

Optimally Sticky Prices: Foundations*

Jean-Paul L'Huillier[†] William R. Zame[‡]

May 31, 2017

Abstract We propose a strategic microfoundation for sticky prices. We model an environment in which a firm has better information than its consumers and show that, when many consumers are uninformed, it is *optimal* for the firm to offer sticky contracts or sticky prices. We establish this result in a general mechanism design framework that allows for non-linear pricing and screening. We then discuss the implications of this microfounded friction for welfare.

JEL E31, E40

Keywords sticky prices, frictions, pricing policies, optimal mechanisms

*The authors are grateful to Nemanja Antic, Marco Bassetto, Paul Beaudry, Pierpaolo Benigno, Francesco Lippi, Morten Ravn, Joel Sobel, Daniele Terlizzese, participants at the Southwest Economic Theory Conference and seminar audiences at BI, Cambridge University, EIEF, Norges Bank, Oxford University, UCL, and University of Mannheim for comments, and to UCLA (L'Huillier) and to EIEF (Zame) for hospitality and support during the preparation of this paper. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any funding agency.

[†]EIEF, Rome, Italy

[‡]Department of Economics, UCLA, Los Angeles, CA

1 Introduction

It is well-documented that prices are *sticky*: they do not always move when fundamentals move. This observed deviation from the predictions of standard models is believed to have significant macroeconomic implications and a substantial literature seeks to provide explanations for it. The explanations given in this literature (typically) posit some sort of friction within the firm that prevents – or at least discourages – adjusting prices in response to changes in fundamentals. Commonly posited frictions include menu costs, including the explicit cost of changing prices: re-marking items on shelves or re-programming pricing software, see Barro (1972), Caplin and Leahy (1991), Golosov and Lucas (2007), and others; or information processing costs that prevent the firm from learning the state; see Reis (2006), Mackowiak and Wiederholt (2009), and others.

In contrast to this popular modeling approach, a longstanding (but informal) idea¹ in economics maintains that frictions originating in the relationship between firms and consumers prevent price adjustments. Motivated by this intuitive idea, our goal is to formalize it in a standard strategic environment. We consider a setting in which there are no frictions *within the firm* – the firm can freely adjust prices (so has no menu costs) and is perfectly informed (so has no information processing costs); the friction, rather, is *between the firm and the consumers*. This friction arises because the firm has more information than (some) of its consumers. The firm would be happy to reveal this information – but may be unable to commit to revealing the information truthfully. We build models that embody the firm’s inability to commit; in each of these models, our conclusion is that price stickiness can arise as *optimal behavior* for the firm when this information asymmetry is severe (many consumers are uninformed).

In our (stylized) model, we consider the interaction between a monopolistic firm that produces a single good using a constant returns to scale technology and continuum of consumers who demand both the good produced by the firm and a single other aggregate consumption good. (In a

¹This idea goes back to (at least) Hall and Hitch (1939), and has been mentioned by many other authors, including Okun (1981), Kahneman, Knetsch, and Thaler (1986). In a highly-cited book, Blinder et al. (1998) notes that when asked to explain their reluctance to increase prices, firms’ managers most common response is that “price increases cause difficulties with customers.”

general equilibrium version of the model, this aggregate good is produced endogenously.) The economy is subject to a shock. We model the shock (which might be interpreted as the result of a nominal disturbance) as the (nominal) price for the aggregate consumption good, which can take on only two values: High or Low.² The shock follows a known distribution; the firm knows the realization of the shock, a fraction of the consumers are informed and so also know the realization of the shock, but the remaining fraction of consumers are uninformed and do not know the realization of the shock. The firm would be happy to reveal the shock – because it would then be able to extract monopoly rents in each state – but (for some values of the parameters) the firm cannot do so *credibly*: when the shock is Low the firm would have an incentive to misrepresent it as High.

We formalize this idea in two settings, distinguished by the objects of choice. In the first, the objects of firm choice are *contracts*. In this setting, we find it convenient and attractive to take an approach via mechanism design, abstracting the interaction between the firm and consumers in terms of mechanisms whose outputs are contracts that specify consumption produced by the firm for the consumer and a money transfer from the consumer to the firm. As usual, it suffices to consider direct mechanisms, in which the firm and consumers report their private information. Among direct mechanisms we restrict attention to those that satisfy incentive compatibility and a weak notion of ex-post individual rationality: agents can refuse to participate *after* they learn the assigned contract and draw whatever inferences are possible from understanding the mechanism and learning the assigned contract but *before* they learn the true state (if they did not already know it). A direct mechanism is *pooling* if it assigns the same contract to uninformed consumers independently of the report of the firm (the state of the world); otherwise it is *separating*. We show that if the fraction of informed consumers is close to 0 the firm optimal incentive compatible, individually rational mechanism is pooling but if the fraction is close to 1 the firm optimal incentive compatible, individually rational mechanism is separating. The mechanism design approach is useful because it offers simplicity, clarity and elegance and also because it assures us that the firm cannot do better even if we admit much

²Our choice to model the shock in this way is purely for convenience and analytic tractability; similar conclusions would obtain for other kinds of shocks. The crucial requirements are that the firm should know more than some consumers about the realization of a parameter that matters to both the firm and the consumers.

more general interactions. We then go on to show that the conclusions of the mechanism design approach remain when we model the behavior of the firm and consumer as a contract-setting game in which the firm can offer a single contract which the consumer can accept or reject, or a menu of contracts from which the consumer can choose. In particular, the optimal mechanisms can be supported as perfect Bayesian equilibria of the contract-setting game. In both formulations, we see that when the fraction of informed consumers is small, contracts – and *a fortiori*, prices – are sticky.

In the second setting, the objects of firm choice are *prices*. In this setting, we focus our attention directly on the obvious price-setting game in which the firm quotes a price and consumers choose quantities to purchase at the quoted price. Here too, the conclusion is that when the fraction of informed consumers is close to 0, the firm-optimal (perfect Bayesian) equilibrium is pooling – the firm quotes the same price in both states of the world, so the price is sticky – while when the fraction of informed consumers is close to 1 the firm-optimal (perfect Bayesian) equilibrium is separating – the firm quotes different prices in the two states of the world, and the price is not sticky.

Our formulation is quite tractable and can be embedded into a fairly standard general equilibrium dynamic model with money. This allows to address the neutrality of money and the welfare effects. For both of these, the implications are most clearly seen by contrasting the conclusions when all consumers are uninformed with the conclusions when all consumers are informed. When all consumers are uninformed, it is optimal for the firm to pool, consumers learn nothing about the true state and prices and quantities are the same in the two states. That is, even though prices are sticky, the state of the world is *neutral* (for these variables). When all consumers are informed, it is optimal for the firm to separate, consumers know the true state, prices are different in the two states, but quantities are the same. Here, as usual, prices are flexible and again the state of the world is *neutral*. Although the firm behaves quite differently in the two extreme cases, the firm is *ex ante* indifferent: it makes the same expected profits whether all consumers are uninformed or all consumers are informed. Similarly, (expected) *social welfare* is the same when all consumers are uninformed and when all consumers are informed. These welfare conclusions are especially striking because assuming that all consumers are informed is equivalent to assuming that the firm could *credibly* reveal the true state. Hence the conclusion is

that, when all consumers are uninformed, the firm’s inability to reveal the true state alters the transfers/prices without altering the welfare of either the firm or the consumers. (Note however, that when the firm sets contracts it is able to extract *all* the social surplus, leaving the consumers with none, but producing the socially efficient quantity and adjusting the transfers. When the firm sets prices, it cannot extract the entire social surplus, but must share it with the consumers.)

Having said this, we must also note that if most, but not all, consumers are uninformed, it remains optimal for the firm to pool, and then pooling does create profit losses for the firm *and* welfare losses for consumers. These losses are small when the fraction of informed consumers is small, but grow as the fraction of informed consumers grows (at least so long as the fraction is small enough to guarantee that the firm pools). Thus our models *do not* imply (large) welfare losses – as do most benchmark monetary models. The reason is that, in the presence of asymmetric information, pooling imposes distortions on the firm but avoids imposing distortions on consumers.

In the presence of additional assumptions about preferences (e.g., quadratic), it would be possible to derive qualitative but testable implications of our models. For instance, as can be seen in the Example, higher per unit cost implies that the region in which the firm prefers to pool is smaller, which would suggest that prices might be stickier in firms or industries in which markups are high. Testing such a prediction would seem possible if sufficient data on markups were available.

Related Literature Our paper is part of a large literature that uses information tools to address macroeconomic outcomes. Seminal contributions by Lucas (1972) and Mankiw and Reis (2002) showed that information frictions are useful to address money non-neutrality, and the persistence of macroeconomic variables. Bacchetta and Van Wincoop (2006), Amador and Weill (2010) and Golosov, Lorenzoni, and Tsyvinski (2014) is an inexhaustive list of recent contributions addressing other issues, but using similar tools. Among the papers addressing the non-neutrality of money, the main novelty is the analysis of the opposite information structure of the one usually considered. That is, we analyze an economy where firms are perfectly informed, but consumers imperfectly informed (starting with Lucas 1972, the literature had focused on the easier-to-handle case of imperfectly informed

firms, and perfectly informed consumers/buyers.³) Another novelty in our paper is methodological: the use of mechanism design to analyze the optimal behavior of strategic agents in the face of these frictions.⁴

A companion paper (L’Huillier 2017) is an application of the price stickiness model derived in this paper. L’Huillier (2017) studies the propagation of monetary shocks in a dynamic economy subject to this friction. L’Huillier (2017) is less ambitious theoretically. Instead, ours is a full, in depth, theoretical analysis of the foundations leading to the optimally-sticky-prices friction.

In IO, a few other papers have also derived a sort of rigidity in prices using different strategic models (see Maskin and Tirole 1988, Nakamura and Steinsson 2011, Cabral and Fishman 2012, among others.) None of these allow for a general equilibrium formulation such as ours. Nakamura and Steinsson (2011) emphasize, as we do, information frictions. However, their mechanism is different. In their model, a repeated game with habit formation on the side of the consumer leads to a lock-in situation. This provides a motive for firms to commit to sticky prices. We rely on a standard, static, mechanism design problem in which the firm-optimal solution is to post sticky prices, even in the absence of repeated purchases and habit formation.

Our work is also related to the theoretical literature studying the implications of commitment problems using mechanism design. Amador, Werning, and Angeletos (2006) introduce techniques for the analysis of time-inconsistency problems. Our mechanism is somewhat simpler to handle, but it also delivers the result that bunching of types is optimal for some ranges of parameters. A similar result is obtained by Athey, Atkeson, and Kehoe (2005), where no discretion (a form of full bunching) is obtained in cases of severe time-inconsistency.

Following this Introduction, Section 2 presents a motivating example, Section 3 describes the environment, Section 4 presents the mechanism design model, Section 5 presents the contract-setting game and Section 6 presents the price-setting game. Section 7 provides results about the effect of money and details the welfare comparisons in each of the three models. Section 8 collects a few remarks about modeling choices and extensions. The Appendix

³An noticeable exception is Mackowiak and Wiederholt (2015), where both households and firms are rationally inattentive.

⁴A closely related paper to ours is by Jovanovic and Ueda (1997), who derives nominal effects in a moral hazard contracting problem.

collects all proofs and the general equilibrium framework.

2 Example

To give a preview of and insight into our results, we begin with a very simple example. We consider an environment with two goods: a special good x and an aggregate consumption good y . There is a single firm that can produce x from y using a constant returns to scale technology $x = Ay$. There is a continuum of unit mass of identical consumers who are endowed with a nominal income M that they trade inelastically for consumption goods. Consumer utility for consumption of the two goods x, y is

$$v(x, y) = (x - x^2/2) + y$$

We assume in what follows that consumers always choose $x, y > 0$.

