

NBER WORKING PAPER SERIES

A STATE-DEPENDENT MODEL OF INTERMEDIATE GOODS PRICING

Brent Neiman

Working Paper 16283

<http://www.nber.org/papers/w16283>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

August 2010

I thank Gita Gopinath, Elhanan Helpman, Ken Rogoff, Ariel Burstein, Steve Davis, Uli Doraszelski, Dan Frisch, Oleg Itskhoki, Virgiliu Midrigan, Emi Nakamura, Roberto Rigobon, Jon Steinsson, and Aleh Tsyvinski for very helpful comments and suggestions. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Brent Neiman. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A State-Dependent Model of Intermediate Goods Pricing
Brent Neiman
NBER Working Paper No. 16283
August 2010
JEL No. E3,F1,L11

ABSTRACT

Recent analyses of transaction-level datasets have generated new stylized facts on price setting and greatly influenced the empirical open- and closed-economy macroeconomics literatures. This work has uncovered marked heterogeneity in price stickiness, demonstrated that even non-zero price changes do not fully "pass through" exchange rate shocks, and offered evidence of synchronization in the timing of price changes. Further, intrafirm prices have been shown to differ from arm's length prices in each of these characteristics. This paper develops a state-dependent model of intermediate goods pricing, which allows for arm's length and intrafirm transactions, and is capable of generating these empirical pricing patterns.

Brent Neiman
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
brent.neiman@chicagobooth.edu

1 Introduction

Recent analyses of transaction-level datasets have generated new stylized facts on price setting and greatly influenced the empirical open- and closed-economy macroeconomics literatures. This work has uncovered marked heterogeneity in price stickiness, demonstrated that even non-zero price changes do not fully "pass through" exchange rate shocks, and offered evidence of synchronization in the timing of price changes. Further, intrafirm prices have been shown to differ from arm's length prices in each of these characteristics. This paper develops a state-dependent model of intermediate goods pricing, which allows for arm's length and intrafirm transactions, and is capable of generating these empirical pricing patterns.

Macroeconomists have long analyzed the impact of price adjustment costs using time-dependent pricing models in which firms cannot control the timing with which they change prices. These models, such as those in Taylor (1980) or Calvo (1983), offer the important advantage that they allow for analytical solutions. They generally cannot, however, match many of the patterns uncovered in the new micro datasets.

For example, there have been several studies of item-level price adjustment underlying the CPI and import price index, such as Bils and Klenow (2004), Nakamura and Steinsson (2008), and Gopinath and Rigobon (2008). These papers document significant cross-sectional differences in price stickiness and further, differences are often correlated with economic fundamentals. Items with a higher elasticity of demand change prices more frequently. Time-dependent models cannot endogenously generate this because they exogenously define the period, or average period, during which a price remains fixed. Cavallo (2009) and Midrigan (2006) find synchronization in the timing of changes in retail prices on the internet and in grocery stores, respectively. Similar to the case of heterogeneous stickiness, in the absence of a synchronized cost shock, time-dependent models have little scope for generating bunching in the timing of price changes.

State-dependent models allow firms to optimally decide when it is worth paying an adjustment cost in order to change prices. These models typically feature monopolistic competition, so strategic responses of firms need not be considered. For instance, with a continuum of competitors with infinitesimal market shares, knowledge of the aggregate price index is often a sufficient statistic summarizing the actions of other firms. As such, firms need not consider the

response of any given competitor when choosing to change its own prices. Like time-dependent setups, state-dependent models with monopolistic competition also cannot generally match these new empirical pricing patterns.^{1,2}

This is particularly important for work on international trade, as papers such as Gopinath et al. (forthcoming), Burstein and Jaimovich (2009), and Fitzgerald and Haller (2009) have all documented the degree to which exchange rate passthrough or pricing to market – equivalent for the purpose of this paper – is incomplete. This is not simply a reflection of nominal rigidity, but rather, holds true even after prices have changed. Most state-dependent models with monopolistic competition, however, use CES preferences that generate constant-markup pricing and counterfactually imply complete cost passthrough.³

Finally, Bernard et al. (2006), Hellerstein and Villas-Boas (forthcoming), and Neiman (2010) use micro-data to examine differences in these dynamic pricing characteristics between arm’s length and intrafirm trade transactions. All three papers find that intrafirm prices exhibit higher exchange rate passthrough, and Neiman (2010) additionally finds that they exhibit less stickiness and synchronization.⁴ Though a state-dependent model without monopolistic competition is likely required to generate these facts, additional structure is also needed to model the difference between arm’s length and intrafirm trade.

Below, I consider a two-firm game in a partial equilibrium environment with trade in intermediate inputs. Upstream firms sell the inputs downstream to either an unrelated party or a wholly owned subsidiary. Each manufacturer’s pricing strategy is a function of the other firm’s pricing strategy. The model is capable of delivering all of the empirical facts described above: (1a) Arm’s length price duration is heterogenous, decreases as goods become less differentiated, and (1b) is smaller for intrafirm prices; (2a) Passthrough is incomplete even after prices change and (2b) higher for intrafirm trades; (3a) Price changes exhibit synchronization, but (3b) less so for intrafirm prices. The empirical motivation for the paper comes largely from papers analyzing international trade data, where it is easiest to measure

¹For example, Midrigan (2006) has to introduce increasing returns in the production function of price changes in order to generate synchronization.

²See Klenow and Kryvtsov (2008) for additional empirical features which present difficulties for standard time- and state-dependent models.

³The use of Kimball (1995) preferences in Gopinath and Itskhoki (forthcoming) or the use of translog preferences in Bergin and Feenstra (2000) are exceptions.

⁴The results in Neiman (2010) are for the set of differentiated traded goods and exclude, for example, commodities.

cost shocks (since the exchange rate acts as an observable cost shifter) and observe intrafirm transactions. I note, however, that the mechanisms in the model and importance of the dynamics are just as salient for domestic transactions.

I start with the unrelated party, or arm's length, case. These firms face idiosyncratic production cost shocks and decide whether to keep their existing price or pay an adjustment cost to change it. The degree to which the cost shock renders the firm's current price suboptimal depends on the elasticity of demand. This generates heterogeneity in stickiness. Further, the inputs are substitutes, so a price change by the competitor firm will also change the profitability of an existing price and may induce a response. This generates both incomplete cost passthrough and synchronization in the timing of price changes.

The dynamics are somewhat different when the downstream input purchaser is a related party. In this case, the upstream firm attempts to avoid double marginalization and sets trade prices to approximately follow marginal cost. Accordingly, intrafirm price setting is primarily inward looking and responds less to competitors' prices, which leads to less price synchronization and greater passthrough of marginal cost shocks. Further, all other things equal, price duration (a measure of stickiness) is positively related to a firm's market share, and negatively to the cost of goods sold and the constancy of its markups. Conditional on market share, the related party cost of goods sold will be higher (and conditional on the cost of goods sold, related party market share will be smaller). Related party markups are also less variable. On average, this will result in shorter related party price duration.

