

PRICE STICKINESS IN Ss MODELS: BASIC PROPERTIES

Ricardo J. Caballero
MIT and NBER

Eduardo M.R.A. Engel
Yale University and NBER

October 12, 2006¹

Abstract

What is the relation between infrequent price adjustment and the dynamic response of the aggregate price level to monetary shocks? Caplin and Spulber (1987) provide a stark example where the answer is “none.” It is well known that by relaxing their limit assumptions some price stickiness is regained but, to our knowledge, there are no general results on this issue. In this paper we study the relation between the frequency of microeconomic adjustment and aggregate price flexibility in a generalized Ss setup. We show that for a wide class of Ss models, the aggregate price level is approximately *three times* as flexible as the frequency of microeconomic price adjustment. This rule of thumb carries over to the cyclical variation in aggregate flexibility: The degree of price flexibility varies three times as much as the frequency of microeconomic adjustment over the business cycle. We also show that in generalized Ss models, strategic complementarities reduce aggregate price flexibility for any given frequency of microeconomic price adjustment, but less so than in Calvo-type models.

Keywords: Aggregate price stickiness, adjustment hazard, adjustment frequency, generalized Ss model, strategic complementarities.

JEL classification: E32, E62.

¹Respectively: caball@mit.edu; eduardo.engel@yale.edu. This paper is being prepared for the Conference on “Microeconomic Adjustment and Macroeconomic Dynamics” to be held at the Study Center Gerzensee on October 20-21, 2006. We thank Emi Nakamura and Jon Steinsson for providing the moments we use to estimate our main model, and Virgiliu Midrigan for providing data from his model to evaluate our rules of thumb. We also thank Ruediger Bachmann, Olivier Blanchard and Bill Brainard for their comments. First draft: September, 2006.

1 Introduction

Understanding the response of the aggregate price level to monetary shocks is among the central questions in monetary economics. Since the origin of almost any aggregate nominal rigidity is a microeconomic rigidity, there have been many studies documenting microeconomic pricing behavior, in particular the frequency of microeconomic price adjustment. The most recent and impressive effort is the set of country studies sponsored by the ECB's Inflation Persistence Network, summarized in Dhyne et al. (2006). Their main finding is that the average duration of price spells in the euro area ranges from four to five quarters, which is similar to the average duration found in the U.S. by Nakamura and Steinsson (2006).²

But what is the mapping from infrequent price adjustment to aggregate price stickiness? We already know that the answer to this question can be surprising. Caplin and Spulber (1987) construct an insightful example where there is no relation between both concepts. They combine a one-sided *Ss* model of microeconomic price adjustment with a specific form of asynchronous adjustment of individual prices (the assumption of a uniform cross-section), and obtain an aggregate price level that responds one-for-one to monetary shocks. Thus there is no aggregate price stickiness in their model—the impulse response is one upon impact and zero thereafter—despite the fact that the frequency of microeconomic price adjustments can take any value. Caballero and Engel (1991) extend this result to show that monetary neutrality holds, *on average*, even if the cross section distribution of firms is not of the specific form assumed by Caplin and Spulber. Again, there is no relation between the the frequency of microeconomic price adjustments and the impulse response function.

In a related recent result, Golosov and Lucas (2006) show that the sluggishness of the aggregate price response to monetary shocks is overestimated when approximating a menu-cost model with a Calvo model, where adjustment is infrequent but uncorrelated with the size of price imbalances. That is, the frequency of microeconomic price adjustments underestimates the flexibility of the aggregate price level in *Ss* models. Similarly, Bils and Klenow (2004) report that the flexibility of aggregated price series in U.S. retail data is significantly higher than suggested by the frequency of price adjustments observed in microeconomic data: they estimate a median monthly frequency of price adjustments of 0.21, while one minus the first-order autocorrelation of the aggregate inflation series—a natural measure of aggregate

²Bils and Klenow (2004) find that average price spells in the U.S. are only half as long as in the euro area. Henceforth we use the estimates in Nakamura and Steinsson (2006) as our reference case, since these estimates are consistent with most of the existing empirical evidence documenting the frequency of microeconomic price adjustments in the U.S. (see, e.g., Kashyap (1995)) and their correction for sales is more precise than the one in Bils and Klenow (2004).

price flexibility which in their Calvo setting should equal the adjustment frequency—is 0.80.

Aside from these illustrative examples, is there anything more general that can be said about the connection between the frequency of microeconomic adjustment and the degree of flexibility of the aggregate price level? We argue in this paper that the answer is yes, and that there is a surprising (at least to us) “universal constant” relating these two concepts. For a wide class of generalized *Ss* models, *the aggregate price level is three times as flexible as the frequency of microeconomic price adjustments.*

The core of the paper derives this result and the intermediate steps to obtain it. We show why, except for the Calvo model, the frequency of microeconomic adjustment is a downward biased estimator of the degree of price flexibility. More precisely, we characterize price flexibility in state dependent models and explain why it is larger than the flexibility of a Calvo model with the same adjustment frequency. We also show that while Calvo overestimates the degree of price stickiness for a given adjustment frequency, using *Ss* models with abrupt adjustment thresholds overestimates the degree of price flexibility relative to empirically more sound models with smoother adjustment. Away from these largely theoretical limits, the relation between the adjustment frequency and the response of aggregate prices to monetary shocks becomes more robust, yielding the rule of thumb described earlier.

We also show that our rule of thumb extends to the *time-variation* in the degree of aggregate price flexibility. That is, given information about the time path of the frequency of microeconomics adjustment, one can obtain an accurate approximation of the path of price flexibility, without having to model the complex dynamics of cross section distributions in state dependent models. In particular, the degree of price flexibility is approximately three times as volatile as the frequency of microeconomic adjustments.

Finally, we show that adding strategic complementarities reduces aggregate price flexibility for any given frequency of price adjustments, and that this effect is stronger in Calvo type models than in generalized *Ss* models.

Section 2 revisits the Caplin and Spulber model and motivates, in a particularly simple setting, the themes we cover later in the paper. Section 3 begins our study of the relation of aggregate price flexibility and the frequency of microeconomic adjustment, considering a simple extension of the Caplin and Spulber model that includes the Calvo model as a limiting case. Section 4 is the core of the paper and describes the key results in the context of a generalized *Ss* model. Section 5 characterizes the time variation in aggregate price flexibility, with an application to U.S. prices. Section 6 adds strategic complementarities, and Section 7 concludes.

2 Caplin and Spulber Revisited

In this section we recreate the Caplin and Spulber (henceforth CS) result and use it to motivate many of the topics we cover in later sections.

2.1 The Model

Let us focus on the aspects of the model which are relevant to our concerns, skipping the derivation of the underlying microeconomic rules or a discussion of general equilibrium aspects, which are largely orthogonal to the issues we address (see, e.g., Stokey (2002) and Dotsey, King and Wolman (1999) for useful references on the steps we skip).

Let p_{it} and p_{it}^* denote the (log of the) actual and target price, respectively, both for firm i at time t .³ There exists a continuum of firms indexed by $i \in [0, 1]$. In CS there are no idiosyncratic shocks and, leaving aside inessential constants:

$$p_{it}^* = m_t \tag{1}$$

where m_t denotes the (log of the) money stock. The sample paths of m_t are continuous and increasing.

Aggregate (log of) output, y_t , is proportional to (the log of) real balances:

$$y_t = m_t - p_t,$$

with the aggregate (log of the) price level, p_t , defined as

$$p_t \equiv \int p_{it} di.$$

If there are no frictions in microeconomic price adjustment, $p_{it} = p_{it}^* = m_t$, so that $p_t = m_t$ and money is neutral. Suppose instead that there is a fixed cost of adjusting individual prices and hence firms adopt Ss rules in setting their prices. As usual, it is convenient to define a state variable:

$$x_{it} \equiv p_{it} - p_{it}^*.$$


The adjustment rule is such that when x_{it} reaches $s - S$, the firm increases the price by $S - s$. This large adjustment catches up with the accumulated monetary expansion since the

³In a model with stochastic adjustment costs, p_{it}^* is defined formally as the price the agent would choose, conditional on the current state of the economy, if its *current* adjustment cost draw is equal to zero.

previous adjustment and anticipates some of the expansion that will take place before the next adjustment (recall that p_{it}^* is equal to m_t plus some constants, which we have dropped for expositional convenience; also, following the convention used in generalized Ss models, we have $x = 0$ immediately after firms adjust their prices).

Firms' adjustments are not perfectly synchronized in CS, because the initial cross-section distribution of actual prices is non-degenerate. In particular, CS assume that the initial distribution of x is uniform over the entire $(s - S, 0]$ interval. It turns out that under the monotonicity and continuity assumptions for the sample paths of money, this uniform distribution is invariant: While the position of individual firms in state space changes over time, the cross-section distribution remains unchanged and uniform over $(s - S, 0]$ (see Figure 1, taken from CS, that illustrates the variation over time of $r = x + S$, for an agent i ; note that in the absence of idiosyncratic shocks, the distance between agents on the circle remains unchanged).

Figure 1: The Caplin and Spulber Model



2.2 Main Result

The main result in CS is that in this context a small monetary expansion has no effect on aggregate output, despite the fact that at any given instant most microeconomic units do not adjust their prices. To see this result, note that a monetary expansion of Δm triggers the adjustment of $\Delta m / (S - s)$ firms, and each of these firms increases its price by $(S - s)$.

The change in the aggregate price level is simply the product of these two terms:

$$\Delta p = \frac{\Delta m}{S - s}(S - s) = \Delta m \quad (2)$$

and hence

$$\Delta y = 0.$$

The standard interpretation of this result is that aggregation erases the impact of microeconomic stickiness. Our first point is that this interpretation depends crucially on the concept of price stickiness we have in mind. If, at the micro level, this concept is defined in terms of the frequency of price adjustments, then the statement is correct, since money is neutral at the aggregate level despite the fact that most firms do not adjust their prices at the micro level. Yet if we have in mind a more standard definition of stickiness, in terms of the impulse response to monetary shocks, then the usual interpretation of Caplin and Spulber's money neutrality result changes dramatically, as now it reflects the absence of *microeconomic* stickiness.