There are two possible states of the world H, L (High, Low), that occur with probabilities ρ_H, ρ_L . The firm is informed of the state of the world, the consumer is not. The state of the world ω represents the (nominal) price level of the aggregate good y ; we assume $p_H > p_L$. For convenience, we define the *harmonic mean price*

$$p_0 = [\rho_H/p_H + \rho_L/p_L]^{-1}$$

We assume the firm is a monopolist, so observes the true state and offers a price, and the consumers maximize (expected) utility given the price and their information. Suppose for the moment that the firm does not condition its price on the state of the world, so offers a fixed price q_0 (independent of the true state). The consumers (who do not know and cannot infer the true state) maximize expected utility of consumption; because income is nominal and utility is quasi-linear in the aggregate good, the assumption that $x, y > 0$ means the consumers maximize

$$E[x - x^2/2 - (q/p_\omega)x] = x - x^2/2 - E(q/p_\omega)x = x - x^2/2 - (q/p_0)x$$

It follows that consumer demand is

$$X_0(q) = 1 - q/p_0$$

and that the firm's expected profit (expressed in real terms; i.e. in units of the aggregate good y) is

$$\Pi_0(q) = (q/p_0 - k)(1 - q/p_0)$$

where $k = 1/A$. The firm maximizes profit by choosing the price $q_0 = [(1 + k)/2] p_0$, and optimal profit is

$$\Pi_0^* = (1 - k)^2/4$$

Now suppose that the firm *does* condition its price on the state of the world, so offers a price q_H when the state is High and q_L when the state is Low. The consumers observe the price offered and so infer the state and maximize utility in the inferred state. Hence in each state $\omega \in \{H, L\}$ the consumers maximize

$$x - x^2/2 - (q/p_\omega)x$$

It follows that consumer demand is

$$X_\omega(q) = 1 - q/p_\omega$$

and that the firm's expected profit is

$$\Pi_\omega(q) = (q/p_\omega - k)(1 - q/p_\omega)$$

The firm maximizes profit by choosing the price $q_\omega = [(1 + k)/2] p_\omega$, and optimal profit is

$$\Pi_\omega^* = (1 - k)^2/4$$

This calculation would seem to suggest that, *ex ante*, the firm is indifferent between the policies of conditioning price on the state of the world or not conditioning price on the state of the world – or, in other words, that the firm would be perfectly willing to simply announce the true state. However, this is not quite right: the firm would be perfectly willing to announce the true state *provided that it could commit to doing so truthfully*. If – as surely might be the case in reality – the firm *cannot* commit to announcing the true state truthfully, we must take into account the incentives of the firm to lie, in particular to offer the price q_H that the consumers expect to see when the

state is High even though the true state is actually Low. If it does so, the firm will realize profit

$$\begin{aligned}
(q_H/p_L - k)X_H(q_H) &> (q_H/p_H - k)X_H(q_H) \\
&= \Pi_H^* \\
&= \Pi_L^* \\
&= (q_L/p_L - k)X_L(q_L)
\end{aligned}$$

Thus the firm would strictly prefer to offer the price q_H even when the true state is L . Hence, we conclude that the firm *cannot credibly commit* to offering the full information monopoly optimal price in each state – it is not *incentive compatible* for the firm to do so. Taking incentive compatibility into account, it would therefore seem that the firm would strictly prefer *not* to condition the price on the true state.

This is still not the full story, however, because it does not take into account the incentives of the firm when it *does not* condition the price on the true state. Rational behavior by the consumers requires that, whatever price the consumers observe, they should then form beliefs about the true state and optimize on the basis of those beliefs. If the consumers observe the anticipated price q_0 their beliefs about the true state should remain the same as its priors ρ_H, ρ_L and they should optimize as above; however, if the consumers observe a price $q \neq q_0$ they should update their priors and optimize with respect to the updated priors. Anticipating this, the firm, having observed the true state, chooses whether to offer the price q_0 or some price $q \neq q_0$. If the firm observes that the true state is H and offers a price $q \neq q_0$, the worst outcome for the firm (the lowest profit) will occur when the consumers update their beliefs to assign probability 1 that the state is L , in which case the consumers will demand the quantity $X_L(q)$ and the firm will realize the profit $(q/p_H - k)X_L(q)$. Hence it will be incentive compatible for the firm to offer the price q_0 when the state is H if and only if

$$(q/p_H - k)X_L(q) \leq (q_0/p_H - k)X_0(q_0)$$

Given our previous calculations, we see that it will be incentive compatible for the firm to offer the price q_0 when the state is H if and only if

$$\max_q (q/p_H - k)(1 - q/p_L) \leq ((1 + k)/2] p_0/p_H - k)(1 - k)/2 \quad (1)$$

Solving the inequality (1) yields a range of values of the marginal cost k for which it will be incentive compatible for the firm to offer the price q_0 when the state is H , but the calculation is entirely unenlightening. Instead, let us observe that if $k = 0$ then the left hand side of (1) will be maximized when $q = p_L/2$ and the maximum will be $p_L/4p_H$, while the right hand side will reduce to $p_0/4p_H$. Since $p_0 > p_L$, it follows that, for k sufficiently small, the left hand side is again strictly less than the right hand side.

We conclude that, for k sufficiently small, a firm that faces uninformed consumers and cannot commit to truthful revelation will strictly prefer *not* to condition price on the true state but rather to offer the same price in both states. Thus, given the informational asymmetry, optimal behavior by the firm leads to *sticky prices*.

Several predictions of our model seem important. The first is that (because sticky prices are predicted only when k is small), sticky prices are more likely to be observed when firms (or industries) are – or become – more efficient. The second is that, in the situations in which our model predicts that prices do not depend on the true state, it also predicts that quantities do not depend on the true state; this is a prediction quite different from that of other sticky price models.

The situation we have analyzed here is special in a number of ways. We have assumed that the firm sets prices, we have assumed that all consumers are uninformed, and we have assumed that consumer preferences are quadratic. In the remainder of the paper we show formally that similar conclusions also obtain for other trading mechanisms, in environments in which some consumers are informed of the true state, and for general consumer utility functions.

3 The Environment

We consider a model with two consumption goods: a special consumption good x and an aggregate consumption good y .

To make the intuition clear, we offer here a partial equilibrium framework in which the price of the aggregate good y is exogenous. However, the results obtained here in the partial equilibrium framework are entirely compatible with a dynamic general equilibrium framework with many firms,

labor, money and financial markets, and in which the supply (and the price) of this aggregate good is obtained endogenously at equilibrium. Because the general equilibrium framework is elaborate and might obscure the informational insight on which we want to focus, we relegate the details to the Appendix and focus here on the partial equilibrium framework.

There is a unit mass of consumers indexed by $c \in [0, 1]$. Of these, a subset $I \subset [0, 1]$ are *informed* and the remainder $D = [0, 1] \setminus I$ are *uninformed*; we frequently write $i \in I, d \in D$ to emphasize that the consumer in question is informed or uninformed (respectively). We write α for the proportion of informed consumers. In the benchmark settings $\alpha = 0, 1$ no (respectively, all) consumers are informed. Consumers value both consumption goods; their common utility function is

$$v(x, y) = u(x) + y$$

For convenience, we assume that u is smooth (twice continuously differentiable), strictly increasing and strictly concave (at least in the relevant region), and that $u(0) = 0$.⁵ Note that consumption is quasi-linear in the aggregate good. Consumers are endowed with a nominal income M that they trade inelastically for consumption goods; that is, they maximize utility for consumption goods given prices and information, subject to the constraint that expenditure equals nominal income.

There is a single firm, which uses the aggregate good y to produce the special good x according to a constant returns technology $x = Ay$. We write $k = 1/A$ for the constant marginal cost of the firm; we assume fixed costs are zero.

There is uncertainty about the state of the world ω which represents the nominal price p_ω of the aggregate consumption good y . We assume there are two possible states, H, L ; without loss assume $p_H > p_L$ (so we refer to H as the High state and L as the Low state).⁶ The firm and all consumers share

⁵These assumptions are all standard. Given that $u(0)$ is finite, the assumption that $u(0) = 0$ is just a convenient normalization. However, the assumption that $u(0)$ is finite has bite: If we allowed $u(0) = -\infty$ then consumers could not survive without consumption of the special good; if the firm can make contract offers then it can extract arbitrarily large transfers in return for arbitrarily small quantities of the special good, and the firm's optimization problem would have no solution.

⁶We restrict attention to two states only for simplicity of exposition; similar conclusions would obtain if there were many states.

a common prior ρ_H, ρ_L about the distribution of states of the world. The firm and the informed consumers know the realized state of the world; the uninformed consumers do not.

It is convenient to write

$$p_0 = \left[\frac{\rho_H}{p_H} + \frac{\rho_L}{p_L} \right]^{-1}$$

for the *harmonic mean* of the prices p_H, p_L with respect to the true probabilities. We define quantities x^*, x_* by the equations

$$\begin{aligned} u'(x^*) &= k \\ u'(x_*) &= \left(\frac{p_H}{p_0} \right) k \end{aligned}$$

Because we have assumed consumer utility to be quasi-linear in the common consumption good, x^* is the *socially optimal* quantity. As we shall see later x_* represents the quantity produced at a particular optimal deviation.

Suppose the firm offers to sell the x good at the nominal price q when the nominal price of the y good is p (known to the consumer). If q is the nominal price of good x and p is the nominal price of good y then consumers choose x, y to maximize utility $u(x)+y$ subject to the budget constraint $qx+py = M$. We assume M is sufficiently large that the non-negativity constraint on y does not bind, so maximizing utility subject to the budget constraint is equivalent to maximizing $u(x) - (q/p)x$; consumer demand $X(q; p)$ for good x is the solution to this problem. It is evident that demand is strictly decreasing in q so $X(q; p)$ is the unique solution to the equation $u'(x) = q/p$. The implicit function theorem guarantees that X is smooth and that $\partial X/\partial q < 0$ and $\partial X/\partial p > 0$.⁷ Note that demand $X(q; p)$ is positively homogeneous of degree 0: if $t > 0$ then $X(tq; tp) = X(q; p)$. Since the real marginal cost of production is k , the firm will only offer prices $q \geq pk$; for such prices, real profit (i.e. profit expressed in terms of the aggregate consumption good) is $\Pi(q; p) = (q/p - k)X(q; p)$. Since $X(q; p)$ is positively homogeneous of degree 0 (and so depends only on real prices) the same is true of profit $\Pi(tq; tp) = \Pi(q; p)$. In particular, demand and profit depend only on real prices.

⁷Demand might be zero for some prices q , but such prices will never occur at the firm optimum so we will be sloppy and ignore this possibility.

We assume that

$$\lim_{q \rightarrow \infty} qX(q; p) = 0$$

This assumption is made to guarantee that for every $k > 0$ there is a (not-necessarily unique) profit-maximizing price.

If the firm offers the nominal price q and the true state $\omega \in \{H, L\}$ is known to the consumer, then the consumer knows the nominal price $p = p_\omega$ and so demands $X(q; p_\omega)$. It is convenient to write $X_\omega(q) = X(q; p_\omega)$. Note that real profit is $(q/p_\omega - k)X_\omega(q)$. We write

$$\begin{aligned} Q_\omega &= \operatorname{argmax} (q/p_\omega - k)X_\omega(q) \\ q_\omega &= \min Q_\omega \\ \Pi_\omega^* &= \max (q/p_\omega - k)X_\omega(q) \end{aligned}$$

Note that homogeneity of degree 0 implies that

$$q \in Q_H \iff (p_L/p_H)q \in Q_L$$

and in particular that $q_L = (p_L/p_H)q_H$, and that $\Pi_L^* = \Pi_H^*$.

