Confidence Signaling Games

Alex McLeod

April 8, 2016

Abstract

In many buyer-seller environments, the seller has private information about the product's quality, the buyer inspects the product before they bargain, and the seller observes how intensely the buyer inspected. While the seller doesn't know the outcome of the inspection, a seller with a high quality product may bargain more confidently, and thereby signal his type. The signaling game is endogenous, because the buyer could choose to perfectly inspect at no cost, but in equilibrium he does not. The two-type model reveals that the high type sellers that would drop out of the maket in a standard market for lemons model can in fact be more likely to sell their product than the low type sellers; further, the buyers can be better off with a greater presence of low type sellers. In the application to the labour market, with three types of workers and technological routinization over time, there can be a simultaneous substantial decrease in the model's lowest wage and substantial increase in its residual wage inequality followed by reversals, which matches the US data from the 1973-2006 period.

1 Introduction

There are many examples of buyer-seller environments where at first the quality of the product is the seller's private information, but the buyer can do an inspection before they bargain: employers screen job applicants, companies do credit checks of people who want loans and risk assessments of people who want insurance, financial institutions acquire information on available assets, people do inspections of houses and cars that are for sale, and governments investigate matters relevant to lobbyists. These environments are complicated by the fact that the buyer's inspection affects not only his evaluation of the product but also the seller's behaviour when they bargain. In turn, the seller's behaviour further affects the buyer's evaluation to the extent the seller reveals his private information, and of course affects how they split the surplus.

To better understand these environments, I analyze a model where the seller knows the quality of the product while the buyer learns the quality with some probability. The seller can observe how intensely the buyer inspects (that is, the probability that the buyer chooses), but cannot observe whether the buyer has learned the quality or not. Finally, nature decides whether the buyer or the seller makes a take-it-or-leave-it offer and the other responds. I restrict the environment in two ways: first, trade generates a positive surplus for all seller types; second, the high type's outside option is weakly greater than the low type's.

When the buyer inspects imperfectly and the seller gets to make the offer, the result is an endogenous signaling game. Further, it belongs to a new class that I've termed *confidence signaling games*. These are signaling games where (1) at some point before the receiver takes his action, he receives a signal about the sender's type apart from the message sent by the sender, and (2) if it weren't for this additional signal, the game would not be a signaling game. Importantly, the sender must not know the content of the additional signal when he sends his message; if he does, the additional signal would in effect just alter the players' priors. Therefore, what drives the sender types to separate is that the high type is more confident that the receiver will respond positively to his message, which has a different flavour than the usual mechanisms that drive separation: that messages are costly and either the high type has a cost advantage or a greater willingness to pay.

In this particular model, if the buyer's inspection (the additional signal) doesn't reveal the quality of the product, then he has to form beliefs, which may depend on the seller's offer. The buyer is correct to think that the seller's private information may be reflected in his offer, because the seller types don't know whether or not the buyer has become informed; to the extent they think he has, a seller with a high quality (low quality) product has an incentive to capitalize on (concede to) that fact and set a high price (low price). Effectively, the seller may signal his type to a completely uninformed buyer just by how confidently he bargains. Notice that the fact that the buyer has stayed uninformed is what makes him need to infer the product's quality from the price, and at the same time the fact that he could have become informed is what makes the inference work.

The reason that the signaling game is endogenous is that the buyer could choose to inspect perfectly, and then their information would be symmetric. To see why he might not, notice that when (1) the probability that the buyer makes the offer is sufficiently small, or (2) the high type's outside option is not much greater than the low type's, then the buyer decides how intensely to inspect largely based on the effects in the case that the seller makes the offer. Since the seller observes the buyer's choice, if the buyer perfectly inspected, then the seller would just demand the entire surplus. Similarly, if the buyer did not inspect at all, then the seller would just demand the entire expected surplus. Only when the buyer inspects with an intermediate intensity can he hope to make a profit.

However, even when the buyer inspects imperfectly, if the intensity exceeds some threshold, then there is full separation, and the seller still demands the entire surplus in the unique modified undefeated equilibrium. Intuitively, the buyer wants to maintain the possibility that he is uninformed to allow the low type to mimic and thereby force the other types' offers down. In a two-type model this just implies that the sellers pool; a variant of the two-type model, however, reveals that the high type sellers that would drop out of the maket in a standard market for lemons model instead stay in the market, which generates price dispersion, and can be more likely to sell their product than the low type sellers. In a three types model, of course, there remain the questions of which type(s) the low type seller will pool with and whether the high and middle type seller will also pool.

The main result there is that when the middle type's quality is either sufficiently high or sufficiently low, all three seller types will pool. The simple part of the explanation is that if the middle type's quality is above the mean, then the high type prefers to pool with both the middle and low types (full pooling) over pooling just with the low type (high-low pooling) simply because it raises his profit. If instead the middle type's quality is below the mean, then the high type would of course prefer high-low pooling to full pooling except for the fact that under high-low pooling the uninformed buyer may have to reject with some probability to deter the middle type from mimicking. If the middle type's quality is very low, then his incentive to mimic will be high, and the uninformed buyer would have to reject with a high probability. On the other hand, if his skill level is very close to the mean, then his incentive to mimic will be low, and the informed buyer may not have to reject with a positive probability at all– the possibility that the buyer is informed will be sufficient. Effectively, the middle type uses a go big (tell a big lie) or go home (tell the truth) strategy. These explanations about the high type's preferences over equilibria are important, because they are a dominant factor in which equilibrium will be played. The more he prefers an equilibrium, the more intensely the buyer can inspect without the high type preferring to separate completely.

In the employer-worker environment, these results provide insight into some of the most mysterious recent trends in wage inequality. In particular, while Autor et al. (2003) argue that one of the major effects of computer capital was to substitute for workers doing routine tasks– which ought to cause a monotonic compression in the productivities of low wage workers– the 10th percentile male (female) wage substantially decreased then receovered, and at the same time the 50/10 residual male (female) wage ratio substantially increased then recovered. With the middle type's go big or go home strategy in mind, this makes perfect sense. If the middle type's initial skill level is sufficiently low, then the compression decreases his incentive to mimic and can lead to a transition from full pooling to high-low pooling. The transition followed by the continued routinization will generate a simultaneous substantial decline in the middle type's wage (the lowest wage in the model) and increase in residual wage inequality followed by reversals.

The principal contribution of this paper is that it identifies the class of confidence signaling games, which to my knowledge there have been only a few examples of. In

Judd and Riordan (1994), each consumer views the firm's price as a signal of how much he values their product but also draws his personal experience with the product. In Cotton (2009, 2015), a politician views the size of each lobbyist's bid for access as a signal of the strength of their case but also draws on the information presented to him by the lobbyist who actually wins. In Barbos (2013), an evalutor views which tier the agent applied to as a signal of the project's quality but also draws on his own evaluation.

More broadly, this paper contributes to the literature on bilateral contracts with information acquisition[1]. These papers do not address the situation where the party making the offer is informed and the party responding to the offer is uninformed but can acquire private information before the contract is offered. This situation is common, in several markets especially, and importantly generates a different kind of strategic information acquisition where the information acquisition decision can affect the informational content of the offer.

Lastly, it contributes to the literature on wage inequality. Acemoglu (1999), Caselli (1999)[2], Galor and Moav (2000), and Shi (2002) generate a simultaneous decrease in the wages of low skilled workers and increase in residual wage inequality. Galor and Tsiddon (1997), Greenwood and Yorukoglu (1997)[3], and Aghion, Howitt, and Violante (2002) generate increases in overall or residual inequality followed by a reversal. However, Greenwood and Yorukoglu generate the reversal only with a roughly simultaneous reversal of the productivity growth slowdown, which is yet to happen, and Aghion, Howitt, and Violante generate only a very small decrease relative to the increase when their model is calibrated. As mentioned above, these papers do not address the differential behaviour of the upper and lower halves of the residual inequality distribution.

The rest of the paper is organized as follows. In section 2, I present the model where there are two types and solve for the unique undefeated perfect Bayesian equilibrium, which makes the three type model in section 3 much easier to understand. In section 4, I illustrate the effects of technological routinization in the three type model with a simulation, and relate the results to the data. Section 5 concludes.

2 The Two Type Model and Equilibrium Analysis

2.1 Model

The sequence of events is as follows:

[1]See Sobel (1993), Lewis and Sappington (1993), Shavell (1994), Cremer and Khalil (1994), Cremer (1995), Kessler (1998), Nosal (2006), Dang (2008), Kaya (2010), Lester, Postlewaite, and Wright (2012).

[2]Caselli (1999) generates an overall wage inequality reversal driven entirely by compositional effects.

[3]Greenwood and Yorukoglu (1997) also generate a very slight decline in the wages of unskiled workers but only with a counterfactual decline in the wages of skilled workers.

1. The seller's type T, his product's quality, is drawn from the distribution $P(H) = \mu = 1 - P(L)$ where $L < H$.
2. The seller observes his type.
3. The buyer chooses an inspection intensity x.
4. The seller observes the buyer's inspection intensity.
5. The buyer observes the seller's type with probability x and observes nothing with probability $1 - x$.
6. Who makes the take-it-or-leave-it offer is drawn from the distribution $P(buyer) = G = 1 - P(seller)$.
7. That person makes a price offer p.
8. The other person accepts or rejects. If he accepts, then the buyer gets a payoff of $\Pi = T - p$ and the seller gets a payoff of p. If he rejects, then the buyer gets a payoff of u_B and the high (low) type seller gets a payoff of u_H (u_L).

Four assumptions should be highlighted. First, trade generates a positive surplus for all seller types: $u_B + u_L < L$ and $u_B + u_H < H$. Second, the high type's outside option is weakly greater than the low type's: $u_L \leq u_H$. These two assumptions imply that u_B and u_L can be normalized to zero while u_H is weakly greater than zero, which leaves the vector of parameters $z = \{u_H, \mu, H, L, G\}$.

Third, the inspection costs are zero to highlight the strategic effects of the buyer's inspection intensity; however, they are simple to incorporate, and I return to this point later.

Fourth, I assume that $G = 0$. To some extent, discussed in more detail later, the equilibrium is not sensitive to the exact value of G: (1) the buyer's offer is independent of G, and the seller's offer depends on G only through the buyer's inspection intensity, which is the subject of the next two points; (2) when $u_H = 0$, the inspection intensity equals $x^*(z) < 1$ for all $G < 1$; (3) when $u_H > 0$, the inspection intensity equals $x^*(z) < 1$ for $G \leq \hat{G}(z)$, and otherwise equals one.

As the final note before the analysis, the reason that the price is determined by a take-it-or-leave-it offer is that the buyer and the seller both have private information: the seller knows the product's quality and the buyer knows whether or not he knows the product's quality; so any protocol where they both take actions will involve signaling both directions, which would be complicated to solve. That said, the mechanisms at play here should be at play under any protocol where the seller has some power– this one is the just the simplest.

2.2 Equilibrium Analysis

The solution concept is perfect Bayesian equilibrium (PBE).

The solution method is as follows. For each x, I find the equilibria of the second stage of the game, and apply the undefeated equilibrium refinement. Then I check which x the buyer would actually choose.

2.2.1 Pooling Equilibria

In a pure strategy pooling equilibrium, by definition, all players play pure strategies, and the type H and type L seller set the same price.

Recall that a PBE specifies beliefs that satify Bayes' rule on the equilibrium path, but that are otherwise unrestricted, and strategies that are optimal given the beliefs. Although, in principle, the off equilibrium path beliefs could be complicated, any equilibrium that can be supported with some beliefs can be supported with the beliefs that minimize the payoff to a deviation. More precisely, if there is a pooling equilibrium with equilibrium price $\hat{p}$ and some beliefs of the uninformed buyer, then there is a pooling equilibrium with equilibrium price $\hat{p}$ and the beliefs of the uninformed buyer that, if the price is $p = \hat{p}$, then the probability that the seller is a type H seller is $q = \mu$ and otherwise $q = 0$. Thus, finding the set of pooling equilibria only requires finding the set of pooling equilibria that can be supported with these beliefs.