We support this claim in two steps. Consider first the price-response $\Delta p_i(\Delta m, x)$ of a representative firm i with state variable x to a small monetary shock of size Δm . A shift of x by Δm leads to no adjustment if the firm is at a distance larger than Δm from the trigger barrier. Only if x is close enough to $s - S$ does the firm adjust, from approximately $s - S$ to 0.⁴

Therefore:

$$\Delta p_i(\Delta m, x) = \begin{cases} 0, & \text{if } x > s - S + \Delta m, \\ S - s, & \text{otherwise,} \end{cases}$$

and it follows that:

$$\frac{\Delta p_i(\Delta m, x)}{\Delta m} = \begin{cases} 0, & \text{if } x > s - S + \Delta m, \\ (S - s)/\Delta m, & \text{otherwise.} \end{cases}$$

To obtain a measure of *microeconomic* flexibility in terms of the impulse response function, we average this expression over all possible values of x . The obvious candidate to weigh

⁴Strictly speaking, the adjustment is from $x - \Delta m$ to 0, but $s - S - x + \Delta m$ is sufficiently small for adjusting firms that $s - S$ is a good approximation and simplifies the expressions. Of course, the limit as $\Delta m \rightarrow 0$ does not depend on this approximation.

the different values of x is the average time-distribution (ergodic distribution) of the state variable x for a given firm. Denoting this density by $h_E(x)$ we have:

$$\frac{\Delta p_i}{\Delta m} \equiv \int_{s-S}^0 \frac{\Delta p_i(\Delta m, x)}{\Delta m} h_E(x) dx, \quad (3)$$

It is intuitively obvious (for a proof see Section 3) that in CS the ergodic density for a given firm is uniform on $(s - S, 0]$, so that $h_E(x)$ is constant and equal to $1/(S - s)$.

The impulse response of Ss (and Calvo) type models typically decreases monotonically and adds up to one because of long run neutrality. Thus, much of the persistence of the IRF to monetary shocks is summarized by the initial response to such a shock. We therefore define our main measure of microeconomic price flexibility as:

$$\mathcal{F}^{\text{micro}} \equiv \frac{\Delta p_i}{\Delta m}.$$

It follows that:

$$\mathcal{F}^{\text{micro}} = \int_{s-S}^{s-S+\Delta m} \frac{S-s}{\Delta m} \times \frac{1}{S-s} dx = 1.$$

That is, if we keep track of a firm over time and draw the histogram of its marginal responses to a monetary shock, $\Delta p_{it}/\Delta m$, most of the observations pile up at zero. Yet a small fraction of observations pile up at $(S - s)/\Delta m$, corresponding to times where the response to a monetary shock is much larger than one-for-one. The average value over time of $\Delta p_{it}/\Delta m$ is equal to one.

Property 1 (The IRF in Caplin-Spulber has no *microeconomic* price-stickiness)

There is no price stickiness at the microeconomic level in Caplin and Spulber. That is, $\mathcal{F}^{\text{micro}} = 1$ in this case.

It is easy to extend the above result to the complete impulse response function (IRF), beyond its first element. Denote by $\text{IRF}_k^{\text{micro}}$ the average price response of a firm at time k to a small monetary shock Δm in period zero, normalized by the size of the shock. We then have:

$$\text{IRF}_k^{\text{CS,micro}} = \begin{cases} 1, & \text{for } k = 0, \\ 0, & \text{for } k \geq 1. \end{cases}$$

Thus, not only $\mathcal{F}^{\text{micro}}$, but any reasonable definition of price-flexibility based on the entire IRF, assigns no *microeconomic* stickiness in the CS context.

The second step in our study of the IRF in Caplin and Spulber connects microeconomic and aggregate price stickiness. It follows from our derivation of Caplin and Spulber’s result—see equation (2)—that the aggregate response of the price level to monetary shocks is also one-for-one upon impact and zero thereafter. This is the well known CS result. Our point, however, is that the actual cross-section distribution has little to do with this result. Once the model has no microeconomic stickiness, the macro result follows *regardless* of what the cross-section distributions look like. We turn to this issue next.

2.3 Relation between Micro and Macro Stickiness

The result that micro and macro stickiness are the same is quite general and follows from the Ergodic Theorem (see, e.g., Walters (1982)). We sketch the proof of this result next.

Given a cross-section $f(x)$ for the state variable, we have that the aggregate price response to a monetary shock Δm is equal to (the superscript f indexes the cross-section distribution):

$$\frac{\Delta p^f}{\Delta m} \equiv \int \frac{\Delta p_i(\Delta m, x)}{\Delta m} f(x) dx. \quad (4)$$

To obtain an aggregate measure of price flexibility, we need to average the above expression over all possible cross-sections $f(x)$: $f_1(x), f_2(x), \dots, f_n(x)$.⁵

$$\frac{\Delta p}{\Delta m} = \sum_{k=1}^n w_k \left[\int \frac{\Delta p_i(x, \Delta m)}{\Delta m} f_k(x) dx \right] = \int \frac{\Delta p_i(x, \Delta m)}{\Delta m} \left[\sum_{k=1}^n w_k f_k(x) \right] dx, \quad (5)$$

where w_k denotes the weight of the k -th cross section, with $w_k > 0$ and $\sum_{k=1}^n w_k = 1$. Denoting the weighted average of all cross-sections by $f_A(x)$, we have:

$$\frac{\Delta p}{\Delta m} = \int \frac{\Delta p_i(\Delta m, x)}{\Delta m} f_A(x) dx. \quad (6)$$

The measure of aggregate price flexibility analogous to $\mathcal{F}^{\text{micro}}$ is defined as:

$$\mathcal{F}^{\text{macro}} \equiv \int \frac{\Delta p_i(\Delta m, x)}{\Delta m} f_A(x) dx. \quad (7)$$

⁵The actual number of cross-sections is infinite and not countable, thus measure theory is required for a formal statement and proof. We assume a finite number for illustrative purposes.

Of course, no averaging is needed in the case of CS, since all cross sections in (5) are the same in this case (and uniform on $(s - S, 0]$). More generally, however, such an average exists under rather weak conditions and, by the Ergodic Theorem, is equal to the individual firm's ergodic density, $h_E(x)$, considered in (3). Since $f_A(x) = h_E(x)$, comparing (3) and (6) yields the following property:

Property 2 (Macro and micro price flexibility are always the same)

Macro- and microeconomic price flexibility, as measured by $\mathcal{F}^{\text{macro}}$ and $\mathcal{F}^{\text{micro}}$, are the same in any (stationary) macroeconomic model:

$$\mathcal{F}^{\text{macro}} = \mathcal{F}^{\text{micro}}.$$

Furthermore, a straightforward extension of the derivation of this result shows that, at all lags, the macro impulse response function (averaged over all possible cross sections) is equal to the micro impulse response function (averaged over all individual time-series) to a monetary shock.

Our interpretation claim now follows from Properties 1 and 2:

Property 3 (The source of aggregate price flexibility in Caplin and Spulber)


When price stickiness is defined in terms of the impulse response function, the source of aggregate price flexibility in Caplin and Spulber is the absence of stickiness at the microeconomic level.

2.4 The Role of the Cross-Section

It is apparent from the previous property that the uniform cross-section distribution in CS has nothing to do with the absence of aggregate price stickiness in their model. So what is the role played by this assumption? The answer is that by choosing a distribution that does not vary over time, CS ensured that the response of the economy to aggregate shocks is the same at all moments in time. For most Ss -type models this is not the case, since the cross-section varies endogenously over time, and so does any reasonable measure of price flexibility (for example, the one defined in (4)). To illustrate this point, let us develop a simple generalization of CS (see Caballero and Engel (1991) for more details).

Consider the CS setting, except for the initial distribution of firms' state variable, which now covers only half the inaction range: the S s bands are normalized to $S - s = 2$ and the initial distribution is uniform on $[-3/2, -1/2]$.

Figure 2: p and m in an extension of Caplin-Spulber



The continuous line in Figure 2 depicts the evolution of the price level under the additional assumption that m grows linearly over time. Initially, there is a period where no firm adjusts its price and the aggregate price level does not change. Eventually, firms reach the trigger level $s - S$ and the aggregate price level rises twice as fast as money. After the last firm adjusts, a new period without price adjustment begins, and so on.

We first note that the ergodic density for a single firm continues to be uniform on the entire inaction range, $(-2, 0]$. Thus our measure of micro price-flexibility, $\mathcal{F}^{\text{micro}}$, is equal to one as in the standard version of Caplin and Spulber, which from Property 2 implies that $\mathcal{F}^{\text{macro}} = 1$ as well.

To evaluate macro price-flexibility *conditional* on a given cross-section distribution, we first define a conditional flexibility index for a cross-section $f(x)$ as:

$$\mathcal{F}^f \equiv \int \frac{\Delta p_i(\Delta m, x)}{\Delta m} f(x) dx.$$

Since in this case f varies over time, so does the conditional price-flexibility measure. In this

example it only takes two values, each one of them half the time:

$$\mathcal{F}^f = \begin{cases} 0 & \text{when no firm is adjusting } (f(-2^+) = 0), \\ 2 & \text{when some firms are adjusting } (f(-2^+) = 1). \end{cases}$$

2.5 Strategic Complementarities

The derivation above can be extended to the case with strategic complementarities. Equation (1) becomes:

$$p_{it}^* = (1 - a)m_t + ap_t, \quad (8)$$

where the parameter $a \in [0, 1/2)$ captures the extent to which firms wish to coordinate their prices. The larger a , the larger the incentives for firms to keep their prices in line with those of other firms.