If the firm offers the nominal price q and the true state $\omega \in \{H, L\}$ is not known to the consumer, but the consumer maintains beliefs equal to the priors ρ_H, ρ_L , then the consumer chooses quantity x to maximize the expected value of net trade which, in view of the definition of the harmonic mean price p_0 is

$$\rho_H[u(x) - qx/p_H] + \rho_L[u(x) - qx/p_L] = u(x) - qx/p_0 \quad (2)$$

Hence consumer demand $X_0(q)$ in this case is the unique solution to the first order condition $u'(x) = q/p_0$. Real profit in this case is

$$\Pi_0(q) = (q/p_0 - k)X_0(q)$$

so the expected optimal real profit is

$$\Pi_0^* = \Pi_0(q_0) = (q_0/p_0 - k)X_0(q_0)$$

We write

$$\begin{aligned} Q_0 &= \operatorname{argmax} (q/p_0 - k)X_0(q) \\ q_0 &= \min Q_0 \end{aligned}$$

As before

$$q \in Q_H \iff (p_0/p_H)q \in Q_0$$

and in particular, $q_0 = (p_0/p_H)q_H$, and $\Pi_0^* = \Pi_H^* = \Pi_L^*$. Write Π^* for the common value of maximum real profit.

3.1 Parameters and their Interpretation

We view the probabilities ρ_H, ρ_L , the prices p_H, p_L , the firm's marginal cost k and the consumer's utility function u as parameters of the environment. We shall require that these parameters lie in the region of the parameter space in which the following three inequalities hold.

$$\frac{u(x^*)}{kx^*} > \frac{p_H}{p_0} \quad (3)$$

$$-kx^* + \left[\frac{p_0}{p_H} \right] u(x^*) > -kx_* + \left[\frac{p_L}{p_H} \right] u(x_*) \quad (4)$$

$$\frac{p_0}{p_L} > \frac{q_0 - kp_0}{q_0 - kp_H} \quad (5)$$

At first glance, it might seem that, given the parameters $p_L, p_H, \rho_L, \rho_H, u$, each of these inequalities would be satisfied whenever k was sufficiently small (the firm is sufficiently efficient). Although we believe this is the correct intuition, the truth is a bit more complicated since the inequalities involve the quantities x^*, x_*, q_0 – all of which are *derived* and depend on k (as well as on $p_L, p_H, \rho_L, \rho_H, u$). However, for the special case in which the utility function is quadratic $u(x) = x - x^2/2$, the intuition is precisely correct. To see this, calculate the derived quantities, obtaining $x^* = 1 - k$, $x_* = 1 - (p_H/p_0)k$, $q_0 = p_0(1 + k)/2$, and then and perform the requisite algebra to see that the

inequalities (3) (4) and (5) reduce to

$$\frac{1+k}{2k} > \frac{p_H}{p_0} \quad (6)$$

$$-k(1-k) + \left(\frac{p_0}{2p_H}\right)(1-k^2) > -k \left[1 - \left(\frac{p_H}{p_0}\right)k\right] + \left(\frac{p_L}{2p_H}\right) \left[1 - \left(\frac{p_H}{p_0}\right)^2 k^2\right] \quad (7)$$

$$\frac{p_0}{p_L} > \frac{p_0(1-k)}{p_0(1-k) - 2(p_H - p_0)k} \quad (8)$$

Because $p_L < p_0 < p_H$ it is clear that that the inequalities (6) (7), (8) are valid for $k = 0$; because the inequalities are strict, they are valid for sufficiently small $k > 0$. However, when utility is not assumed to be quadratic, it seems the most we can say is that there is an open set of parameters $p_H, p_L, \rho_H, \rho_L, u, k$ for which the required inequalities (3) (4), (5) obtain.

The role played by the inequalities (3) (4), (5) will become clear later but it may be useful to say a little bit now. When the true state is High, production is more expensive for the firm (in nominal terms) than it is in expectation and even more expensive than when the true state is Low. This creates difficulties in satisfying individual rationality and/or incentive compatibility constraints for the firm when the state is High. The inequality (3) solves this problem for the pooling mechanism we construct in Section 4; the inequality (4) solves this problem for the Contract-Setting Game; the inequality (5) solves this problem for the Price-Setting Game.

4 Mechanism Design

We begin by formulating the problem in terms of contracts and mechanisms; in the next Section we show that the mechanisms we identify can be implemented by a natural contract-setting game. As usual, it suffices to consider direct mechanisms; in fact we consider only direct mechanisms with some additional properties that correspond to what firms and consumers might actually do in the world. We look for the firm-optimal mechanism. We show that when the fraction of informed consumers is low the firm-optimal

mechanism is pooling; when the fraction of informed consumers is high the firm-optimal mechanism is separating. (In the intermediate range, it seems the firm-optimal mechanism may depend in a complicated way on the parameters, including the utility function of the consumer.)

Before continuing, we should address a small point. We have in mind the interaction between a single firm and many consumers, under the assumption that no consumer can observe the action of any other consumer. (This assumption preserves the asymmetry between informed and uninformed consumers.) For the purposes of mechanism design it is necessary to think how to formalize this interaction. In a direct mechanism all consumers report their types – their information. For an informed consumer, the true type is the fact that the consumer is informed and the true state. If a positive fraction of consumers are informed then (because we only consider mis-representation by at most one agent) the mechanism will always “know” the true state and hence can make assignments that depend on the true state and are independent of mis-reporting by a single consumer or by the firm. To avoid this difficulty we view the firm as interacting with each consumer independently, so that neither consumers nor the mechanism can “observe” the actions/reports of other consumers. Formally, therefore, we consider a direct mechanism with only two agents: a firm and a consumer, the latter of which could be one of several types: uninformed, informed that the state is High, informed that the state is Low. (We discuss several other issues after presenting the mechanism design framework.)

4.1 Direct Mechanisms

Formally we consider a direct mechanism with two agents: a firm and a consumer. There are two states of the world H, L with probabilities $\rho_H, \rho_L = 1 - \rho_H$. The firm is informed of the true state so may be one of two types H, L ; the consumer may be informed or uninformed and hence may be one of three types H, L, D (where being of type D means being uninformed). We use $\omega \in \{H, L\}$ for (true) states, $r, r' \in \{H, L\}$ for (true or false) reports of the firm, and $s, s' \in \{H, L, D\}$ for (true or false) reports of the consumer. Note that the common prior joint distribution of firm and consumer types (always writing the firm type/report as the first variable and the consumer

type/report as the second variable) is:

$$\begin{aligned}\rho(r, D) &= (1 - \alpha)\rho_r && \text{if } r \in \{H, L\} \\ \rho(r, s) &= \alpha\rho_r = \alpha\pi_s && \text{if } r = s \in \{H, L\} \\ \rho(r, s) &= 0 && \text{if } r \neq s \in \{H, L\}\end{aligned}$$

Of course, the true types of the firm and the informed consumer are perfectly correlated.

As usual in a direct mechanism, the firm and consumer report their types and the mechanism returns an outcome, which in this setting is a *contract* $\langle x, t \rangle$ consisting of a quantity x produced by the firm and sold to the consumer in return for the nominal transfer t from the consumer to the firm. We assume the firm cannot be forced to offer contracts that would be certain to lose money in *both* states of the world; given that profits in state ω are $-kx + t/p_\omega$ and that $p_H > p_L$ this means the set of contracts under consideration is

$$\mathcal{C} = \{\langle x, t \rangle : -kx + t/p_L \geq 0\}$$

Thus a *direct mechanism* is a function

$$\mu = (r, s) \mapsto \langle x(r, s), t(r, s) \rangle : \{H, L\} \times \{H, L, D\} \rightarrow \mathcal{C}$$

We consider only deterministic budget-balanced mechanisms. A contract $\langle x, t \rangle$ represents a trade between the consumer and the firm; if the true state is ω , firm profit is

$$\Pi(x, t|\omega) = -kx + t/p_\omega$$

and consumer utility for the trade is

$$U_c(x, t|\omega) = u(x) - t/p_\omega$$

(Of course income M does not enter into the utility for net trade.) It is important to keep in mind that, *conditional on the contract and the true state*, the informed and uninformed consumers obtain the same utility.

We compute the profit of the firm and the utility of the consumer as a function of reports of both the firm and consumer, assuming as usual that agent in question may mis-report but that the counterparty always reports truthfully. In computing the profit of the firm we must keep in mind that it

meets an informed consumer with probability α and an uninformed consumer with probability $1 - \alpha$, so if the true type of the firm (which is the true state) is $r \in \{H, L\}$ and it reports $r' \in \{H, L\}$ its expected profit is

$$\Pi(r'|r) = \alpha [t(r', r)/p_r - kx(r', r)] + (1 - \alpha) [t(r', D)/p_r - kx(r', D)]$$

(The first term is expected profit deriving from a meeting between the firm and an informed consumer who reports the true state r ; the second term is expected profit deriving from a meeting between the firm and an uninformed consumer, who truthfully reports D .) The informed consumer knows the state so if the true type of the informed consumer (which is the true state) is $s \in \{H, L\}$ and it reports $s' \in \{H, L, D\}$ its utility is

$$U_i(s'|s) = u(x(s', s)) - t(s', s)/p_s$$

The uninformed consumer does not know the state (has no private information) so when it reports $s' \in \{H, L, D\}$ its (expected) utility is :

$$\begin{aligned} U_d(s'|D) &= \rho_H [u(x(H, s')) - t(H, s')/p_H] \\ &\quad + \rho_L [u(x(L, s')) - t(L, s')/p_L] \end{aligned}$$

It follows that the Incentive Compatibility constraints for the Firm whose true type is $r \in \{H, L\}$, for the informed consumer whose true type is $s \in \{H, L\}$ and for the uninformed consumer are:

- IC-Fr** $\Pi(r|r) \geq \Pi(r'|r)$ for $r, r' \in \{H, L\}$
- IC-Is** $U_i(s|s) \geq U_i(s'|s)$ for $s \in \{H, L\}, s' \in \{H, L, D\}$
- IC-D** $U_d(D|D) \geq U_d(s'|D)$ for $s' \in \{H, L, D\}$

We are interested in mechanisms with the property that agents can refuse to participate *after* they learn the assigned contract and draw whatever inferences are possible from understanding the mechanism and learning the assigned contract but *before* they learn the true state if they did not already know it; we refer to this property as *weak ex post individual rationality*. (Ex post individual rationality in the usual sense means that all agents can refuse to participate after they learn the assigned contract *and* the true state.) Because the firm and the informed consumer *know* the true state and the mechanism, weak ex post individual rationality and ex post individual rationality both reduce to the usual interim individual rationality constraint for the firm and the informed consumer.

WIR-Fr $\Pi(r|r) \geq 0$ for $r \in \{H, L\}$

WIR-Is $U_i(s|s) \geq 0$ for $s \in \{H, L\}$

The uninformed consumer does *not know* the true state but can draw an inference about the true state from learning the contract *if* the mechanism assigns different contracts to the uninformed consumer when the state is H and when the state is L , but can draw no inference otherwise. Hence the individual rationality constraint for the uninformed consumer is

WIR-D $\mu(H, D) = \mu(L, D) \Rightarrow U_d(D|D) \geq 0$
 $\mu(H, D) \neq \mu(L, D) \Rightarrow u(q(\omega, D)) - t/p_\omega \geq 0$ for $\omega \in \{H, L\}$

We are interested in mechanisms that are incentive compatible and weakly ex post individually rational, which we abbreviate IC+WIR. Note that the incentive compatibility constraints for the firm depend on α (although the incentive compatibility constraints for the consumer(s) and the individual rationality constraints for all agents do not), so that a given mechanism may be IC+WIR for some values of α and not others.