In sum, many of the new facts on import, export, producer, and retail prices suggest the need for a dynamic model of price adjustment with at least four features: firms with positive market shares, price adjustment costs, different vertical structures, and state-dependent pricing. I now describe a partial equilibrium model with these features that is capable of producing the salient facts on arm's length price setting – and the comparison along these dimensions with intrafirm price setting – found across a large set of empirical studies.

2 A Partial Equilibrium Model of Trade in Intermediate Goods

The model is a nested constant elasticity of substitution (CES) structure closely related to that used in Yang (1997) and more recently in Atkeson and Burstein (2008). An infinitely lived

representative consumer buys a continuum of final goods that are assembled by distributors from two inputs. Under one structure studied, both inputs are produced at arm's length. I also consider a structure in which one input is produced at arm's length and the other is produced by a wholly-owned subsidiary.

The cost of production for these inputs varies over time due to idiosyncratic cost shocks. Distributor pricing is completely flexible, while manufacturers must pay a fixed adjustment cost to change their prices.⁵ Consumers maximize their lifetime expected utility and arm's length manufacturers maximize their lifetime expected profits. Integrated firms maximize the lifetime combined profits of their manufacturing plants and distributors.

I consider an exogenously imposed vertical ownership structure for any given industry. Due to double marginalization, there is always an incentive in the model to vertically integrate. Hence, to generate the existence of both structures, I am implicitly assuming a firm-specific integration cost that varies randomly, even conditional on all the variables in the model. If simulated in a multi-firm setting, this would produce a combination of integrated and arm's length firms in most sectors. The starkest example of such a cost in the international setting would be a ban on foreign ownership in certain countries and sectors.⁶ To the extent integration costs are fixed in nature, the firm-specific reason (or lack thereof) for integration should not impact pricing.

2.1 Consumers

Consumers maximize their expected lifetime utility of consumption, a discounted consumption stream at times t , $E_t \sum_{t=0}^{\infty} \beta^t U(C^t)$, where they exhibit a CES love of variety $C^t = \left[\int_0^1 c^t(z)^{\frac{\eta-1}{\eta}} dz \right]^{\frac{\eta}{\eta-1}}$ over a continuum of final goods c that are indexed by $z \in [0, 1]$. As is standard in this setup, consumer demand for good $c(z)$ is $c^t(z) = C^t (p^t(z))^{-\eta} (P^t)^{\eta}$, where the price index is defined as: $P^t = \left[\int_0^1 p^t(z)^{1-\eta} dz \right]^{\frac{1}{1-\eta}}$.

⁵The assumption of greater flexibility in downstream prices is supported in the data. See, for instance, Shoenele (2009) and Gopinath and Rigobon (2008).

⁶For example, foreign direct investment, a necessary precursor to intrafirm trade, was highly restricted in China and India during most of the dataset.

2.2 Distributors

There is a continuum of distributors that costlessly assemble each final good using a CES production technology that combines two product-specific manufactured intermediate inputs:

$$c^t(z) = \left[\gamma(z) c_1^t(z)^{\frac{\rho(z)-1}{\rho(z)}} + (1 - \gamma(z)) c_2^t(z)^{\frac{\rho(z)-1}{\rho(z)}} \right]^{\frac{\rho(z)}{\rho(z)-1}},$$

where $\eta < \rho(z) < \infty$ and $\gamma(z) \in (0, 1)$ for all z . Sectors with higher values of ρ are less differentiated as the distributor can more easily substitute away from any given input in those sectors. Distributors take input prices as given and solve the problem:

$$\max_{p^t(z)} p^t(z) c^t(z) - p_1^t(z) c_1^t(z) - p_2^t(z) c_2^t(z), \quad (1)$$

which results in demand for the first manufactured input (expression for the second input, not shown, is symmetric) of:

$$c_1^t(z) = c^t(z) (p_1^t(z))^{-\rho(z)} (\gamma(z) x^t(z))^{\rho(z)},$$

where

$$x^t(z) = \left[\gamma(z)^{\rho(z)} p_1^t(z)^{1-\rho(z)} + (1 - \gamma(z))^{\rho(z)} p_2^t(z)^{1-\rho(z)} \right]^{\frac{1}{1-\rho(z)}}$$

is the total unit production cost of the final good. Distributors then set price at a constant markup over this marginal cost of production, $p^t(z) = (\eta/(\eta - 1)) x^t(z)$.

2.3 Manufacturers

Intermediate good manufacturers use a linear technology to produce $c_j^t(z)$ at a constant marginal cost for each firm j , which I write in logs for convenience of notation as: $\ln[m_j^t(z)] = \text{constant} + \lambda e_j^t(z)$. $e_j^t(z)$ shifts the marginal costs of firm j supplying inputs for final good z at time t . In an open-economy setting, it can alternatively be thought of as an exchange rate. In a closed-economy setting, it can be thought of as an idiosyncratic productivity term. In describing the model, I will focus on this the case where the two firms' shocks are uncorrelated. I do this to emphasize that the model generates synchronized price changes even when cost

shocks are not synchronized, an assumption that is perhaps most natural in the international setting. The framework, however, can handle any correlation structure.

A share of the total production costs, $(1 - \lambda)$, is impacted by this shock. This captures the case when productivity gains only impact certain production processes, or in the open-economy case, when the exchange rate does not fully impact the unit cost because the manufacturer itself imports intermediate inputs from abroad. Though I do not focus on any particular set of quantitative estimates, I introduce λ to come closer to matching the highly incomplete rates of exchange rate passthrough seen in the empirical literature. I model the shock process as an AR(1):

$$e_j^t(z) = \delta e_j^{t-1}(z) + \mu_j^t(z),$$

where $\mu_j^t(z)$ is normally distributed with cumulative distribution function $F_j(\mu_j^t(z))$. This allows for shocks that are strongly mean-reverting as well as those arbitrarily close to fully persistent (as $\delta \rightarrow 1$).

2.4 Price Setting: Arm's Length Trades

I start by considering the case in which both input suppliers are unrelated to their customer and trade at arm's length. Unlike the distributors, the manufacturers pay a fixed cost to change their nominal prices. These trades are business-to-business transactions, and hence, this fixed cost is more typically thought to reflect the cost of changing processes, communicating, and negotiating with customers than the retail price interpretation as "menu" costs (See Zbaracki et al., 2004).

