Rational inattention in dynamic adverse selection

Sander Heinsalu [*]

September 18, 2016

Abstract

I analyze a market with asymmetric information, interdependent values, multiple trading opportunities and trade frictions. The frictions can be reduced at a cost, e.g. by increasing attention, search or computing power. Raising the difference between the values of buyers and sellers can delay trade, despite the greater gains from trade. Rejecting an initial offer is a stronger signal when the offer is more attractive, so a larger difference in values raises the signalling motive, which may overwhelm the increased incentive to accept the better offer. As a result, a subsidy on trade may have the unintended consequence of freezing the market.

Keywords: Lemons market, signalling, asymmetric information, dynamic pricing, rational inattention.

JEL classification: D82, D83, C72.

Markets with asymmetric information (mortgages, health insurance, used cars) may feature inefficiently few transactions. One solution used in practice is to subsidize trade, e.g. by government guarantees, purchases of troubled assets, or a lower interest rate. By increasing the gap between buyer and seller values, such interventions may reduce the lemons problem of Akerlof (1970) enough to restore the efficient level of trade. This insight holds for some one shot markets, but may fail in a dynamic environment with trading frictions, as the current work will show. With multiple opportunities for exchange, a larger difference between buyer and seller values may delay trade, reducing its volume initially. The reason is that greater gains from trade raise the benefit of signalling high quality by initial rejection of offers. The signalling motive may outweigh the increased incentive to trade earlier when there is a larger surplus to share. More initial rejection lowers welfare.

In more detail, there is a privately informed buyer facing a perfectly competitive market of sellers (the situation with an informed seller and competitive buyers is similar and omitted). The buyer is one of two types and knows which, but the sellers only have a common prior belief about the type. Selling to the

high-value type buyer is more costly for the sellers, but gains from trade are positive for any seller matched with any buyer type.

First the buyer privately invests in reducing noise in her later action. Second, each seller makes a price offer. Third, the buyer decides whether to accept or reject the best price among the sellers. The buyer can accept at most one offer. Noise means that a buyer intending to accept sometimes rejects and vice versa (the interpretation is discussed later). If the buyer accepts, the game ends. Fourth, after rejection, the sellers make another offer. Fifth, the buyer again decides to accept or reject. Noise in this last stage, or noise in seller actions does not affect the qualitative results, so these choices are assumed noiseless for simplicity.

Choosing the noise level can be interpreted as rational inattention—the buyer decides how frequently to check advertisements, email or post and how carefully to examine the offers received before clicking 'accept' or 'reject'. The more costly attention is directed to these activities, the lower the chance of missing a trading opportunity or misreading an incoming offer as acceptable when it is in fact unacceptable. The task of paying attention can be delegated to more or less competent workers. For example, on 8 Dec 2005 a trader at Mizuho Securities sold 610,000 shares of J-Com Co. at 1 yen apiece when intending to sell 1 share at 610,000 yen. Mizuho lost 225 million USD. A Deutsche Bank forex dealer mistakenly transferred 6 billion USD to a hedge fund in June 2015. The recipient returned the money the next day.

Investment in reducing noise can take the form of avoiding or fixing technical problems with phone or email which can result in missed offers. Or checking computer code to prevent erroneous trading, e.g. Knight Capital Group went bankrupt on 1 Aug 2012 after losing 450 million USD due to mistaken computerized asset trades. The second interpretation is costly search effort in search and matching situations, which increases the probability of having the opportunity to trade. High frequency trading in financial markets requires investment in rapidly identifying and responding to trading opportunities. Presumably greater investment leads to a higher chance of making the correct decision. The probability is never 1 or 0, regardless of the importance of the decision, e.g. a broker in Japan entered mistaken stock orders worth 617 billion USD on 1 Oct 2014 (the orders were later cancelled).

The solution concept is perfect Bayesian equilibrium (PBE). Because gains from trade are positive for any seller-buyer match, in a one shot interaction all types would trade, which is efficient. With multiperiod trading, reservation values are endogenous and some types may delay trade. Delay is inefficient due to discounting. For a nonempty open set of parameters, there is a unique PBE in which raising the value of the high-value buyer type lowers the probability of initial acceptance for both types. This means the buyer switches from attempting to accept to trying to reject, or lowers the investment in reducing noise when attempting to accept, or increases this investment when trying to reject. Counterintuitively, a higher expected trading surplus delays trade. Raising the value of the low-value buyer increases the probability of acceptance for both types. Raising the values of both types proportionately may still reduce acceptance of

2

the sellers' initial offer.

The reason greater gains from trade may lower welfare is that rejecting an initial offer is a stronger signal when the benefit of accepting that offer is higher. The improvement in the belief of the uninformed sellers after an initial rejection may outweigh the loss of surplus. A better belief leads to a higher offer later, justifying rejecting an offer at the start. More signalling lowers welfare, as is usual in costly signalling.

The gains from trade do not have to rise uniformly across time periods for the result to hold. A higher surplus from exchange in the initial period has the same effect as a greater surplus every period. More gains from trade in the last period reduce the signalling motive. This is somewhat surprising, because signalling is like investment—a cost paid at the start and a benefit reaped later. Despite this, there is less signalling with higher later surplus in the current work, because the greater later gains from trade motivate the low-value type to imitate the high-value more. The increased similarity between the actions of the types leads to belief responding less to signals, lowering the signalling motive for all.

The results are similar when the game is modified so the sellers also have a type—high or low cost, which is publicly realized after the buyer chooses the hidden effort of noise reduction. In this game, for a nonempty open set of parameters, there is a unique equilibrium in which the low-value buyer accepts the initial high price (that the sellers offer when their cost is high) and rejects the initial low price (offered in case of low seller cost). The high-value buyer accepts both prices.

Literature

This paper builds on the literature on markets with adverse selection that started from Akerlof (1970), adding to it rational inattention proposed in Sims (2003). Dynamic adverse selection (in two periods) is studied in Fuchs et al. (2016), who focus on the comparison between public and private offers without trading frictions. They find private offers to be welfare-enhancing. Hörner and Vieille (2009) derive a similar result with an infinite horizon and a different monopolist making an offer each period. With public offers, trade may stop forever after a rejection. The current paper imposes frictions that can be reduced at a cost and considers a different question, namely the comparative statics when the gap between the values of the buyer and sellers increases.