Among IC+WIR mechanisms, we seek the *firm optimal mechanism*; i.e. the mechanism μ that maximizes the firm's *ex ante* expected real profit, which is

$$E_\rho \Pi(\mu) = E_\rho \left[\alpha [(-kx(\omega, \omega) + t(\omega, \omega)/p_\omega)] \right. \\ \left. + (1 - \alpha) [-kx(\omega, D) + t(\omega, D)/p_\omega] \right]$$

(The expectation is taken with respect to ρ , the probability distribution on states of the world.) We say a mechanism is *pooling* if

$$\mu(H, D) = \mu(L, D)$$

and *separating* otherwise. Note that we identify mechanisms as pooling or separating on the basis of the contracts assigned to the *uninformed* consumer only. As will become clear in the next section when we turn to implementation of the mechanisms we identify, this allows for the possibility that the firm *screens* by offering *menus of contracts* rather than a single contract. Screening equilibria can be pooling in the sense that the firm offers the same menu in different states, the uninformed consumers choose the same contract in each state but the informed consumers choose different contracts in different states.

4.2 Benchmark Mechanisms

It is useful to begin by introducing two benchmark mechanisms μ^0, μ^1 that are optimal for the extreme parameter values $\alpha = 0, 1$. To this end, define contracts

$$\begin{aligned}\langle x_0, t_0 \rangle &= \langle x^*, p_0 u(x^*) \rangle \\ \langle x_H, t_H \rangle &= \langle x^*, p_H u(x^*) \rangle \\ \langle x_L, t_L \rangle &= \langle x^*, p_L u(x^*) \rangle\end{aligned}$$

Then define μ^0 by

$$\begin{aligned}\mu^0(r, D) &= \langle x_0, t_0 \rangle \text{ for } r \in \{H, L\} \\ \mu^0(r, H) &= \langle x_0, t_0 \rangle \text{ for } r \in \{H, L\} \\ \mu^0(r, L) &= \langle 0, 0 \rangle \text{ for } r \in \{H, L\}\end{aligned}\tag{9}$$

and define μ^1 by

$$\begin{aligned}\mu^1(r, D) &= \langle x_r, t_r \rangle \text{ for } r \in \{H, L\} \\ \mu^1(r, s) &= \langle x_r, t_r \rangle \text{ for } r = s \in \{H, L\} \\ \mu^1(L, H) &= \langle x_L, t_L \rangle \\ \mu^1(H, L) &= \langle 0, 0 \rangle\end{aligned}\tag{10}$$

The basic facts about these benchmark mechanisms are contained in the following propositions. The proofs (and all other proofs) are deferred to the Appendix.

Proposition 1 *For every $\alpha \in [0, 1]$:*

- (i) *the mechanism μ^0 is IC+WIR*
- (ii) *the firm's expected profit is*

$$\Pi(\mu^0, \alpha) = [\rho_H + (1 - \alpha)\rho_L] [-kx^* + u(x^*)]$$

- (iii) *μ^0 extracts all the surplus from the uninformed consumer in expectation*

Note that if $\alpha = 0$ then μ^0 is the optimal IC+WIR mechanism. Indeed, it is optimal in the wider class of mechanisms that are *ex ante* individually rational for the uninformed consumer, since it maximizes profits subject to yielding the uninformed consumer expected utility at least 0. But it is *not* necessarily the optimal IC+WIR mechanism for $\alpha > 0$: the firm might prefer a mechanism that yields a contract carefully tailored to α .

Proposition 2 *There is an $\alpha^1 \in (0, 1)$ such that for every $\alpha \in (\alpha^1, 1]$:*

- (i) μ^1 is IC+WIR
- (ii) the firm's expected profit is $\Pi(\mu^1, \alpha) = -kx^* + u(x^*)$
- (iii) μ^1 extracts all the surplus from both the informed and uninformed consumer in each state
- (iv) μ^1 is the optimal IC+WIR mechanism

4.3 Optimal Mechanisms

With the preliminaries in hand we can now state the main results of this Section: when the fraction of informed consumers is small, pooling mechanisms are optimal; when the fraction of informed consumers are large, separating mechanisms are optimal.

Theorem 1 *There is an $\alpha^0 \in (0, 1)$ such that for every $\alpha \in [0, \alpha^0)$*

- (i) the firm strictly prefers the benchmark mechanism μ^0 to every separating IC+WIR mechanism
- (ii) the optimal IC+WIR mechanism is pooling

Theorem 2 *For every $\alpha \in (\alpha^1, 1]$, the separating mechanism μ^1 is the optimal IC+WIR mechanism.*

5 Contract-Setting

Because a direct mechanism is a Bayesian game, it is a tautology to say that the optimal mechanisms we have identified can be implemented as Bayesian Nash Equilibria (BNE) of *some* game form. However, we can show more than this: they can be implemented as Perfect Bayesian Nash Equilibria (PBE) of the game in which firms make take-it-or-leave-it offers of contracts or menus and consumers either Accept or Reject the given contract or select a contract from the offered menu. Note that these game forms give the firm enormous power – but even this enormous power is not enough to overcome the incentive problem leading to contract stickiness.

Contract-Setting Game The game unfolds as follows:

1. the firm and the informed consumers learn the true state $\omega \in \{H, L\}$
2. the firm offers a finite menu of contracts $\mathbb{M} = \langle x^1, t^1 \rangle, \dots, \langle x^n, t^n \rangle$
3. consumers choose a single contract from the offered menu or else Reject the entire menu – which means choosing the contract $\langle 0, 0 \rangle$

In this game a *strategy* for the firm is a map $\sigma_F : \{H, L\} \rightarrow \mathcal{M}$ (the set of finite menus of contracts), a strategy for the uninformed consumers is a map $\sigma_D : \mathcal{M} \rightarrow \mathcal{C}$ such that $\sigma_D(\mathbb{M}) \in \mathbb{M} \cup \{\langle 0, 0 \rangle\}$ and a strategy for the informed consumers is a map $\sigma_I : \mathcal{M} \times \{H, L\} \rightarrow \mathcal{C}$ such that $\sigma_I(\mathbb{M}) \in \mathbb{M} \cup \{\langle 0, 0 \rangle\}$.

As usual, a strategy profile $\sigma = (\sigma_F, \sigma_D, \sigma_I)$ is a BNE if each agent is optimizing given its own information and the strategy of other agents. It is a PBE if in addition for each offer \mathbb{M} the informed and uninformed agents hold beliefs that are consistent with Bayes' Rule and the equilibrium strategy of the firm and choose actions that are optimal with respect to those beliefs. An equilibrium is *pooling* if $\sigma_F(H) = \sigma_F(L)$ and *separating* otherwise. As noted earlier, in a pooling equilibrium it will necessarily be the case that the uninformed consumers obtain the same contracts in both states, but the informed consumers may obtain different contracts in different states: the firm may successfully screen consumers.

Theorem 3 *There is an $\alpha_*^0 \in (0, \alpha^0]$ such that if $\alpha \in [0, \alpha_*^0]$ then the mechanism μ^0 can be implemented as a pooling PBE of the Contract-Setting Game.*

We have already noted that when $\alpha > 0$, μ^0 is *not* the firm-optimal IC+WIR mechanism. Similarly, when $\alpha > 0$, in the firm-optimal PBE the firm does not offer the contract $\langle x_0, t_0 \rangle$; it can do better by offering a contract carefully tailored to the precise value of α – and may do still better by offering a menu of contracts and thereby screening the consumers. However in view of Theorem 1, when α is small, in the firm-optimal BNE the uninformed consumer must choose the same contract in both states and hence must be unable to infer the true state. Hence the firm may find it useful to screen, but it is useful to do so only by offering *the same menu in both states*.

We have seen that the mechanism μ^0 can be implemented as a PBE when α is close to 0; now we show that the mechanism μ^1 can be implemented as a PBE when α is close to 1.

Theorem 4 *If $\alpha \in (\alpha^1, 1]$ then the mechanism μ^1 can be implemented as a separating PBE of the the Contract-Setting Game.*

6 Price-Setting

Firms are often not able to offer take-it-or-leave-it contracts. In this Section we show that quite similar results obtain in the perhaps more realistic setting in which the firm is a monopolist and sets prices – but allows consumers to purchase any desired quantity at the offered price. In this setting the obvious strategic form seems sufficiently compelling that we will formulate it directly in terms of a game form and not bother with a mechanism design formulation. Note however that the firm’s problem does not reduce to that of an ordinary price-setting monopolist because some consumers are uninformed and will draw inferences from the price offered.

6.1 The Price-Setting Game

Price-Setting Game The game unfolds as follows:

1. the firm and the informed consumers learn the true state $\omega \in \{H, L\}$
2. the firm offers a nominal price $q \in [0, \infty)$

3. consumers choose and purchase a quantity x at the price q

Thus a strategy for the firm is a map $\sigma_F : \{H, L\} \rightarrow [0, \infty)$, a strategy for the uninformed consumers is a quantity choice $\sigma_D : [0, \infty) \rightarrow [0, \infty)$ that satisfies the budget constraint, and a strategy for the informed consumers is a quantity choice $\sigma_I : [0, \infty) \times \{H, L\} \rightarrow [0, \infty)$ that satisfies the budget constraint. A strategy profile $\sigma = (\sigma_F, \sigma_D, \sigma_I)$ is a BNE if all agents are optimizing (i.e. the firm maximizes expected profits and consumers maximize expected utility) given their information and the strategies of other agents. This entails that, following a price offer q *on the equilibrium path*, both the informed and uninformed consumers choose quantities that are optimal (equal to demand), with respect to their information. It is a PBE if in addition for *every price offer* q whether on or off the equilibrium path, the informed and uninformed agents hold beliefs that are consistent with Bayes' Rule and the equilibrium strategy of the firm and choose quantities that are optimal (equal to demand) with respect to those beliefs. An equilibrium is *pooling* if $\sigma_F(H) = \sigma_F(L)$ and *separating* otherwise.

We distinguish two candidate equilibria which parallel the pooling and separating mechanisms of the previous Section:

- σ^0
 - the firm offers the price q_0
 - after observing any price q and the state ω the informed consumer chooses the quantity $X_\omega(q)$
 - after observing the price q_0 , the uninformed consumer chooses the quantity $X_0(q_0)$; after observing any price $q \neq q_0$, the uninformed consumer chooses the quantity $X_L(q)$

The informed consumer knows the true state, so in a PBE whatever price q is offered, the informed consumer simply optimizes on the basis of its knowledge. The uninformed consumer does not know the true state, but must form beliefs on the basis of the price q and then optimize on the basis of those beliefs. In this case, after observing the price q_0 (on the equilibrium path) the uninformed consumer's beliefs coincide with its priors (as must be the case in a BNE), but after observing any other price $q \neq q_0$ (off the equilibrium path) the uninformed consumer believes that the state is Low with probability 1.

- σ^1
 - after observing the state ω the firm offers the price q_ω

- after observing the state ω and any price q the informed consumer chooses the quantity $X_\omega(q)$
- after observing the price q_ω , the uninformed consumer chooses the quantity $X_\omega(q_\omega)$; after observing any price $q \neq q_H, q_L$, the uninformed consumer chooses the quantity $X_L(q)$

As before, the informed consumer knows the true state, so in a PBE whatever price q is offered, the informed consumer simply optimizes on the basis of its knowledge. The uninformed consumer does not know the true state, but must form beliefs on the basis of the price q and then optimize on the basis of those beliefs. In this case, after observing a price q_ω (on the equilibrium path) the uninformed consumer believes that the state is ω (as must be the case in a BNE), but after observing any other price $q \neq q_L, q_H$ (off the equilibrium path) the uninformed consumer believes that the state is Low with probability 1.

Note that σ^0 is pooling and σ^1 is separating.