Each period, the manufacturer that provides the first input (the setup is symmetric, so we focus on this manufacturer without loss of generality) earns operating profits $\pi_1 = p_1 c_1 - m_1 c_1$, which are defined to exclude the cost of price adjustment. For notational convenience, I drop the sector and time indices, z and t , when they are not needed, and re-write operating profits:

$$\pi_1 = CP^\eta \left(\frac{\eta}{\eta - 1} \right)^{-\eta} [\gamma^\rho p_1^{1-\rho} + (1 - \gamma)^\rho p_2^{1-\rho}]^{\frac{\rho - \eta}{1 - \rho}} \gamma^\rho p_1^{-\rho} (p_1 - m_1).$$

Arm's length manufacturers maximize the present value of real profits, less real adjustment

costs ϕ_{AL}/P^t , by solving:

$$\max_{p_1(s^t)} E_t \sum_{t=0}^{\infty} \beta^t \left[\frac{\pi_1}{P^t} - \frac{\phi_{AL}}{P^t} \chi_1^t \right]. \quad (2)$$

χ_j^t is an indicator function equalling 1 when $p_j^t \neq p_j^{t-1}$ and 0 otherwise.

2.5 Price Setting: Intrafirm Trades

We can also use the model to consider the case in which one product is assembled from a related party, which sells its input to a wholly owned subsidiary (or parent). We do not consider the case in which both firms are related parties as this would render the setup, in which manufacturers do not coordinate price-setting with each other, unrealistic. Distributors that purchase from a related party also purchase from arm's length suppliers, a feature that is supported in the data.⁷

Vertically integrated firms aim to maximize overall profits – the sum of its profits at the manufacturer and distributor levels – as follows. The manufacturing firm (or a separate headquarters division) instructs the distributor to take input prices as given, and to purchase from the arm's length or related party manufacturer in whatever way maximizes distributor profits. This should not be interpreted as if the distributor is naive of the ownership structure or acts myopically, but rather, is simply following the pricing mechanism designed by the integrated firm. As part of this mechanism, the manufacturer knows how the distributors will act and thus chooses prices in order to maximize the expected present value of all future integrated profits, after subtracting price adjustment costs.⁸ Anecdotal evidence suggests that the essence of this pricing mechanism is used by actual companies.⁹

Without loss of generality, I assume the related party manufacturer in this case supplies

⁷Bernard et al. (2007) shows that the vast majority of firms that import from related parties also do so from arm's length suppliers.

⁸I assume the manufacturer would incur the adjustment cost if it communicated its desired price to the distributor, even if it does not actually change its price to that level. This is consistent with the above interpretation of an adjustment cost and rules out a potentially more profitable mechanism whereby the distributor changes its price even if its suppliers do not change theirs.

⁹The managing director of a consulting firm specializing in transfer pricing told me of the case of a large multinational company that evaluates upstream manufacturing managers on their ability to minimize production costs, without any link to the upstream unit's profits (the company delegates the determination of transfer prices, but not retail prices, to a separate group that aims to maximize overall firm profitability). The consultant described another integrated relationship in which the downstream unit, by design, made purchases without even knowing which suppliers were related parties and which were arm's length firms. Both anecdotes support the idea in the model that transfer prices may be both allocative and designed to maximize the sum of upstream and downstream profits.

the second manufactured input. The integrated firm's operating profits are:

$$\pi_2 = [\pi_2^{Distributor}] + [\pi_2^{Manufacturer}] = pc - p_1c_1 - m_2c_2, \quad (3)$$

and can be re-written as:

$$\begin{aligned} \pi_2 = & CP^\eta \left(\frac{\eta}{\eta - 1} \right)^{1-\eta} [\gamma^\rho p_1^{1-\rho} + (1 - \gamma)^\rho p_2^{1-\rho}]^{\frac{1-\eta}{1-\rho}} \\ & - CP^\eta \left(\frac{\eta}{\eta - 1} \right)^{-\eta} [\gamma^\rho p_1^{1-\rho} + (1 - \gamma)^\rho p_2^{1-\rho}]^{\frac{\rho-\eta}{1-\rho}} [\gamma^\rho p_1^{1-\rho} + (1 - \gamma)^\rho p_2^{-\rho} m_2]. \end{aligned}$$

Vertically integrated firms maximize the present value of real profits, less real adjustment costs ϕ_{RP}/P^t , and maximize an expression equivalent to (2). The related party manufacturer pays an adjustment cost because coordination, communication, and process changes between business units of the same firm are also costly.

3 Determinants of Pricing Patterns

In this section, before proceeding to the full dynamic model's solution and simulation, I try to build intuition for the model's ability to match characteristics in the data. We start with the case without nominal rigidities ($\phi_{AL} = \phi_{RP} = 0$). In this flexible price setting, I focus on synchronization and passthrough. I next add an adjustment cost, take an approximation of the firm's profit function, and run some simple one-period numerical examples. These are designed to generate intuition for the determinants of price duration. These exercises suggest the model will produce the patterns on duration, passthrough, and synchronization found in the data.

3.1 Flexible Price Passthrough and Synchronization: Arm's Length Prices

The arm's length firm's optimal price is set by taking its competitor's price as given and pricing at a variable markup over marginal cost, $p_j = \frac{\varepsilon(s_j)}{\varepsilon(s_j)-1} m_j$. The market share of arm's length input manufacturer j in that sector, s_j , can be expressed as:

$$s_j = \frac{p_j c_j}{x c} = \left(1 + (\gamma_j / \gamma_{-j})^{-\rho} (p_j / p_{-j})^{\rho-1} \right)^{-1},$$

and the elasticity of demand, ε_j , is the market-share weighted average of the elasticities of substitution for final goods and for the sector's intermediate inputs:

$$\varepsilon_j(s_j) = \eta s_j + \rho(1 - s_j).$$

The optimal price depends on both the firm's own cost and, through its impact on market share, the competitor's price. This strategic complementarity is often assumed away in setups with monopolistic competition. Arm's length markups decrease with the elasticity of demand, and firms with given market shares will charge lower markups for more substitutable goods.

Totally differentiating the markup, elasticity, and market share definitions, I approximate the change in arm's length price as a weighted average of the shocks to a firm's own cost and its competitor's price:

$$\widehat{p}_j = \alpha_j \widehat{m}_j + (1 - \alpha_j) \widehat{p}_{-j} = \lambda \alpha_j \mu_j + (1 - \alpha_j) \widehat{p}_{-j}, \quad (4)$$

where:

$$\alpha_j = \frac{\varepsilon_j(\varepsilon_j - 1)}{\varepsilon_j(\varepsilon_j - 1) + (\rho - \eta)(\rho - 1)s_j(1 - s_j)}, \quad (5)$$

and where $\widehat{x} = dx/x$ denotes the size (in percent) of a change in a variable x . Expression (4) measures the responsiveness of the arm's length flexible price to a percentage change in marginal cost or to its competitor's price. It assumes that the competitor does not subsequently change its price further. $\lambda \alpha_j$ approximates passthrough of the cost shock for arm's length firms. Given that $\lambda, \alpha_j \in (0, 1)$, arm's length passthrough will be incomplete, even after a price changes, consistent with the data. As the elasticity of demand ε_j changes with market share, the markup $\varepsilon_j/(\varepsilon_j - 1)$ will change, and a varying amount of the cost shock will be absorbed, rather than passed through.

