclosure about Product Quality,” *Journal of Law and Economics*, 24(4), 461–483.
- GUERRIERI, V., R. SHIMER, AND R. WRIGHT (2009): “Adverse Selection in Competitive Search Equilibrium,” University of Chicago and University of Pennsylvania.
- HARDY, G. H., J. E. LITTLEWOOD, AND G. PÓLYA (1952): “Inequalities,” Cambridge, UK. Cambridge University Press.
- HOPKINS, E. (forthcoming): “Job Market Signalling of Relative Position or Becker Married to Spence,” *Journal of the European Economic Association*.
- HOPPE, H., B. MOLDOVANU, AND A. SELA (2009): “The Theory of Assortative Matching Based on Costly Signals,” *Review of Economic Studies*, 76(1), 253–281.

- JOHNSTONE, D. B. (2001): "The Economics and Politics of Income Contingent Repayment Plans," Mimeo, Center for Comparative and Global Studies in Education, State University of New York at Buffalo.
- MAILATH, G. J. (1987): "Incentive Compatibility in Signaling Games with a Continuum of Types," *Econometrica*, 55(6), 1349–1365.
- MILGROM, P. R. (1981): "An Axiomatic Characterization of Common Knowledge," *Econometrica*, 49(1), 219–222.
- MILGROM, P. R., AND I. SEGAL (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70(2), 583–602.
- NICHOLSON, S. (2003): "Barriers to Entering Medical Specialties," Working Paper 9649, National Bureau of Economic Research.
- OKUNO-FUJIWARA, M., A. POSTLEWAITE, AND K. SUZUMURA (1990): "Strategic Information Revelation," *Review of Economic Studies*, 57(1), 25–47.
- OSTROVSKY, M., AND M. SCHWARZ (forthcoming): "Equilibrium Information Disclosure and Unraveling in Matching Markets," *American Economic Journal: Microeconomics*.
- PETERS, M. (2009): "Noncontractible Heterogeneity in Directed Search," University of British Columbia.
- ROYDEN, H. L. (1988): *Real Analysis*. Prentice Hall, Englewood Cliffs, NJ.