The work closest to the current one is Fuchs and Skrzypacz (2015). They show that government interventions in a dynamic lemons market without frictions may increase or decrease welfare. Fuchs and Skrzypacz exhibit a condition ensuring that an initial subsidy of trade followed by a tax that shuts the market is optimal. The present paper, using different assumptions, derives the opposite result: an initial subsidy with a later tax may freeze the market. Contrary to Fuchs and Skrzypacz (2015), there are frictions and no static adverse selection (all types would trade in a one shot situation). The market freeze is entirely due to dynamic incentives, which may change in either direction when the gap

3

between buyer and seller values increases.

Intervention in a static market with adverse selection has been studied in Philippon and Skreta (2012); Tirole (2012), where a round of government financing of privately informed firms is followed by one shot competitive trade. Government intervention affects this later trade and the expectation of this effect influences the firms' response to the intervention. Both Philippon and Skreta (2012) and Tirole (2012) show that intervention cannot increase welfare in their static frictionless context. The current paper shows it is possible to structure subsidies and taxes to raise welfare in a dynamic environment with frictions. The focus is on the dynamics, which drive the novel results.

Adverse selection is combined with maturity mismatch in Bolton et al. (2011); Heider et al. (2009). The current paper does not model maturity mismatch directly, but this could be one reason for the gains from trade and asymmetric information between the buyer and sellers.

Dynamic competitive markets with adverse selection are studied in Janssen and Roy (2002) who show that equilibrium prices increase over time and eventually all types trade. Unlike Janssen and Roy (2002), the current paper focusses on the comparative statics of interventions, not on the price path. An increasing price is also found in Camargo and Lester (2014); Chiu and Koeppl (2016) in a search context. Camargo and Lester argue that sunset provisions can improve welfare, because the expectation of a future subsidy can delay trade. The present work qualifies this finding—a current or future subsidy targeted to the high-value type may have the opposite effect to targeting the low-value buyer. Chiu and Koeppl present an argument for an increase in total surplus from delaying asset purchasing programs. The current paper shows present and future subsidies have opposite effects for any parameter configuration, but the welfare impact of a future subsidy may be positive or negative.

1   Model and preliminaries

The players are a buyer and a competitive market of sellers. The buyer has a type $\theta \in \{H, L\}$. The buyer knows the type, but the market only has a common prior belief $\mu_0 = \Pr(H) \in (0, 1)$. Type θ buyer values the good at v_θ, with $v_H > v_L$. Selling to type θ costs the market c_θ, normalized to $c_H = 1$, $c_L = 0$. Gains from trade are assumed large enough for all types to trade in a one shot interaction: $v_L > c_H$.

The timing of actions is as follows.

1. Buyer type θ chooses the probability $\rho_\theta \in \left[\frac{1}{2}, 1\right]$ of correctly responding to the initial offer.

2. The market makes a price offer $P_1 \in [0, 1]$ without observing ρ_θ.

3. The buyer observes P_1 and chooses accept or reject, denoted $1, 0$ respectively.

4

4. With probability ρ_θ, the choice is implemented. With probability $1 - \rho_\theta$, action 1 is switched to 0 and vice versa. Realization of 1 ends the game.

5. If action 0 realized, then the market makes a second price offer $P_2 \in [0, 1]$, observing P_1, but not ρ_θ or the action the buyer tried to choose.

6. The buyer observes P_2 and chooses 1 or 0, which is implemented noiselessly.

Restricting prices to $[0, 1] = [c_L, c_H]$ is w.l.o.g. as shown below. Making the final choice of the buyer noisy either exogenously or endogenously (controlled by ρ_θ or by a second investment $\rho_{\theta 2}$ made initially or after 0) does not change the qualitative results. Similarly, a noisy offer by the market has little effect. For simplicity, the market's offer and the buyer's final action are assumed noiseless. The first accept-reject choice of the buyer is subject to controlled noise, which does affect the results. The benchmarks of noiseless actions and exogenous noise are discussed in Section 4. Noise ensures all information sets are on the equilibrium path and all beliefs are determined by Bayes' rule. Controlled noise is realistic (its interpretation is discussed in the Introduction) and removes bang-bang equilibria, which have trivial and extreme comparative statics.

The buyer's strategy is a function $s : \{H, L\} \to \left[\frac{1}{2}, 1\right] \times [0, 1]^{\{1,0\}} \times [0, 1]^{\{1,0\}}$. Type θ's component of the strategy is denoted $s(\theta) = (\rho_\theta, s_1(\theta), s_2(\theta))$. Here, $s_t(\theta) \in \{1, 0\}$ for $t = 1, 2$ is the buyer's attempted choice, which for $s_1(\theta)$ need not realize. The realized choice is denoted $\sigma_t(\theta) \in [0, 1]$, which incorporates both possible intentional mixing and the noise. The strategy that the market expects the buyer to use in equilibrium is denoted s^*.

If the market expects type θ's realized acceptance decision for the first offer to be 1 with probability $\sigma_1^*(\theta)$, then the probability of H conditional on acceptance is

$$\Pr(G|1) := \frac{\sigma_1^*(H)\mu_0}{\sigma_1^*(H)\mu_0 + \sigma_1^*(L)(1 - \mu_0)}. \tag{1}$$

The price is then $P_1 = \Pr(G|1)c_H + (1 - \Pr(G|1))c_L = \Pr(G|1)$. The probability of H conditional on rejection is $\Pr(G|0) := \frac{(1-\sigma_1^*(H))\mu_0}{(1-\sigma_1^*(H))\mu_0 + (1-\sigma_1^*(L))(1-\mu_0)}$. The second offered price is derived similarly to the first. If the market expects type θ to accept the second offer with probability $\sigma_2^*(\theta)$, then

$$P_2 = \Pr(G|01) := \frac{\sigma_2^*(H)\Pr(G|0)}{\sigma_2^*(H)\Pr(G|0) + \sigma_2^*(L)(1 - \Pr(G|0))}. \tag{2}$$

The market's price offers P_1, P_2 are determined by a zero profit condition (the market breaks even, given the types that the market expects to accept). The zero profit price can only be in $[c_L, c_H] = [0, 1]$, so restricting price offers to $[0, 1]$ is w.l.o.g.