6.2 Pooling and Separating Equilibria

We prove two results in the price-setting environment that parallel our results in the contract-setting environment. The first shows that pooling is optimal for the firm – i.e. maximizes (expected) profits – when the fraction of uninformed consumers is low; the second shows that separating is optimal for the firm when the fraction of uninformed consumers is high.

Theorem 5 *There is a cut-off $\tilde{\alpha}^0 \in (0, 1)$ such that if $\alpha \in [0, \tilde{\alpha}^0)$ then*

- (i) σ^0 is a PBE of the Price-Setting Game;
- (ii) σ^0 yields higher firm profit than any separating PBE of the Price-Setting Game.

Theorem 6 *There is a cut-off $\tilde{\alpha}_1 \in (0, 1)$ such that if $\alpha \in (\tilde{\alpha}_1, 1]$ then*

- (i) σ^1 is a PBE of the Price-Setting Game
- (ii) σ^1 maximizes firm profit among all PBE of the Price-Setting Game.

7 Implications for Output and Welfare

In this brief Section, we analyze the implications of our model(s) for output and welfare. We compare produced quantities (output) and social welfare when all (or most) consumers are uninformed with quantities and social welfare when all (or most) consumers are informed.⁸ We find that, in both the mechanism design framework and the price-setting game, quantities and social welfare in the firm optimal solution are the same when all consumers are uninformed as when all consumers are informed. (Because the contract-setting game implements the firm-optimal mechanisms, the conclusions are the same as for the mechanism design framework.) Because assuming that all consumers are informed is equivalent to assuming that the firm can credibly reveal the true state, our conclusion is, surprisingly, that the firm's inability to credibly reveal the true state does not lead to either different produced quantity or to a welfare loss. When a small but strictly positive fraction of consumers are informed, the firm's inability to credibly reveal the true state does lead to a different produced quantity and a welfare loss. The extent to which quantities are different and welfare losses are larger is parametrized by α : so long as the firm pools, distortions are increasing in α (the fraction of informed consumers).

We first consider the mechanism design framework. (For the definitions of the benchmark pooling mechanism μ^0 and the benchmark separating mechanism μ^1 see Section 4.) The following proposition expresses formally our conclusions about quantities and welfare.

Proposition 3 *Total output and social welfare in the pooling mechanism μ^0 when $\alpha = 0$ coincide with total output and social welfare in the separating mechanism μ^1 when $\alpha = 1$.*

We now show that parallel conclusions obtain in the price-setting game. (For the definitions of the benchmark pooling equilibrium σ^0 and the benchmark separating equilibrium σ^1 see Section 6.)

⁸The reader acquainted with the general equilibrium model in the Appendix will note that statements here apply for all τ and $\varsigma = 1$.

Proposition 4 *Total output and social welfare in the pooling equilibrium σ^0 when $\alpha = 0$ coincide with total output and social welfare in the separating equilibrium σ^1 when $\alpha = 1$.*

Although in both the contract-setting framework and the price-setting framework, total output and social welfare are the same when $\alpha = 0$ and when $\alpha = 1$, total output and social welfare are not the same in the contract-setting framework and in the price-setting framework. In the former, the firm produces the socially optimal quantity (in each state) and extracts the full surplus from the consumer (in expectation when the consumers are uninformed and in each state separately when the consumers are informed). In the price-setting case, the firm produces less, so social welfare is lower and of course the firm must share the surplus with the consumers, so makes less profit.

Finally, we note that when $\alpha > 0$, the benchmark mechanism μ^0 assigns the quantity x^* to all consumers when the state is L and to the uninformed consumer when the state is H , but assigns 0 to the informed consumer when the state is H – so the expected produced quantity and expected social welfare are less when $\alpha > 0$. In the benchmark pooling equilibrium σ^0 , the uninformed consumers buy $X_0(q_0)$ but the informed consumers buy different quantities in the H, L states and there is a welfare loss in the H state. Thus the state is not neutral in this case.

8 Conclusion

In this paper, we have developed a strategic microfoundation for price stickiness. Our first approach is based on mechanism design; this approach makes it easier for us to solve for the firm optimal mechanism (the firm optimal outcome) and guarantees that the firm cannot do better than in our solution even in a very general contracting environment. Our conclusion is that, when many (but not all) consumers are uninformed, the firm optimal mechanism requires that contracts – and *a fortiori* prices – be sticky. The mechanism design solution can be implemented in a natural setting in which the firm offers contracts. In a more familiar setting in which the firm quotes prices we show that the firm optimal (perfect Bayesian) equilibrium again requires that prices be sticky.

In this investigation we have decided to focus extensively on attempting to provide a general model of price stickiness, i.e. analyzing several types of firm-consumer interaction. We feel this effort is worthwhile in order to deliver a solid microfoundation for the friction.

We have shown that it can be optimal for the firm to choose contracts or prices that are sticky with respect to imperfectly observed changes in the aggregate state. However, the firm might wish to choose contracts or prices that adjust perfectly to anticipated and observable changes in the aggregate state. For example, if money increases on average at a constant and commonly known rate, the pooling prices can grow at the same rate but be sticky with respect to un-anticipated and imperfectly observable variations around the average. This distinguishes our theory from a subset of existing theories such as menu cost models.

Price stickiness is a central friction in many macroeconomic models, so a microfounded model of price rigidity seems useful as a way to analyze how the friction is modified by the environment and, most importantly, how economic policy affects the friction. For instance, in work in progress we use this approach to show that *inflation targeting* increases the amount of nominal rigidity, leading to a flattening of the Phillips curve that is welfare reducing.

As it stands, our model does not provide implications that are literally testable. However, as we have observed, pooling is predicted for some ranges of parameters and not for others. This observation might provide the basis for tests across firms/industries and over time.

Appendix: Proofs

Proof of Proposition 1 (i) The proof requires checking the various IC and WIR constraints. The only one that requires a little care is the WIR constraint for the firm when the state is H . When the state is H the firm “sells” to both the informed and uninformed consumer and so its profit is $\Pi(\mu^0, \alpha|H) = -kx^* + p_0u(x^*)/p_H$; we must show that this is non-negative. Collecting terms shows that

$$-kx^* + p_0u(x^*)/p_H \geq 0 \Leftrightarrow u(x^*)/kx^* \geq p_H/p_0$$

which is precisely (3).

(ii) To see that the firm’s expected profit is indeed given by (ii), simply note that when the state is H the firm “sells” to both the informed and uninformed consumer but when the state is L the firm “sells” only to the uninformed consumer; this is (ii). That μ^0 extracts all the surplus from the uninformed consumer in expectation follows immediately definition of the contract $\langle x_0, t_0 \rangle$; this is (iii). ■

Proof of Proposition 2 (i) We must (as in the proof of Proposition 1) check a collection of IC and WIR constraints. The only one that is not trivial is the IC constraint for the firm when the state is L . If the state is L and the firm reports L then it will sell to both the informed and uninformed consumer so its expected profit will be

$$\Pi(L|L) = -kx^* + p_Lu(x^*)/p_L = -kx^* + u(x^*)$$

If the firm reports H then it will sell *only* to the uninformed consumer so its expected profit will be

$$\Pi(H|L) = (1 - \alpha)[-kx^* + p_Hu(x^*)/p_L]$$

In order that the IC constraint be satisfied, we require $\Pi(L|L) \geq \Pi(H|L)$; doing the requisite algebra we see that this will be true exactly when

$$\alpha \geq \frac{(p_H/p_L)u(x^*) - u(x^*)}{(p_H/p_L)u(x^*) - kx^*}$$

By definition, $u'(x^*) = k$; since u' is strictly decreasing and $u(0) = 0$ it follows that $u(x^*) > kx^*$. Hence the denominator of the fraction on the

right-hand side is strictly greater than the numerator and both are strictly positive. Setting α^1 equal to this fraction, we see that $\alpha^1 \in (0, 1)$ and the IC constraint of the firm in the L state is satisfied exactly when $\alpha \in [\alpha^1, 1]$, so we obtain (i).

(ii) and (iv) follow simply by plugging in the definitions. To see (iii), note that the mechanism yields the firm the largest possible profit in each state, consistent with the requirements that the mechanism be weakly ex post incentive compatible, which requires yielding the consumer non-negative utility in each state. ■

Proof of Theorem 1 If $\alpha = 0$ then (expected) firm profit in the mechanism μ^0 is $-kx^* + u(x^*) = \widehat{\Pi}$; if α is small then firm profit is almost $\widehat{\Pi}$. To establish (i) we show that if α is small then firm profit in any separating IC+WIR mechanism μ is bounded away from $\widehat{\Pi}$. To do this we first show that firm profit when the state is L is bounded by $\widehat{\Pi}$, use the IC constraint when the state is L to find a bound for firm profit when the state is H , and use this bound to find the cutoff α^0 .

Fix α and a separating IC+WIR mechanism μ . Write $\Pi_\omega(\mu)$ for (expected) firm profit in the mechanism μ when the state is ω and all agents report truthfully, and $\Pi(\mu) = \rho_H \Pi_H(\mu) + \rho_L \Pi_L(\mu)$ for expected firm profit in the mechanism μ .

We first estimate the expected profit from the benchmark mechanism μ^0 . By definition of the contract space \mathcal{C} , the mechanism never assigns any contract that yields the firm negative profits, so

$$\Pi(\mu^0) \geq (1 - \alpha)\widehat{\Pi}$$

We now turn to the separating mechanism μ . Because μ is separating, weak ex post individual rationality guarantees that both the informed and uninformed agents obtain utility at least 0 from the contracts they are assigned in the mechanism. By definition this means that, for each state ω and each consumer the contract $\langle x, t \rangle$ assigned when the state is ω must satisfy the utility constraint $u(x) - t/p_\omega \geq 0$. The contract $\langle x, t \rangle$ yields the firm a per-contract profit of $-kx + t/p_\omega$. Subject to the utility constraint, the unique firm-optimal contract is $\langle x^*, p_\omega u(x^*) \rangle = \langle x_\omega, t_\omega \rangle$ and yields firm profit $\widehat{\Pi}$ in state ω so the mechanism μ cannot assign a contract that yields per-contract profit more than $\widehat{\Pi}$ from either consumer in either state. In

particular, $\Pi_L(\mu) \leq \widehat{\Pi}$.

Let $\mu(H, D) = \langle \bar{x}, \bar{t} \rangle$ be the contract assigned to the uninformed agents when the firm reports H . Because the contracts assigned by the mechanism never yield per-contract profit greater than $\widehat{\Pi}$ in either state we must have

$$\Pi_H(\mu) = (1 - \alpha)[-k\bar{x} + \bar{t}/p_H] + \alpha\widehat{\Pi}$$

and hence

$$\bar{t} \geq p_H \left[\frac{\Pi_H(\mu) - \alpha\widehat{\Pi}}{1 - \alpha} \right]$$

Now consider the IC constraint for the firm when the true state is L . If the firm reports L it obtains profit $\Pi_L(\mu) \leq \widehat{\Pi}$; if it misreports H then it obtains profit at least $(1 - \alpha)[-k\bar{x} + \bar{t}/p_L] + \alpha \cdot 0 = (1 - \alpha)[-k\bar{x} + \bar{t}/p_L]$. Hence the IC constraint guarantees that

$$\begin{aligned} \widehat{\Pi} &\geq (1 - \alpha)[-k\bar{x} + \bar{t}/p_L] \\ &= (1 - \alpha)[-k\bar{x} + \bar{t}/p_H] + (1 - \alpha)\bar{t} [(1/p_L) - (1/p_H)] \\ &\geq \Pi_H(\mu) - \alpha\widehat{\Pi} + (1 - \alpha)p_H \left[\frac{\Pi_H(\mu) - \alpha\widehat{\Pi}}{1 - \alpha} \right] [(1/p_L) - (1/p_H)] \\ &\geq \Pi_H(\mu) - \alpha\widehat{\Pi} + [\Pi_H(\mu) - \alpha\widehat{\Pi}][(p_H/p_L) - 1] \end{aligned}$$

Rearranging and collecting terms yields

$$\Pi_H(\mu) \leq [(p_L/p_H) + \alpha]\widehat{\Pi}$$

and hence

$$\Pi(\mu) \leq \rho_L\widehat{\Pi} + \rho_H[(p_L/p_H) + \alpha]\widehat{\Pi} \quad (11)$$

We want to find α^0 so that $\Pi(\mu^0) > \Pi(\mu)$ when $\alpha < \alpha^0$. In view of the inequalities (8) and (11), it suffices to have

$$(1 - \alpha)\widehat{\Pi} > \rho_L\widehat{\Pi} + \rho_H[(p_L/p_H) + \alpha]\widehat{\Pi}$$

or equivalently to have

$$(1 - \alpha) > \rho_L + \rho_H[(p_L/p_H) + \alpha]$$

Solving for α , we see that this obtains provided that

$$\alpha < \alpha^0 = \frac{\rho_H [1 - (p_L/p_H)]}{1 + \rho_H}$$

which yields the desired result.