Having specified the beliefs, what remains to be specified are the conditions under which pooling strategies are optimal given those beliefs. Of course, the strategy of the informed (uninformed) buyer is to accept non-negative payoffs (non-negative expected payoffs) and reject negative payoffs (non-negative expected payoffs). For the sellers, it cannot be optimal to pool on a price greater than $\mu H + (1-\mu)L$, because the uninformed buyer would reject, and then the type H seller and the type L seller would be better off appealing only to the informed buyer and setting the separating prices H and L respectively. Thus, condition 1 is $\hat{p} \leq \mu H + (1-\mu)L$. The next condition relies on the fact that, if the pooling price were too low, then the type H seller would be better off appealing only to the informed buyer and setting the price H. Thus, condition 2 is $xH + (1-x)u_H \leq \hat{p}$. The last condition depends on whether the sellers are pooling on a price greater than L or equal to L, because the type L seller faces risk in the first case but not in the second. In the first case, similarly to the type H seller's condition, if the pooling price were too low, then the type L seller would be better off appealing only to the informed buyer and setting the price L. In the second case, the type L seller has no potentially profitable deviations. Thus, condition 3 is either $L \leq \hat{p}(1-x)$ or $L = \hat{p}$. The following proposition summarizes.

Proposition 2.1. *In a pooling equilibrium, the type H and type L sellers set the price $\hat{p}$ where (i) $\hat{p} \leq \mu H + (1-\mu)L$, (ii) $xH + (1-x)u_H \leq \hat{p}$, and (iii) either $L \leq \hat{p}(1-x)$*

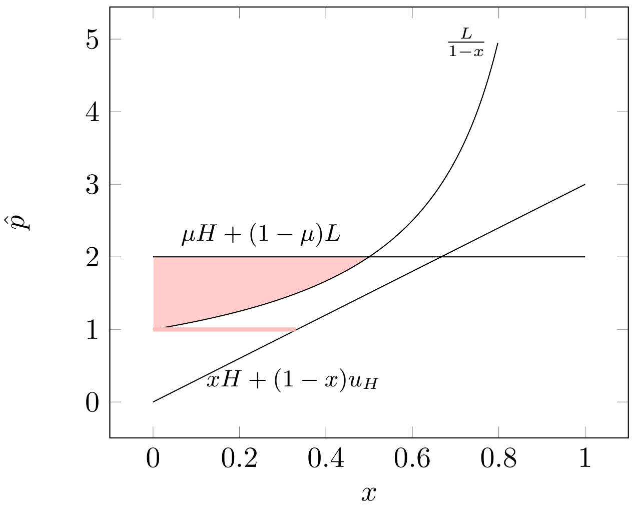

Figure 1: The correspondence maps from the inspection intensity to the set of pooling equilibrium prices. The figure depicts the case where $L/(1-x)$ intersects $\mu H + (1-\mu)L$ before $xH + (1-x)u_H$ does: $u_H < \mu H + (1-\mu)L - \sqrt{\mu H(H-L)}$.

or $L = \hat{p}$. Any pooling equilibrium can be supported by the uninformed buyer's beliefs that if the price is $\hat{p}$, then the probability that the seller is a type H seller is $q = \mu$ and otherwise $q = 0$.

Figure 1 (Figure 2) depicts the set of pooling equilibrium prices for the case where $L/(1-x)$ intersects $\mu H + (1-\mu)L$ before (after) $xH + (1-x)u_H$ does. Figures 3 and 4 (Figures 5 and 6) repeat for separating (mixed strategy) equilibrium prices. When $xH + (1-x)u_H$ does not intersect $\mu H + (1-\mu)L$, there are no pooling equilibria, and in fact the seller just demands the entire surplus. In figures 1-7, the parameter u_H is equal to zero for aesthetic reasons; it's easy to visualize the differences if instead $u_H > 0$.

2.2.2 Separating Equilibria

In a pure strategy separating equilibrium, by definition, all players play pure strategies, and the type H and type L sellers set different prices. Through reasoning similar to that for pooling equilibria, we have the following proposition.

Proposition 2.2. *In a separating equilibrium, the type L seller sets the price $p = L$ and type H seller sets the price $\hat{p}$ where (i) $\hat{p} \leq H$, (ii) $xH + (1-x)u_H \leq \hat{p}$, and (iii) $\hat{p}(1-x) \leq L$. Any separating equilibrium can be supported by the uninformed buyer's beliefs that if the price is $\hat{p}$, then the probability that the seller is a type H seller is $q = 1$ and otherwise $q = 0$.*

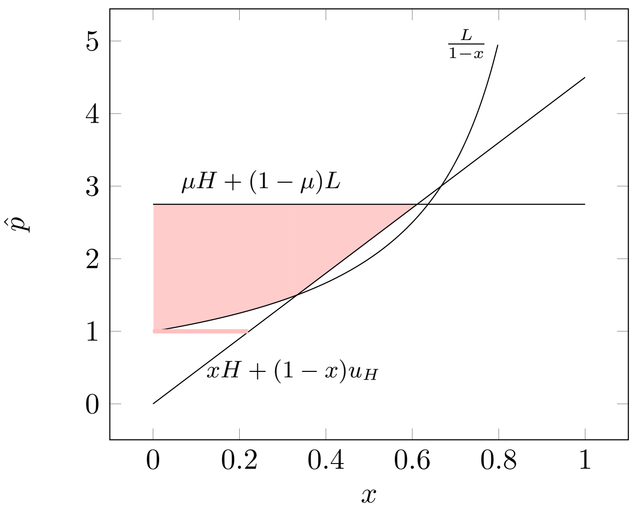

Figure 2: The correspondence maps from the inspection intensity to the set of pooling equilibrium prices. The figure depicts the case where $xH + (1-x)u_H$ intersects $\mu H + (1-\mu)L$ before $L/(1-x)$ does: $u_H > \mu H + (1-\mu)L - \sqrt{\mu H(H-L)}$

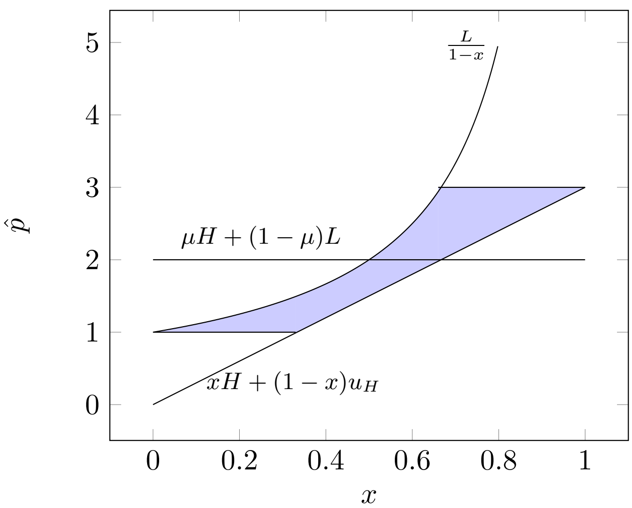

Figure 3: The correspondence maps from the inspection intensity to the set of separating equilibrium prices. The figure depicts the case where $L/(1-x)$ intersects $\mu H + (1-\mu)L$ before $xH + (1-x)u_H$ does: $u_H < \mu H + (1-\mu)L - \sqrt{\mu H(H-L)}$

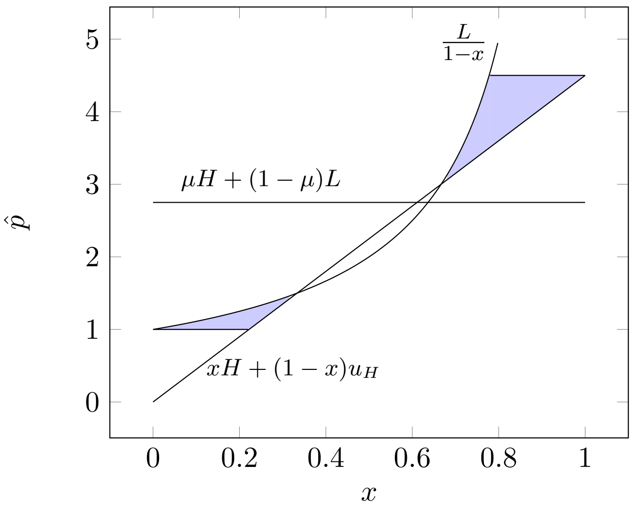

Figure 4: The correspondence maps from the inspection intensity to the set of separating equilibrium prices. The figure depicts the case where $xH + (1-x)u_H$ intersects $\mu H + (1-\mu)L$ before $L/(1-x)$ does: $u_H > \mu H + (1-\mu)L - \sqrt{\mu H(H-L)}$

Here the reason for the two cases becomes clear. When $u_H > \mu H + (1-\mu)L - \sqrt{\mu H(H-L)}$, there is an interval where $\mu H + (1-\mu)L < xH + (1-x)u_H$ and $L/(1-x) < xH + (1-x)u_H$. In that interval, no pure strategy pooling equilibrium exists and no pure strategy separating equilibrium exists. This makes sense– if there were pooling on a price greater than $xH + (1-x)u_H$, then the uninformed buyer would reject and the type L seller would prefer to separate; if there were separating and the type H seller were to set a price greater than $xH + (1-x)u_H$ and less than H, then the uninformed buyer would accept and the type L seller would prefer to pool. This intuition is suggestive, and indeed there are mixed strategy equilibria in the interval where no pure strategy equilibria exist. There are also mixed strategy equilibria outside that interval, and unlike most other signaling games, these mixed strategy equilibria can be important in the overall game.

2.2.3 Mixed Strategy Equilibria

In a mixed strategy equilibrium, by definition, one or more of the players plays a mixed strategy. I assume that the informed buyer will always accept when indifferent, which eliminates the possibility of him playing a mixed strategy, and that the uninformed buyer will always accept the price L. For the analogous proposition and proof, see the appendix.

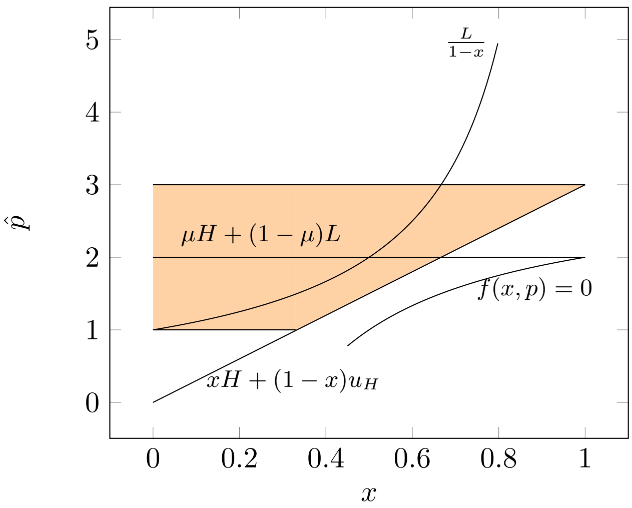

Figure 5: The correspondence maps from the inspection intensity to the set of mixed strategy equilibrium prices. The figure depicts the case where where $L/(1-x)$ intersects $\mu H + (1-\mu)L$ before $xH + (1-x)u_H$ does: $u_H < \mu H + (1-\mu)L - \sqrt{\mu H(H-L)}$. The function $f(x,p) = u_H L + p[x(H-p) - L]$.