The buyer pays a cost $\hat{\kappa}(\rho)$ for choosing ρ, with $\hat{\kappa} \in C^2([\frac{1}{2}, 1))$, $\hat{\kappa}''(\rho) > 0$ for $\rho < 1$, $\hat{\kappa}'\left(\frac{1}{2}\right) = 0$ and $\lim_{\rho \to 1-} \hat{\kappa}'(\rho) > v_H$. Accepting P_1 yields $v_\theta - P_1$ to type θ. There is discounting between offers: accepting P_2 provides the buyer

$\delta(v_\theta - P_2)$, with $\delta \in (0, 1)$. Rejecting P_2 gives zero. The accept-reject decisions choose $\max\{v_\theta - P_1,\ \delta(v_\theta - P_2)\}$ and $\max\{\delta(v_\theta - P_2),\ 0\}$. The total expected payoff when choosing ρ, assuming later choices are optimal is

$$\rho \max\{v_\theta - P_1,\ \max\{\delta(v_\theta - P_2),\ 0\}\} \tag{3}$$
$$+ (1 - \rho) \min\{v_\theta - P_1,\ \max\{\delta(v_\theta - P_2),\ 0\}\} - \hat{\kappa}(\rho).$$

The equilibrium concept is perfect Bayesian equilibrium, hereafter simply called equilibrium.

Definition 1. A buyer strategy s^* and market price offers P_1, P_2 are an equilibrium if

(a) $\delta(v_\theta - P_2) \gtrless 0 \Rightarrow s_2^*(\theta) = \sigma_2^*(\theta) = {1 \atop 0}$,

(b) $v_\theta - P_1 \gtrless \max\{\delta(v_\theta - P_2),\ 0\} \Rightarrow s_1^*(\theta) = {1 \atop 0}$,

(c) ρ_θ^* maximizes (3),

(d) $\sigma_1^*(\theta) = \rho_\theta^* s_1^*(\theta) + (1 - \rho_\theta^*)(1 - s_1^*(\theta))$,

(e) P_1 is given by (1) and P_2 by (2).

The assumption $v_L > 1$ implies $P_t < v_L$ for $t = 1, 2$, so the buyer's final choice is $s_2^*(\theta) = 1$ in any equilibrium. The final price is then $P_2 = \Pr(G|0)$. The buyer's response to the first offer always satisfies $\sigma_1^*(\theta) \in \{1 - \rho_\theta^*,\ \rho_\theta^*\}$, because otherwise the initial cost $\hat{\kappa}(\rho^*)$ could be reduced, keeping the noise the same. Thus the buyer's action sequence $(\rho_\theta, s_1(\theta), s_2(\theta))$ can w.l.o.g. be reduced to just choosing $\sigma_1(\theta) \in [0, 1]$ at cost $\kappa(\sigma_1(\theta)) := \hat{\kappa}\left(\left|\sigma_1(\theta) - \frac{1}{2}\right| + \frac{1}{2}\right)$. The subscript 1 is dropped for simplicity in what follows. The new $\kappa : [0, 1] \to \mathbb{R}$ satisfies $\kappa \in C^2((0, 1))$, $\kappa'' > 0$, $\kappa(\sigma) = \kappa(1 - \sigma)$ and $\sigma \in \left[0, \frac{1}{2}\right) \Rightarrow \kappa'(\sigma) < 0$. A choice $\sigma_\theta < \frac{1}{2}$ is interpreted as intending to reject P_1 and choosing the noise accordingly.

2   Results

After reducing buyer type θ's maximization problem to choosing $\sigma(\theta)$, each type's payoff function is continuous, strictly concave in own action, the strategy set is compact and convex and the market's best response continuous in buyer strategy, so an equilibrium exists by standard arguments. The first order condition (FOC) for type θ is $v_\theta - P_1 - \delta(v_\theta - P_2) - \kappa'(\sigma(\theta)) = 0$. The assumptions on κ and μ_0 ensure an interior solution. Substituting the prices (1), (2) into the FOC and rearranging yields

$$(1 - \delta)v_\theta - \frac{\sigma(H)\mu_0}{\sigma(H)\mu_0 + \sigma(L)(1 - \mu_0)} \tag{4}$$
$$+ \delta \frac{(1 - \sigma(H))\mu_0}{(1 - \sigma(H))\mu_0 + (1 - \sigma(L))(1 - \mu_0)} = \kappa'(\sigma(\theta)).$$
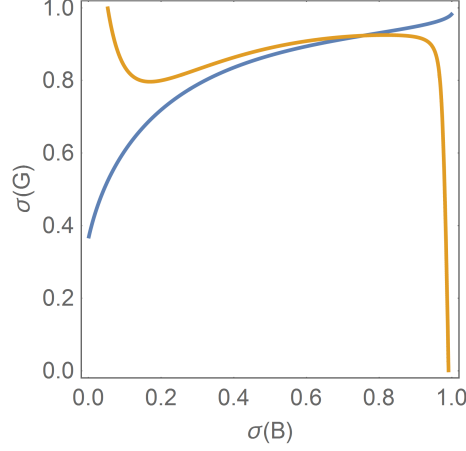
Figure 1: Solutions to first order conditions. Blue: H, orange: L. Parameters $v_H = 1.5$, $v_L = 1.1$, $\delta = \mu_0 = \frac{1}{2}$, $\kappa(\sigma) = -\frac{1}{5}\sqrt{\frac{1}{4} - (\sigma - \frac{1}{2})^2}$.

Solving the system of FOCs, one for each type, is not possible in general. When a closed form exists, it is complicated. A special case is $\delta = 1$ when the unique equilibrium is $\sigma^*(\theta) = \frac{1}{2}$ for $\theta = 1, 2$, regardless of the other parameters.