(ii) follows immediately since μ^0 is certainly no better for the firm than the *best* pooling IC+WIR mechanism. ■

Proof of Theorem 2 This is immediate from Proposition 2. ■

Proof of Theorem 3 The strategy of the firm is to offer the single contract $\langle x_0, t_0 \rangle$ in both states. The strategy of the informed consumer is to accept any contract that yields non-negative utility given the (known) state and to reject all others. The strategy of the uninformed consumer is to accept the offered contract, to accept any other contract that yields non-negative utility in the L state and to reject all others. (If the firm offers menus rather than single contracts, the strategy of the informed consumer is to accept the best contract among those that yield non-negative utility in the known state and to reject if no such contract is offered; the strategy of the uninformed consumer is to accept the best contract among those that yield non-negative utility in the L state and to reject if no such contract is offered.)

It is evident that these strategies constitute a BNE equilibrium of the Contract-Setting Game; to see that they constitute a PBE we have to consider what might happen off the equilibrium path. It is easy to see that the only issue is what happens when the true state is H and the firm deviates and offers a contract $\langle x, t \rangle \neq \langle x_0, t_0 \rangle$. This contract will be rejected by the uninformed consumers unless $u(x) - t/p_L \geq 0$; to simplify the analysis, suppose for the moment that $\alpha = 0$ so all consumers are uninformed. In this case, the largest profit the firm could realize from a deviation comes from choosing x, t to maximize profit $-kx + t/p_H$ subject to the constraint $u(x) - t/p_L \geq 0$. The maximum occurs when $u(x) - t/p_L = 0$ and $t = p_L u(x)$, so the profit-maximizing quantity solves $u'(x) = (p_H/p_L)k$; this is the quantity we have called x_* in equation (2). The optimal deviation is therefore the contract $\langle x_*, p_L u(x_*) \rangle$ and the profit resulting from the optimal deviation is therefore $-kx_* + (p_L/p_H)u(x_*)$. However, we have assumed in (4) that

$$-kx^* + (p_0/p_H)u(x^*) > -kx_* + (p_L/p_H)u(x_*)$$

Since the left-hand side is the profit the firm realizes from *not deviating* we see that deviation is *strictly worse* for the firm than following the prescribed

strategy. But if deviation is strictly worse when $\alpha = 0$ it will also be strictly worse when α sufficiently small, which is the assertion of the theorem.⁹ ■

Proof of Theorem 4 The strategy of the firm is to offer the single contract $\langle x_H, t_H \rangle$ when the state is H and the single contract $\langle x_L, t_L \rangle$ when the state is L . The strategy of the informed consumer is to accept any contract that yields non-negative utility given the (known) state and to reject all others. The strategy of the uninformed consumer is to accept the offered contracts, to accept any other contract that yields non-negative utility in the L state and to reject all others. (If the firm offers menus rather than single contracts, the strategy of the informed consumer is to accept the best contract among those that yield non-negative utility in the known state and to reject if no such contract is offered; the strategy of the uninformed consumer is to accept the best contract among those that yield non-negative utility in the L state and to reject if no such contract is offered.) These strategies form a PBE. ■

Proof of Theorem 5 (i) It is evident that consumers optimize following every price offer so it suffices to show that, if α is small enough, the firm does not wish to deviate.

- Suppose the true state is L . If the firm follows σ_F^0 and offers the price q_0 then its profit will be

$$\begin{aligned}\Pi_\alpha &= (1 - \alpha)(q_0/p_L - k)X_0(q_0) + \alpha(q_0/p_L - k)X_L(q_0) \\ &> (1 - \alpha)(q_0/p_L - q_0/p_0 + q_0/p_0 - k)X_0(q_0) \\ &= (1 - \alpha)[(q_0/p_0 - k)X_0(q_0) + (q_0/p_L - q_0/p_0)X_0(q_0)] \\ &= (1 - \alpha)[\Pi_0^* + (q_0/p_L - q_0/p_0)X_0(q_0)]\end{aligned}$$

If the firm deviates and offers the price $q \neq q_0$ the uninformed consumers will choose $X_L(q)$ so the firm's profit will be

$$\Pi'_\alpha = (1 - \alpha)(q/p_L - k)X_L(q) + \alpha(q/p_L - k)X_L(q) = \Pi_L^*$$

As we have noted, $\Pi_0^* = \Pi_L^*$; since $(q_0/p_L - q_0/p_0)X_0(q_0) > 0$ we see that $\Pi_\alpha > \Pi'_\alpha$ when $\alpha = 0$ and hence also when $\alpha > 0$ is smaller than some $\hat{\alpha}_0 > 0$.

- Suppose the true state is H . If the firm follows σ_F^0 and offers the price q_0 then its profit will be

$$\Pi_\alpha = (1 - \alpha)(q_0/p_H - k)X_0(q_0) + \alpha(q_0/p_H - k)X_H(q_0)$$

⁹We leave it to the interested reader to compute an explicit estimate for α_*^0 .

If the firm deviates and offers the price $q \neq q_0$ the uninformed consumers will choose $X_L(q)$ so the firm's profit will be

$$\Pi'_\alpha = (1 - \alpha)(q/p_H - k)X_L(q) + \alpha(q/p_H - k)X_H(q)$$

We need to show $\Pi_\alpha > \Pi'_\alpha$ (for all q); as before, if this is true for $\alpha = 0$ it will necessarily be true for small α .

To check that $\Pi_0 > \Pi'_0$, multiply and divide Π_0 by $q_0/p_0 - k$ and rearrange:

$$\begin{aligned} \Pi_0 &= \left(\frac{q_0/p_0 - k}{q_0/p_0 - k} \right) (q_0/p_H - k)X_0(q_0) \\ &= \left(\frac{q_0/p_H - k}{q_0/p_0 - k} \right) (q_0/p_0 - k)X_0(q_0) \\ &= (p_0/p_H) \left(\frac{q_0 - kp_H}{q_0 - kp_0} \right) \Pi_0^* \end{aligned}$$

Similarly, multiply and divide Π'_0 by $q/p_L - k$ and rearrange

$$\begin{aligned} \Pi'_0 &= \left(\frac{q/p_L - k}{q/p_L - k} \right) (q/p_H - k)X_L(q) \\ &= \left(\frac{q/p_H - k}{q/p_L - k} \right) (q/p_L - k)X_L(q) \\ &\leq \left(\frac{q/p_H - k}{q/p_L - k} \right) \Pi_L^* \\ &= (p_L)(p_H) \left(\frac{q - kp_H}{q - kp_L} \right) \Pi_L^* \\ &\leq (p_L/p_H)\Pi_L^* \end{aligned}$$

Since $\Pi_0^* = \Pi_L^*$ it suffices to show that

$$(p_0/p_H) \left(\frac{q_0 - kp_H}{q_0 - kp_0} \right) > (p_L/p_H)$$

or equivalently that

$$\frac{p_0}{p_L} > \frac{q_0 - kp_0}{q_0 - kp_H}$$

which is (5). Hence $\Pi_0 > \Pi'_0$, whence $\Pi_\alpha > \Pi'_\alpha$ for α sufficiently small. We conclude that, for α sufficiently small, the strategy profile σ^0 is a PBE of the Price-Setting Game, as asserted.¹⁰

(ii) The intuition is simple. When $\alpha = 0$ firm profit in the equilibrium σ^0 is Π^* , so when α is small then firm profit in the equilibrium σ^0 is close to Π^* . We show that if α is small then firm profit $\Pi(\tilde{\sigma})$ in any separating PBE $\tilde{\sigma}$ is bounded away from Π^* by supposing otherwise and deriving a violation of the IC constraint for the firm.

Fix $\alpha \in [0, 1]$ and a separating PBE $\tilde{\sigma}$. For $\omega \in \{H, L\}$ let $\tilde{q}_\omega = \tilde{\sigma}_F(\omega)$ be the price offered by the firm. The expected profit of the firm in the equilibrium σ^0 is

$$\begin{aligned} \Pi(\sigma^0) &= (1 - \alpha)(q_0/p_0 - k)X_0(q_0) \\ &\quad + \alpha[\rho_H(q_0/p_H - k)X_H(q_0) + \rho_L(q_0/p_L - k)X_L(q_0)] \end{aligned}$$

Note that $(q_0/p_0 - k)X_0(q_0) = \Pi_0^* = \Pi^*$ so the first term, which is the firm's profit from sales to the uninformed consumers, is $(1 - \alpha)\Pi^*$. The second term is the firm's profit from sales to informed consumers, which might be negative. However, for small α this term is at least $-\alpha\Pi^*$ so for small α we conclude that

$$\Pi(\sigma^0) \geq (1 - 2\alpha)\Pi^*$$

The expected profit to the firm in the separating equilibrium $\tilde{\sigma}$ is

$$\Pi(\tilde{\sigma}) = \rho_H(\tilde{q}_H/p_H - k)X_H(\tilde{q}_H) + \rho_L(\tilde{q}_L/p_L - k)X_L(\tilde{q}_L)$$

Suppose $\Pi(\tilde{\sigma}) \geq \Pi(\sigma^0) \geq (1 - 2\alpha)\Pi^*$. Because each of the profit terms $(\tilde{q}_H/p_H - k)X_H(\tilde{q}_H)$, $(\tilde{q}_L/p_L - k)X_L(\tilde{q}_L)$ is no greater than Π^* , it is necessary that

$$(\tilde{q}_H/p_H - k)X_H(\tilde{q}_H) > \Pi^* - 2\alpha\Pi^*/\rho_H \quad \text{and} \quad (\tilde{q}_L/p_L - k)X_L(\tilde{q}_L) > \Pi^* - 2\alpha\Pi^*/\rho_L$$

In order that $\tilde{\sigma}$ be a BNE, incentive compatibility when the state is L requires that the firm weakly prefers to offer the price \tilde{q}_L rather than the price \tilde{q}_H . When the state is L and the firm offers the price \tilde{q}_L its profit is not

¹⁰A specific estimate for how small α must be could be obtained by keeping careful track of terms as in the proof of part (i) of Theorem 1.

greater than Π^* . When the state is L and the firm offers the price \tilde{q}_H the uninformed consumers demand $X_H(\tilde{q}_H)$ and the informed consumers demand $X_L(\tilde{q}_H)$ so we must have

$$\begin{aligned}
\Pi^* &\geq (1 - \alpha)(\tilde{q}_H/p_L - k)X_H(\tilde{q}_H) + \alpha(\tilde{q}_H/p_L - k)X_L(\tilde{q}_H) \\
&\geq (1 - \alpha)(\tilde{q}_H/p_L - k)X_H(\tilde{q}_H) \\
&= (1 - \alpha)(\tilde{q}_H/p_L + \tilde{q}_H/p_H - \tilde{q}_H/p_H - k)X_H(\tilde{q}_H) \\
&= (1 - \alpha)(\tilde{q}_H/p_H - k)X_H(\tilde{q}_H) + [(1/p_L) - (1/p_H)]\tilde{q}_H X_H(\tilde{q}_H) \\
&\geq (1 - \alpha)(\tilde{q}_H/p_H - k)X_H(\tilde{q}_H) + [(1/p_L) - (1/p_H)]\Pi^* \\
&\geq (1 - \alpha)(\Pi^* - 2\alpha\Pi^*/\rho_H) + [(1/p_L) - (1/p_H)]\Pi^*
\end{aligned}$$

Simplifying and re-arranging yields

$$(1 - \alpha)2\alpha/\rho_H \geq [(1/p_L) - (1/p_H) - \alpha]$$

By assumption, $p_L < p_H$ so this inequality cannot obtain if α is sufficiently small. This is a contradiction so the proof is complete. (We leave it to the reader to find an explicit expression for the cutoff $\tilde{\alpha}^0$.) ■

Proof of Theorem 6 (i) It is clear that the strategies of the informed and uninformed consumer are optimal following any price offer, so to show that σ^1 is a PBE we have to examine three potential deviations

- (a) the true state is H and the firm offers a price $q \neq q_H$
- (b) the true state is L and the firm offers some price $q \neq q, q_H$
- (c) the true state is L and firm offers the price q_H

(We examine these assertions in the order given because only for (c) will we need to know anything about α .)