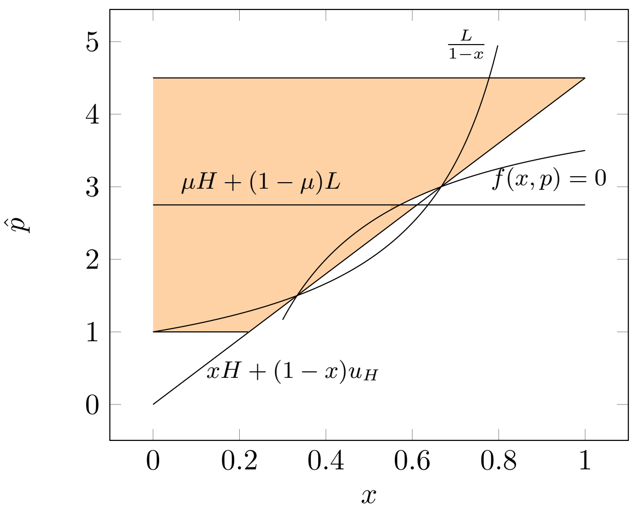

Figure 6: The correspondence maps from the inspection intensity to the set of mixed strategy equilibrium prices. The figure depicts the case where $xH+(1-x)u_H$ intersects $\mu H+(1-\mu)L$ before $L/(1-x)$ does: $u_H > \mu H+(1-\mu)L-\sqrt{\mu H(H-L)}$. The function $f(x,p) = u_H L + p[x(H-p) - L]$.

2.2.4 The Equilibrium Refinement

What may be the most familiar PBE refinement, the intuitive criterion, is insufficient here– it does not refine the set of equilibria with prices above $L/(1-x)$. Moreover, en route to their improvement, Mailath, Okuno-Fujiwara, and Postlewaite (1993) argue that the usual motivation for the intuitive criterion, as well as for several other refinements, is illogical. In particular, after some type(s) adjust their strategies, the motivation goes that the receiver should reconsider his beliefs starting with the idea that "All players were guaranteed at least their equilibrium payoffs," and then sometimes proceed to adjust his beliefs in a way that makes a deviation payoff increasing for some type(s). However, the receiver knows that he follows that line of thinking and that the other players know that. Thus, following an action that is not a deviation, the receiver knows that the other player cannot be a type for whom a deviation is payoff increasing. It is then not necessarily true that "All players were guaranteed at least their equilibrium payoffs." Instead, after some type(s) adjust their strategies, the receiver ought to adjust his beliefs using Bayes' rule while keeping in mind that all other players, not just the ones for whom a deviation is payoff increasing, might respond by adjusting their strategies. But beliefs that follow Bayes rule given strategies and strategies that are optimal given beliefs is exactly the PBE concept. It is that line of thinking– that deviations should be to an equilibrium not to an arbitrary message– that leads to the authors' "undefeated equilibrium refinement."

An equilibrium (E') ***defeats*** another equilibrium (E) when (i) there exist some E' equilibrium price (p) for some type(s) (T) that is not an E equilibrium price, (ii) at least one of the type(s) in T strictly prefers and all of the type(s) in T weakly prefer their payoffs in the E' equilibrium when setting the price p to their payoffs in the E equilibrium, and (iii) the uninformed buyer's beliefs that support the E equilibrium are inconsistent with (a) the strictly preferring type(s) in T setting the price p with the probability one and (b) the weakly preferring type(s) in T setting the price p with any probability.

An equilibrium is ***undefeated*** if there is no equilibrium that defeats it.

Through direct application of the definition of defeats, the set of pooling equilibria is refined to the one where the sellers set the price $\mu H + (1-\mu)L$, and the set of separating equilibria is refined to the one where the type H seller sets the price $\min\{L/(1-x), H\}$. However, as mentioned in their paper, the refinement is not written to apply to mixed strategy equilibria, because they were not important in the signaling games the authors had in mind. Thus, it too is insufficient– it refines the set mixed strategy equilibria where the type H seller sets a price greater than $\mu H + (1-\mu)L$ (the other mixed strategy equilibria are defeated by the undefeated pooling equilibrium) to the subset where the type H seller sets the price H, which the uninformed buyer accepts with a probability $\lambda \leq \hat{\lambda} = min\{\frac{L}{(1-x)H}, 1\}$, and the type L seller sets the price L. Happily, it can be modified in a very reasonable way to apply to mixed strategy equilibria, that is, to apply when deviating to an equilibrium doesn't necessarily mean deviating to a particular price. At the same time, since the requirement that "all of the type(s) in T weakly prefer their payoffs in the E' equilibrium when setting the price p to their payoffs in the E equilibrium" is completely unmotivated and has real consequences in the three type model, I simply remove it.

In the modified definition, an equilibrium (E') ***defeats*** another equilibrium (E) when (i) there exists some E' equilibrium price (p) for some type(s) (T) that is not an E equilibrium price, (ii) at least one of the type(s) in T strictly prefers their payoff in the E' equilibrium when setting the price p to their payoffs in the E equilibrium, and (iii) the uninformed buyer's beliefs that support the E equilibrium are inconsistent with (a) the strictly preferring type(s) in T setting the price p with the probability that they do in the E' equilibrium and (b) all other type(s) in T setting the price p with any probability less than or equal to the probability with which they do in the E' equilibrium.

With the modified definition of defeats, the set of remaining mixed strategy equilibria is refined to the one where the type H seller sets the price H, which the uninformed buyer accepts with probability $\hat{p}$, and the type L seller sets the price L. Each of the equilibria where the uninformed buyer accepts with probability $\lambda < \hat{\lambda}$ is now defeated by an equilibrium where the type H seller sets the price $H - \epsilon(\lambda)$, which the uninformed buyer accepts with probability $min\{\frac{L}{(1-x)(H-\epsilon(\lambda))}, 1\}$. Notice that, without the modification, the buyer could believe that the type L seller deviated to the price $H - \epsilon(\lambda)$ with probability one despite the fact that the type L seller sets that price with a very small probability in those equilibria. I use only the modified definition of "defeats" throughout the rest of the paper.

For each x, which of the within group undefeated equilibria– pooling, separating, or mixed strategy– is the undefeated equilibrium of the second stage of the game[4] depends on which the type H seller prefers, that is, which of $\mu H + (1-\mu)L$, $\min\{L/(1-x), H\}$, and $\min\{xH + L + \frac{(1-x)H-L}{H}u_H, H\}$ is the greatest. Since the third is always greater than or equal to the second, it is only the first and the third in tension. Of course, the first does not exist for some x, but the third would be greater for those x anyways. The following proposition summarizes.

Proposition 2.3. *The unique undefeated equilibrium of the second stage of the game is as follows: (1) For* $x \leq \hat{x} = \frac{H-L}{H-u_H}[\mu - u_H/H]$*, the type* H *and type* L *seller set the price* $\mu H + (1-\mu)L$*, which the informed buyer accepts if and only if the seller is a type* H *and the uninformed buyer accepts. (2) For* $x > \hat{x}$*, the type* H *and type* L *sellers set their prices equal to* H *and* L *respectively. The informed buyer accepts, while the uninformed buyer accepts* H *with probability* $\hat{\lambda} = min\{\frac{L}{(1-x)H}, 1\}$ *and accepts* L*.*

Since the undefeated equilibrium of the second stage of the game is a mixed strategy equilibrium for some values of x, it's worth the detour to think about the famous critique

[4]The two-step refinement process works, because if an equilibrium that is refined away from group A in the first step can defeat the within-group undefeated equilibrium from group B, then the undefeated equilibrium from group A also defeats it.

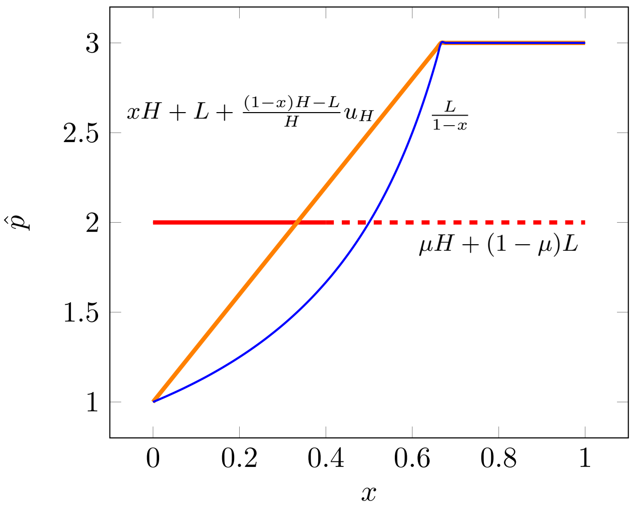

Figure 7: The type H seller's expected payoff is $\mu H + (1-\mu)L$ in the undefeated pooling equilibrium, $min\{\frac{L}{1-x}, H\}$ in the undefeated separating equilibrium, and $\min\{xH + L + \frac{(1-x)H-L}{H}u_H, H\}$ in the undefeated mixed strategy equilibrium.

of mixed strategy equilibria that a player who is indifferent between two actions would be unlikely to play them with exactly the probability that supports the equilibrium, and also about the famous response, due to Harsanyi (1973), that a mixed strategy equilibrium where player i plays a mixed strategy can be interpretated as the limit of a sequence of pure strategy equilibria, where there is a continuum of player i types, as the difference between the types goes to zero. In the context of this game, the interpretation is very compelling– if the undefeated mixed strategy equilibrium is taken as the limit of a sequence in n where the type H seller sets the price $H - \epsilon_n$, which is in line with the refinement process above, then the interpretation is just that even the uninformed buyer (who is the player who plays a mixed strategy) receives some very noisy signal from a continuous distribution. Then uninformed buyers with sufficiently postive noisy signals accept and the others reject.

2.2.5 The Buyer's Inspection Intensity

To begin, notice that if $\mu H < u_H$, then $\hat{x} < 0$. In that case, the buyer's inspection intensity is indeterminate and the seller just demands the entire suplus. Therefore, I assume $u_H < \mu H$ in what follows.

That aside, the optimal inspection problem is simple. For x less than $\hat{x} = \frac{H-L}{H-u_H}[\mu - u_H/H]$, the type H and type L sellers set the price $\mu H + (1-\mu)L$, and otherwise set the prices H and L respectively. The buyer's expected surplus increases with x up to the point $\hat{x}$, since prices do not change and the probability that he's informed increases, but beyond that his surplus drops to zero. Thus, despite the fact that inspection with any intensity is completely costless, the buyer just chooses $x = \hat{x}$. Intuitively, since the type H seller either wants to set a safe medium price or a risky high price, the buyer needs to keep the high price sufficiently risky. The following proposition summarizes.

Proposition 2.4. *The unique undefeated equilibrium of the game is that the buyer chooses $\hat{x} = \frac{H-L}{H-u_H}[\mu - u_H/H]$, the type H and type L sellers set their price equal to $\mu H + (1-\mu)L$, which the informed buyer accepts if and only if the seller is a type H and the uninformed buyer accepts.*

2.2.6 Comparative Statics and Robustness

For comparative statics under the restriction $u_H = 0$, the buyer's expected surplus equals $\mu^2(1-\mu)\frac{(H-L)^2}{H}$, which increases with μ for $\mu < 2/3$, increases with H, and decreases with L; the type H seller's expected profit equals $\mu H + (1-\mu)L$, which increases with μ, H, and L; lastly, the type L seller's expected profit equals $[1 - \mu(H - L)/H][\mu H + (1-\mu)L]$, which increases with μ for $\mu < 1/2$, and increases with H and L. Intuitively, everyone wants the expected quality to be high, but the buyer also wants some uncertainty about the quality to capture an informational rent, and the low type seller doesn't want the buyer to inspect intensely.

In the model section, I mentioned that the results were robust to the addition of inspection costs and to a positive probability that the buyer makes the offer ($G > 0$). First, inspection costs are easy to incorporate. In particular, the buyer would never inspect more intensely than the threshold intensity whether or not there are costs, and depending on the cost function he may choose to inspect with any lesser intensity. The analysis just involves a comparison of the marginal value of information to the marginal inspection cost, since the sellers' prices do not change below the threshold.