The solutions to the FOCs of the types are characterized in the following lemma and illustrated in Fig. 1. The dependence of P_t on $\sigma(\theta)$ is written explicitly as $P_t(\sigma(L), \sigma(H))$ when needed.

Lemma 1. *There exists a continuous strictly increasing function $h_H : [0, 1] \to [0, 1]$ s.t. for all $\sigma(L)$,*

$$(1 - \delta)v_H - P_1(\sigma(L), h_H(\sigma(L))) + \delta P_2(\sigma(L), h_H(\sigma(L))) = \kappa'(h_H(\sigma(L))). \quad (5)$$

There exist $S \subseteq [0, 1]$ and a continuous onto function $h_L : S \to [0, 1]$ s.t. S is a union of nonempty closed intervals and for all $\sigma(L) \in S$,

$$(1 - \delta)v_L - P_1(\sigma(L), h_L(\sigma(L))) + \delta P_2(\sigma(L), h_L(\sigma(L))) = \kappa'(\sigma(L)). \quad (6)$$

Further, $h_L(\min S) = 1$, $h_L(\max S) = 0$.

Proof. Fixing $\sigma^*(H)$, the L FOC (4) maps $\sigma^*(L)$ into $\sigma(L)$ continuously. A continuous function from $[0, 1]$ to $[0, 1]$ has a fixed point, so setting $\sigma(L) = \sigma^*(L)$ in the FOC, for every $\sigma^*(H)$ there exists $\sigma^*(L)$ satisfying the FOC. This shows h_L satisfying (6) is defined on $S \neq \emptyset$, $S \subseteq [0, 1]$.

Fixing $\sigma^*(L)$, the H FOC maps $\sigma^*(H)$ into $\sigma(H)$ continuously. So for each $\sigma^*(L)$, there exists a fixed point $\sigma(H) = \sigma^*(H)$ of the FOC. This shows $h_H : [0, 1] \to [0, 1]$ satisfying (5) is defined for all $\sigma(L) \in [0, 1]$.

The left hand side (LHS) of (4) strictly decreases in $\sigma(H)$ and strictly increases in $\sigma(L)$, which is easy to show by taking derivatives. The right hand

7

side (RHS) of (4) for type H is constant in $\sigma(L)$ and vice versa. The RHS is strictly increasing in $\sigma(\theta)$ for type θ, due to $\kappa'' > 0$. For each θ and $\sigma(L)$, the LHS of (4) is strictly decreasing in $\sigma(H)$ and the RHS weakly increasing, so there is at most one $\sigma(H)$ that solves (4).

Both sides of (4) depend continuously on $\sigma(L)$, so both h_H and h_L are continuous on their respective domain. Given θ and $\hat{\sigma}(L)$, if $\sigma(H) \in (0, 1)$ solves (4), then by the continuity of the LHS and the RHS, there exists $\epsilon > 0$ s.t. for $\sigma(L) \in [\hat{\sigma}(L) - \epsilon, \hat{\sigma}(L) + \epsilon]$, the solution to (4) satisfies $\sigma(H) \in [0, 1]$. The union of maximal intervals on which h_L takes values in $[0, 1]$ is denoted S.

The rest of the proof is in the spirit of the monotone comparative statics of Milgrom and Roberts (1994); Milgrom and Shannon (1994). If $\sigma(H)$ solves (4) for H and a given $\sigma(L)$, then increasing $\sigma(L)$ leaves the RHS constant, but strictly raises the LHS. To restore equality of LHS and RHS, $\sigma(H)$ must rise (LHS falls, RHS increases in $\sigma(H)$). Therefore h_H is strictly increasing.

Due to $\kappa'(0) < -v_H$, we have for $\sigma(L) = 0$ and any $\sigma(H)$,

$$(1 - \delta)v_B - \frac{\sigma(H)\mu_0}{\sigma(H)\mu_0} + \delta\frac{(1 - \sigma(H))\mu_0}{(1 - \sigma(H))\mu_0 + 1 - \mu_0} > \kappa'(0),$$

so at $\sigma(L) = \min S$, the LHS crosses the RHS from above. If $\sigma(H) < 1$ when $\sigma(L) = \min S$, then the LHS could be reduced by raising $\sigma(H)$. Then the LHS would equal the RHS at a lower $\sigma(L)$, which contradicts $\sigma(L) = \min S$.

Due to $\kappa'(1) > v_H$, we have for $\sigma(L) = 1$ and any $\sigma(H)$,

$$(1 - \delta)v_B - \frac{\sigma(H)\mu_0}{\sigma(H)\mu_0 + 1 - \mu_0} + \delta\frac{(1 - \sigma(H))\mu_0}{(1 - \sigma(H))\mu_0} < \kappa'(1).$$

so at $\sigma(L) = \max S$, the LHS crosses the RHS from below. If $\sigma(H) > 0$ when $\sigma(L) = \max S$, then the LHS could be increased by reducing $\sigma(H)$. Then the LHS would equal the RHS at a higher $\sigma(L)$, which contradicts $\sigma(L) = \max S$. □

Although derivatives were used in the proof of Lemma 1, the underlying reasoning is that of monotone comparative statics (Milgrom and Roberts, 1994; Milgrom and Shannon, 1994) and extends to more general environments, provided the appropriate monotonicity is preserved. In preparation for describing the comparative statics of equilibria, the next Lemma discusses the comparative statics of the solutions to the FOCs.