Now, we consider the case where each firm does respond to any change in the other firm's price. Substituting $\widehat{p}_{-j} = \alpha_{-j} \widehat{m}_{-j} + (1 - \alpha_{-j}) \widehat{p}_j$ into (4), I write :

$$\widehat{p}_j = \beta_j \widehat{m}_j + (1 - \beta_j) \widehat{m}_{-j},$$

where j and $-j$ are arm's length firms and:

$$\beta_j = \frac{\alpha_j}{\alpha_j + \alpha_{-j} - \alpha_j \alpha_{-j}} \in (0, 1)$$

is now the equivalent expression to (5). $\lambda\beta_j$ is now an approximation to cost passthrough. Note that $\beta_j > \alpha_j$, implying that an arm's length firm with a given market share will have higher passthrough when competing against a more responsive firm than otherwise. Here, we see analytically that this implication of strategic complementarity, which is often abstracted from in state-dependent models, can be important, particularly at low levels of stickiness.

3.2 Flexible Price Passthrough and Synchronization: Intrafirm Prices

In this setting with zero adjustment costs, pricing for related parties is simple. Comparing the expression for distributor per-period profits in equation (1) with that for the integrated firm in equation (3), it is clear that in order for the solution to the distributor's problem to always equal the solution to the integrated firm's problem, related party manufacturers should charge their marginal cost: $p_j = m_j$ if $j = RP$. As discussed in Hershleifer (1956), the transfer occurs at marginal cost because the firm wants to use inputs as efficiently as possible in generating the final good, since the final good consumer is the only real customer. Above, heterogeneous good arm's length firms were shown to charge higher markups than homogenous good arm's length firms. Combined with marginal cost transfer pricing, this implies that intrafirm prices of equivalent goods will be lower than arm's length prices, and the difference should be larger for heterogeneous goods. This is precisely the result found empirically in Bernard et al. (2006).

With no adjustment costs, related parties will fully pass through the portion of the shock μ_j that changes its unit cost. In particular:

$$\widehat{p}_j = \widehat{m}_j = \lambda de_j \approx \lambda \mu_j \text{ if } j = RP, \quad (6)$$

where the approximation becomes an equality as $\delta \rightarrow 1$. Hence, intrafirm passthrough equals λ . Incomplete passthrough in the related party flexible price case is entirely due to the existence of some share of the marginal cost being unaffected by the cost shocks.

In this sense, the related party manufacturer is less concerned with the arm's length

firm and is focused entirely inward, on its own marginal cost. In the dynamic model with adjustment costs, related parties will not strictly price at marginal cost because the firm must weigh whether it prefers to be slightly above or below its ideal flexible price in future periods where a price change is not warranted. It will remain true in the model, however, that related party passthrough is very close to λ . This implies, consistent with the empirical results, that intrafirm passthrough will be higher than arm's length passthrough. Further, the competitor firm's price is absent from the pricing equation (6), so this model will generate less intrafirm synchronization, also consistent with the data.

3.3 Static One-Period Game: Arm's Length Case

I now return to the environment with positive adjustment costs and consider the model's ability to match the empirical findings that price duration is larger for more heterogenous products and that prices change with significant synchronization. This model will be able to generate both of these comparative statics.

As seen in equation (4), there are two shocks that could lead a firm to change its price – a shock to its own production cost and a change in its competitor's price – and a host of conditions and parameters, such as the market share and the size of the adjustment cost, that influence this decision. To build intuition, I start by considering a one-period game where there is no price response from competitors and firms start in their flexible price equilibrium, with initial profits denoted by $\pi_j^+ = \pi_j^+(m^+, p_{-j}^+)$. From this point, if firm j foregoes price adjustment in the face of higher production costs, there is no change in revenue or demand, and the firm's profits will decline by exactly this cost change times the number of units:

$$d\pi_j^N = \pi_j(m_j^+ + dm_j, p_{-j}^+) - \pi_j^+ = -c_j dm_j = -c_j m_j \widehat{m}_j,$$

where the superscript "N" stands for "non-adjustment." This expression holds equally for both related parties and the arm's length firms. To consider the change in profits that would occur under adjustment (represented with "A") to this shock, I write the second-order approximation around the flexible price equilibrium just prior to a cost shock:

$$d\pi_j^A = \pi_j(m_j^+ + dm_j, p_{-j}^+) - \pi_j^+ \approx \frac{\partial \pi_j^+}{\partial m_j} dm_j + \frac{1}{2} \frac{\partial^2 \pi_j^+}{\partial m_j^2} dm_j^2,$$

where the expressions will differ for related parties and arm's length firms. The overall incentive to change prices, an object that implies shorter price durations as it gets bigger, is approximated as the difference between the two: $d\pi_j^A - d\pi_j^N$.

I show in Appendix A that $\partial\pi_j^+/\partial m_j = -c_j$ for both types of firms, and hence the first order terms for the change in profit with and without adjustment cancel. As a result, the approximated second-order incentive is only the second-order term $\frac{1}{2}\frac{\partial^2\pi_j^A}{\partial m_j^2}dm_j^2 = \frac{1}{2}\Omega_j\widehat{m}_j^2$. The arm's length markup structure leads to an expression for the adjustment incentive:

$$\Omega_{AL,j} = (\varepsilon_j - 1) s_j \alpha_j cx, \quad (7)$$

where cx denotes total spending on the sector's inputs. After fixing manufacturer revenues, $\Omega_{AL,j}$ can be written as the product of $(\varepsilon_j - 1)$ and α_j , as is the focus of Gopinath and Itskhoki (forthcoming), which first derived such an expression for the arm's length case in a similar model with monopolistic competition.

Unfortunately, it cannot be shown analytically that in the arm's length case, more heterogeneous good prices will always be stickier because $d\Omega_{AL,j}/d\rho$ cannot be unambiguously signed. Hence, to get a sense for the comparative static of duration with respect to degree of heterogeneity, I consider the following numerical exercise, plotted in Figure 1. I set initial productivity levels equal, $m_{AL,j} = m_{AL,-j}$, and pick a uniform value for the adjustment cost ϕ . Under this configuration, each firm starts with equal market share. Starting from equilibrium in the flexible price model (noted with the black plus sign), a firm observes its own cost shock and its competitor's price change and determines if adjustment merits payment of the fixed cost.

The left plot is drawn from the perspective of an arm's length manufacturer in a highly differentiated sector, where shocks to its competitor's price and its own cost are represented with the horizontal and vertical axes, respectively. The right plot is the exact same, but for an arm's length manufacturer in a less differentiated sector (with higher ρ). The red regions are then defined as the portions of the state space where a firm does not adjust prices and the boundaries can be thought of as s-S bands.