Second, the buyer's offer is independent of G: he offers u_H to the high type and zero to the low type if he's informed, and offers $p^*(z) = \{u_H, 0\}$ if he's uninformed, and the seller's offer only depends on G indirectly via the buyer's choice of inspection intensity. That is, the purely mechanical effect aside, the price won't change with G unless the optimal inspection intensity does. With respect to that, the buyer's expected surplus weakly increases with x if he makes the offer while, if the seller makes the offer, his expected surplus increases until $\hat{x}$ then falls to zero. Therefore, if the buyer doesn't pick $x = \hat{x}$, then he'll pick $x = 1$. But the value of the extra information is zero (small) when $G = 0$ (G is small) or when $u_H = 0$ (u_H is small). More precisely, there exist $\hat{G}(z)$ and $\hat{u}_H(z)$ such that (1) for any $u_H > 0$, the inspection intensity equals $\hat{x}$ for all $G \leq \hat{G}(z)$, and (2) for any $G > 0$, the inspection intensity equals $\hat{x}$ for all $u_H \leq \hat{u}_H(z)$.

2.2.7 The Market for Lemons

The results from the two-type model sharply contrasts those from Akerlof's (1970) famous paper *The Market for Lemons*[5]. To see this, consider a variant of the two type model where a fraction α of type H sellers have an outside option $l = 0$ and the rest have an outside option $h > \mu H + (1 - \mu)L$. In Akerlof's paper, the equilibrium is that the type H_h sellers drop out of the market, while the type H_l and the type L sellers set their price equal to $\alpha\mu H + (1 - \mu)L$ and successfully sell their cars.

The present paper shows what a difference the buyers' inspections can make. In particular, (1) the type H_h sellers stay in the market, set their price equal to H, and successfully sell their cars with probability $\hat{x} + L/H$ where $\hat{x} = \frac{\alpha\mu}{1-(1-\alpha)\mu}\frac{H-L}{H}$, and (2) the type L sellers only sell their cars with probability $1 - \hat{x}$. Thus, far from Akerlof's results, when $\alpha\mu > 1 - \mu$ both the type H_h and type H_l sellers are in fact more likely to make a sale than the type L sellers, and when $\alpha\mu > 2(1 - \mu)$ the buyers would actually be better off with a greater presence of type L sellers as shown in the comparative statics subsection above.

In addition, this version matches several key features of the used car market: (1) customers elect to take test drives, (2) some customers buy a car that they know could be a lemon, and some refuse to buy a car that they know to be a lemon, (3) some customers refuse to buy a car that they know isn't a lemon just because the price is too much, and (4) some sellers demand what the car is actually worth, while others set a lower price and have a better chance to make a sale.

2.2.8 Discussion

The solution of the two type model is characterized by simplicity. In buyer-seller models where the inspection intensity is exogenous, sellers will pool on the expected quality for intensities below the threshold and separate to their true quality otherwise. In models where it is endogenous, the buyer chooses the threshold intensity and the sellers pool. The results are robust to different outside options for seller types, inspection costs, and the probability with which the seller makes the offer. All these factors likely make the two type model easy to embed into models where the analysis is already complicated by other components.

However, the simplicity derives from the fact that the buyer cannot make any profit when there is any separation, which does not carry over when there are more than two types. In fact, if a middle type's quality were just below the mean, then the buyer would have to inspect with a very low intensity for all the sellers to pool, and then he would

[5]Levin (2001) studies the market for lemons in the presence of an exogenous public signal, which differs from the present paper in that the accuracy of the signal isn't the buyer's choice and the seller's price isn't a signal of quality.

make a very small expected profit. Instead, he could abandon that plan, inspect with a higher intensity, and have only some types pool. Effectively, with more than two types, the buyer faces a tradeoff between the inspection intensity and what kind of equilibrium will be played in the second stage of the game. A complete treatment of confidence signaling games should take this tradeoff into account, and of course the application to wage inequality requires a model where the workers do not always pool on a wage.

In a model with infinite types each of measure zero, however, there is no undefeated equilibrium of the second stage where a positive measures of sellers pool. The proof is in the appendix, but intuitively there is a free rider problem, because if positive measures were to pool, then each above average type in the pool would prefer to separate and leave the other above average types to "babysit" the below average ones. But despite the fact that no two products have exactly the same quality, there are reasons to doubt the infinite type model prediction. First, with infinite types of measure zero, equilibria where positive measure of sellers pool can Pareto dominate the mixed strategy equilibrium where each type separates. Second, we think that sometimes sellers actually do bargain more confidently to appear to have a high quality product, that is, sellers actually do pool. Third, since one equilibrium defeats another when a seller credibly informs the buyer that the first equilibrium is being played, we should wonder when a seller sets a price whether, in reality, a buyer would believe that seller alone plays that strategy or that similar sellers play similar strategies. If we think the latter, then the free rider problem may disappear.

All that considered, to have the most realistic results likely requires a number of types that is greater than two and less than infinity. I find that the three type model captures the tradeoff between the inspection intensity and what kind of equilibrium will be played in the second stage of the game, and also with technological routinization can match the important trends in wage inequality mentioned in the introduction.

3 Three Type Model and Equilibrium Analysis

The three type model is identical to the two type model except that (1) the seller's type T is drawn from the distribution $P(H) = P(M) = P(L) = \frac{1}{3}$ where $H > M > L$ where L is normalized to one in all calculations, and (2) the seller types have equal outside options $u_H = u_M = u_L$.

The approach for the two type model was to find all the equilibria of the second stage of the game. In the three type case, this task would be much more involved and some of them would be impossible to depict on a 2D graph. Instead, I can skip that step and just find the within-group undefeated equilibria for each x. Then I can proceed as before to find the between group undefeated equilibrium for each x, and finally find which x the buyer will actually choose. It turns out that it is easier to split the analysis into two cases. I present the case one analysis below and, since it is similar but more complicated, put the case two analysis in the appendix.

3.1 Case I: $M > \frac{1}{2}(H + L)$

3.1.1 Within-Group Undefeated Equilibria

Full Pooling: In the set of full pooling equilibria, where all three types of sellers set the same price, the undefeated equilibrium is the one where the sellers set the price $\frac{1}{3}(H + M + L)$ which the uninformed buyer accepts. The conditions under which this equilibrium exists are that $xH \leq \frac{1}{3}(H + M + L)$ and that $L \leq (1 - x)\frac{1}{3}(H + M + L)$.

HL Pooling Equilibria: In the set of HL pooling equilibria, where the type H and the type L seller set the same price but there is not full pooling, there is no undefeated equilibrium.[6] However, there can be important HL pooling equilibria where the type M seller sets the price M with a very small probability, that is, where the sellers' expected payoffs are different by only a very small amount from the full pooling equilibrium. Of course, the conditions under which some of these equilibria exist are the same as the conditions for the undefeated full pooling equilibrium with weak inequalities replaced with strict inequalities.

HM Pooling Equilibria: In the set of HM pooling equilibria, where the type H and the type M seller set the same price but there is not full pooling, each equilibrium is defeated by the equilibrium where each seller sets his price equal to his type and the uninformed buyer accepts H and M with the highest probabilities such that the type L seller does not strictly prefer to mimic either one (henceforth the "mixed separating equilibrium").[7]

ML Pooling Equilibria: In the set of ML pooling equilibria, where the type M and the type L seller set the same price but there is not full pooling, the undefeated equilibrium is the one where the type H seller sets the price H which the uninformed buyer accepts with the highest probability such that the type L seller does not strictly prefer to mimic the type H seller and the type M and type L sellers set the price $\frac{1}{2}(M+L)$

[6] To see this, notice that, in the equilibrium where the type H and the type L sellers set the price $\frac{1}{2}(H + L)$ and the type M seller sets the price M, the uninformed buyer accepts the price $\frac{1}{2}(H + L)$ with probability one and accepts the price M with the probability such that the type H and the type M sellers are indifferent between the two prices. As the type M seller decreases the probability with which he sets the price M and increases the probability with which he pools with the other types, the HL pooling price increases, the uninformed buyer accepts the price M with a higher probability, and all the sellers' expected payoffs increase. Thus, the sequence of defeating HL pooling equilibria converges to the undefeated full pooling equilibrium, which is outside the set of HL pooling equilibria.

[7] To see this, notice that the type H seller prefers the price H in the mixed separating equilibrium over the price $\frac{1}{2}(H + M)$ in his most preferred HM pooling equilibrium, and the type M seller is indifferent between the price M in the mixed separating or in his most preferred HM pooling equilibrium.

Equilibrium	**Type H's Expected Payoff**	**Type M's Expected Payoff**
Full Pooling	$\frac{1}{3}(H+M+L)$	$\frac{1}{3}(H+M+L)$
HL Pooling	$\frac{1}{3}(H+M+L)$	$\frac{1}{3}(H+M+L)$
ML Pooling	$xH+(1-x)\frac{1}{2}(M+L)$	$\frac{1}{2}(M+L)$
Mixed Separating	$min\{xH+L, H\}$	$min\{xM+L, M\}$
HL-ML Pooling	$< xH+(1-x)\frac{1}{2}(M+L)$	$< \frac{1}{3}(H+M+L)$

Table 1: The type H and type M sellers' expected payoffs in each within-group undefeated equilibrium for Case I.

which the uninformed buyer accepts. The conditions under which this equilibrium exists are that $xM \leq \frac{1}{2}(M+L)$ and that $L \leq (1-x)\frac{1}{2}(M+L)$.

Remaining Equilibria: First, if the type L seller sets the price L with positive probability, then the undefeated equilibrium is the mixed separating equilibrium, which always exists. Second, if the type L seller does not set the price L with positive probability, then the undefeated equilibrim is one where the type H seller sets a price greater than M, the type M seller sets a price less than M, and the type L seller sets both their prices with positive probability, that is, a sort of HL-ML pooling equilibrium. To see this, notice that each equilibrium where the type H and the type L seller pool on a price and the type H seller also sets another price is defeated by an equilibrium where he does not. As he shifts probability from the other price to the pooling price, the pooling price increases, and so does his expected payoff. Then, since there is not full or HL pooling, the type M and type L sellers must also pool on a price, and there the same statements are then true with type M replacing type H. The conditions under which these equilibria exists are that $xH \leq xp_{HL} + (1-x)p_{ML}$ and that $L \leq (1-x)p_{ML}$.

3.1.2 Between-Group Undefeated Equilibrium

Moving forward, through finding the values of x where the type H or the type M seller's preferences over within-group undefeated equilibria change, there is immediately a problem– there can be no between-group undefeated equilibrium for some values of x. In Figure 8, which depicts a not-at-all atypical situation, there is a shaded interval where the ML pooling equilibrium defeats every other equilibrium, but is itself defeated by the HL pooling equilibrium. Intuitively, the type H seller prefers to deviate from HL to ML pooling and the type M seller prefers to deviate from ML to HL pooling; that is, each type prefers that the other pool with the low types.