(a) If the true state is H and the firm offers the price q_H it will derive the same profit $(q_H/p_H - k)X_H(q_H) = \Pi_H^* = \Pi^*$ from each informed consumer and each uninformed consumer and so its total profit will be $\Pi_H^* = \Pi^*$. Suppose instead that the firm deviates and offers the price $q \neq q_H$ and consider separately the profit derived from the informed consumers and uninformed consumers. The informed consumers know the true state so demand $X_H(q)$, and the firm will therefore derive profit $(q/p_H - k)X_H(q)$ from each informed

consumer; this is no greater than Π_H^* , since the latter is the maximal per consumer profit when the state is known to be H . The uninformed consumers do not know the true state but observe a price different from q_H and hence believe the state to be L and so demand $X_L(q)$, and the firm will therefore derive profit $(q/p_H - k)X_L(q)$ from each informed consumer. Since $p_H > p_L$ this profit is strictly less than $(q/p_L - k)X_L(q)$, which in turn is no greater than $\Pi_L^* = \Pi^*$ since the latter is the maximal per-consumer profit when the state is known to be L . Thus this deviation does not yield higher profit from the informed consumers and strictly lower profit from the uninformed consumers, so the firm does not gain from this deviation (and indeed loses if any consumers are uninformed).

(b) If the true state is L and the firm offers a price $q \neq q_H, q_L$ then both consumers will demand $X_L(q)$ and the firm's profit will be $(q/p_L - k)X_L(q)$; since this profit is maximized when $q = q_L$ the firm cannot gain from this deviation.

(c) If the true state is L and the firm offers the price q_L both consumers will demand $X_L(q_L)$ and the firm's profit will be $(q_L/p_L - k)X_L(q_L) = \Pi_L^* = \Pi^*$. If the firm offers the price q_H the informed consumers will demand $X_L(q_H)$ and the uninformed consumers will demand $X_H(q_H)$ so the firm's profit will be

$$\Pi' = \alpha [(q_H/p_L - k)X_L(q_H)] + (1 - \alpha) [(q_H/p_L - k)X_H(q_H)]$$

Note that, because Π_L^* is the optimal profit when the state is known to be L , the first term on the right is less than $\alpha\Pi^*$ but the second is greater than $(1 - \alpha)\Pi^*$. We require that $\Pi^* \geq \Pi'$. Doing the requisite algebra and keeping in mind the signs of the various terms shows that this will be the case when

$$\alpha \geq \frac{(q_H/p_L - k)X_H(q_H) - (q_L/p_L - k)X_L(q_L)}{(q_H/p_L - k)X_H(q_H) - (q_H/p_L - k)X_L(q_H)}$$

The numerator and denominator of the fraction on the right hand side are both positive and the numerator is strictly smaller than the denominator because $(q_L/p_L - k)X_L(q_L) = \Pi^*$ and $(q_H/p_L - k)X_L(q_H) < \Pi^*$, so if we set $\tilde{\alpha}_1$ equal to the right hand side it follows that σ^1 is a PBE when $\alpha \in [\tilde{\alpha}_1, 1]$.

(ii) To see that σ^1 maximizes expected profit among all PBE – indeed among all BNE – of the Price-Setting Game, consider an alternative BNE σ . If σ is separating then all consumers know the state so the firm cannot make profit greater than Π_H^* when the state is H and Π_L^* when the state is

L ; since $\Pi_H^* = \Pi_L^* = \Pi^*$ the firm cannot make a greater profit following σ than following σ^1 . If σ is pooling then the firm offers the same price q in both states; the expected profit it makes from the uninformed consumers is no greater than Π_0^* and the profit it makes from the informed consumers in a given state ω is no greater than Π_ω^* . Since $\Pi_0^* = \Pi_\omega^* = \Pi^*$, the total expected profit it makes is no greater than Π^* , so again the firm cannot make a greater profit following σ than following σ^1 . (Note that, because there may be many profit-maximizing prices, there may be many PBE that yield the *same* profit at σ^1 .) ■

Proof of Proposition 3 Suppose first that $\alpha = 0$. By construction, the mechanism μ^0 assigns the contract $\langle x^*, p_0 u(x^*) \rangle$ to all consumers so the quantity produced is x^* and social welfare (which is the sum of firm profit and consumer welfare) is $[-kx^* + p_0 u(x^*)] + [u(x^*) - p_0 u(x^*)] = u(x^*) - kx^*$. (Of course, firm revenue and consumer expenditure cancel.) Now suppose that $\alpha = 1$. By construction, the mechanism assigns the contract $\langle -kx^*, p_\omega u(x^*) \rangle$ to all consumers when the state is ω , so the quantity produced in either state is x^* and social welfare in state ω is $[-kx^* + p_\omega u(x^*)] + [u(x^*) - p_\omega u(x^*)] = -kx^* + u(x^*)$. Thus quantity produced and social welfare are identical in the two settings, as asserted. ■

Proof of Proposition 4 Suppose first that $\alpha = 0$. By construction, the price charged in the equilibrium σ^0 is q_0 , the quantity produced is $X_0(q_0)$ and social welfare is $-kX_0(q_0) + u(X_0(q_0))$. Now suppose that $\alpha = 1$. By construction, when the state is ω the price charged in the equilibrium σ^1 is q_ω , the quantity produced is $X_\omega(q_\omega)$ and (keeping in mind that, as above, firm revenue and consumer expenditure cancel), social welfare is $-kX_\omega(q_\omega) + u(X_\omega(q_\omega))$. As noted earlier, demand depends only on the real price and $q_0/p_0 = q_\omega/p_\omega$, so $X_\omega(q_\omega) = X_0(q_0)$ and $-kX_\omega(q_\omega) + u(X_\omega(q_\omega)) = -kX_0(q_0) + u(X_0(q_0))$. Thus quantity produced and social welfare are identical in the two settings, as asserted. ■

References

- Amador, M. and P.-O. Weill (2010). Learning from prices: public communication and welfare. *Journal of Political Economy* 118(5), 866–907.
- Amador, M., I. Werning, and G.-M. Angeletos (2006). Commitment vs.

- flexibility. *Econometrica* 74(2), 365–396.
- Athey, S., A. Atkeson, and P. J. Kehoe (2005). The optimal degree of discretion in monetary policy. *Econometrica* 73(5), 1431–1475.
- Bacchetta, P. and E. Van Wincoop (2006). Can information heterogeneity explain exchange rate determination puzzle? *American Economic Review* 96(3), 552–576.
- Barro, R. J. (1972). A theory of monopolistic price adjustment. *Review of Economic Studies* 39(1), 17–26.
- Blinder, A., E. Canetti, D. Lebow, and J. Rudd (1998). *Asking About Prices: A New Approach to Understanding Price Stickiness*. New York. Russell Sage Foundation.
- Cabral, L. and A. Fishman (2012). Business as usual: A consumer search theory of sticky prices and asymmetric price adjustment? *International Journal of Industrial Organization* 30, 371–376.
- Caplin, A. and J. Leahy (1991). State-dependent pricing and the dynamics of money and output. *Quarterly Journal of Economics* 106(3), 683–708.
- Golosov, M., G. Lorenzoni, and A. Tsyvinski (2014). Decentralized trading with private information. *Econometrica* 82(3), 1055–1091.
- Golosov, M. and R. E. Lucas (2007). Menu costs and Phillips curves. *Journal of Political Economy* 115(2), 171–199.
- Hall, R. and C. Hitch (1939). Price theory and business behavior. *Oxford Economic Papers* 2, 12–45.
- Jovanovic, B. and M. Ueda (1997). Contracts and money. *Journal of Political Economy* 105(4), 700–708.
- Kahneman, D., J. L. Knetsch, and R. H. Thaler (1986). Fairness as a constraint on profit seeking: Entitlements in the market. *American Economic Review* 76(4), 728–741.
- Lagos, R. and R. Wright (2005). A unified framework for monetary theory and policy analysis. *Journal of Political Economy* 113(31), 463–484.
- L’Huillier, J.-P. (2017). Consumers’ imperfect information and price rigidities. *Mimeo*.

- Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory* 4(2), 103–124.
- Lucas, R. E. and N. L. Stokey (1987). Money and interest in a cash-in-advance economy. *Econometrica* 55(3), 491–513.
- Mackowiak, B. and M. Wiederholt (2009). Optimal sticky prices under rational inattention. *American Economic Review* 99(3), 769–803.
- Mackowiak, B. and M. Wiederholt (2015). Business cycle dynamics under rational inattention. *Review of Economic Studies* 82(4), 1502–1532.
- Mankiw, G. and R. Reis (2002). Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve. *Quarterly Journal of Economics* 117(4), 1295–1328.
- Maskin, E. and J. Tirole (1988). A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica* 56(3), 571–599.
- Nakamura, E. and J. Steinsson (2011). Price setting in a forward-looking customer markets. *Journal of Monetary Economics* 58(3), 220–233.
- Okun, A. M. (1981). *Prices and Quantities: A Macroeconomic Analysis*. Brookings Institution.
- Reis, R. (2006). Inattentive producers. *Review of Economic Studies* 73, 793–821.

[FOR ONLINE PUBLICATION]

Appendix: General Equilibrium Framework

This section lays down a general equilibrium framework in which our model of price stickiness can be embedded. The setup is fairly standard. However, it is quite involved and therefore we start its description with an overview of the key economic interactions and main technical pieces. Subsequently we fully describe every piece of the model.

Preview The population of the economy is composed by a unit mass of households. These households own a unit mass of firms, which operate in different geographic locations called islands. There is a unit mass of islands, and in each island there is a single firm.

Households are divided into workers and consumer-shoppers. For brevity, we call consumer-shoppers just ‘consumers’.

The aggregate state of the economy is the supply of money m . Firms, by assumption, are informed about this quantity.¹¹ Consumers are imperfectly informed and learn m by looking at firms’ prices.¹² Workers learn m by looking at the wage in a centralized economy-wide labor market.