Lemma 2. *Increasing v_θ raises $h_\theta(\sigma(L))$ for all $\sigma(L)$, but does not affect $h_{-\theta}(\sigma(L))$ for $-\theta \neq \theta$.*

Proof. Fix $\sigma(L)$. Increasing v_θ raises the LHS of (4) for θ, but does not affect the FOC of $-\theta$. To restore equality for θ, either the RHS must rise or the LHS decrease. For H, both effects are achieved by a rise in $\sigma(H) = h_H(\sigma(L))$. For L, the RHS is fixed, because $\sigma(L)$ is. To decrease the LHS, $\sigma(H) = h_L(\sigma(L))$ must increase. □
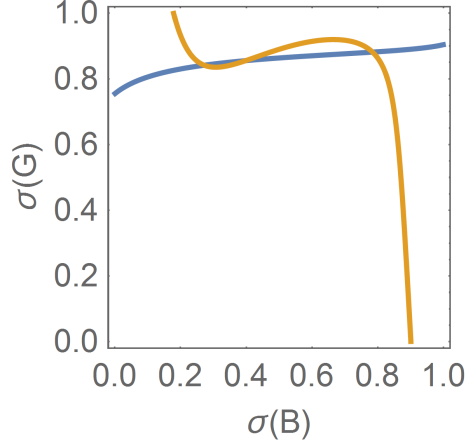
Figure 2: Solutions to first order conditions. Blue: H, orange: L. Parameters $v_H = 2.1$, $v_L = 1.1$, $\delta = \mu_0 = \frac{1}{2}$, $\kappa(\sigma) = \frac{1}{100\sigma(1-\sigma)}$.

There may be multiple equilibria, as illustrated in Fig. 2. Generically, there is an odd number of equilibria (this is a standard result, see e.g. Mas-Colell et al. (1995) p. 598) and stable and unstable equilibria alternate. The minimal and maximal equilibrium (together called *extremal equilibria*) are generically stable. The comparative statics results will focus on the minimal and maximal equilibrium, following Milgrom and Roberts (1994). In any unstable equilibrium, the comparative statics are the opposite to those in the stable equilibria surrounding it. If there are multiple equilibria, they all lie on a strictly increasing function, as shown in Lemma 3. In other words, if one equilibrium has a higher $\sigma(L)$ than another, then it also has a higher $\sigma(H)$. Lemma 3 also shows that $\sigma(H) > \sigma(L)$ in any equilibrium.

Lemma 3. *All equilibria lie on the strictly increasing continuous bijection $f : [0, 1] \to [0, 1]$, $f : \sigma(L) \mapsto \sigma(H)$ that solves*

$$(1 - \delta)v_G - \kappa'(f(\sigma(L))) = (1 - \delta)v_B - \kappa'(\sigma(L)). \qquad (7)$$

Further, $f(\sigma) > \sigma \ \forall \sigma \in (0, 1)$. If $\kappa(\sigma) = \kappa(1 - \sigma) \ \forall \sigma \in [0, 1]$, then $f(\sigma) = 1 - f(1 - \sigma)$.

Proof. Set the FOCs of the types equal and cancel $P_1 - \delta P_2$ to obtain (7) with $\sigma(H)$ instead of $f(\sigma(L))$. Due to $\kappa'' > 0$, the function κ' is strictly increasing. Thus

$$f(\sigma(L)) := (\kappa')^{-1} \left((1 - \delta)(v_H - v_L) + \kappa'(\sigma(L))\right) \qquad (8)$$

is strictly increasing, which implies one-to-one. Because $\kappa'(\sigma(L))$ is defined for all $\sigma(L) \in [0, 1]$, f is also. The same reasoning applies to f^{-1} to show it is

9

defined for all $\sigma(H) \in [0, 1]$, so f is onto. Both sides of (7) are continuous, so f is continuous.

Due to the strictly increasing $\kappa', (\kappa')^{-1}$ and $v_H - v_L > 0$, equation (8) implies $f(\sigma(L)) > \sigma(L)$.

If $\kappa(\sigma) = \kappa(1 - \sigma)$, then $\kappa'(\sigma) = -\kappa'(1 - \sigma)$. Substituting $1 - \sigma(L)$ into (7) yields $(1 - \delta)v_G - \kappa'(f(1 - \sigma(L))) = (1 - \delta)v_B - \kappa'(1 - \sigma(L))$. Using $\kappa'(\sigma) = -\kappa'(1 - \sigma)$ and (7), this becomes $(1 - \delta)v_G + \kappa'(1 - f(1 - \sigma(L))) = (1 - \delta)v_B + \kappa'(\sigma(L)) = (1 - \delta)v_G + \kappa'(f(\sigma(L)))$. This holds for all $\sigma(L)$ and the functions $(1 - \delta)v_\theta + \kappa'(\cdot), \theta = B, G$ are strictly increasing, so $1 - f(1 - \sigma(L)) = f(\sigma(L))$. $\square$

The relationship $f(\sigma) = 1 - f(1 - \sigma)$ means symmetry of the graph of f around the line through $(0, 1)$ and $(1, 0)$. This does not imply the symmetry of the equilibrium set, because equilibria are generically located asymmetrically on the graph of f.

The main result of the paper is Prop. 4 that describes the comparative statics of $\sigma(\theta)$ when v_G or v_L changes. Raising one or both v_θ increases the gap between buyer and seller values. It will be shown that this rise in the trade surplus sometimes results in a delay in trade. The idea of the proof can be seen in Fig. 2, noting that based on Lemma 2, h_H shifts up when v_H increases and h_L shifts up when v_L increases. As a result, $\sigma(L)$ and $\sigma(H)$ can move in any direction, depending on the relative magnitudes of the changes in v_B, v_G. For the parameters in Fig. 2, raising v_H reduces $\sigma(L)$ and increases $\sigma(H)$ in extremal equilibria. Raising v_L increases both $\sigma(L), \sigma(H)$. In the middle equilibrium, the opposite comparative statics hold. In the unique equilibrium in Fig. 1, increasing v_H reduces both $\sigma(L)$ and $\sigma(H)$, but raising v_L increases $\sigma(L)$ and $\sigma(H)$.