What would be useful then is a compelling reason to strengthen the requirements for one equilibrium to defeat another. Happily, the situation in the right subinterval

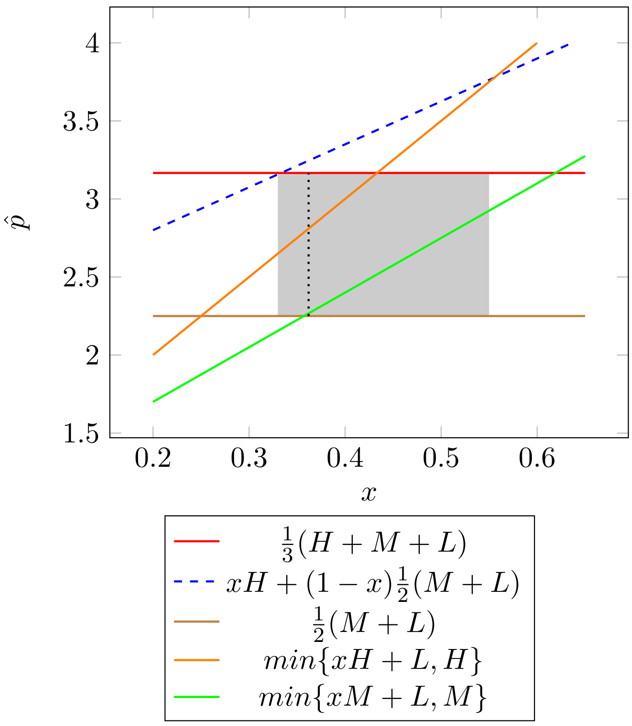

Figure 8: The type H seller's expected payoff is $\frac{1}{3}(H+M+L)$ in the full and HL pooling equilibria, $xH+(1-x)\frac{1}{2}(M+L)$ in the ML pooling equilibrium, and $min\{xH+L,H\}$ in the mixed separating equilibrium. The type M seller's expected payoff is $\frac{1}{3}(H+M+L)$ in the full and HL pooling equilibria, $\frac{1}{2}(M+L)$ in the ML pooling equilibrium, and $min\{xM+L,M\}$ in the mixed separating equilibrium.

where $xM+L > \frac{1}{2}(M+L)$ suggests such a reason. Recall that the motivation for the undefeated equilibrium refinement was that, after one or more types adjust their strategies, the buyer ought to adjust his beliefs through Bayes' rule while keeping in mind that all other players would respond by adjusting their strategies. Thus, in the right subinterval, an HL pooling equilibrium is defeated by the ML pooling equilibrium, because the type H seller adjusts his price to H and the type M and type L sellers respond by adjusting their prices to $\frac{1}{2}(M+L)$. In that subinterval, however, among the equilibria where the type H seller sets the price H, the type M seller prefers the mixed separating equilibrium not the ML pooling equilibrium. A very compelling reason then to strengthen the requirements for one equilibrium to defeat another is that the other sellers ought not just respond but best respond. That is, after the type H adjusts his price to H, the type M seller ought to adjust his price to M and then the type L seller ought to adjust his price to L. Of course, the type H seller will anticipate the adjustments of the other sellers before making his initial adjustment, and not adjust his price to H until he prefers $xH+L$ to $\frac{1}{3}(H+M+L)$.

An equilibrium ($E' \in \mathcal{E}'$) ***iteratively defeats*** another equilibrium (E) when E'

defeats E where $\mathcal{E}'$ is the set of undefeated equilibria of the game where the type(s) T set their E' equilibrium prices.

Using the iteratively undefeated equilibrium refinement, the within-group analysis above continues to hold word for word except with *defeats* and *undefeated* replaced with *iteratively defeats* and *iteratively undefeated*. The differences come in the between-group analysis.

3.1.3 Between-Group Iteratively Undefeated Equilibrium

The iteratively undefeated equilibrium refinement improves matters, but still there can be no between-group iteratively undefeated equilibrium for some x. First, there remains the problem that the ML pooling equilibrium and an HL pooling equilibrium iteratively defeat each other in the left subinterval. Second, in the right subinterval, though the ML pooling equilibrium does not iteratively defeat an HL pooling equilibrium, a similar problem emerges with the HL-ML pooling equilibria in place of the ML pooling equilibrium. To see this, notice that, a small amount after $xM + L$ crosses $\frac{1}{2}(M + L)$, the type H seller can set a price a small amount less than H and thereby make it possible for the uninformed buyer to accept with a positive probability while a fraction of type L sellers set the type H seller's price. Then the remaining fraction of type L sellers set the type M seller's price, which can now be greater than $\frac{1}{2}(M + L)$, and, more importantly, greater than $xM + L$. In effect, the type H and type M sellers go from disagreeing about who should pool with the type L sellers to disagreeing about exactly what fraction of type L sellers the other should pool with.

However, when we work within the PBE structure and find the motivation for the undefeated equilibrium compelling, these problems become fundamental to any bargaining game with at least two types that are better than the worst type. A perfect analogy is two people (the type H and type M sellers) bargaining over how to split a surplus where, after reaching an agreement, one of them is chosen at random to inform a banker (the buyer) about their argeement. Just as there can be multiple equilibria that iteratively defeat each other, there can be multiple ways to split the surplus that are acceptable to both sides, and just as there will be no overall iteratively undefeated equilibrium, there will be no agreement that the banker can believe, because each person would just say that the agreement was one that is the most favourable for himself. Therefore, there is no reasonable prediction for the game in the intervals where no between-group iteratively undefeated equilibrium exists. Since it's not that too many equilibria are reasonable but rather that none are reasonable, I simply assume that the buyer will not choose an x in those intervals.

We are now ready to find the values of x where the type H or type M sellers' preferences over within-group equilibria change. There are three important crossings: [1] $xH + (1 - x)\frac{1}{2}(M + L)$ crosses $\frac{1}{3}(H + M + L)$, [2] $xM + L$ crosses $\frac{1}{2}(M + L)$, and [3]

$xH+L$ crosses $\frac{1}{3}(H+M+L)$. Through simple algebra, it can be shown that [1] occurs at $x=\frac{1}{3}$ and [2] occurs before [3]. The possible orderings are [1][2][3], [2][1][3], and [2][3][1]. Each requires a separate explanation; only the first explanation, the most complicated of the three, has an accompanying graph. It is extremely useful to switch back and forth between the explanation and the graph. Notice throughout the explanations that the type H, type M, and type L sellers' expected payoffs are never less than xH, xM, and $L/(1-x)$, respectively, and thus that there is no problem of the conditions under which the equilibria exist not being satisfied.

[1][2][3]: Before [1], there is full pooling. From [1] to [2], ML and HL pooling iteratively defeat each other. After [2], there is an inteval where HL-ML and HL pooling iteratively defeat each other. In that interval, for HL-ML pooling to iteratively defeat HL pooling, the type H seller's price must be sufficiently low that that type M seller's price can be greater than $xM+L$. If either (i) the type H seller's price falls below M (the M-constrant), or (ii) the type H seller's expected payoff falls below $\frac{1}{3}(H+M+L)$ (the payoff constraint), then after that there would be an interval where there is full pooling. Unfortunately, though the M-constraint cannot bind and later unbind, this is not necessarily true for the payoff constraint. Thus, if the payoff constraint binds, then the full pooling interval may end either where the payoff constraint becomes slack, and then the process may even repeat, or at [3]. If either constraint binds for the final time before [3], then after it does there is full pooling followed by mixed separating. If not, then after [3], the bound on the payoff constraint becomes $xH+L$ rather than $\frac{1}{3}(H+M+L)$ and there is an interval where two HL-ML pooling equilibria iteratively defeat each other– one that the type H seller prefers and one that the type M seller prefers. Anywhere after [3] that either of the constraints exactly binds, there is HL-ML pooling. The reason that there is HL-ML pooling exactly where one of the constraints binds if and only if one binds after [3] is that, before [3], HL pooling still iteratively defeats the HL-ML pooling equilibrium, but after [3], nothing defeats it– it and mixed separating are both iteratively undefeated. After one of the constraints binds for the final time, there is mixed separating. [$M>3$]

[2][1][3]: Before [1], there is full pooling. After [1], there is an interval where there continues to be full pooling. To see this, notice that since [2] occurs before [1], then just after [1] the type H seller's full pooling payoff is still greater than his HL-ML pooling payoff if the type M seller's payoff is greater than or equal to $xM+L$. In effect, the type H seller's HL-ML pooling payoff has to catch up. If it doesn't ever catch up or doesn't before the M-constraint binds, then the full pooling interval continues until [3] and after that there is mixed separating. If it does, then after that, rather than "After [2]" the explanation is the same as it was for the [1][2][3] case. [$2<M<3$]

[2][3][1]: Before [3], there is full pooling. After [3], there is mixed separating. [$M<2$]

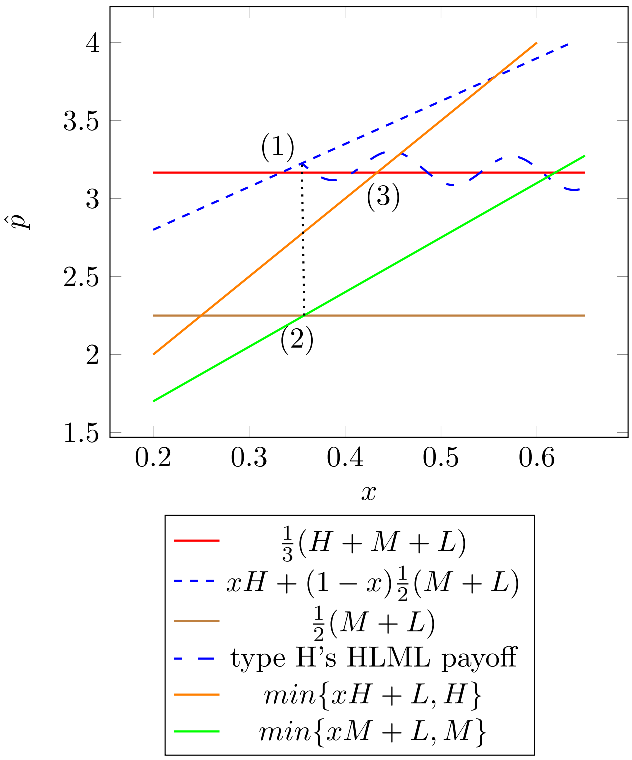

Figure 9: The type H seller's expected payoff is $\frac{1}{3}(H+M+L)$ in the full and HL pooling equilibria, $xH+(1-x)\frac{1}{2}(M+L)$ in the ML pooling equilibrium, and $min\{xH+L,H\}$ in the mixed separating equilibrium. Here the curve that represents the type H seller's HL-ML pooling payoff is purely illustrative. The type M seller's expected payoff is $\frac{1}{3}(H+M+L)$ in the full and HL pooling equilibria, $\frac{1}{2}(M+L)$ in the ML pooling equilibrium, and $min\{xM+L,M\}$ in the mixed separating equilibrium. The graph depicts the case where $M>3$.

3.2 Buyer's Inspection Intensity

Just as in the two-type model, the buyer will never choose to induce full separation, and if he were to inspect with an intensity that induces full pooling, then he would inspect with the hightest intensity that induces full pooling. The important difference in the three-type model is that the buyer has two other options. By similar logic– that his information improves while the prices remain constant– it is also true that if he were to inspect with an intensity that induces HL pooling, then he would inspect with the highest one. It turns out, see appendix, that the same can be said for HL-ML pooling. Thus, for any values of H and M (recall that L is normalized to one) the buyer will have at most three inspection intensities to choose from.

The best way to get a sense of which equilibrium the buyer will choose is to run through a (6 x 1000) grid in (H,M)-space. Each of the eight graphs in Figure 10 corresponds to a different value of H. Along the x-axis, the value of M ranges from L to H. Along the y-axis is the equilibrium of the overall game where "1" means full pooling, "2" means HL-ML pooling, and "3" means HL pooling.

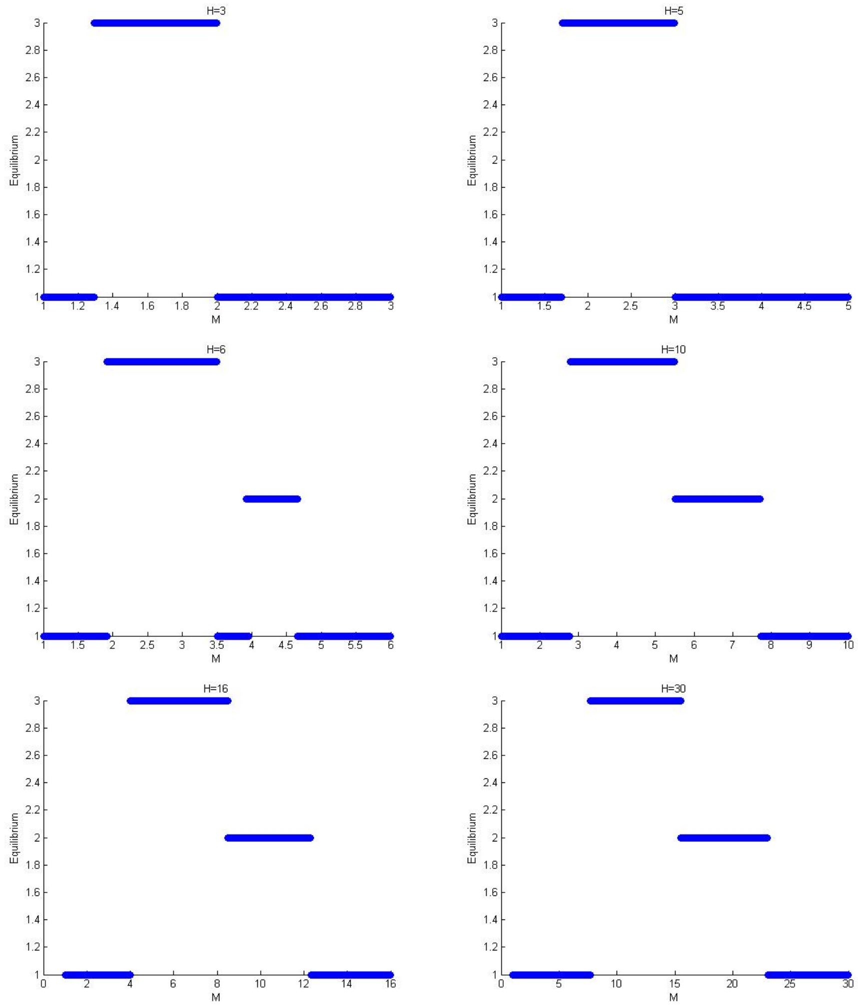

Figure 10: Each of the eight graphs corresponds to a different value of H. Along the x-axis, the value of M ranges from L to H. Along the y-axis is the equilibrium of the overall game where "1" means full pooling, "2" means HL-ML pooling, and "3" means HL pooling.