The aggregate change in prices during a small time unit Δt is given by:

$$\Delta p = \Delta p^*(S - s)f(s - S)\Delta t, \quad (9)$$

where $\Delta p^*f(s - S)\Delta t$ is the fraction of firms that adjust and $S - s$ the size of their adjustments. Equation (9) extends (2) to this more general setting.

Substituting (8) in (9) and solving for Δp leads to:

$$\Delta p = \begin{cases} 0 & \text{when no firm is adjusting } (f(-2^+) = 0), \\ (2 - 2a)\mu\Delta t/(1 - 2a) & \text{when some firm are adjusting } (f(-2^+) = 1). \end{cases}$$

Hence:

$$\mathcal{F}^f = \begin{cases} 0 & \text{when no firm is adjusting } (f(-2^+) = 0), \\ (2 - 2a)/(1 - 2a) & \text{when some firms are adjusting } (f(-2^+) = 1). \end{cases} \quad (10)$$

The dash-dotted line in Figure 2 depicts the evolution of the price level when $a = 0.4$. Compared with the case without strategic complementarities ($a = 0$, continuous line), the aggregate price level remains constant during longer periods of time. The flip side is that when the aggregate price level increases, it does so at a faster rate, since a larger fraction of firms adjust their price in any given time period. The longer periods of inaction and the

shorter but brisker periods with price adjustments cancel each other out so that, on average, the flexibility index \mathcal{F} continues being equal to one.

Thus, what is special about CS within the class of one-sided Ss models is not that there is no aggregate price stickiness. It follows from the Average Neutrality Result in Caballero and Engel (1993a) that this is the case for a broad family of models of this type. Instead, what is special about CS's cross-section distribution is that it is *invariant* and hence the response of the aggregate price level to monetary shocks does not vary over time. This observation and the above extension leads to a claim we develop in Section 5: In most Ss models the cross-section distribution is not invariant, and hence the degree of aggregate price flexibility fluctuates over time. The robustness of our main conclusions (which are derived for the $a = 0$ case) to strategic complementarities is considered in Section 6.

3 From Caplin-Spulber to Calvo

What relation should we expect between the price adjustment frequency and aggregate price flexibility in menu-cost models? The one-sided Ss models described in the preceding section are at one extreme, with total price flexibility at the aggregate level regardless of the adjustment frequency. We argue next that the other extreme is the Calvo (1983) model, where both concepts coincide and aggregate flexibility is equal to the adjustment frequency. More generally, however, the answer is in between these two extremes.

In the CS model, the fraction of firms that adjust in one time period, henceforth the *frequency of adjustment index*, is equal to

$$\mathcal{A}^{\text{CS}} = \int_{s-S}^{s-S+\mu} \frac{1}{S-s} dx = \frac{\mu}{S-s},$$

where μ denotes the money growth rate (assumed constant) and $(s-S, 0]$ the inaction range. The expression above assumes that the choice of units in which time is measured is such that there are always firms that do not adjust within a given period (i.e., $\mu < S-s$). By contrast, as derived in Section 2, the price flexibility index \mathcal{F}^{CS} is one, to imply:

$$\mathcal{F}^{\text{CS}} > \mathcal{A}^{\text{CS}}.$$

Let us modify this model and assume that in addition to the trigger threshold $s-S$, there is a strictly positive hazard λ that a firm adjusts at any point in time, regardless of its price

imbalance x . Thus, we have a model that nests both Calvo and one-sided Ss models: As s tends to minus infinity we obtain the Calvo model, while if we take $\lambda = 0$ we are back to CS.

If $f(x, t)$ denotes the cross-section density at time t , then

$$f(x, t + \Delta t) = (1 - \lambda\Delta t)f(x + \mu\Delta t, t), \quad s - S < x \leq 0.$$

This follows from the fact that a necessary condition for a firm to have a price imbalance x at time $t + \Delta t$ is to have a price imbalance $x + \mu\Delta t$ at time t and that the fraction of firms at $x + \mu\Delta t$ at time t that reaches x at time $t + \Delta t$ is $1 - \lambda\Delta t$ —the remaining firms adjust because of a Calvo-type adjustment shock. From the derivation of Property 2 we know that the time-average of all possible cross-sections, $f_A(x)$, is equal to the ergodic distribution of an individual price setter. Let us calculate this average, as it is the concept we need to compute \mathcal{A} and \mathcal{F} .

Setting $f(\cdot, t + \Delta t) = f(\cdot, t) \equiv h_E(\cdot) = f_A(\cdot)$ in the expression above, using a first-order Taylor expansion and letting $\Delta t \rightarrow 0$ leads to:

$$f'_A(x) = \alpha f_A(x),$$

with $\alpha = \lambda/\mu$. Imposing that the integral of f_A over the inaction range is one, then yields:

$$f_A(x) = \frac{\alpha e^{\alpha(x+S-s)}}{e^{\alpha(S-s)} - 1}, \quad s - S \leq x \leq 0. \quad (11)$$

Choosing the unit with which we measure time small enough so that the probability of two Poisson-shocks for the same firm in a given time-period is negligible, we have that the fraction of firms that adjust in one time period is:

$$\text{Fraction of adjusters} = \lambda + (1 - \lambda)F_A(s - S + \mu) \quad (12)$$

where F_A denotes the c.d.f. for f_A . The first term on the right hand side is the fraction of firms that adjust because of a Poisson shock. The second term considers, among those that did not receive such a shock, the fraction that adjusted because their state variable reached the trigger $s - S$. It follows that:

$$\text{Fraction of adjusters} \cong \lambda + (1 - \lambda)f_A(s - S)\mu, \quad (13)$$

and from (11) we have that:

$$\mathcal{A}^\lambda = \lambda \left(1 + \frac{1 - \lambda}{e^{\alpha(S-s)} - 1} \right). \quad (14)$$

Hence, since the fraction of agents that adjust before reaching the CS-trigger $s - S$ grows with λ , the frequency of adjustment index \mathcal{A}^λ increases monotonically with λ .

Once this Poisson term is introduced, when calculating the price-flexibility index \mathcal{F} it becomes useful to turn to discrete time as well. We then have (as the impulse Δm tends to zero):

$$\frac{\Delta p^\lambda}{\Delta m} \cong \lambda + (1 - \lambda) f_A(s - S)(S - s).$$

The first term on the right hand side is the marginal price increase because of the monetary shock, for those firms with a Poisson-induced adjustment. The second term is the additional contribution to inflation from firms that adjust because they reach the trigger barrier $s - S$. It follows that:

$$\mathcal{F}^\lambda \equiv \frac{\Delta p^\lambda}{\Delta m} = \lambda + (1 - \lambda) \frac{\alpha(S - s)}{e^{\alpha(S-s)} - 1}. \quad (15)$$

Since both terms on the right hand side are strictly positive when $0 < \lambda < 1$, we have that


$$\mathcal{F}^\lambda < 1 = \mathcal{F}^{\text{CS}}, \quad 0 < \lambda < 1.$$

Furthermore, \mathcal{F}^λ does not vary monotonically with λ . It is one for $\lambda = 0$ (the CS case) and again one for $\lambda = 1$ (no micro frictions). It is decreasing for small values of λ and increasing for larger values (see Figure 3).

This stark example illustrates the complex connection between aggregate flexibility and the price adjustment frequency. Whether a decrease in the frequency of price adjustments generates more price stickiness depends a great deal on where the decline in adjustments is coming from. Or, using the Golosov and Lucas (2006) terminology, on the strength of the *selection effect*. In the stark extreme of the CS model, where only the firms with largest deviations adjust and only do so upwards, the selection effect is large enough to fully undo any price stickiness that might have emerged as the frequency of price adjustments decreases.

Once we incorporate Poisson shocks into the CS environment, we have that, as expected, the flexibility and frequency of price adjustment indices move in the same direction for large values of λ : more firms adjusting means that the aggregate price level responds more to shocks. Yet, as shown in Figure 3, which plots \mathcal{F}^λ and \mathcal{A}^λ as a function of λ for a given

Figure 3: \mathcal{F} and \mathcal{A} for a Caplin-Spulber model with Poisson Shocks



$(s - S, 0]$ interval, for small values of λ both indices move in opposite directions. The Poisson shocks weaken the selection effect, since now some of the firms that adjust are not among those who benefit the most, but instead are chosen at random. For small values of λ the weakening of the selection effect dominates over the standard positive relation between adjustment frequency and flexibility.

This example also helps motivate a point we develop in a more general context in the following section: The frequency of adjustment index is a lower bound for flexibility in Ss models, and this lower bound is achieved only by the Calvo model.⁶ To see this, compare (14) and (15), and recall that our choice of time period ensures that $\mu < (S - s)$ (also recall that $\alpha = \lambda/\mu$). It follows that, for $\lambda < 1$,

$$\mathcal{F}^\lambda > \mathcal{A}^\lambda.$$

This example illustrates that one of the determinants of the difference between \mathcal{A}^λ and \mathcal{F}^λ is the extent to which the selection effect applies. The larger the selection effect, the larger the difference. As mentioned above, the Calvo model is obtained by taking the limit as s

⁶Recall that the Calvo model can be viewed as a particular case of a (generalized) Ss model (see the following subsection), where the distribution of adjustment costs has mass λ at zero and $1 - \lambda$ at infinity.

goes to minus infinity, which yields (for all λ):

$$\lim_{s \rightarrow -\infty} \mathcal{A}^\lambda = \lim_{s \rightarrow -\infty} \mathcal{F}^\lambda = \lambda.$$

4 Generalized Ss Models

Let us generalize the model of the previous sections to consider a broader set of shocks and adjustment rules. These extensions are more easily implemented in discrete time. Shocks to the growth rate of money are i.i.d. with mean μ and variance σ_A^2 , and firms experience idiosyncratic (productivity and demand) shocks v_{it} which are i.i.d. with zero mean and variance σ_f^2 . These shocks are independent across agents and from the aggregate shock.⁷ With these assumptions, the target price follows the process:

$$\Delta p_{it}^* = \Delta m_t + v_{it}.$$

With no further changes, and preserving the fixed cost of adjusting prices at the microeconomic level, this model typically yields a two-sided Ss policy (see, e.g., Barro (1972)). We generalize it further and assume that there are i.i.d. idiosyncratic shocks to adjustment costs as well, drawn from a distribution $G(\omega)$. Integrating over all possible realizations of these adjustment costs, we obtain an *adjustment hazard*, $\Lambda(x)$, defined as the probability of adjusting—prior to knowing the current adjustment cost draw—by a firm that would adjust by x if its adjustment cost draw were zero.⁸ Of course,

$$0 \leq \Lambda(x) \leq 1, \quad \forall x.$$

It follows that for non-degenerate distributions $G(\omega)$, $\Lambda(x)$ is decreasing for $x < 0$ and increasing for $x > 0$: the cost of deviating from the target price is increasing with respect to the distance from this price and therefore adjustment is more likely when $|x|$ is larger. This is the *increasing hazard* property.