The scenario where a firm's production cost increases by 5 percent and the competitor raises prices by 10 percent is represented by a move upward from the black plus sign by 0.05

and to the right by 0.10. If such a move does not exit the red region, it means that given these shocks, a firm would not change its price. If such a move crosses the upper boundary into the "raise" region, it means the shocks are sufficiently large to warrant a price increase, even if facing an adjustment cost.

The first key observation is that the no-adjust region for both arm's length firms has negative slope. If a change in the other firm's price is large enough, it can induce the first firm to change prices, even if the first firm does not incur a shock to its marginal cost. This is the visual manifestation of the strategic complementarity in the model and is the force generating synchronization in the timing of price changes. Secondly, the vertical width of the band is wider for the more differentiated, or heterogenous, case. Given the degree of stickiness in the data, own-cost shocks are far more prevalent than competitor-price shocks and hence, the vertical width is the crucial determinant of stickiness. It is clear that any given cost shock is more likely to exit the red region, up or down, for the less differentiated good arm's length firm. Though one can find places in the parameter space where these results do not hold, they are far away from the most natural benchmarks such as symmetry and generally require significantly skewed productivity distributions in the sector.

3.4 Static One-Period Game: Intrafirm Case

I now consider the same exercise but fixing the sector's elasticity of substitution and instead focusing on the comparison of arm's length and related party trades. I show that the model can match the empirical results showing that intrafirm prices change more frequently and with less synchronization. The related party pricing structure leads to an expressions for the adjustment incentive $\Omega_{RP,j}$:

$$\Omega_{RP,j} = \varepsilon_j s_j cx, \quad (8)$$

where cx denotes total spending on the sector's inputs.

The difference between the related party expression in (8) and the arm's length expression in (7) reflects the fact that a firm's cost of goods sold, $COGS_j = c_j m_j$, scales each firm's incentive to change prices for a given percentage cost shock. Since arm's length firms charge a markup and related parties do not, the cost of goods sold is related differently to market shares and elasticities for the two firms. Substituting $COGS_{RP,j} = s_{RP,j} cx$ and $COGS_{AL,j} =$

$s_{AL,j}cx^{\frac{\varepsilon_j-1}{\varepsilon_j}}$ into expressions (8) and (7), we can write the incentives as $\Omega_{AL,j} = \varepsilon_j\alpha_jCOGS_{AL,j}$ and $\Omega_{RP,j} = \varepsilon_jCOGS_{RP,j}$. This gives the intuition for why related party duration will be shorter, conditional on the market share, and all other things equal. The market share uniquely determines the demand elasticity ε_j , and given the related party charges no markup, its cost of goods sold must be higher. The variable markup component of passthrough, α_j , is strictly less than one, so $\Omega_{RP,j} > \Omega_{AL,j}$.

In Appendix A, I demonstrate for the two-firm case that $s_{AL,j} < \eta/(2\eta - 1) = \overline{s_{AL}}$ is a sufficient, though not necessary, condition for $\Omega_{RP} > \Omega_{AL}$. Note that as $\eta \rightarrow 1$, $\overline{s_{AL}} \rightarrow 1$, and there is no portion of the parameter space where the approximation suggests stickier related parties, regardless of initial productivities. In the model's other extreme, as $\eta \rightarrow \infty$, $\overline{s_{AL}} \rightarrow 1/2$. Given arm's length markups exceed those of related parties, this implies that with equal productivities, related parties are less sticky everywhere in the parameter space. Numerical exercises suggest that for any given η , an increase in ρ increases the maximum arm's length market share below which its prices will be stickier. For plausible parameter values in this model, the threshold is at least two-thirds, and often much higher. This absolute level will of course decrease in a multifirm model, but the requirement that arm's length firms hold a significantly larger market share in order to be less sticky will generally hold, regardless of the number of firms. Hence, this static model generally predicts less sticky intrafirm prices.

Further, Figure 2 shows s-S bands similar to those shown for the arm's length case, but instead of comparing across elasticities of substitution, it compares the pricing decision of an arm's length firm (left) to that of a related party (right). Again, I set initial productivity levels equal, $m_{AL} = m_{RP}$, and pick a uniform value for the adjustment cost ϕ . This implies market shares will differ, but plots from the case of equal market shares are qualitatively the same.

First, note that the no-adjust region for the related party is essentially flat. This means that, when integrated firm prices are close enough to their flexible price target, there is no price change from the competitor (arm's length) firm that could induce a firm to change its own price. Only as one moves vertically away from the horizontal line $p_{RP} = 1$ does the region begin to have any curvature. This follows because the result that related party price setting is inwardly focused is only strictly true when at the flexible price equilibrium. In this sense, Figure 2 helps one visualize why the model is able to produce greater synchronization among

arm's length trades than intrafirm trades. Secondly, the vertical width of the band is smaller for the related party case, indicating less price stickiness and corroborating the results from the second order approximation.

4 Recursive Formulation and Solution

Some of the above intuitions on what drives incomplete exchange rate passthrough and synchronization in arm's length firms have been shown previously in models with flexible prices, such as Atkeson and Burstein (2008). Adding price stickiness is important to ensure that these intuitions from the flexible price literature are preserved when realistic frictions are added and when simulated data better resembles the highly sticky price data actually tested by empiricists. Further, the previous sections' derivations rely on several simplifying assumptions or approximations, abstract from option value, and consider the occurrence of each shock and each firm's pricing decision only one at a time. In reality, firms have expectations about each other's responses to shocks and typically start periods away from their flexible price equilibrium. In this section, we move to a dynamic setting in order to address these shortcomings.

The monetary authority maintains a constant retail price level, $P^t = 1$, and thus fixes aggregate consumption, $C^t = 1$. This leaves four principal state variables in the system – the two manufacturing prices from the previous period and the two marginal costs in the current period. I bundle these four dimensions of the state space as $\Theta^t = \{p_1^{t-1}, p_2^{t-1}, m_1^t, m_2^t\}$. Most dynamics are generated by the fully observable shocks to the marginal cost of production for each firm.

The other source of dynamics follows from the random adjustment cost, ϕ_j^t , drawn identically and independently each period from the distributions $G_j(\cdot)$. This follows Dotsey, King, and Wolman (1999) and renders the problem more tractable. Though firms know the distribution of their competitor's adjustment cost, they only observe their own realized cost. This assumption of random and private adjustment costs is helpful because it rules out certain cases in which there would be multiple (or no) equilibria. It allows for a game in pure strategies, but where each player treats the other as if she were playing a mixed strategy due to uncertainty about the other's state. As implemented, this assumption does not impact any of the qualitative results.

Firms follow pure strategies in price setting. For a given state $\{\Theta^t, \phi_j^t\}$, each firm j simultaneously chooses a unique price. As emphasized in Doraszelski and Satterthwaite (forthcoming), due to the uncertainty about the competitor's adjustment cost, a firm generally does not know with certainty what strategy its competitor will play. Hence, from the perspective of firm $-j$, the probability that firm j changes prices in a given period is $\xi_j(\Theta) = \int \chi_j(\Theta, \phi_j) dG_j(\phi_j)$. A Markov Perfect Equilibrium is defined as a set of pricing policies for each firm j , $p_j^t = p_j(\Theta^t, \phi_j^t)$, where p_j^t maximizes expected firm profits, consistent with consumer demand, and where each firm has correct expectations about the distribution of its competitor's prices across realizations of the competitor's adjustment cost.