There are some features that require explanation. The focus will be mostly on the type H seller's preferences and then on the buyer's profit. The type H seller's preferences are a dominant factor in which equilibrium will be played, because the more he prefers an equilibrium, the more intensely the buyer can inspect without him preferring to separate.

First, for $M < \frac{1}{2}(H + L)$, there is always full pooling followed by HL pooling. This follows from the type H seller's preference for full pooling over HL pooling when M is sufficiently small relative to H but not when it is more intermediate. The reason for this preference is that, as M increases relative to H, the type M seller's incentive to mimic the HL pooling price rather than separate decreases, and so does the probability with which the uninformed buyer has to reject the HL pooling price to deter mimicking. Though there exists a between group undefeated HL-ML pooling equilibrium for some values of x, it happens that the buyer never chooses it. Second, for $M > \frac{1}{2}(H+L)$, there is never HL pooling, which this time follows from the type H seller's preference for full pooling over HL pooling when M is sufficiently large relative to H. The reason for this preference is just that the addition of the type M seller increases the full pooling price.

For $M > \frac{1}{2}(H + L)$, what remains to be explained is the choice between full pooling and HL-ML pooling. It turns out that there is HL-ML pooling wherever there exists a between group undefeated HL-ML pooling equilibrium, and there is full pooling otherwise. There is effectively an interval of M values where neither the payoff constraint nor the M constraint binds for the final time before $xH + L$ crosses $\frac{1}{3}(H + M + L)$. For small values of H, the interval is empty; for intermediate values, the interval is non-empty; for large values, the interval extends all the way down to $M = \frac{1}{2}(H + L)$.

Applied to the labour market, if M begins sufficiently small relative to H and technological routinization for routine tasks increased the values of H, M, and L and decreased the ratios of H/M, H/L, and M/L, then the type M seller's wage would fall in the transition from full pooling to HL pooling but in general rise, and residual inequality would rise in the transition from full pooling to HL pooling but in general fall. That is, the model has the potential to match some of the most mysterious trends in the wage inequality literature.

4 Applications to Trends in Wage Inequality

In this section, I compare the prediction of the repeated static model while the parameters change to the data. Alternatively, I could create a simple dynamic model, and assume that the workers' skill levels were random across matches with firms. This assumption would ensure that the worker's outside option does not depend on his type, and also that the distribution of types in the unemployment pool does not change as the types' relative probabilities of acceptance changes. However, since the only important endogenous variable is newly bargained wages, which the repeated static model and the dynamic model would capture just as well, I chose the simpler approach.

To make the connection between the model and wage dynamics, I need to know how the workers' relative productivities change over the period, which arguably was in large part determined by technological advances, in particular computer capital. As mentioned in the introduction, Autor et al. (2003) argue that one of the major effects of computer capital was to substitute for workers who do routine tasks. Since workers who perform routine tasks tend to have low levels of education (Table 5b Acemoglu and Autor 2012) and thus low wages, we should expect a compression in the relative productivities of low wage workers. Lastly, figure one in their paper suggests that these changes started to take effect sometime in the mid-seventies.

In the simulation, the fraction of each type of worker is equal to a third. The initial productivities are set to $H = 10$, $M = 2.5$, $L = 1$, and to impose the technological routinization hypothesis, the sequences are set to $H_t = 1.0015H_{t-1}$, $M_t = 1.0035M_{t-1}$, and $L_t = 1.0055L_{t-1}$, which if each period is a quarter impies an average annual growth rate of 1.14% for low skilled workers' productivities.

The comparisons will be between (1) the lowest wage in the simulation and the 10th percentile male and female wages over the period 1973-2006, and (2) the ratio of the highest wage to the lowest wage in the simulation and the 50-10 ratio of residual wages over the period.

4.1 Episodic Fall in 10th Percentile Wages

Figures 11 and 12 depict the cumulative log change in the smoothed wage of the middle skilled workers from the simulation and the cumulative log change in the wage of the 10th percentile worker for males and females (Acemoglu and Author 2012), respectively. The dynamics of all three graphs are qualitatively similar. In all three graphs, the wage initially increases, then it experiences an episodic decrease, and finally it increases again. These dynamics occur in the model because the middle type's wage falls in the transition from full to HL pooling but in general rises due to technological progress.

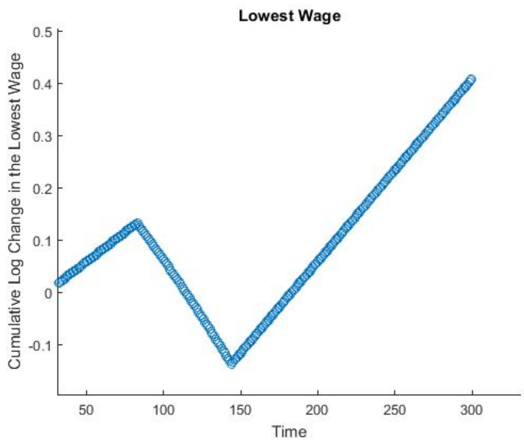

Figure 11: The graph depicts the cumulative log change in the smoothed wage of the middle skilled workers from the simulation. The smoothed wage of middle skilled workers in a period is the mean of their simulated wages in that period, the thirty periods before, and the thirty periods after.

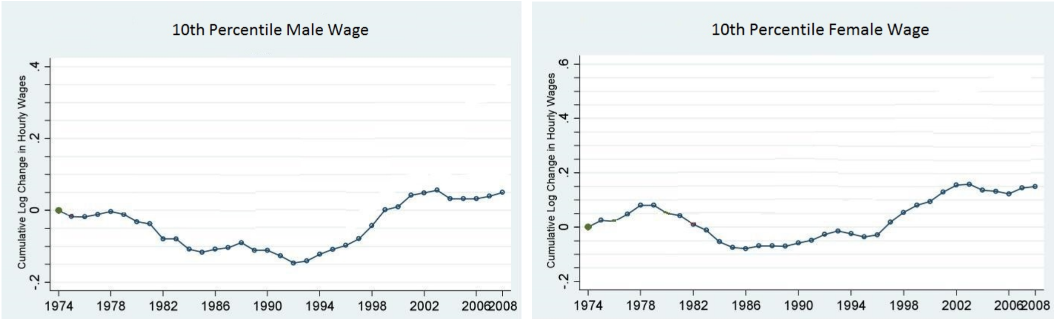

Figure 12: Cumulative log change in real weekly earnings at the 10th wage percentile for males and females respectively.

4.2 Episodic Rise in Residual Wage Inequality

Figures 13 and 14 depict the residual wage inequality measured as the smoothed ratio of the highest wage to the lowest wage from the simulation and of the 50-10 residual wage inequality for males and females from the data (AKK 2008), respectively. To read the graphs from the data requires a little explanation. The least flat line is the actual 50-10 residual wage inequality and the three flatter lines are measures of the composition effect with different base years. To read the composition adjusted 50-10 residual wage inequality, you have to try to subtract the change in one of the flatter lines from the least flat line. That is, since the flatter lines are upward sloping, you have to try to see the least flat line being more downward sloping. While doing that, the dynamics in all three graphs are qualitatively similar. In all three graphs, the residual wage inequality is initially stagnant, then it experiences an episodic increase, and finally it slowly decreases again. These dynamics occur in the model because the ratio of the highest wage to the lowest wage rises in the transition from full to HL pooling but in general is falling due to technological routinization.

4.3 Further Evidence

4.3.1 The Education and Residual Inequality Relationship

The simulation comparisons were motivated by three facts: (1) routinization compresses the productivities of workers performing routine tasks, (2) these workers tend to have low levels of education (Table 5b Acemoglu and Autor 2012), and (3) workers with low levels of education tend to have lower wages, and thus be in the lower half of the residual wage inequality distribution. However, we can look directly at the behaviour of residual wage inequality for each education level separately. If the residual wage inequalities for the lower levels of education, specifically high-school dropout, high-school graduate, and some college exhibit an increase followed by a decrease and the higher

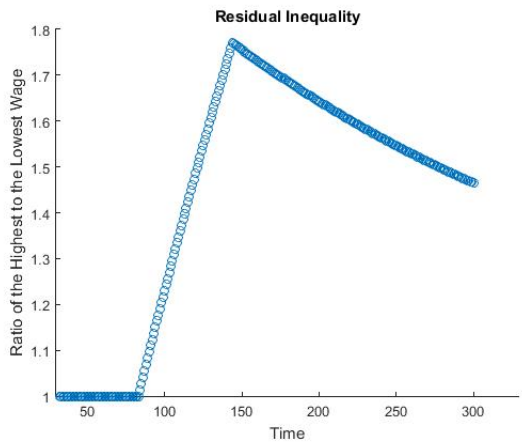

Figure 13: The graph depicts the smoothed ratio of the highest wage to the lowest wage from the simulation; it was smoothed the same way the lowest wage was.

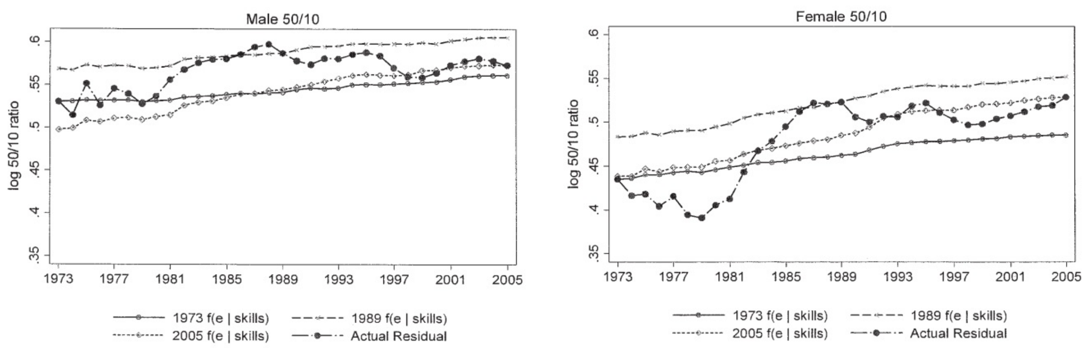

Figure 14: The least flat line in each of the graphs is the ratio of the 50th residual wage to the 10th percentile residual wage for males and females respectively. The three flatter lines are measures of the composition effect with different base years. To read the composition adjusted 50-10 residual wage inequality, you have to try to subtract the change in one of the flatter lines from the least flat line. For details, see the reference.

levels do not, then the theory does not somehow rely on the third relationship. Figure 15 depicts residual wage inequality by education for males and females (Lemieux 2006). As expected, the increasing then decreasing behaviour is only there for the lower levels of education.