Denoting by $f(x, t)$ a cross section immediately before adjustments take place at time t ,

⁷When simulating the model we assume that the distribution of aggregate and idiosyncratic shocks are normal, but we could use some other distribution as well, as in Midrigan (2006).

⁸See Caballero and Engel (1999) for a detailed discussion of such a model, Dotsey, King and Wolman (1999) for an application to prices in a dynamic general equilibrium context, and Caballero and Engel (1993b) for an estimation of a generalized hazard model for prices.

we have:

$$\Delta p_t = - \int x \Lambda(x) f(x, t) dx.$$

The smoothness properties of this adjustment hazard, which is not present for conventional *Ss* models, is quite useful in our derivations below.

Let $h_E(x)$ denote the ergodic distribution of x for an individual firm; then the frequency of adjustment index is:

$$\mathcal{A} = \int \Lambda(x) h_E(x) dx.$$

In what follows we study the relation between the adjustment frequency and flexibility as we vary the underlying parameters (μ , σ_A , σ_I and $G(\omega)$) and therefore the shape of the adjustment hazard and the ergodic density.

4.1 A Basic Inequality

Let $\Delta p_0(\Delta m^d)$ denote the average (over all possible cross-section distributions) inflation response to a monetary deviation of Δm^d from its average growth rate. It follows that the first element of the impulse response function with respect to this shock is our flexibility index:

$$\mathcal{F} \equiv \Delta p'_0(\Delta m^d = 0).$$

To obtain a useful expression for \mathcal{F} we note that:

$$\Delta p_0(\Delta m^d) = - \int x \Lambda(x) f_A(x + \Delta m^d) dx = - \int (x - \Delta m^d) \Lambda(x - \Delta m^d) f_A(x) dx. \quad (16)$$

Differentiating this expression with respect to Δm^d and evaluating at $\Delta m^d = 0$ yields:⁹

$$\mathcal{F} = \int \Lambda(x) f_A(x) dx + \int x \Lambda'(x) f_A(x) dx \quad (17)$$

and therefore (recall that $f_A(x) = h_E(x)$)

$$\mathcal{F} = \mathcal{A} + \int x \Lambda'(x) f_A(x) dx. \quad (18)$$

It follows that:

⁹Differentiating under the integral requires that $x \Lambda(x) f(x)$ be continuous, see the Appendix for details.

Property 4 (Flexibility and adjustment frequency are the same in the Calvo model)

In the Calvo model, where $\Lambda'(x) = 0$, we have

$$\mathcal{F}^{\text{Calvo}} = \mathcal{A}^{\text{Calvo}}.$$

More importantly, it also follows that:


Property 5 (The adjustment frequency is a lower bound for flexibility)

In any increasing hazard model:

$$\mathcal{F} > \mathcal{A}.$$

Proof The increasing hazard property states that $\Lambda'(x) > 0$ for $x < 0$ and $\Lambda'(x) < 0$ for $x > 0$.¹⁰ It follows that $x\Lambda'(x) > 0$ for all x , and therefore $\int x\Lambda'(x)f(x)dx > 0$ and $\mathcal{F} > \mathcal{A}$.

Figure 4: \mathcal{F} and \mathcal{A} and the steepness of $\Lambda(x)$

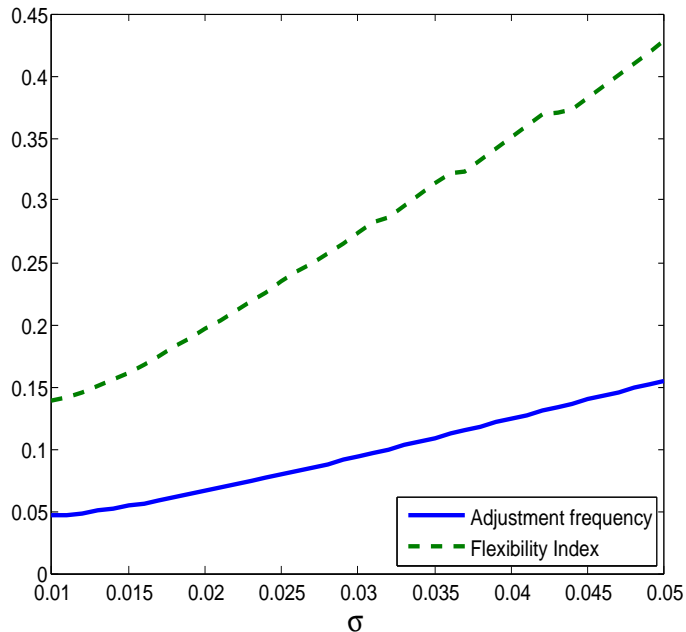


The term $\int x\Lambda'(x)f_A(x)dx$ quantifies the importance of the selection effect. It depends on how increasing the hazard is (captured by the term $x\Lambda'(x)$), and on how much weight the cross section gives to the larger values of $|x|$ (captured by $f_A(x)$).

¹⁰This result also holds if we work with the weaker concept of “increasing hazard property”, according to which $\Lambda(x)$ satisfies this property if $\Lambda'(x) \leq 0$ for $x \leq 0$ and $\Lambda'(x) \geq 0$ for $x \geq 0$, with strict inequality on a set with positive measure (under the model’s ergodic density).

Figure 4 shows how \mathcal{A} and \mathcal{F} vary when the steepness λ_2 of a quadratic hazard, $\Lambda(x) = \lambda_2 x^2$, increases.¹¹ Larger values of λ_2 imply that firms are more likely to adjust for a given imbalance x , and therefore the frequency of adjustment index increases as λ_2 grows. Figure 5 illustrates the impact of increasing the standard deviation of shocks while keeping the hazard fixed.¹² As $\sigma \equiv \sqrt{\sigma_A^2 + \sigma_I^2}$ increases, the fraction of firms adjusting increases as well, since firms move faster to regions with higher values of $\Lambda(x)$; this explains why \mathcal{A} increases with σ .

Figure 5: \mathcal{F} and \mathcal{A} as function of the volatility of shocks



In addition to the specific patterns of these figures, there is a common feature worth highlighting: Not only does flexibility \mathcal{F} increase together with \mathcal{A} , it also grows much faster than \mathcal{A} . In fact, in the next section we show that there is a sort of “universal constant” that relates \mathcal{F} and \mathcal{A} .

¹¹Parameter values correspond to Model 1 estimated in Section 4.3.

¹²Parameter values are those from Model 3 estimated in Section 4.3. Also note that in our setting it does not matter whether we vary the idiosyncratic or aggregate variance of innovations. The relevant measure is $\sigma \equiv \sqrt{\sigma_A^2 + \sigma_I^2}$. This follows from Property 2, since σ is the relevant measure of volatility when calculating the individual firm’s ergodic density.

4.2 A Rule of Thumb

Knowing that for all Ss models $\mathcal{F} > \mathcal{A}$ is of limited use without an estimate of the difference, that is, without a way to gauge the importance of the selection effect. Furthermore, a priori it would seem that this difference could depend on intricate properties of the particular Ss model under consideration, making a robust estimation of the difference between both indices rather difficult. Contrary to our priors, we find a simple rule-of-thumb that works well in a wide variety of realistic scenarios. This rule can be used to gauge price flexibility based on little more than an estimate for the frequency of price adjustments. This is good news, since recent work has provided estimates for the latter based on detailed microeconomic information.¹³

To motivate our rule of thumb, we note that from (17) we have:

$$\mathcal{F} = \int \Lambda(x) \left[1 + \frac{x\Lambda'(x)}{\Lambda(x)} \right] f_A(x) dx = \int \Lambda(x) [1 + \eta(x)] f_A(x) dx, \quad (19)$$

where

$$\eta(x) = \frac{x\Lambda'(x)}{\Lambda(x)}$$

denotes the elasticity of the adjustment hazard with respect to price imbalances as summarized by x . This elasticity determines how the fraction of adjusters varies with x . Large absolute values indicate that a firm's adjustment probability responds strongly to a marginal increase in its price imbalance. By contrast, a value of $\eta(x)$ close to zero suggest a Calvo-type behavior, where the probability of adjusting does not depend on x .