Let $V_j(\Theta, \phi_j)$ denote the conditional values of the firm, after each has observed its own price adjustment cost. I define these value functions recursively as:

$$V_j(\Theta^t, \phi_j^t) = \max_{\tilde{p}_j} \xi_{-j} \pi_j(\tilde{p}_j, \tilde{p}_{-j}) + (1 - \xi_{-j}) \pi_j(\tilde{p}_j, p_{-j}^{t-1}) - \chi_j(\Theta^t, \phi_j^t) \phi_j + \beta \int \int_{u_j u_{-j}} V_j(\Theta^{t+1}) dF_j dF_{-j}, \quad (9)$$

for firm j . Here, it is easy to see the difficulty in modeling this type of strategic behavior – it requires solving a coupled system of Belman equations where each firm's optimal policy depends on the other's. $V_j(\Theta') = \int V_j(\Theta', \phi_j) dG_j(\phi_j)$ is the expected value function of firm j , conditional on being in state Θ' , but before observing its adjustment cost (expectations here are taken only over uncertainty about the realization of this cost).

Following Doraszelski and Satterthwaite (forthcoming), I integrate both sides of these Belman equations over all realizations of their respective adjustment costs and re-write the value function in equation (9), which is a function of five variables, as the expected value function, which is no longer a function of the adjustment cost:

$$V_j(\Theta^t) = \max_{\substack{\xi_j \in [0,1] \\ \tilde{p}_j \in (0,\infty)}} E[\pi_j] - \int_{\phi_j < G_j^{-1}(\xi_j(\Theta^t))} \phi_j dG_j(\phi_j) + \beta \int \int_{u_j u_{-j}} E[V_j(\Theta^{t+1})] dF_j dF_{-j}. \quad (10)$$

Expected profit, $E[\pi_j]$, is the probability weighted average across the four combinations of $\{\text{adjust, no-adjust}\} \times \{\text{adjust, no-adjust}\}$, and the transition of the first two state variables is

similarly defined in the expected continuation value. Formally:

$$E[\pi_j] = \sum_{\alpha_1=0}^1 \sum_{\alpha_2=0}^1 \xi_1^{\alpha_1} (1 - \xi_1)^{1-\alpha_1} \xi_2^{\alpha_2} (1 - \xi_2)^{1-\alpha_2} \pi_j \left((\tilde{p}_1)^{\alpha_1} (p_1^{t-1})^{1-\alpha_1}, (\tilde{p}_2)^{\alpha_2} (p_2^{t-1})^{1-\alpha_2}, m_j \right),$$

and:

$$E[V_j(\Theta^{t+1})] = \sum_{\alpha_1=0}^1 \sum_{\alpha_2=0}^1 \xi_1^{\alpha_1} (1 - \xi_1)^{1-\alpha_1} \xi_2^{\alpha_2} (1 - \xi_2)^{1-\alpha_2} V_j \left((\tilde{p}_1)^{\alpha_1} (p_1^{t-1})^{1-\alpha_1}, (\tilde{p}_2)^{\alpha_2} (p_2^{t-1})^{1-\alpha_2}, \cdot, \cdot \right).$$

Subject to the above system of demand, production, and cost shocks, the two firms play a non-cooperative dynamic game in pure Markov pricing strategies. I follow Midrigan (2006, forthcoming) and Miranda and Vedenov (2001) and use projection methods (collocation, specifically) to approximate the solution to this coupled system of Belman equations.¹⁰ A detailed description of the solution algorithm is given in Appendix B. Figure 3 shows a sample plot (holding fixed the values for the competitor's previous price and current cost) of a policy function from the solution of the model. The vertical axis gives the conditional probability of a price change before observing the menu cost realization and the x- and y-axes give the firm's previous price and current cost. This plot makes clear that, despite the time-dependency added by the stochastic menu cost, the model preserves its state-dependent flavor. The probability of a price change fluctuates dramatically across the state space, even if it transitions more smoothly than the zero to one fluctuations in a standard state-dependent model.

5 Simulation Results

To compare the model's predictions for price duration, passthrough, and synchronization with the data, I take the approximated policy functions and generate series of costs and prices for various ranges of the parameter space. The two-input structure of my model rules out treatment of the simulation as a true calibration exercise – I do not focus on comparing precise quantitative levels of variables in the simulation to the data, but rather, show that I can match the key features of the data on arm's length and intrafirm duration, passthrough,

¹⁰The techniques used are described in-depth in Miranda and Fackler (2002), which also provides an accompanying MATLAB toolbox (CompEcon) that was used extensively for this paper.

and synchronization.

I solve two versions of the model. First, I simulate a sector with two arm's length firms to test if the model generates price durations that increase with heterogeneity, incomplete passthrough even after a price change, and greater than random synchronization in price changes. Next, I simulate a sector with one arm's length and one related party firm to test if the model generates less intrafirm stickiness, higher passthrough, and lower synchronization. I simulate these two structures for four sectors with varying elasticities of substitution, ρ . The period length is intended to represent one month and the discount factor is set at $\beta = 0.99$. I set a normal distribution for the monthly shock process, μ , with a standard deviation of 2.5 percent for both manufacturers, roughly that of the U.S. dollar to Euro exchange rate. Identical uniform distributions (with limited support) are used for each firm's adjustment costs such that the median duration magnitude roughly fits the level of stickiness in the international trade micro-data and results in spending on adjustment as a share of annual manufacturing revenues of about 0.2 percent.

I consider three cases: In the first, I set the firms' steady state market shares equal ($s_j = s_{-j}$); in the second, I set productivities equal ($m_j = m_{-j}$); and in the third, I set the firms' steady state cost of goods sold to be equal ($c_j m_j = c_{-j} m_{-j}$). For the case with related parties, these scenarios imply the related party's market share will be equal, larger, and smaller, respectively, than that of the arm's length firm. When simulating the sector with two arm's length firms, I choose productivities such that the first has both equal demand and market share compared to the arm's length firm in the hybrid sector.¹¹ As shown in Section 3, price stickiness is proportional to the spending on a sector, so for each set of parameter values, I vary aggregate consumption to equalize steady state spending on the manufacturing sector.

I set $\lambda = 0.75$, which of course will scale down passthrough levels for both firms.¹² This and other parameter values are summarized below:

¹¹This, of course, implies that the second firm, whose statistics are not reported, does not have the same productivity as the related party in the hybrid case.

¹²In the open-economy interpretation of the model, this parameter is consistent with a typical import share from OECD input-output tables.