4.3.2 The Rise in On-The-Job Wage Growth

In reality, a person's productivity tends to increase while in a job and so too may his wage. To compare the model to panal data on income, I assume that outside the

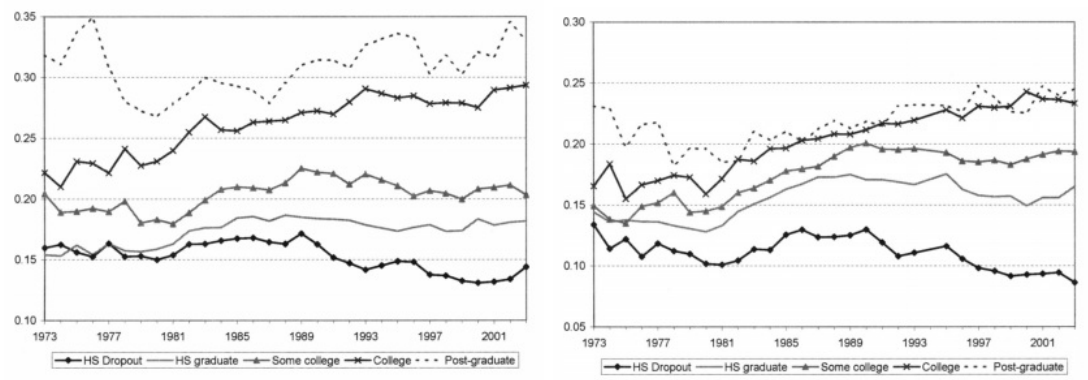

Figure 15: Variance of wages by education for males and females respectively

Wage Growth Within Job

Years	Wage Growth
1970-75	0.031
1975-80	0.041
1981-86	0.050
1986-91	0.051

Table 2: The table reports the annual wage growth for stayers. The standard error is 0.002 for each value. For details, see the reference.

model a worker will receive at least the wage he bargained for plus some fraction of his on-the-job productivity gains when his productivity starts to exceed his wage. I think this is a reasonable assumption, and since each type's productivity growth doesn't change over the transition, the sluggishness of each type's on-the-job wage growth in the model will be determined by the extent to which their bargained wage exceeds their true productivity. Thus, after the transition from full pooling to HL pooling, there will be a small negative change in new hires' on-the-job wage growth for low types, a big positive change for middle types, and no change for high types. That is, over the period, there will be a positive change in on-the-job wage growth. Violante (2002) finds with many different measures that this is true. I report the results for his baseline measure in Table 2.

Further, the model in combination with the wage growth assumption predicts that the wage growth should be concentrated among workers (1) with low levels of education, (2) with low levels of income (the middle types), and (3) who changed jobs during the period. Gottschalk and Moffitt (1994) find that, to the extent that wage growth over a group of workers can be proxied by the variance of their transitory earnings[8], these more specific facts are also borne out by the data.

[8]A worker's transitory earnings is defined as his deviation from his mean earnings.

Variance of Transitory Earnings by Education

Years of Completed Education	1970-1978	1979-1987	Percent Change
Fewer than 12	0.106	0.208	96
12 or more	0.081	0.123	52
16 or more	0.065	0.093	43

Table 3: The table reports the percent change in the variance of transitory earnings by years of completed education. For details, see the reference.

Variance of Transitory Earnings by Permanent Earnings Percentile

Permanent Earnings Percentile	1970-1978	1979-1987	Percent Change
Bottom 25	0.229	0.337	47
25-75	0.075	0.097	29
Top 25	0.045	0.054	20

Table 4: The table reports the percent change in the variance of transitory earnings by permanent earnings percentile. For details, see the reference.

Variance of Transitory Earnings by Job Status and Education

Years of Completed Education	Job Stayer Change	Job Mover Change
Fewer than 12	0.012	0.132
12 or more	0.016	0.040
16 or more	0.010	0.016

Table 5: The table reports the change between 1970-78 and 1979-87 in the variance of transitory earnings for people who did not change jobs within the period and for people who changed jobs at least once by years of completed education. For details, see the reference.

4.4 Effects of the Simulation Parameters

The intial values of H and M have been chosen so that their ratio is high, which allows for the equilibrium to transition from full pooling to HL pooling. Of course, it is impossible to know the joint skill distribution for all the skills that are valuable to employers, although we might well think that there is more weight near the bottom than near the top.

The other simulation parameters are the growth rates of H, M, and L. First, if the growth rates were too high, then any effect of the wage bargaining game would be swamped, but those growth rates would be unrealistic. Second, since H grows slower than M which grows slower than L, the inequalities $H > M > L$ will eventually be violated, but that can be avoided with more realistic functional forms. Third, there is a force for upward transitions because H grows slower than M but there is also a force for downward transitions because M grows slower than L. Thus, there can be no transitions, one or more transitions, and transitions that are followed by their reversals. That is, the one transition outcome of the simulation above is not the only possibility, but across initial value of H, M, and L near the threshold and reasonable growth rates, it is by far the most common.

5 Conclusion

This paper studies a model in which the seller has private information about the product's quality, the buyer inspects the product before they bargain, and the seller observes how intensely the buyer inspected. The result is an endogenous signaling game deriving from the fact that sellers with a high quality product can bargain more confidently.

With two types of sellers, the equilibrium is simple– the sellers pool for inspection intensities below a threshold and separate otherwise, so the buyer chooses to inspect just up to the threshold. With three types of sellers, however, there is effectively a bargaining game within a bargaining game, because the two better types bargain about which equilibrium should be played. A fundamental problem emerges for inspection intensities with multiple equilibria such that (i) the best response of one seller playing an equilibrium is for the other seller to play it as well, and (ii) the sellers have different rankings over the equilibria. For those inspection intensities, there is no reasonable prediction of the game.

That aside, there is a definite pattern in which equilibrium is played as the sellers' relative qualities change. Most importantly, when the middle type is sufficiently dissimilar to the high type, there will be full pooling. But if he becomes more similar to the high type, then his incentive to mimic decreases, and there will be high-low pooling.

Imposing the routinization hypothesis– that there has been a compression in skill levels of workers performing routine tasks and an expansion in the skill levels of workers performing non-routine tasks– immediately predicts the substantial increase in the upper half residual wage inequality. However, only through the model does it further predict the simultaneous substantial decrease in the lowest wages and the substantial increase in the lower half residual wage inequality followed by reversals. These results, plus the additional on-the-job wage growth tests, set this model apart.

6 Appendix

6.1 Mixed Strategy Equilibria (In the Two Type Model)

Proposition 6.1. *In a fully separating mixed strategy equilibrium, the type L seller sets the price L and the type H seller either sets the price H with probability one or mixes over $\hat{p}$ and H where (i) $xH + (1-x)u_H \leq \hat{p}$ and (ii) $L \geq \hat{p}(1-x)$. Any fully separating mixed strategy equilibrium can be supported by the uninformed buyer's beliefs that if the price is H, in the first case, or if the price is $\hat{p}$ or H, in the second, then the probability that the seller is a type H seller is $q = 1$ and otherwise $q = 0$.*

Proof. If the type H seller puts positive probability on only a single price, then it must be H. If it were less than H, then the uninformed buyer would accept with probability one, and it would not be a mixed strategy equilibrium.

If the type H seller puts positive probability on more than a single price, then he must set the price H with a positive probability and set a single other price $\hat{p} < H$ with the complementary probability. The proof is simply that all prices less than H will be accepted, and it is impossible for the type H seller to be indifferent between two prices that are accepted with the same probability. □

Proposition 6.2. *Let $x > 0$. In a semi-separating mixed strategy equilibrium, the sellers pool on exactly one price $\hat{p}$. A semi-separating equilibrium can be constructed with pooling price $\hat{p}$ if and only if (1) $xH + (1-x)u_H \leq \hat{p}$, (2) $L/(1-x) \leq \hat{p}$, and either (3) $\hat{p} \leq \mu H + (1-\mu)L$ or (iii)' $\mu H + (1-\mu)L < \hat{p}$ and $u_H L + \hat{p}[x(H-\hat{p}) - L] \leq 0$. Any semi-separating mixed strategy equilibrium can be supported by the uninformed buyer's beliefs that if the price is $\hat{p}$, then the probability that the seller is a type H seller is such that the uninformed buyer is indifferent between accepting and rejecting; if the price is H, then $q = 1$, and otherwise $q = 0$.*

Proof. If the sellers were to pool on two prices $p_1 > p_2 > L$, then $xp_1 + C_1 + (1-x)(1-\lambda_1)u_H = xp_2 + C_2 + (1-x)(1-\lambda_2)u_H$ where $\lambda_1 < \lambda_2$ and $C_1 = C_2$– a contradiction. If the sellers were to pool on two prices $p_1 > p_2 = L$, then it must be that $xp_1 + C_1 + (1-x)(1-\lambda_1)u_H = L$ and $C_1 = L$– a contradiction.

The prices can be sorted as follows. If the uninformed buyer were to continuously decrease the probability with which he accepts a given price λ whether the type H or the type L seller becomes indifferent first between that price and H or L respectively. This defines a line $u_H L + p[x(H-p) - L] = 0$ in $\{x, p\}$-space below which the type H seller becomes indifferent first and above which the type L seller does.

If $\hat{p} < \mu H + (1-\mu)L$, then a mixed strategy equilibrium can be constructed with pooling price $\hat{p}$ through continuously decreasing λ until the type H or the type L seller becomes indifferent who then plays a mixed strategy over that price and H or L respectively. If $\hat{p} > \mu H + (1-\mu)L$ and $u_H L + p[x(H-p) - L] < 0$, a mixed strategy equilibrium can be constructed at a price through decreasing λ until the type L seller becomes indifferent who then plays a mixed strategy over that price and L. □

6.2 Infinite Types Each of Measure Zero

Proposition 6.3. *With infinite types each of measure zero, there is no second stage undefeated equilibrium where a positive measure of seller types pool. However, an equilibrium where a positive measure of seller types pool can Pareto dominate the mixed strategy equilibrium where each type separates for some values of x.*

Proof. Suppose, to reach a contradiction, that there were an undefeated equilibrium where a positive measure of seller types pool. If a seller's type K is above the pooling price and he sets the pooling price, then his expected payoff can be written $xp+C$ where C may be less than $(1-x)p$ if the buyer is deterring some types from mimicking. If he were to deviate to the equilibrium where everything else is the same except that he sets his price equal to his type, then his expected payoff could be written $xK + C > xp + C$. Thus, the initial undefeated pooling equilibrium is defeated– a contradiction.

The proof that an equilibrium where a positive measure of seller types pool can Pareto dominate the mixed strategy equilibrium where each type separates for some x is just that both (i) $(1-x)\mathbb{E}(T) > x\mathbb{E}(T) + L$ and (ii) $\mathbb{E}(T) > xH + L$ for sufficiently small x. □

6.3 The Type H Seller's Undefeated HL-ML Pooling Expected Payoff

Proposition 6.4. *Let L be normalized to one. In the undefeated HL-ML pooling equilibrium, the type H seller's expected payoff is $x\frac{H+z}{1+z} + (1-x)(\frac{M+(1-z)}{1+(1-z)})$ where z is such that $\frac{M+(1-z)}{1+(1-z)} = xM + 1$.*

Proof. In the undefeated HL-ML pooling equilibrium, the type H and type M sellers' prices can be written $\frac{H+z}{1+z}$ and $\frac{M+(1-z)}{1+(1-z)}$ respectively where z is the probability with which the type L seller sets the price equal to the type H seller's price. The type H seller's expected payoff can be written $F = x\frac{H+z}{1+z} + (1-x)\frac{M+(1-z)}{1+(1-z)}$, the second derivative of which with respect to z is positive.