When comparing (19) with:

$$\mathcal{A} = \int \Lambda(x) f_A(x) dx$$

it becomes apparent that for a hazard with constant elasticity η we have

$$\mathcal{F} = (1 + \eta)\mathcal{A}.$$

The main idea behind our rule of thumb is that $\eta = 2$ provides a good approximation for a broad class of Ss models. To see why, note that almost any Ss type model can be

¹³For the U.S. see Bils and Klenow (2004), Klenow and Kryvtsov (2005), Midrigan (2006) and Nakamura and Steinsson (2006); for Europe see Dhyne et al. (2006) and Fabiani et al. (2006).

approximated with the following hazard:

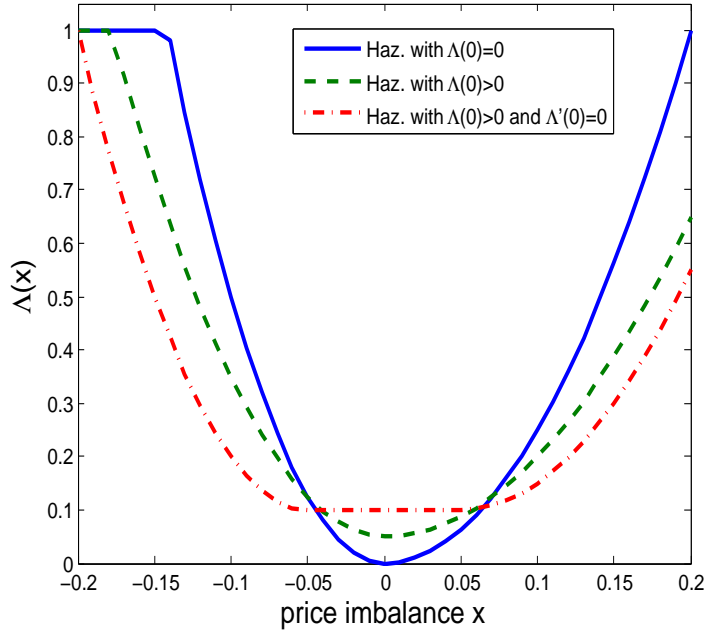
$$\Lambda(x) = \begin{cases} 1, & x \leq x_1, \\ \lambda_2^p(x - x_2)^2, & x_1 < x < x_2, \\ 0, & x_2 \leq x \leq x_3, \\ \lambda_2^n(x - x_3)^2, & x_3 < x < x_4, \\ 1, & x \geq x_4, \end{cases} \quad (20)$$

where $x_1 \leq x_2 \leq x_3 \leq x_4$. Also, λ_2^p and λ_2^n are such that $\Lambda(x)$ is non-increasing for $x < 0$ and non-decreasing for $x > 0$ and $0 \leq \Lambda(x) \leq 1$. It turns out that:

$$\mathcal{F} \cong 3\mathcal{A} \quad (21)$$

is a good approximation for most hazards within this broad class (see the appendix for a formal proposition and proof). In fact, no approximation is involved when the hazard is

Figure 6: Examples of adjustment hazards $\Lambda(x)$



quadratic with no gap around zero (for an example, see the hazard depicted with a continuous

line in Figure 6):

$$\Lambda(x) = \begin{cases} \lambda_2^p x^2, & x \leq 0, \\ \lambda_2^n x^2, & x \geq 0. \end{cases}$$

In this case $\eta(x) \equiv 2$ and $\mathcal{F} = 3\mathcal{A}$.

There are two reasons why (21) might not be a good approximation within the class of hazard described by (20). First, when a significant number of firms adjusts from values of x where the adjustment probability is one, since $\eta(x)$ is equal to zero in this range, not to 2. Second, when $x\Lambda(x)f(x)$ has jumps, as in the Caplin and Spulber model or any other stark Ss model with abrupt thresholds. The derivation of (17) assumes no such jumps, since otherwise we cannot exchange differentiation and integration (see the Appendix for details). However, these sharp scenarios are theoretical abstractions rather than accurate representations of reality, since in practice there is much more heterogeneity in agents' adjustment costs, or some other parameter, than assumed by strict Ss models. These additional sources of "noise" lead to smooth adjustment hazard, where no jumps are present and firms are likely to adjust long before they reach the region where adjustment probabilities approach one.

A more substantive reason for why (21) may not hold is that the hazard (20) rules out the possibility that $\Lambda(0) > 0$, as is the case for some Ss models with multiple-goods, such as Midrigan (2006), see Figure 7. It turns out that adding this extension still yields a simple rule of thumb.

Defining

$$\tilde{\Lambda}(x) = \Lambda(x) - \Lambda(0)$$

we have that (19) becomes


$$\mathcal{F} = \Lambda(0) + \int \tilde{\Lambda}(x)f_A(x)dx + \int x\tilde{\Lambda}'(x)f_A(x)dx,$$

and the earlier derivations carry through with $\tilde{\Lambda}(x)$ in the place of $\Lambda(x)$, leading to:

$$\mathcal{F} \cong 3\mathcal{A} - 2\Lambda(0). \tag{22}$$

As shown in the Appendix, (22) also is a good approximation when the adjustment hazard is constant in a region around zero and quadratic to the left and right (see, for example, the dashed hazard in Figure 6 and the hazard from Midrigan's model, depicted in Figure 7). Again, the approximation is good if a negligible fraction of agents adjusts from the region

Figure 7: Adjustment hazard $\Lambda(x)$ and Ergodic Density $f_A(x)$ from Midrigan (2006)



where the hazard is equal to one.

Of course, (22) can be rewritten as:

$$\mathcal{F} = \mathcal{A} + 2[\mathcal{A} - \Lambda(0)]. \quad (23)$$

If there is no selection effect, as in the Calvo model, the probability of adjusting does not depend on the extent to which the agent benefits from adjusting, that is, on x . In this case $\mathcal{A} = \Lambda(0)$ and $\mathcal{F} = \mathcal{A}$. The term $2[\mathcal{A} - \Lambda(0)]$ therefore captures the importance of the selection effect. Large values of this measure mean that the average probability that a firm adjusts, \mathcal{A} , is much larger than the probability of adjusting when a firm does not benefit at all from adjusting, $\Lambda(0)$.

Let us take stock:

Property 6 (A Useful Rule-of-Thumb)

In generalized Ss models, the following expression provides a good approximation:

$$\mathcal{F} \cong \mathcal{A} + 2[\mathcal{A} - \Lambda(0)]. \quad (24)$$

In particular, if $\Lambda(0) = 0$, as often is the case, we have:

$$\mathcal{F} \cong 3\mathcal{A}. \tag{25}$$

An interesting implication of the rule of thumb, already hinted by Figures 4 and 5, is that whatever factor increases the frequency of price adjustments, such as an increase in idiosyncratic uncertainty or a narrowing or shift of the adjustment hazard, increases flexibility three times faster than it increases the fraction of firms adjusting.

Table 1 evaluates the tightness of our rule-of-thumb in two important cases. The first row considers the model in Dotsey, King and Wolman (1999).¹⁴ In this case $\Lambda(0) = 0$ and applying our approximation does not require information beyond the average fraction of firms adjusting. The second row in Table 1 considers the model in Midrigan (2006), where multiproduct firms face economies of scale when adjusting their prices.¹⁵ This weakens the selection effect, since some prices are adjusted when they are close to their target level more often than in Dotsey, King and Wolman (1999) (a larger $\Lambda(0)$ for any given level of \mathcal{A}), thereby leading to less price flexibility. Indeed, the selection effect, as measured by $2[\mathcal{A} - \Lambda(0)]$, is equal to 0.394 under Dotsey, King and Wolman (1999), while in Midrigan (2006) it is only 0.218.

Table 1: EVALUATING THE RULE-OF-THUMB

Model	$\Lambda(0)$	\mathcal{A}	\mathcal{F}	Rule-of-thumb \mathcal{F}
Dotsey-King-Wolman (1999):	0	0.197	0.590	0.591
Midrigan (2006):	0.106	0.215	0.444	0.432

Even though Dotsey, King and Wolman (1999) and Midrigan (2006) develop considerably richer models than our stripped down version, in both cases the rule-of-thumb is as good as one could hope for.

Finally, note that the rule of thumb is conditional on the cross section of price imbalances: it says that by just knowing how much adjustment this cross-sections generate (plus one

¹⁴We read off the value of \mathcal{F} from Figure IV.B, the value for \mathcal{A} is mentioned in the main text.

¹⁵Virgiliu Midrigan provided us with data for p and p^* for 1000 firms, each producing 2 products, over 100 periods. Based on this data we estimated $\Lambda(x)$ and $f(x)$ using 101 bins of equal size, with their centers ranging from -0.4 to 0.4 (almost all price imbalances are in this range: see Figure 7, which shows the hazard and ergodic density for this model). Next we estimated $\Lambda'(x_i)$ by $(\Lambda(x_{i+1}) - \Lambda(x_{i-1}))/2h$, where the x_i denote the center of the bins and h their width. Next we calculated $\int \Lambda(x)f(x)dx$ and $\int x\Lambda'(x)f(x)dx$ based on our estimates for $\Lambda(x)$, $\Lambda'(x)$ and $f(x)$, obtaining in this way the “true” values of \mathcal{A} and \mathcal{F} . To calculate our approximation, we estimated $\Lambda(0)$ by averaging $\Lambda(x_i)$ over values of x_i between -0.04 and 0.04 .

parameter in the case of Midrigan’s model), one can obtain a fairly accurate estimate of the degree of aggregate price flexibility.

4.3 An Application to US Consumer Prices

In this section we use some moments from Nakamura and Steinsson (2006) to estimate an adjustment hazard for US consumer prices. We then use the model to determine its implied aggregate stickiness.

Our goodness-of-fit criterion considers the square root of the average absolute log deviation (henceforth, RMS) of the following four statistics calculated in Nakamura and Steinsson (2006) based on the BLS data that underlie the CPI. The first two statistics are the median frequency of upward and downward price changes (6.1% and 2.6%). The third and fourth statistics are the median size of upward and downward price adjustments (7.3% and 10.5%). All statistics exclude sales and cover monthly data over the 1998-2005 period.