Proposition 4. *In extremal equilibria, if v_H increases, then $\sigma^*(L)$ decreases and there are both an open set of parameters in which $\sigma^*(H)$ decreases and an open set of parameters in which $\sigma^*(H)$ increases. If v_L increases, then in the extremal equilibria, both $\sigma^*(L), \sigma^*(H)$ increase.*

Proof. Equilibria are intersections of h_H and h_L in $(\sigma(L), \sigma(H))$-space. By Lemma 1, h_H and h_L are continuous, h_H strictly increasing, $h_L(\min S) = 1$ and $h_L(\max S) = 0$. Due to h_H strictly increasing, $h_H(\sigma(L)) < 1$ for all $\sigma(L) < 1$ and $h_H(\sigma(L)) > 0$ for all $\sigma(L) > 0$. Therefore at extremal equilibria, h_L crosses h_H from above. By Lemma 2, raising v_H increases h_H and leaves h_L constant, so by Milgrom and Roberts (1994) Lemma 1, the extremal intersections of h_H and h_L move left (this is also obvious geometrically, see Figs 1, 2). Thus $\sigma^*(L)$ decreases. By Lemma 2, raising v_L increases h_L and leaves h_H constant, so by Milgrom and Roberts (1994) Lemma 1, the extremal intersections of h_H and h_L move right, so $\sigma^*(L)$ increases. Due to h_H constant, equilibria move along h_H, which is strictly increasing, so if $\sigma^*(L)$ increases, then $\sigma^*(H)$ increases.

Examples showing $\sigma^*(H)$ may increase or decrease in extremal equilibria when v_H increases are in Figs 1, 2. The h_θ are continuous in the parameters by Lemma 1. By continuity, there exists $\epsilon > 0$ s.t. for parameters in an open

ϵ-ball around the values in Fig. 1, $\sigma^*(H)$ increases in extremal equilibria and in an open ϵ-ball around the values in Fig. 2, $\sigma^*(H)$ decreases. $\qquad\square$

The reason why $\sigma^*(L), \sigma^*(H)$ move together when h_L is raised with h_H constant is that the intersections of h_L, h_H move along the strictly increasing h_H. The same reasoning cannot be applied when h_H is raised with h_L constant, because h_L need not be monotone. If h_L is increasing when h_H crosses it from below, then $\sigma^*(H)$ falls when h_H is raised, otherwise $\sigma^*(H)$ increases. A closed form for h_L can be found by solving the L type FOC (4) for $\sigma(H)$. The parameter region where h_L is increasing is easy to compute numerically, but difficult to characterize analytically.

Combining shifts in v_H and v_L combines their effects on the equilibrium strategy. Interventions cannot usually target only one type, because types are unobservable. The effect of an intervention on the values of the types is likely unequal. It depends on the specific environment and intervention whether the good or bad type's value increases more. The consequences of interventions may be in the desired direction or the opposite and are difficult to predict based on observable data.

Between any two distinct equilibria with the comparative statics in Prop. 4, there is generically an equilibrium with the exact opposite comparative statics for $\sigma(L)$. This is a special case of the well known result that stable and unstable fixed points alternate and have opposite comparative statics. The geometric intuition (see Fig. 2) is that generically h_L crosses h_H from below at any equilibrium adjacent to another in which h_L crosses h_H from above. If h_L crosses h_H from below, then the converse reasoning to that in the proof of Prop. 4 holds.

3 Extensions

3.1 Values vary over time

If the value of the buyer can differ between the times of receiving offers, then the reasoning at the end of Section 1 still holds. Denote the value at the time of the first offer of the sellers by $v_{\theta 1}$ and at the time of the last offer $v_{\theta 2}$. The FOC changes from (4) to

$$v_{\theta 1} - \delta v_{\theta 2} - \frac{\sigma(H)\mu_0}{\sigma(H)\mu_0 + \sigma(L)(1-\mu_0)} \tag{9}$$
$$+ \delta \frac{(1-\sigma(H))\mu_0}{(1-\sigma(H))\mu_0 + (1-\sigma(L))(1-\mu_0)} = \kappa'(\sigma(\theta)).$$

The best response characterization in Lemma 1 is unchanged, as is the equilibrium ordering in Lemma 3. The comparative statics of h_θ in Lemma 2 hold for $v_{\theta 1}$, but are the opposite for $v_{\theta 2}$. This is because the effect of $v_{\theta 1}$ in (9) is the same as that of v_θ in (4), but $v_{\theta 2}$ has the reverse effect.

The comparative statics of equilibria are derived from those of best responses as in Prop. 4, so the influence of $v_{\theta 1}$ on $\sigma^*(L), \sigma^*(H)$ is the same as that of v_θ

earlier, but $v_{\theta 2}$ has the opposite effect. Subsidizing trade early may delay trade, but the expectation of later subsidies may increase the probability of trading at the first opportunity. This is the opposite of the result of Fuchs and Skrzypacz (2015) that to encourage trade, there should be an early subsidy and a later tax.

3.2 Uncertain seller cost

Instead of costs $0, 1$ for serving buyer type B, G, the sellers have cost types $\eta = 1, 2$ with $0 \le c_B^\eta < c_G^\eta < v_L$ and $c_\theta^1 < c_\theta^2$. The cost types can be interpreted as levels of market supply. The common prior probability of c_θ^η is p^η. When the buyer chooses the precision of the noise $\rho_\theta \in \left[\frac{1}{2}, 1\right]$, the cost type η is unknown. Otherwise the game would reduce to the baseline in Section 2 for each η. The cost type becomes public before the market makes the first price offer. The results are unchanged if the buyer does not learn η directly, because the offer of the market reveals η.

The definition of equilibrium is modified. Conditional on each cost type, the market makes a price offer that results in zero profit, given the type mix the market expects to accept this offer. Conditional on each cost type and price offer, the buyer chooses accept or reject.

The final offer of the market is still accepted by both types for all η, due to $c_G^2 < v_L$. Unlike in Section 1, the buyer need not choose the minimal or maximal acceptance probability in response to every first offer. However, an interior probability being optimal implies indifference, which can be replaced by the minimal or maximal probability when the incentives for choosing the noise are considered. Denote the realized probability of buyer type θ accepting price P_1^η by $\sigma_1^\eta(\theta)$ and the market's expectation of this by $\sigma_1^{\eta*}(\theta)$. The FOC for the noise choice is $\sum_{\eta=1}^{2} p^\eta |(1-\delta)v_\theta - P_1^\eta + \delta P_2^\eta| = \hat{\kappa}'(\rho_\theta)$, with

$$P_1^\eta := \frac{c_H^\eta \sigma_1^{\eta*}(H)\mu_0 + c_L^\eta \sigma_1^{\eta*}(L)(1-\mu_0)}{\sigma_1^{\eta*}(H)\mu_0 + \sigma_1^{\eta*}(L)(1-\mu_0)}, \tag{10}$$

$$P_2^\eta := \frac{c_H^\eta(1-\sigma_1^{\eta*}(H))\mu_0 + c_L^\eta(1-\sigma_1^{\eta*}(L))(1-\mu_0)}{(1-\sigma_1^{\eta*}(H))\mu_0 + (1-\sigma_1^{\eta*}(L))(1-\mu_0)}.$$

Type θ tries to accept P_1^η if $(1-\delta)v_\theta - P_1^\eta + \delta P_2^\eta > 0$ and tries to reject it if the reverse inequality holds. Due to $v_H > v_L$, if L tries to accept, then H also: $\sigma_1^{\eta*}(H) \ge \sigma_1^{\eta*}(L)$.