Notice that in order to allow a situation in which firms have better information than consumers about the nominal aggregate state, we need to move away from monopolistic competition (or other forms of centralized goods markets). The reason is that in such market structures typically all consumers observe all prices in the economy, and therefore they would learn the aggregate state right away. In our environment, instead, consumers observe one price at a time, which limits learning and allows for price stickiness of the type addressed in this paper. Consumers become informed by seeing a

¹¹It is possible to relax this assumption and letting firms learn m from their interactions with consumers, as long as an arbitrary small proportion of consumers know m and—in contrast to Lucas (1972)—each firm sells to a representative sample of consumers. To simplify the exposition, here we assume that firms are informed right from the start.

¹²One can think about this assumption as representing the fact that—for at least a portion of the consumer population—gathering precise information directly about money supply is a costly and complex process. But prices may convey this information more readily, as it is the case in our model.

price that has adjusted to the aggregate amount of money m . Firms adjust prices as a function of how many consumers are informed.

The setup is based on two tools drawn from the literature: Lagos and Wright (2005) and Lucas and Stokey (1987). As Lagos and Wright (2005), we exploit quasilinearity and periods that are divided in two subperiods to be able to handle heterogeneity. As Lucas and Stokey (1987), we use a cash-in-advance model with credit and cash goods. The quasilinearity of preferences in our model, together with a time structure including periods and subperiods, allow us to handle the heterogeneity implied by dispersed information in a simple way, and to model game theory interactions preserving compatibility with general equilibrium. Every period is divided in two subperiods. Specifically, in the first subperiod, trade happens in a decentralized market for goods. In the second subperiod, the market for goods is centralized. Importantly, the focus is on the first subperiod, which is when price stickiness can occur. In the second subperiod trade takes place under perfect information. The infinite recurrence of periods in the model is only used as a technical device to introduce money in a standard cash-in-advance framework. Regarding the use of both credit and cash goods, we will focus on the transactions of credit goods, which will allow consumers to buy from firms without knowing the supply of money in the decentralized market. Trade of the cash good happens at the end of each period, in the centralized market, and is used simply as a way of “closing” the model.

Population and Geography There is a unit mass of households indexed by i . Each of these households is divided into a worker and a consumer-shopper, called for brevity ‘consumer’. There is a unit mass of islands, indexed by j .

Time Structure Similar to Lagos and Wright (2005), periods are divided in two subperiods. Time is discrete. Periods are indexed by τ and run from $\tau = 0$ to infinity. Subperiods are indexed by ς , and run from $\varsigma = 1$ to $\varsigma = 2$.

Money Shocks Money supply evolves as

$$\log m_\tau = \log m_{\tau-1} + \nu_\tau \tag{12}$$

where ν_τ is a monetary shock that hits *at the beginning* of period τ . ν_τ is drawn from a binary probability distribution over $\mathcal{V} = \{\nu_h, \nu_l\}$, with $\nu_h > 0$ and $\nu_l < 0$. We refer to $\nu_\tau = \nu_h$ as the High state, and to $\nu_\tau = \nu_l$ as the Low state. Both states are equally likely: $Pr(\nu_\tau = \nu_h) = Pr(\nu_\tau = \nu_l) = 1/2$. (Within this framework it is possible to specify other money supply processes

than (12).)

We impose the following assumption regarding \mathcal{V} : The space of realizations of monetary shocks \mathcal{V} is such that

$$E [e^{-\nu_\tau}] = 1 \tag{13}$$

This centering assumption implies that the the inverse of the money supply, i.e. the real value of a 1 dollar bill, is a martingale:

$$E \left[\frac{1}{m_\tau} \right] = E \left[\frac{e^{-\nu_\tau}}{m_{\tau-1}} \right] = \frac{1}{m_{\tau-1}}$$

Restriction (13) ensures that, when a firm does not make its price contingent on m_τ , it posts the same price as in the previous period. However, this assumption is not essential for any of the results of the paper.

Notice of course that (12) implies that the amount of money is the same within a period.

Information Structure Firms are informed about the state of the world, i.e. they know the realization of ν_τ from the beginning of period τ , and the implied value of m_τ . At the beginning of every period, there is an exogenous proportion α_τ of consumers who are informed. Workers become informed when they supply labor in the centralized economy-wide labor market, to be fully described below.

Goods Markets We start by describing how trade of goods happens in the decentralized market. These goods are bought on credit. Specifically, every consumer is sent randomly to an island. The sampling of consumers is such that every island receives a representative unit mass of the population of consumers.

On island j there is a firm. This firm is a monopolist and sets terms of trade for a good x . The terms of trade can be to set a price, or more generally to offer contracts. (The formulation of the general equilibrium admits both cases.) Throughout this section we refer to this firm as “firm j ” or “monopolist j ” interchangeably.

At the end of period τ ($\varsigma = 2$), consumers go to a centralized competitive markets to buy a good y on cash from a competitive firm. The price of this good is p . We now comment on the role of good x in the model. This good

is simply a way of “closing” the model, because in equilibrium, the price p will be proportional to the money supply m (to be shown later.)

Labor and Financial Markets At the end of every period τ , a number of events happen together with the opening of the centralized market for good y described earlier. First, workers sell labor in an economy-wide competitive labor market at a wage w_τ . At this point, production of all goods bought in the period takes place, and these goods are delivered to households and consumed.¹³ Moreover, as in Lucas and Stokey (1987), workers bring home labor income, credit goods are paid, and profits from firms are received. Only then financial markets open and bonds and cash for period $\tau + 1$ are traded.

Households’ Preferences Having described the environment together with the timing and information assumptions, we will now present household i ’s preferences. This household faces the problem

$$\max E_{i\zeta\tau} \left[\sum_{\tau=0}^{\infty} \beta^\tau (u(x_{i\tau}) + v(y_{i\tau}) - l_{i\tau}) \right] \quad (14)$$

where $x_{i\tau}$ is consumption of the credit good x at subperiod 1 time τ , produced by a randomly matched firm \hat{j} of island \hat{j} , $y_{i\tau}$ is consumption of the cash good, and $L_{i\tau}$ is labor supplied by the worker. $E_{i\zeta\tau}$ is the expectation operator at the relevant stage of each period, taking into account the household’s information set. This maximization is subject to the budget constraint

$$t_{\hat{j}(i,1,\tau)\tau} + p_\tau y_{i\tau} + m_{i\tau} + b_{i\tau} = (1 + r_\tau)b_{i\tau-1} + m_{i\tau-1} + \iota_\tau + w_\tau l_{i\tau} + \pi_{i\tau} \quad (15)$$

where $\hat{j}(i, 1, \tau)$ is a function that designates firm \hat{j} that is randomly matched to household i at subperiod 1 time τ , and $t_{\hat{j}(i,1,\tau)\tau}$ is the transfer from household i to firm \hat{j} (in both the mechanism design or contract-setting game formulations, to be written as price times quantity demanded in the price-setting formulation.) $b_{i\tau}$ are bond holdings at the end of the period, r_τ is the interest rate paid by bonds from the previous period, ι_τ is a lump-sum money transfer from the monetary authority, and $\pi_{i\tau}$ are total profits.¹⁴

¹³We could have avoided production taking place at the end of the period by introducing another type of labor supplied within the period. With the intention of not making the environment even more involved, we use here only one type of labor which is supplied at the end of every period.

¹⁴ ι_τ is such that $\iota_\tau = m_\tau - m_{\tau-1}$. Due to quasilinearity, all agents have the same money holdings and therefore we can write this transfer in this way.

The cash-in-advance constraint for good y is

$$p_\tau y_{i\tau} \leq m_{i\tau-1} + l_\tau \tag{16}$$

A salient feature of households' preferences is the quasilinearity in labor. It implies an absence of income effects in the demand of goods x and y which is the key for tractability in the model. We develop the reasons fully when we solve the model further down in this section.

The utility functions $u(\cdot)$ and $v(\cdot)$ are assumed to be twice continuously differentiable on \mathbb{R}^{++} , strictly increasing, and strictly concave.

Production All firms in the economy have a linear technology and produce using only labor. Within every period, monopolist j of the decentralized market produces according to the production function

$$x_{j\tau} = Al_{j\tau}$$

For simplicity, we assume that $A \equiv 1/k$ is common knowledge. In general equilibrium (derived below), the specification of production of the credit or special good x is exactly the same in the body of the paper. The equilibrium of the game played between consumers and firms is the same as in the body. Below we shall prove that, in this setup, any (game theory) equilibrium between firms and consumers is compatible with general equilibrium.

The competitive firm produces y according to the production function

$$y = l$$

where productivity has been normalized to one.

Definition of Equilibrium for the Economy A general equilibrium of this economy is given by allocations $\{x_{i\tau}, y_{i\tau}\}$, labor supply $\{l_{i\tau}\}$, labor demand $\{l_{j\tau}, l_\tau\}$, bond holdings $\{b_{i\tau}\}$, profits $\{\pi_{i\tau}\}$, nominal transfers $\{t_{j\tau}\}$, nominal prices $\{y_\tau\}$, nominal wages $\{w_\tau\}$, nominal interest rates $\{1 + r_\tau\}$, for all i, j, ς, τ , s.t.

1. Households' conditions for optimality and corresponding constraints are satisfied;
2. The mechanism design problem, contract-setting game, or price-setting game is solved as specified above;

3. The representative firm maximizes profits taking the price as given;
4. Goods, labor, bonds, and money markets clear.

Households' Optimality Conditions The condition for optimality for good x is obtained as in the body. Notice that if the shopper is uninformed, he faces uncertainty regarding w_τ . Thus, here, w_τ plays the role of p in the body of the paper, but in GE both are equal to m_τ (see below.)

When the shopper buys the cash good y , he computes a first order condition for consumption of this cash good after observing its price. This good is sold in a centralized market, and therefore its price reveals the realization of the monetary shock to the shopper in case he did not know it already. Therefore, at this point the shopper does not face any uncertainty, and the first order condition is:

$$\beta^\tau v'(y_{i\tau}) = p_\tau(\lambda_{i\tau} + \psi_{i\tau}) \quad (17)$$

The worker computes a first order condition for labor supply after observing the equilibrium wage. This is a centralized market, and therefore this wage reveals the realization of the monetary shock to the worker. Therefore, the worker does not face any uncertainty, and the first order condition is:

$$\beta^\tau = w_\tau \lambda_{i\tau} \quad (18)$$

The first order condition for money holdings is computed at a financial market at the end of every period, and therefore under perfect information:

$$\lambda_{i\tau} = E_\tau[\lambda_{i\tau+1} + \psi_{i\tau+1}] \quad (19)$$

The first order condition for bond holdings is (for the same reason) computed under perfect information:

$$\lambda_{i\tau} = (1 + r_{\tau+1})E_\tau[\lambda_{i\tau+1}] \quad (20)$$

General Equilibrium First, we conjecture that y is constant in equilibrium. If so, then the price of this good is pinned down by the cash in advance constraint, and therefore it is proportional to money supply. Optimality of production for the representative firm immediately implies that the wage w_τ

is also proportional to money supply m_τ . After the normalization of productivity of the competitive firm, we have that all of these three quantities are equal:

$$p_\tau = w_\tau = m_\tau \tag{21}$$

Then, (18) gives the value of the multiplier $\lambda_{i\tau}$. Manipulating expressions (17), (19) and (20) and using (13) gives the other equilibrium values for choices of the household as $v'(y) = 1/\beta$, $m_{i\tau} = m_\tau$, and $b_\tau = 0$ (which verifies our conjecture about y), and $r_\tau = 1/\beta - 1$. Notice that because of quasilinearity none of these depend on subperiods' choices.

It remains to check that the labor market clears. Because of quasilinearity, labor supply is set to satisfy the budget constraint. Aggregating the budget constraint gives the economy's resource constraint, and from this one can establish that the labor market clears. This implies that any solution to the mechanism design problem (or game played between firms and consumers) is compatible with GE.

This completes the characterization of the GE.