Table 2: ESTIMATION OF HAZARD MODELS

	NS Data 1998-2005	Sym. haz. model 1	Asym. haz. model 2	Asym. haz. model 3
Frac. pos. adj.:	0.061	0.0504	0.0650	0.0631
Frac. neg. adj.:	0.026	0.0296	0.0254	0.0255
$E[\Delta p \Delta p > 0]$:	0.073	0.0895	0.0778	0.0757
$E[\Delta p \Delta p < 0]$:	0.105	0.0811	0.1017	0.1031
RMS-log-deviation:	—	0.2010	0.0491	0.0280
λ_0 :	—	0	0.039	0
λ_2^p :	—	28.8	51.2	61.0
λ_2^n :	—	28.8	0.00	12.4
σ_{tot} :	—	0.0260	0.0368	0.0282
\mathcal{A} :	0.087	0.0800	0.0903	0.0886
\mathcal{F} :	—	0.2392	0.1821	0.2583
\mathcal{F}/\mathcal{A} :	—	2.99	2.02	2.92
$(\mathcal{F} + 2\Lambda(0))/\mathcal{A}$:	—	2.99	2.88	2.92

We generate the ergodic densities of adjustment hazard models as in Caballero and Engel (1993b) and compute the four statistics mentioned earlier. We assume that the mean of monetary shocks is that reported in Nakamura and Steinsson (2006) for average monthly inflation (0.21%) and that idiosyncratic shocks are normal.

Table 2 summarizes our results. Model 1 reports the symmetric ($\lambda_2^p = \lambda_2^n$) quadratic hazard model that best fits the Nakamura-Steinsson statistics. We estimate three parameters: λ_0 , λ_2 and $\sigma \equiv \sqrt{\sigma_A^2 + \sigma_I^2}$, but the best non-negative estimate of λ_0 is zero. The RMS of our best model is relatively large, at around 20%. The main reason for the large RMS is that a symmetric hazard, combined with a positive average monetary growth rate, cannot generate downward adjustments that are larger than upward adjustments, as is required by the data.

Model 2 considers the downward rigidity found in Caballero and Engel (1993a), where the hazard is increasing for price imbalances that lead to price increases and constant (yet different from zero) for price imbalances that lead to price reductions. The fit improves significantly, to an RMS of 4.9%. Finally, the fit improves further if we consider asymmetric (and increasing) hazards, as in Model 3: the RMS for our best model now is only 2.8%.

Figure 8: Adjustment hazard $\Lambda(x)$ and Ergodic Density $f_A(x)$ for Model 3




Figure 8 shows the hazard function $\Lambda(x)$ and the ergodic density $f_A(x)$ for our preferred Model 3. The asymmetry in the hazard is evident: for a price imbalance of magnitude $|x|$, firms are less likely to adjust when their price is too high ($x < 0$) than when it is too low ($x > 0$). This may reflect the fact that the option value of waiting when $x < 0$ is higher than when $x > 0$, since a positive underlying inflationary process implies that future shocks are more likely to undo the current price imbalance when x is positive.

The last two rows of Table 2 test the precision of the rule of thumb for the estimated

model. The simple rule (25) does a good job for Models 1 and 3. Recall that if it worked exactly, the ratio reported in the next to last row should be 3. However, when $\Lambda(0) > 0$ as in Model 2, we need to resort to the modified rule of thumb in (24). The last row of the table reports $(\mathcal{F} + 2\Lambda(0))/\mathcal{A}$, which is equal to 3 if the modified rule holds exactly.

We see that all the estimated models have similar frequency of (monthly) price adjustments, around 10%. Flexibility, on the other hand, varies across these models and is about 26% in our preferred model. That is, while the estimated price adjustment frequency suggests that on average microeconomic prices are adjusted every 11 months, the implied aggregate flexibility is closer to 4 months. An implication of this finding is that if a researcher were to use a Calvo model as an approximation for a more realistic but complex Ss model, she should use 4 rather than 11 months in calibrating the frequency of price adjustments.

5 Time Variation in the Aggregate IRF

As we showed in Section 2, time variation in the cross-section distribution leads to price-flexibility indices that vary over time. In this section we characterize this time variation for the generalized Ss model.

5.1 Conditional price flexibility

In generalized Ss models, the cross section distribution of price imbalances is influenced by the sequence of monetary shocks hitting the economy. To quantify the extent to which the IRF fluctuates over time, we note that our rule-of-thumb applies as well to time-varying measures of price-adjustments and price-flexibility.

Denote by \mathcal{A}_t the fraction of firms adjusting their price in period t and by \mathcal{F}_t the first element of period t 's impulse response function. We then have that:

$$\mathcal{F}_t \cong 3\mathcal{A}_t - 2\Lambda(0), \tag{26}$$

with (almost) equality for hazards that belong to the family of piecewise quadratic hazards considered in Section 4. This leads to the following useful approximation:

Property 7 (A Rule-of-Thumb for the Volatility of Flexibility) *The volatility of the flexibility index \mathcal{F} is approximately three times the volatility of the fraction of firms adjusting.*

Let us now turn to a concrete application.

5.2 An application to US prices

First, note that combining the previous property with the time series for yearly averages of monthly frequencies of price changes reported by Nakamura and Steinsson (2006, Figure 4), suggests that during the 1988-2005 period annual averages of \mathcal{F}_t varied between a minimum of 25.6 percent in 1999 and a maximum of 40.2 percent in 2005; with a standard deviation of 4.7 percent.¹⁶ We also “tested” the rule of thumb with Midrigan’s (2006) model, following 1000 multiproduct firms over 100 periods. The volatility of the flexibility index \mathcal{F}_t , as measured by its standard deviation, was 0.111, compared with 0.114 obtained from our approximation (26). The corresponding means were 0.425 and 0.438 while the correlation between both series was 0.979.¹⁷ Although there are small differences, which are probably due to time variation in $\Lambda(0)$, the rule-of-thumb seems to perform reasonably well for second moments as well.

Table 3: ESTIMATION OF HAZARD MODELS: σ_A AND σ_I

	NS Data	Model 3
IQR yearly fraction adjusters:	0.88%	0.88%
IQR yearly fraction positive adjusters:	1.51%	1.44%
IQR yearly fraction negative adjusters:	0.71%	0.68%
σ_A :	—	0.00692
σ_I :	—	0.02734


Going beyond this rule of thumb, we also estimate the dynamics corresponding to Model 3 in Section 4. The first step in this estimation is to decompose σ into its aggregate and idiosyncratic components by matching the interquartile range in yearly averages of monthly adjustment frequencies reported by Nakamura and Steinsson (2006) for the 1998-2005 period (0.88%). Table 3 reports our parameter estimates. It also shows that our stylized model fits well the interquartile range for data on frequency of upward and downward price adjustments separately, even though we only used the interquartile range of all price adjustments—upward and downward—to decompose σ into its aggregate and idiosyncratic components.

In the second step, we simulated an economy based on the above model for 10,000 time periods, and calculated the flexibility index \mathcal{F}_t at each point in time. Figure 9 shows the

¹⁶The above percentages assume $\Lambda(0) = 0$, the fractions reported above decreases by $2\Lambda(0)$ otherwise.

¹⁷We calculated the “true” flexibility index and our approximation using, in each period, the methodology described in footnote 15. Since we had fewer observations in each period, we considered wider bins (0.02 instead of 0.008).

Figure 9: Histogram of \mathcal{F}_t for a time-series with with 10,000 observations



histogram of the flexibility index. It is apparent that while the bulk of the observations are around 25 percent, there are also episodes when the degree of flexibility grows sharply, even to levels above 40 percent. Which events trigger these sharp drops in stickiness? The following two figures illustrate the answer.

The left panel in Figure 10 shows the path of monthly inflation rates (percentages) for a cumulative aggregate shock, in deviation from its trend, equal to $4\sigma_A$, distributed evenly over 2, 4 and 6 months, respectively. The economy begins at the ergodic density in period 6 in all cases. The right panel shows the corresponding paths with shocks that, in deviation from their trend, are the same but of opposite sign. The difference between both panels is evident. Inflation responds much more to a sequence of positive shocks than it does to a sequence of negative shocks. This asymmetric pattern is even more prominent when we look at the evolution of the flexibility index, as shown in Figure 11. The aggregate price level (and inflation) responds much more to monetary shocks that are above average, than to shocks that are below average, since firms themselves are more responsive to positive shocks. This leads to an economy where price flexibility increases considerably after a sequence of positive shocks, but decreases only slightly after an equivalent sequence of negative shocks, explaining the highly skewed histogram shown in Figure 9. Not surprisingly, the skewness for the fraction of price adjustments (yearly averages of monthly frequencies) over the 1988-2005 period reported in Nakamura and Steinsson has a significantly positive skewness of 0.80.¹⁸

¹⁸In 10,000 bootstrap simulations of their series, skewness was positive 98.4% of the time.