η	β	γ	λ	σ	δ	ρ_{Min}	ρ_{Max}
2	0.99	0.5	0.75	0.025	0.985	4	10

Figure 4 shows an example of the price and cost series generated by the simulation program. The prices and costs are plotted against the left axis, while the probability of adjustment ξ_j is indicated by the shaded bars and is measured on the right axis. In the start of year 5, arm's length firm j increases its price even though its own cost has clearly been declining. This is labeled an "example of complementarity" because the price increase is clearly driven by the (correct) expectation that the other firm, its competitor, would increase its own price. Again, this feature is typically excluded from state-dependent models. Figures 5-8 report the results from these simulations, including median price durations, cost passthrough conditional on a price change, and a measure of synchronization. I now discuss these simulations and compare each in turn to the empirical facts.

5.1 Duration: Empirics and Simulation

The empirical results suggest that more heterogenous arm's length prices change less frequently. For example, Gopinath and Rigobon (2008) show in their Table IV that the mean frequency of price change for reference priced (i.e. homogenous) goods is more than twice that of differentiated goods. "Raw goods", a highly substitutable category, is the least sticky in Bils and Klenow (2004) while Medical care, presumably highly differentiated, is the most sticky. Nakamura and Steinsson (2008) show that less differentiated goods like "unprocessed food", "vehicle fuel", or "transportation goods" change prices far more often than more differentiated products like "processed food" or "services". Figure 5 plots the median arms length duration in the three market share configurations against the elasticity of substitution (ρ) and shows that in this model, for sensible parameter values, price duration or stickiness decreases as goods become less differentiated.

Further, Neiman (2010) shows that related party prices are stickier than arm's length prices in the same sector. Figure 6 plots the median for arm's length and intrafirm prices in sectors with both types of firms and for the three market share configurations and four values of substitution elasticities. The solid lines in the figure report price statistics from the arm's length firms while the dashed lines do so for the intrafirm prices. Colors and labels indicate

which market share configuration was used to generate the median duration series. Again, the model is able to match the empirical observation that median price durations are shorter for intrafirm prices.¹³

5.2 Passthrough: Empirics and Simulation

Many recent papers have demonstrated that, even conditional on price adjustment, cost passthrough is less than 1, including Gopinath et al. (forthcoming), Burstein and Jaimovich (2009), and Fitzgerald and Haller (2009). Others, such as Bernard et al. (2006), Hellerstein and Villas-Boas (forthcoming), and Neiman (2010) have shown that such passthrough measures are lower for arm's length than for intrafirm price changes. To capture this concept in the simulated data, I consider the β coefficient from the pooled regression:

$$\Delta \ln p_j^{t_j - t_j^{-1}} = \alpha + \beta \Delta \ln e_j^{t_j - t_j^{-1}} + \varepsilon_j^{t_j}$$

where t_j and t_j^{-1} are good specific and respectively denote the times of the most recent and penultimate price changes. Only non-zero price changes are included in the regression, and $\Delta \ln p_j^{t_j - t_j^{-1}} = \ln(p_j^{t_j} / p_j^{t_j^{-1}})$ denotes (in percentage terms) the size of the most recent price change and $\Delta \ln e_j^{t_j - t_j^{-1}} = \ln(e_j^{t_j} / e_j^{t_j^{-1}})$ denotes the accumulated change in the cost shock from the time of previous price change to the time of the most recent change. Figure 7 plots this passthrough coefficient (which, given it is run on simulated data, is very precisely estimated) for arm's length prices in the two arm's length firm structure and for intrafirm prices in the hybrid structure. As in the data, arm's length conditional passthrough is clearly incomplete and is below that of related parties.¹⁴

¹³Another natural assumption might be that intrafirm adjustment costs are lower than arm's length adjustment costs. This could certainly replicate the finding of lower intrafirm duration, but could not explain the results on passthrough and synchronization. Further, as mentioned in Neiman (2010), this would imply counterfactual differences in the size distribution of price changes.

¹⁴Passthrough estimates in the micro-data literature are quite small and range from about 10 percent to about 50 percent. As with most of the passthrough literature, this model's average arm's length rate of passthrough of 55 percent is thus too high. I acknowledge this, and focus on the model's ability to match the comparative statics of passthrough in the data.

5.3 Synchronization: Empirics and Simulation

Finally, Cavallo (2009) and Midrigan (2006) demonstrate that price changes are synchronized, and Neiman (2010) additionally finds that this synchronization is larger for arm's length transactions than for intrafirm transactions. There is no standard measure used to quantify price change synchronization. Here, I observe the percentage of simulated months in which both manufacturers' prices change and compare it to the percentage that would be randomly generated. For instance, if firm 1 changes its prices every d_1 months, and firm 2 does so every d_2 months, zero synchronization would imply the existence of months with two price changes about $100/(d_1 d_2)$ percent of the time. Hence, I measure synchronization in the simulated data as a ratio ("synchronization ratio") of the frequency of months with two price changes to the frequency that would be expected with randomly timed changes. The vast majority of time-dependent models would, for example, generate ratio values of 1. Values greater than 1 suggest synchronization in the data.

Figure 8 shows the synchronization ratio for the sectors with both arm's length firms as well as for the hybrid sectors. With only one exception, the ratios are greater than 1 and demonstrate that the model produces price change synchronization. While the analytics and static exercise in Section 3 indicate that, all things equal, we expect less synchronization in hybrid sectors with a related party, it is impossible to generate hybrid and fully non-integrated sectors with all things equal. Nonetheless, the hybrid sector exhibits less synchronization in 8 of the 12 simulations, consistent with the evidence that related party price changes are less synchronized.

6 Conclusion

A large number of recent empirical studies have documented new facts on stickiness, cost passthrough, and synchronization in final good and traded intermediate prices. Arm's length price stickiness is heterogenous and decreases with the elasticity of demand for a good. Incomplete cost passthrough is not simply a function of nominal rigidities and persists even after prices are changed. There is evidence of bunching in the timing of price changes. Further, studies that consider transactions between related parties have found that intrafirm stickiness and synchronization are lower and passthrough is higher. These facts present challenges

to traditional pricing models in the open- and closed-economy macroeconomics literature. I write a model of intermediate good pricing that can be used to describe both arm's length and intrafirm pricing strategies and is capable of delivering all these empirical patterns.