There exist parameter values and equilibria in which H accepts both prices, but the L type buyer accepts the high price and rejects the low. Example parameters at which this occurs are $v_H = 2$, $v_L = 1$, $\mu_0 = \delta = c_L^2 = \frac{1}{2}$, $c_L^1 = 0$, $c_H^2 = \frac{61}{64}$ and $c_H^1 = \frac{239}{256}$, with $\hat{\kappa}$ and p^1 chosen so that $\rho_L = \frac{3}{4}$, $\rho_H = \frac{5}{8}$ satisfy the FOC.

It is unusual for a buyer to accept a high and reject a low price. The reason here is that rejecting a low price signals to the sellers that the buyer has a low value. This leads to a lower price next period. Rejecting a high price would not be as strong a signal and the price would not fall enough after that to justify

the delay. Also, the high-value type is more tempted to imitate the low-value by rejecting a high price, because the gain from accepting (that is lost upon rejection) is not as large.

In light of the findings in Section 2, the existence of an equilibrium where one type accepts the high and rejects the low price is not surprising. In some sense, this is merely pasting together the equilibria in Prop. 4 before and after the value change. However, mathematically the model in this section is substantially more complicated, because it does not reduce to a one-dimensional problem for each type and the FOC contains absolute values, so the best responses are kinked.

4 Benchmarks

4.1 Noiseless or exogenously noisy

The action that the buyer chooses is implemented exactly and there is no choice of noise. Type θ accepts the first offer P_1 only if $v_\theta - P_1 \geq \delta v_\theta - \delta P_2$ and rejects it only if the reverse inequality holds. If L accepts, then H also, because $v_H > v_L$.

If the market expects both types to accept the first offer, then $P_1 = \mu_0$ and P_2 is off-path and undefined. If the market expects both types to reject the first offer, then $P_2 = \mu_0$ and P_1 is undefined. Impose the refinement that a deviation to accepting the undefined price comes from the type for whom this deviation is more profitable. This sets $P_2 = 0$ when the market expects both types to accept and $P_1 = 1$ when it expects both to reject. These prices are obtained as limits when H accepts and L mixes with probability of acceptance approaching 1, or when L rejects and H mixes with acceptance probability going to 0. These limits respect the order of best responses that if L accepts, then H also.

The equilibrium where both types accept exists iff $(1 - \delta)v_B \geq \mu_0$. Both rejecting is an equilibrium iff $(1-\delta)v_G \leq 1-\delta\mu_0$, which may hold simultaneously with the preceding condition. The equilibrium where H accepts and L rejects the first offer exists iff $v_G - 1 \geq \delta v_G$ and $v_B - 1 \leq \delta v_B$, equivalently $(1-\delta)v_L \leq 1 \leq (1 - \delta)v_H$. This cannot exist simultaneously with the equilibrium where both types reject. Mixed equilibria may also exist, but the focus is on pure.

The comparative statics are intuitive: raising the gains from trade leads to earlier trading. Increasing v_L may create the equilibrium where both types accept and destroy the equilibrium where H accepts and L rejects. Increasing v_H may create the equilibrium where H accepts and L rejects and destroy the equilibrium where both reject.

Suppose there is exogenous noise such that the buyer's acceptance or rejection is implemented with probability $\rho > \frac{1}{2}$ and switched with probability $1 - \rho$. Then type θ tries to accept the first offer P_1 only if $\rho(v_\theta - P_1) + (1 - \rho)(\delta v_\theta - \delta P_2) \geq (1 - \rho)(v_\theta - P_1) + \rho(\delta v_\theta - \delta P_2)$, which reduces to the condition for the noiseless case.

If the market expects both types to accept or both to reject the first offer, then $P_1 = P_2 = \mu_0$. With $\delta < 1$, it is not possible that both types reject.

The equilibrium where both accept always exists. If the market expects H to accept and L to reject, then $P_1 = \frac{\rho\mu_0}{\rho\mu_0 + (1-\rho)(1-\mu_0)} > P_2 = \frac{(1-\rho)\mu_0}{(1-\rho)\mu_0 + \rho(1-\mu_0)}$. This equilibrium exists iff $(1-\delta)v_L \leq P_1 - \delta P_2 \leq (1-\delta)v_H$. Mixed equilibria may also exist.

Increasing both v_θ does not affect the all-accept equilibrium, but may create or destroy the equilibrium where H accepts and L rejects the initial offer.

4.2 Noiseless with a continuum of types

Buyer types $\theta \in [0,1]$ are distributed according to the atomless prior cdf FThe values of the buyer types are given by a strictly increasing function $v : [0,1] \to (1,\infty)$. The prior expected type is $\bar{\theta} := \int_0^1 \theta dF(\theta)$.

Due to $v(0) > 1$, all types accept the last offer of the market conditional on reaching that point in the game. Let $\sigma^*(\theta)$ denote the probability that the market assigns to buyer type θ accepting the first offer. Then

$$P_1 := \frac{\int_0^1 \theta \sigma^*(\theta) dF(\theta)}{\int_0^1 \sigma^*(\theta) dF(\theta)}, \qquad P_2 := \frac{\int_0^1 \theta (1 - \sigma^*(\theta)) dF(\theta)}{\int_0^1 (1 - \sigma^*(\theta)) dF(\theta)}. \tag{11}$$

If the market expects no type to accept the first offer, then P_1 is off-path and undefined and if the market expects all types to accept, then P_2 is undefined.