Figure 10: Monthly inflation and the time-distribution of an aggregate shock

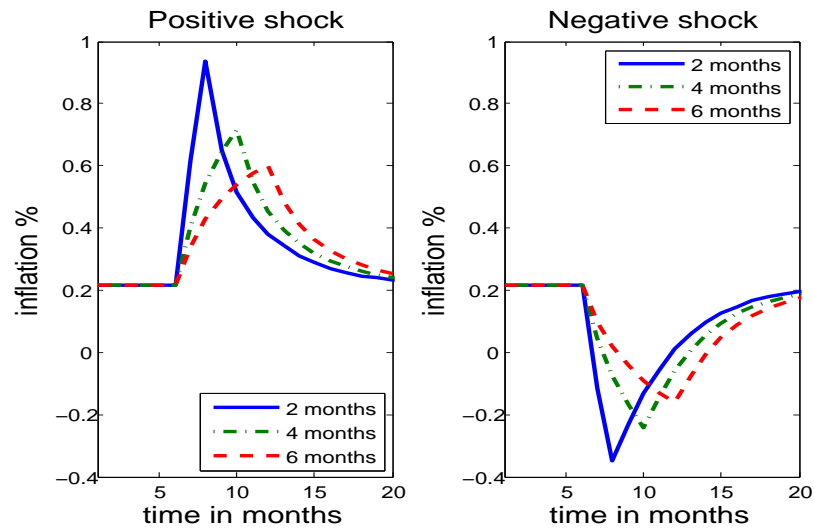
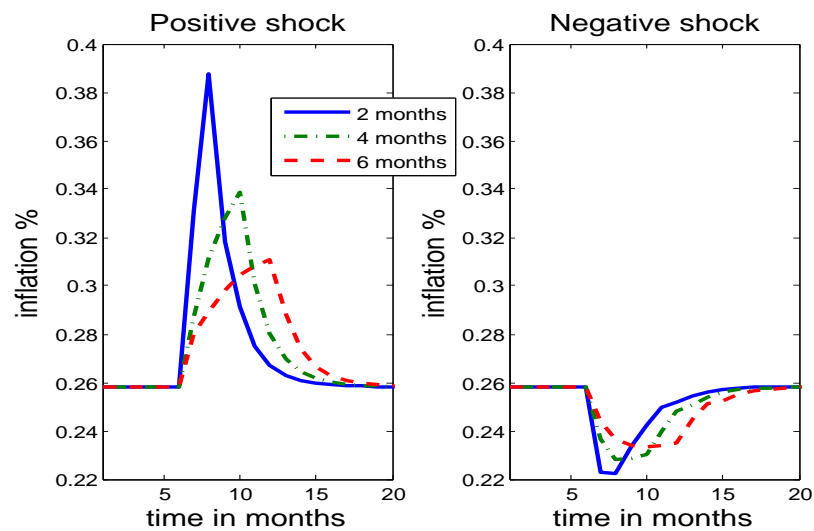


Figure 11: The \mathcal{F} index and the time-distribution of an aggregate shock



6 Strategic Complementarities

It is well known that strategic complementarities have the potential to slowdown significantly the adjustment of prices to monetary shocks (see, e.g., Ball and Romer (1990)). How do these affect our conclusions? In a nutshell, complementarities further widen the gap between price flexibility inferred under the *Ss* and Calvo model assumptions, for an observed price adjustment frequency.

The presence of strategic complementarities affect both the hazard and the dynamics of target prices. Let us isolate the latter effect and describe a firm's target price process as:

$$\Delta p_{it}^* = (1 - a)\Delta m_t + a\Delta p_t + v_{it}.$$

If $a = 0$, we recover the previous model, while when $a > 0$, there are strategic complementarities; that is, the target price for firm i depends on the aggregate price level. We assume that strategic complementarities, while present, are not strong enough to generate multiple equilibria.

Let Δm be the actual monetary shock (at time t) and v the marginal monetary shock (we will let $v \rightarrow 0$) also at time t . Denoting $\Phi(x) \equiv x\Lambda(x)$, we can write:

$$\Delta p(\Delta m + v) = \int \Phi(x + (1 - a)\Delta m + (1 - a)v + a\Delta p(\Delta m + v))f(x, t)dx.$$

Differentiating both sides with respect to v and evaluating at $v = 0$ leads to:

$$\text{IRF}_0 = [(1 - a) + a\text{IRF}_0]J,$$

where J plays the role of the integral we have been focussing on throughout the paper—that is, the sum of both terms on the r.h.s. of (17)—which in this case takes the form:

$$J \equiv \int \Phi'(x + (1 - a)\Delta m + a\Delta p(\Delta m))f(x, t)dx.$$

It follows that:

$$\mathcal{F} = \frac{(1 - a)J}{1 - aJ}.$$

If, for a given fraction of adjusters, \mathcal{A} , we fit a Calvo model, we have that $J = \mathcal{A}$ and therefore:

$$\mathcal{F}^{\text{Calvo}} = \frac{(1 - a)\mathcal{A}}{1 - a\mathcal{A}},$$

which implies that flexibility decreases as strategic complementarities rise (a rises). This is the result typically highlighted in the literature. Our main point, however, is that while the effect of strategic complementarities on aggregate price inertia is still present in Ss models, it is somewhat diluted relative to the strength of this mechanism in a Calvo model.


Using the rule of thumb to substitute $3\mathcal{A}$ for J , in an increasing hazard model (with $\Lambda(0) = 0$) we have:

$$\mathcal{F}^{Ss} \cong \frac{3(1-a)\mathcal{A}}{1-3a\mathcal{A}}.$$

It is now easy to see that for $a > 0$, we have $\mathcal{F}^{Ss}/\mathcal{F}^{\text{Calvo}} > 3$. In fact, if $a\mathcal{A} \ll 1$:

$$\frac{\mathcal{F}^{Ss}}{\mathcal{F}^{\text{Calvo}}} \cong 3 + 6a\mathcal{A}.$$

Figure 12: \mathcal{F}^{Ss} and $\mathcal{F}^{\text{Calvo}}/\mathcal{F}^{Ss}$ as a function of a for Model 3



As strategic complementarities rise, both Calvo and Ss models become more sticky, but proportionally less in Ss than in Calvo. The reason is that complementarities depend on \mathcal{F} rather than on \mathcal{A} and, for a given fraction of firms adjusting, $\mathcal{F}^{Ss} > \mathcal{F}^{\text{Calvo}}$.

The left panel in Figure 12 shows how \mathcal{F}^{Ss} varies with a for Model 3 estimated in Section 4. As expected, aggregate price flexibility decreases when strategic complementarities are

more important. However, the right panel in Figure 12 reports the ratio $\mathcal{F}^{\text{Ss}}/\mathcal{F}^{\text{Calvo}}$ as a function of a , clearly showing that strategic complementarities induce more flexibility in a Calvo setting than in an Ss setting. Nonetheless, the ratio between both flexibility indices remains close to 3, as in the case without strategic complementarities.

Let us take stock:

Property 8 (*Strategic complementarities and relative flexibility*) For an observed frequency of price adjustments, strategic complementarities decrease price flexibility, more so in the Calvo than in the Ss model.

7 Final Remarks

The recent work of Golosov-Lucas (2006), Burstein (2002), Klenow and Krystow (2005), Midrigan (2006), Gertler and Leahy (2006), and others has rekindled the interest on menu cost type models. As with many microeconomic variables, the prices of goods and services are seldom adjusted continuously. However, the implications of microeconomic inaction for the stickiness of the aggregate price level is only gradually being understood, usually restricted to important but specific examples. In this article we have tried to take a step further in understanding the conceptual issues involved in the connection between the frequency of price adjustments and aggregate price flexibility, extracting the basic aggregate properties of generalized Ss models. Throughout the paper we listed several properties, the most important of which is that, away from largely theoretical stark Ss models, there is a simple rule of thumb mapping the frequency of price adjustments, \mathcal{A} , to aggregate price flexibility, \mathcal{F} , for a broad class of Ss models:

$$\mathcal{F}_t \cong 3\mathcal{A}_t - 2\Lambda(0),$$

where $\Lambda(0)$ denotes the probability of price adjustments for agents that benefit the least from changing their price, and is equal to zero in many instances.

Of course, there is nothing specific to prices in our results. These are properties about contexts where lumpy microeconomic adjustment is prevalent, of which prices is just one, and probably not the best, application. At the microeconomic level, investment decisions, durable purchases, hiring and firing decisions, inventory accumulation, and many other important economic variables are lumpy in nature and hence fit the essentials of the model in this paper.

Finally, for realistic parameter values, time variation in the impulse response function of aggregate prices does not seem to be as important as in other applications, such as investment (see, e.g., Bachmann et al 2006). Given this observation, in price applications that do not focus on extreme events, a sensible strategy may well be to continue using the simpler Calvo model. However, in such a case it is important to recalibrate the adjustment probability parameter λ to match \mathcal{F} rather than \mathcal{A} . The parameter \mathcal{F} can be obtained from our rule of thumb combined with microeconomic evidence on \mathcal{A} from studies such as Bils and Klenow (2004), Fabiani et al (2006), Nakamura and Steinsson (2006), and Midrigan (2006). For example, Nakamura and Steinsson (2006) find that prices are adjusted approximately every 11 months, which from our rule of thumb implies that the Calvo-equivalent model should be calibrated to firms adjusting their prices on average every 4 months.

References

- [1] Bachmann, R., R.J. Caballero and E.M.R.A. Engel (2006). “Lumpy Investment in Dynamic General Equilibrium,” NBER Working Paper No. 12336.
- [2] Ball, L. and D. Romer (1990). “Real Rigidities and the Non-Neutrality of Money,” *Review of Economic Studies*, **57**, 193–203.
- [3] Barro, R.J. (1972). “A Theory of Monopolistic Price Adjustment” *Review of Economic Studies* **39**, 17–26.
- [4] Bils, M. and P.J. Klenov (2004). “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy* **112**(5), 947–985.
- [5] Burstein, A. (2002) “Inflation and Output Dynamics with State Dependent Pricing”, Mimeo, UCLA.
- [6] Caballero, R.J. and E.M.R.A. Engel (1991). “Dynamic (S, s) Economies,” *Econometrica*, **59**(6), 1659–1686.
- [7] ————— (1993a). “Heterogeneity and Output Fluctuations in a Dynamic Menu-Cost Economy”, *Review of Economic Studies*, **60**, 95–119.
- [8] ————— (1993b). “Microeconomic Rigidities and Aggregate Price Dynamics”, *European Economic Review*, **37** (4), 697–711.
- [9] ————— (1999). “Explaining Investment Dynamics in US Manufacturing: A Generalized (S, s) Approach,” *Econometrica* **67**(4), 741–782.
- [10] Calvo, G. (1983). “Staggered Prices in a Utility Maximizing Framework” *Journal of Monetary Economics*, **12**, 383–398
- [11] Caplin, A. and D. Spulber (1987). “Menu Costs and the Neutrality of Money,” *Quarterly Journal of Economics*, **102**, 703–726.
- [12] Dhyne, E., L.J. Alvarez, H.L. Bhan, G. Veronese, D. Dias, J. Hoffmann, N. Jonker, P. Lunnemann, F. Rumler and J. Vilmunen (2006). “Price Setting in the Euro Area and the United States: Some Facts From Individual Consumer Price Data,” *J. of Economic Perspectives*, **20**(2), 171–192.