In the undefeated HL-ML pooling equilibria, we have the restrictions (1) $xM + 1 \leq \frac{M+(1-z)}{1+(1-z)}$ and (2) $M \leq \frac{H+z}{1+z}$. Thus, since F is convex, the maximum under the

Equilibrium	type H's expected payoff	type M's expected payoff
Full Pooling	$\frac{1}{3}(H+M+L)$	$(1-x)\frac{1}{3}(H+M+L)$
HL Pooling	$min\{x\frac{1}{2}(H+L)+M, \frac{1}{2}(H+L)\}$	M
ML Pooling	$xH+(1-x)\frac{1}{2}(M+L)$	$\frac{1}{2}(M+L)$
Mixed Separating	$min\{xH+L, H\}$	$min\{xM+L, M\}$
HL-ML Pooling	$< xH+(1-x)\frac{1}{2}(M+L)$	$< M$

Table 6: The type H and type M sellers' expected payoffs in each within-group undefeated equilibrium for Case I.

restrictions is achieved either at (a) z_1 where $xM + 1 = \frac{M+(1-z_1)}{1+(1-z_1)}$ or (b) $min\{z_2, 1\}$ where $M = \frac{H+z_2}{1+z_2}$.

However, if the maximum under the restrictions were not achieved at z_1, then the payoff constraint would bind. First, if the maximium were achieved at z_2, then the type H seller's payoff would be even greater at $\frac{1}{3}(H+M+L)$ by the convexity of F. Second, if the maximium were achieved at one, then the type H seller's payoff would be equal to $x\frac{1}{2}(H+L)+M$, which equals or is less than the payoff bound. □

Corollary 6.5. *The M-constraint cannot bind and later become slack.*

6.4 Between Group Undefeated Equilibrium in Case II: $M < \frac{1}{2}(H+L)$

The analysis of Case II is similar to that of Case I. The main difference in the within group analysis is that there exists an iteratively undefeated HL pooling equilibrium where the type H and type L sellers set the price $\frac{1}{2}(H+L)$, which the uninformed buyer accepts with the highest probability such that the type M seller does not strictly prefer to set the price $\frac{1}{2}(H+L)$, and the type M seller sets the price M, which the uninformed buyer accepts.

In Case IIA ($M < \frac{(H+1)(H+2)}{5H+1}$), $xH+L$ crosses $x\frac{1}{2}(H+L)+M$ below $\frac{1}{3}(H+M+L)$. The first subcase is where $M < 2$. Then $xH+L$ crosses $\frac{1}{3}(H+M+L)$ before $xH+(1-x)\frac{1}{2}(M+L)$ does, and the explanation is just that there is full pooling until the crossing and then there is mixed separating. The second subcase is where $M > 2$. The important crossings are [1] $xH+(1-x)\frac{1}{2}(M+L)$ crossing $\frac{1}{3}(H+M+L)$, [2] $xM+L$ crossing $\frac{1}{2}(M+L)$, and [3] $xH+L$ crossing $x\frac{1}{2}(H+L)+M$. All are possible except [1][3][2].

[1][2][3] and [3][1][2]: Before [1], there is full pooling. After [2], if the HL-ML pooling payoff constraint binds for the final time before [3], then after that there is full pooling

followed by mixed separating. If not, then after [3] anywhere that it exactly binds there is HL-ML pooling and besides that there is mixed separating or nothing.

[2][1][3], [2][3][1], and [3][2][1]: Before [1], there is full pooling. After [1], there is an interval where there continues to be full pooling, because the type H seller's HL-ML pooling payoff needs to catch up. If it doesn't, then the interval continues until $xH + L$ crosses $\frac{1}{3}(H + M + L)$ and after that there is mixed separating. If it does, then after that rather than "After [2]", the explanation is the same as it was in the [1][2][3] and [3][1][2] case.

In Case IIB ($\frac{(H+1)(H+2)}{5H+1} < M < (H+L)/4$), $xH+L$ crosses $x\frac{1}{2}(H+L)+M$ above $\frac{1}{3}(H+M+L)$ and $xH+(1-x)\frac{1}{2}(M+L)$ crosses $\frac{1}{3}(H+M+L)$ before $x\frac{1}{2}(H+L)+M$ does. The important crossings are [1] $xH+(1-x)\frac{1}{2}(M+L)$ crossing $\frac{1}{3}(H+M+L)$, [2] $xM+L$ crossing $\frac{1}{2}(M+L)$, and [3] $xH+L$ crossing $min\{x\frac{1}{2}(H+L)+M, \frac{1}{2}(H+L)\}$. By the definition of Case IIB, [1] comes before [3].

[1][2][3] and [1][3][2]: Before [1], there is full pooling. After [2], if the HL-ML pooling payoff constraint binds for the final time before $x\frac{1}{2}(H+L)+M$ crosses $\frac{1}{3}(H+M+L)$, then after that there is full pooling followed by HL pooling followed by mixed separating. If it binds for the final time after that crossing but before [3], then after that there is HL pooling followed by mixed separating. Lastly, if it binds for the final time after [3], then anywhere after that it exactly binds there is HL-ML pooling and besides that there is mixed separating or nothing.

[2][1][3]: Before [1], there is full pooling. After [1], there is an interval where there continues to be full pooling. because the type H seller's HL-ML pooling payoff needs to catch up. If it doesn't, then the interval continues until $x\frac{1}{2}(H + L) + M$ crosses $\frac{1}{3}(H + M + L)$, and after that there is HL pooling followed by mixed separating. If it does, then after that rather than "After [2]", the explanation is the same as it was in the [1][2][3] and [3][1][2] case.

In Case IIC, ($M > \frac{(H+1)(H+2)}{5H+1}$ and $M > (H + L)/4$, the important crossings are [1] $xH + (1 - x)\frac{1}{2}(M + L)$ crossing $min\{x\frac{1}{2}(H + L) + M, \frac{1}{2}(H + L)\}$, [2] $xM + L$ crossing $\frac{1}{2}(M + L)$, and [3] $xH + L$ crossing $min\{x\frac{1}{2}(H + L) + M, \frac{1}{2}(H + L)\}$. There are two subcases. The first is where [3] is before [1], that is, where $H < 2M + 3$ and/or $HM + M + 1 < 3H$. There the explantion is just that there is full pooling followed by HL pooling followed by mixed separating. The second is where [1] is before [3]. It turns out that [1][3][2] is impossible.

[1][2][3]: Before [1], there is full pooling followed by HL pooling. After [2], if the HL-ML pooling payoff constraint binds for the final time before [3], then after that there is HL pooling followed by mixed separating. If not, then after [3] anywhere that it exactly binds there is HL-ML pooling and besides that there is mixed separating or nothing.

[2][1][3]: Before [1], there is full pooling followed by HL pooling. After [1], there is an interval where there continues to be HL pooling, because the type H seller's HL-ML pooling payoff needs to catch up. If it doesn't, then the interval continues until [3], and after that there is mixed separating. If it does, then after that rather than "After [2]", the explanation is the same as it was in the [1][2][3] case.

6.5 The Buyer's Expected Undefeated HL-ML Pooling Profit

Proposition 6.6. *Whenever there is more than one undefeated HL-ML pooling equilibrium, the one with the highest inspection intensity is the most profitable.*

Proof.

$$\begin{aligned}\Pi_{HL-ML} &= \frac{x}{3}[(H - w_H) + (M - w_M)] \\ &= \frac{1}{3}[2x^2HM - xH(M-1) - 3x^2M + 2x(M-1)]\end{aligned}$$

$$\begin{aligned}\frac{d\Pi_{HL-ML}}{dx} &= \frac{1}{3}[4xHM - H(M-1) - 6xM + 2(M-1)] \\ &> \frac{1}{3}[\frac{4}{3}HM - HM + H - 2] > 0\end{aligned}$$

The first inequality derives from $x > 1/3$ and $H > 2$ whenever there are HL-ML pooling equilibria, and the second inequality derives from $H > 2$. □

7 Works Cited

Acemoglu, D. (1999). Changes in unemployment and wage inequality: An alternative theory and some evidence. American Economic Review, 89(5), 1259-1278.

Acemoglu, D., and Autor (2012). Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings. Handbook of Labor Economics, 1043-1171.

Aghion, P., Howitt, P., and Violante, G. L. (2002). General purpose technology and wage inequality. Journal of Economic Growth, 7(4), 315-345.

Bagwell, K., and Riordan, M. H. (1991). High and declining prices signal product quality. American Economic Review, 81(1), 224-239.

Caselli, F. (1999). Technological revolutions. American Economic Review, 89(1), 78-102.

Cooper, R., and Ross, T. W. (1984). Prices, product qualities and asymmetric information: The competitive case. Review of Economic Studies, 51(2), 197-207.

Cremer, J. (1995). Arm's length relationships. Quarterly Journal of Economics, 110(2), 275-295.

Cremer, J., and Khalil, F. (1994). Gathering information before the contract is offered: The case with two states of nature. European Economic Review, 38(3-4), 675-682.

Dang, T. V. (2008). Bargaining with endogenous information. Journal of Economic Theory, 140(1), 339-354.

Daughety, A. F., and Reinganum, J. F. (1994). Asymmetric information acquisition and behavior in role choice models: An endogenously generated signaling game. International Economic Review, 35(4), 795-819.

Delacroix, A., and Shi, S. (2013). Pricing and signaling with frictions. Journal of Economic Theory, 148(4), 1301-1332.

Feltovich, N., Harbaugh, R., and To, T. (2002). Too cool for school? Signaling and countersignaling RAND Journal of Economics 33(4), 630–649.

Galor, O., and Moav, O. (2000). Ability-biased technological transition, wage inequality, and economic growth. Quarterly Journal of Economics, 115(2), 469-497.

Galor, O., and Tsiddon, D. (1997). Technological progress, mobility, and economic growth. American Economic Review, 87(3), 363-382.

Gottschalk, P., and Moffitt, R. (1994). The growth of earnings instability in the U.S.

labor market. Brookings Papers on Economic Activity, (2), 217-254.

Greenwood, J., and Yorukoglu, M. (1997). 1974. Carnegie-Rochester Conference Series on Public Policy, 46, 49-95.

Juhn, C., Murphy, K. M., and Pierce, B. (1993). Wage inequality and the rise in returns to skill. Journal of Political Economy, 101(3), 410-442

Kaya, A. (2010). When does it pay to get informed? International Economic Review, 51(2), 533-551.

Kessler, A. S. (1998). The value of ignorance. RAND Journal of Economics, 29(2), 339-354.

Lemieux, T. (2006). Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill? American Economic Review, 96(3), 461-498.

Lester, B., Postlewaite, A., and Wright, R. (2012). Information, liquidity, asset prices, and monetary policy. Review of Economic Studies, 79(3), 1209-1238.

Levin, J. (2001). Information and the market for lemons. RAND Journal Of Economics (RAND Journal Of Economics), 32(4), 657-666.

Lewis, T. R., and Sappington, D. E. M. (1993). Ignorance in agency problems. Journal of Economic Theory, 61(1), 169-183.

Mailath, G. J., Okuno-Fujiwara, M., and Postlewaite, A. (1993). Belief-based refinements in signalling games. Journal of Economic Theory, 60(2), 241-276.

Milgrom, P., and Roberts, J. (1986). Price and advertising signals of product quality. Journal of Political Economy, 94(4), 796-821.

Nosal, E. (2006). Information gathering by a principal. International Economic Review, 47(4), 1093-1111.

Riordan, M. H. (1986). Monopolistic competition with experience goods. Quarterly Journal of Economics, 101(2), 265-279.

Rocheteau, G. (2011). Payments and liquidity under adverse selection. Journal of Monetary Economics, 58(3), 191-205.

Shavell, S. (1994). Acquisition and disclosure of information prior to sale. RAND Journal of Economics, 25(1), 20-36.

Shi, S. (2002). A directed search model of inequality with heterogeneous skills and skill-based technology. Review of Economic Studies, 69(2), 467-491.

Sobel, J. (1993). Information control in the principal-seller problem. International Economic Review, 34(2), 259-269.

Violante, G. L. (2002). Technological acceleration, skill transferability, and the rise in residual inequality. Quarterly Journal of Economics, 117(1), 297-338.

Wolinsky, A. (1983). Prices as signals of product quality. Review of Economic Studies, 50(4), 647-658.