A best response of buyer type θ to the first offer of the sellers is $\sigma(\theta) = 1$ if $v(\theta) - P_1 \geq \delta v(\theta) - \delta P_2$. If the reverse inequality holds, $\sigma(\theta) = 0$ is a BR. Clearly, $\theta_1 < \theta_2$ implies $\sigma(\theta_1) \leq \sigma(\theta_2)$. At most one type (denoted θ^*) is indifferent between the first and the second offer, so at most one type mixes for any expectation of the market. The prices (11) can thus be redefined in terms of θ^* as

$$P_1(\theta^*) := \frac{\int_{\theta^*}^1 \theta dF(\theta)}{\int_{\theta^*}^1 dF(\theta)}, \qquad P_2(\theta^*) := \frac{\int_0^{\theta^*} \theta dF(\theta)}{\int_0^{\theta^*} dF(\theta)}. \tag{12}$$

Both prices increase in θ^* and if both prices are defined, then $P_1 > \bar{\theta} > P_2$. Impose the refinement requiring the undefined prices to be the limits of on-path prices when θ^* approaches zero or one. Specifically, if all types accept the first offer, then $\theta^* = 0$ and $P_2 = 0$. If all reject, then $\theta^* = 1$ and $P_1 = 1$.

Each equilibrium is characterized by its θ^*, which is a point where the differentiable function $g(\theta) := (1-\delta)v(\theta) - P_1(\theta) + \delta P_2(\theta)$ crosses zero. Using (12), $g(0) = (1-\delta)v(0) - \bar{\theta}$ and $g(1) = (1-\delta)v(1) - 1 + \delta\bar{\theta}$. The equilibrium where all accept the first offer exists if $(1-\delta)v(0) \geq \bar{\theta}$. The equilibrium where all reject exists for $(1-\delta)v(1) \leq 1 + \delta\bar{\theta}$. Interior equilibria where $\theta^* \in (0,1)$ may also exist.

Increasing $v(\theta)$ for some θ may destroy the all-reject equilibrium, because if $g(1) \leq 0$, then raising g by increasing v may lead to $g(1) > 0$. Similarly, raising v may create the all-accept equilibrium by making $g(0) \geq 0$.

The comparative statics of an interior equilibrium θ^* depend on whether g crosses zero from above or below at θ^*. Increasing $v(\theta)$ for some θ raises

14

$g(\theta)$ for that θ. Thus at equilibria where g crosses 0 from above, θ^* weakly increases in $v(\cdot)$ (fewer types accept the first offer) and if g crosses zero from below, then θ^* decreases. If the minimal equilibrium is not all-accept, then in it increasing v increases trade (θ^* falls), because $g(0) < 0$ and continuity of g imply $g'(\min \theta^*) \geq 0$. Similarly, if the maximal equilibrium is not all-reject, then in it increasing v increases trade. Overall, the minimal and maximal equilibrium (whether corner or interior) have the intuitive comparative static that higher gains from trade lead to earlier trade.

5 Conclusion

Markets with adverse selection and multiple opportunities to trade feature an incentive to signal by delay. This signalling may manifest as seemingly counterintuitive behaviour. Subsidizing trade early and taxing it later may delay trade. Similarly, raising the gains from trade in all periods may delay trade. These findings run counter to the previous literature, which uses different assumptions on the possible actions and the valuations of the trading parties. The importance of considering controllable noise (rational inattention) is illustrated by the absence of the effect in noiseless or exogenously noisy environments.

Generalizations to more than two periods and types are left for future research, as is the addition of moral hazard (investing in increasing one's value). An important modification of the model to study is the case where the uninformed party of the transaction is a monopolist. This would connect the current work to the large literature on the Coase conjecture. Another avenue of extension is to endogenize the values of the buyers and sellers by embedding the model in a search and matching framework. There has been recent interest in search markets with adverse selection and the geometric solution method in this paper may prove tractable in that setting.

References

AKERLOF, G. (1970): "The market for lemons: Quality uncertainty and the market mechanism," *Quarterly Journal of Economics*, 84, 488–500.

BOLTON, P., T. SANTOS, AND J. SCHEINKMAN (2011): "Outside and inside liquidity," *Quarterly Journal of Economics*, 126, 259–321.

CAMARGO, B. AND B. LESTER (2014): "Trading dynamics in decentralized markets with adverse selection," *Journal of Economic Theory*, 153, 534–568.

CHIU, J. AND T. V. KOEPPL (2016): "Trading dynamics with adverse selection and search: Market freeze, intervention and recovery," *The Review of Economic Studies*.

FUCHS, W., A. ÖRY, AND A. SKRZYPACZ (2016): "Transparency and distressed sales under asymmetric information," *Theoretical Economics*.

Fuchs, W. and A. Skrzypacz (2015): "Government interventions in a dynamic market with adverse selection," *Journal of Economic Theory*, 158, 371–406.

Heider, F., M. Hoerova, and C. Holthausen (2009): "Liquidity hoarding and interbank market spreads: the role of counterparty risk," Working Paper 1126, European Central Bank.

Hörner, J. and N. Vieille (2009): "Public vs. private offers in the market for lemons," *Econometrica*, 77, 29–69.

Janssen, M. C. W. and S. Roy (2002): "Dynamic trading in a durable good market with asymmetric information," *International Economic Review*, 43, 257–282.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995): *Microeconomic theory*, New York: Oxford University Press.

Milgrom, P. and J. Roberts (1994): "Comparing equilibria," *The American Economic Review*, 84, 441–459.

Milgrom, P. and C. Shannon (1994): "Monotone comparative statics," *Econometrica: Journal of the Econometric Society*, 62, 157–180.

Philippon, T. and V. Skreta (2012): "Optimal interventions in markets with adverse selection," *American Economic Review*, 102, 1–28.

Sims, C. A. (2003): "Implications of rational inattention," *Journal of monetary Economics*, 50, 665–690.

Tirole, J. (2012): "Overcoming adverse selection: how public intervention can restore market functioning," *American Economic Review*, 1, 29–59.