- [13] Dotsey, M., R.G. King, and A.L. Wolman (1999)., “State Dependet Pricing and the General Equilibrium Dynamics of Money and Output” *Quarterly Journal of Economics*, **114**, 655–690
- [14] Fabiani, S., M. Druant, I. Hernando, C. Kwapil, B. Landau, C. Loupias, F. Martins, T. Matha, R. Sabbatini, H. Stahl, and A. Stokman (2006). “The Pricing Behavior of Firms in the Euro Area: New Survey Evidence,” *International Journal of Central Banking*, **2** (3), September, 3–47.
- [15] Gertler, M. and J. Leahy (2006). “A Phillips Curve with an *Ss* Foundation,” NBER Working Papers 11971.
- [16] Golosov, M. and R.E. Lucas (2006). “Menu Costs and the Phillips Curve,” MIT Mimeo.
- [17] Kashyap, A.K. (1995). “Sticky Prices: New Evidence from Retail Catalogs,” *Quarterly Journal of Economics* **110**, 245–274.
- [18] Klenow, P.J. and O. Kryvtsov (2005). “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation,” Working Paper, Stanford Univeristy.
- [19] Midrigan, V. (2006). “Menu Costs, Multi-product Firms, and Aggregate Fluctuations.” Mimeo.
- [20] Nakamura, E. and J. Steinsson (2006). “Five Facts About Prices: A Reevaluation of Menu Cost Models,” mimeo, Harvard University.
- [21] Stokey, N. (2002). *Brownian Models in Economics*. Manuscript.
- [22] Walters, P. (1982). *An Introduction to Ergodic Theory*, Springer-Verlag.

A Rule of Thumb: Formal Proof

The following proposition provides an expression for the size of the selection effect for a rich family of hazards.

Proposition 1 *Consider the following family of adjustment hazards:*

$$\Lambda(x) = \begin{cases} 1, & x \leq x_1, \\ \lambda_0 + \lambda_2^p(x - x_2)^2, & x_1 < x < x_2, \\ \lambda_0, & x_2 \leq x \leq x_3, \\ \lambda_0 + \lambda_2^n(x - x_3)^2, & x_3 < x < x_4, \\ 1, & x \geq x_4, \end{cases} \quad (27)$$

where $x_1 \leq x_2 \leq x_3 \leq x_4$. Also, λ_0 , λ_2^p and λ_2^n are such that $\Lambda(x)$ is non-increasing for $x < 0$ and non-decreasing for $x > 0$ and $0 \leq \Lambda(x) \leq 1, \forall x$.

Denote by π_1 the probability assigned by the cross section density $f(x)$ to values of x with $\Lambda(x) = 1$ and by D_p and D_n the absolute value of the discontinuity of $x\Lambda(x)f(x)$ at x_1 and x_4 , respectively. We then have:

$$\mathcal{F} = \mathcal{A} + (D_p - D_n) + 2[(1 - \pi_1)(1 - \Lambda(0)) + \mathcal{I}] + C, \quad (28)$$

with

$$C = 2\lambda_2^p x_2 \int_{x_1}^{x_2} (x - x_2) f(x) dx + 2\lambda_2^n x_3 \int_{x_3}^{x_4} (x - x_3) f(x) dx.$$

Proof Denote $A_i = (x_{i-1}, x_i)$, $i = 1, 2, 3, 4, 5$; with $x_0 = -\infty$ and $x_5 = \infty$. Defining $\Delta p_0(v)$ as in (16), with v in place of Δm^d ,

$$\Delta p_0(v) = - \int (x - v) \Lambda(x - v) f(x) dx = - \sum_{i=1}^5 \int_{x_{i-1}+v}^{x_i+v} (x - v) \Lambda(x - v) f(x) dx. \quad (29)$$

A basic result from calculus (Leibniz's formula) states that if $a(v)$, $b(v)$ and $g(x, v)$ are differentiable and

$$G(v) \equiv \int_{a(v)}^{b(v)} g(x, v) dx,$$

then

$$G'(v) = g(b(v), v)b'(v) - g(a(v), v)a'(v) + \int_{a(v)}^{b(v)} \frac{\partial g}{\partial v}(x, v) dx.$$

Applying this result to calculate the five integrals in the sum on the r.h.s. of (29), differentiating w.r.t. v and evaluating at $v = 0$:

$$\mathcal{F} = - \sum_{i=1}^5 \{x_i \Lambda(x_i^-) f(x_i^-) - x_{i-1} \Lambda(x_{i-1}^+) f(x_{i-1}^+)\} + \sum_{i=1}^5 \int_{x_{i-1}}^{x_i} [\Lambda(x) + x \Lambda'(x)] f(x) dx, \quad (30)$$

where $f(x_i^+)$ and $f(x_i^-)$ denote the limit of $f(x)$ when x approaches x_i from the right (larger values than x_i) and left, respectively.

The first sum on the r.h.s. of (30) corresponds to the discontinuities of $x \Lambda(x) f(x)$ —in our setting jumps may occur only at x_1 and x_4 . Since we are working with densities $f(x)$ with finite expectations, we have that $\lim_{x \rightarrow \pm\infty} x f(x) = 0$, which combined with (30) implies

$$- \sum_{i=1}^5 \{x_i \Lambda_i(x_i) f(x_i^-) - x_{i-1} \Lambda_i(x_{i-1}) f(x_{i-1}^+)\} = D_p - D_n. \quad (31)$$

Since in A_1 and A_5 we have $\Lambda(x) = 1$ and $\Lambda'(x) = 0$, it follows that

$$\int_{A_1} [\Lambda(x) + x \Lambda'(x)] f(x) dx + \int_{A_5} [\Lambda(x) + x \Lambda'(x)] f(x) dx = \int_{A_1 \cup A_5} f(x) dx = \pi_1. \quad (32)$$

We also have:

$$\begin{aligned} \int_{A_2} [\Lambda_1(x) - \lambda_0 + x \Lambda_1'(x)] f(x) dx &= 3\lambda_2^p \int_{x_1}^{x_2} (x - x_2)^2 f(x) dx + 2\lambda_2^p x_2 \int_{x_1}^{x_2} (x - x_2) f(x) dx, \\ \int_{A_4} [\Lambda_1(x) - \lambda_0 + x \Lambda_1'(x)] f(x) dx &= 3\lambda_2^n \int_{x_3}^{x_4} (x - x_3)^2 f(x) dx + 2\lambda_2^n x_3 \int_{x_3}^{x_4} (x - x_3) f(x) dx, \\ \int_{A_2 \cup A_3 \cup A_4} \lambda_0 f(x) dx &= \lambda_0 (1 - \pi_1). \end{aligned}$$

The sum of the three integrals above and the integral in (32) is equal to the second sum on the r.h.s. of (30). Substituting this sum and (31) in (30):

$$\mathcal{F} = (D_p - D_m) + \pi_1 + \lambda_0 (1 - \pi_1) + 3B_2 + 3B_4 + C, \quad (33)$$

with

$$\begin{aligned}
B_2 &= \lambda_2^p \int_{x_1}^{x_2} (x - x_2)^2 f(x) dx, \\
B_4 &= \lambda_2^n \int_{x_3}^{x_4} (x - x_3)^2 f(x) dx, \\
C &= 2\lambda_2^n x_3 \int_{x_3}^{x_4} (x - x_3) f(x) dx + 2\lambda_2^p x_2 \int_{x_1}^{x_2} (x - x_2) f(x) dx,
\end{aligned}$$

A straightforward calculation shows that the fraction of firms adjusting is:

$$\mathcal{A} \equiv \int \Lambda(x) f(x) dx = \sum_{i=1}^5 \int \Lambda(x) f(x) dx = \pi_1 + \lambda_0(1 - \pi_1) + B_2 + B_4.$$

Using the above expression to get rid of $B_2 + B_4$ in (33) yields the desired result:

$$\mathcal{F} = \mathcal{A} + 2[\mathcal{A} - \pi_1 - \Lambda(0)(1 - \pi_1)] + (D_p - D_N) + C. \tag{34}$$

As mentioned in the main text, many of the terms in (28) are relevant only in limit cases such as strict Ss models. For example, consider the Caplin and Spulber setup. To obtain a discrete-time approximation, we consider $f(x)$ uniform on $(s - S - \mu, -\mu]$, since the relevant cross-section is the one after the aggregate shock (of size μ), immediately before adjustments take place. We also assume $\Lambda(x) = 0$ if $x \in (s - S, 0)$ and $\Lambda(x) = 1$ otherwise. To match the assumptions in the above proposition, we let $x_1 = x_2 = s - S$, $x_3 = x_4 = 0$, and $\lambda_2^p = \lambda_2^n = \lambda_0 = 0$. A few steps of algebra show that the terms on the r.h.s. of (28) take the following values: $\mathcal{A} = \mu/(S - s)$, $D_p = 1$, $D_n = 0$, $(1 - \pi_1)(1 - \Lambda(0)) = 1 - \mu/(S - s)$ and $C = 0$. It follows that:

$$\mathcal{F}^{\text{CS}} = 1 + \frac{\mu}{S - s}. \tag{35}$$

Of course, a uniform distribution on the inaction range is only an approximation to the ergodic density in discrete time. This explains why we obtain even more flexibility than in the continuous time case.

If we adapt (15) to discrete time, we have:

$$\mathcal{F}^\lambda = \lambda + (1 - \lambda) \frac{\alpha(S - s + \mu)}{e^{\alpha(S - s)} - 1}.$$

and we obtain (35) letting $\lambda \rightarrow 0$.