References

- [1] Atkeson, Andrew and Ariel Burstein. 2008. "Pricing to Market, Trade Costs, and International Relative Prices." *American Economic Review*.
- [2] Bergin and Feenstra, 2000. "Staggered Price Setting and Endogenous Persistence," *Journal of Monetary Economics*, 43: 657-680.
- [3] Bernard, Andrew, J. Bradford Jensen, and Peter Schott. 2006. "Transfer Pricing by U.S.-based Multinational Firms." NBER Working Paper No. 12493.
- [4] Bernard, Andrew, J. Bradford Jensen, and Peter Schott. 2007. "Importers, Exporters, and Multinationals: A Portrait of Firms in the U.S. that Trade Goods." Unpublished.
- [5] Bills, Mark, and Peter Klenow. 2004. "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112: 947-985.
- [6] Burstein, Ariel and Nir Jaimovich. 2009. "Understanding Movements in Aggregate and Product-Level Real-Exchange Rates," Working Paper.
- [7] Calvo, Guillermo. 1983. "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12(3): 383-398.
- [8] Cavallo, Alberto. 2009. "Scraped Data and Sticky Prices: Frequency, Hazards, and Synchronization," Working Paper.
- [9] Doraszelski, Ulrich and Mark Satterthwaite. forthcoming. "Computable Markov-Perfect Industry Dynamics," *Rand Journal of Economics*.
- [10] Dotsey, Michael, Robert King, and Alexander Wolman. 1999. "State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output." *Quarterly Journal of Economics*, 114(2): 655-690.
- [11] Fitzgerald, Doireann and Stefanie Haller. 2009. "Pricing-to-Market: Evidence from Producer Prices," Working Paper.
- [12] Gopinath, Gita, Oleg Itskhoki, and Roberto Rigobon. forthcoming. "Currency Choice and Exchange Rate Pass-through." *American Economic Review*.
- [13] Gopinath, Gita and Oleg Itskhoki. forthcoming. "Frequency of Price Adjustment and Passthrough," *Quarterly Journal of Economics*.
- [14] Gopinath, Gita and Roberto Rigobon, 2008. "Sticky Borders." *Quarterly Journal of Economics*, 123(2): 531-575.
- [15] Hellerstein, Rebecca and Sofia Berto Villas-Boas, forthcoming. "Outsourcing and Pass-through." *Journal of International Economics*.
- [16] Hershleifer, Jack. 1956. "On the Economics of Transfer Pricing." *Journal of Business*, 29: 172-183.
- [17] Judd, Kenneth. 1988. *Numerical Methods in Economics*. Cambridge, Massachusetts: MIT Press.
- [18] Kimball, Miles. 1995. "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, 27: 1241-1277.
- [19] Klenow, Peter and Oleksiy Kryvtsov. 2008. "State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?" *Quarterly Journal of Economics* 123: 863-904.

- [20] Midrigan, Virgiliu. 2006. "Menu Costs, Multi-Product Firms and Aggregate Fluctuations." <http://homepages.nyu.edu/~vm50/>.
- [21] Midrigan, Virgiliu. forthcoming. "Is Firm Pricing State or Time-Dependent? Evidence from U.S. Manufacturing." *Review of Economics and Statistics*.
- [22] Miranda, Mario and Paul Fackler (2002). *Applied Computational Economics and Finance*. Cambridge, Massachusetts: MIT Press.
- [23] Miranda, Mario and Dmitry Vedenov. 2001. "Numerical Solution of Dynamic Oligopoly Games with Capital Investment." *Economic Theory*, Vol. 18(1): 237-261.
- [24] Nakamura, Emi, and Jon Steinsson, 2008. "Five Facts About Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics*, 123(4): 1415-1464.
- [25] Nakamura, Emi and Jon Steinsson. forthcoming. "Monetary Non-Neutrality in a Multi-Sector Menu Cost Model," *Quarterly Journal of Economics*.
- [26] Neiman, Brent. 2010. "Stickiness, Synchronization, and Passthrough in Intrafirm Trade Prices," *Journal of Monetary Economics*, forthcoming.
- [27] Schoenle, Raphael. 2009. "International Menu Costs and Price Dynamics," Working Paper.
- [28] Taylor, John. 1980. "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, 88(1): 1-23.
- [29] Yang, Jiawen. 1997. "Exchange Rate Pass-Through in U.S. Manufacturing Industries." *Review of Economics and Statistics*, 79(1): 95-104.
- [30] Zbaracki, Mark J., Ritson, Mark, Levy, Daniel, Dutta, Shantanu, and Mark Bergen. 2004. "Managerial and Customer Dimensions of the Costs of Price Adjustment: Direct Evidence from Industrial Markets," *Review of Economics and Statistics*, 86(2): 514-533.

Figures




Figure 1: Arm's Length No-Adjust Regions by Elasticity of Substitution

Notes: Red regions define s-S bands within which a firm will not change prices. Movement along the vertical axes represents a percentage shock to a firm's own production cost, and movement along the horizontal represents a percentage shock to the competitor's price. No-adjust regions are calculated assuming there is no response from the competitor. The change from a thick band for differentiated sectors to a thin one for less differentiated sectors will generate heterogeneity in stickiness. Shocks will push the black market outside of the red band more frequently in the right plot, leading to less stickiness. Both bands have a slope, indicating strategic complementarities and generating synchronization in price changes.




Figure 2: No-Adjust Regions by Vertical Structure

Notes: Red regions define s-S bands within which a firm will not change prices. Movement along the vertical axes represents a percentage shock to a firm's own production cost, and movement along the horizontal represents a percentage shock to the competitor's price. No-adjust regions are calculated assuming there is no response from the competitor. The change from a thicker band for arm's length goods sectors to a thinner one for related parties will generate shorter intrafirm duration. Shocks will push the black market outside of the red band more frequently in the right plot, leading to less stickiness. The right band is essentially flat, leading to less synchronization among intrafirm price changes.




Figure 3: Sample Simulated Policy Function and Model-Generated Data

Notes: Fixing particular values for the competitor's previous price and current cost, this is a sample policy functions where the vertical axis gives the conditional probability of a price change before observing the menu cost realization and the x- and y-axes give the firm's previous price and current cost. Though the random adjustment cost adds some time-dependence to the problem, this plots shows that the solution retains a state-dependent flavor as the probability of adjustment changes sharply with the state. The constant probability of adjustment in a Calvo model, for instance, would appear above as a flat plane.




Figure 4: Sample Simulated Policy Function and Model-Generated Data

Notes: This is a sample of the simulated price and cost data (lines, left axis) and probability of adjustment (shading, right axis) for a given sector. $p_{AL,j}$ and p_{-j} are the manufacturers' prices, $m_{AL,j}$ is the cost of arm's length firm j , and $\xi_{AL,j}$ is the probability of adjustment immediately prior to observing the month's adjustment cost. In the start of year 5, firm j increases its price even though its own cost has clearly been declining. This is labeled an "Example of Complementarity" because the price increase is clearly driven by the (correct) expectation that the competitor would increase its own price.




Figure 5: Simulated Arm's Length Duration

Notes: Results from simulation detailed in Section 4 and Appendix B. All three simulated configurations generate heterogeneous stickiness that decreases with the elasticity of substitution.




Figure 6: Simulated Arm's Length and Intrafirm Duration

Notes: Results from simulation detailed in Section 4 and Appendix B. All three simulated configurations generate higher median duration for the arm's length firms than for the related parties.

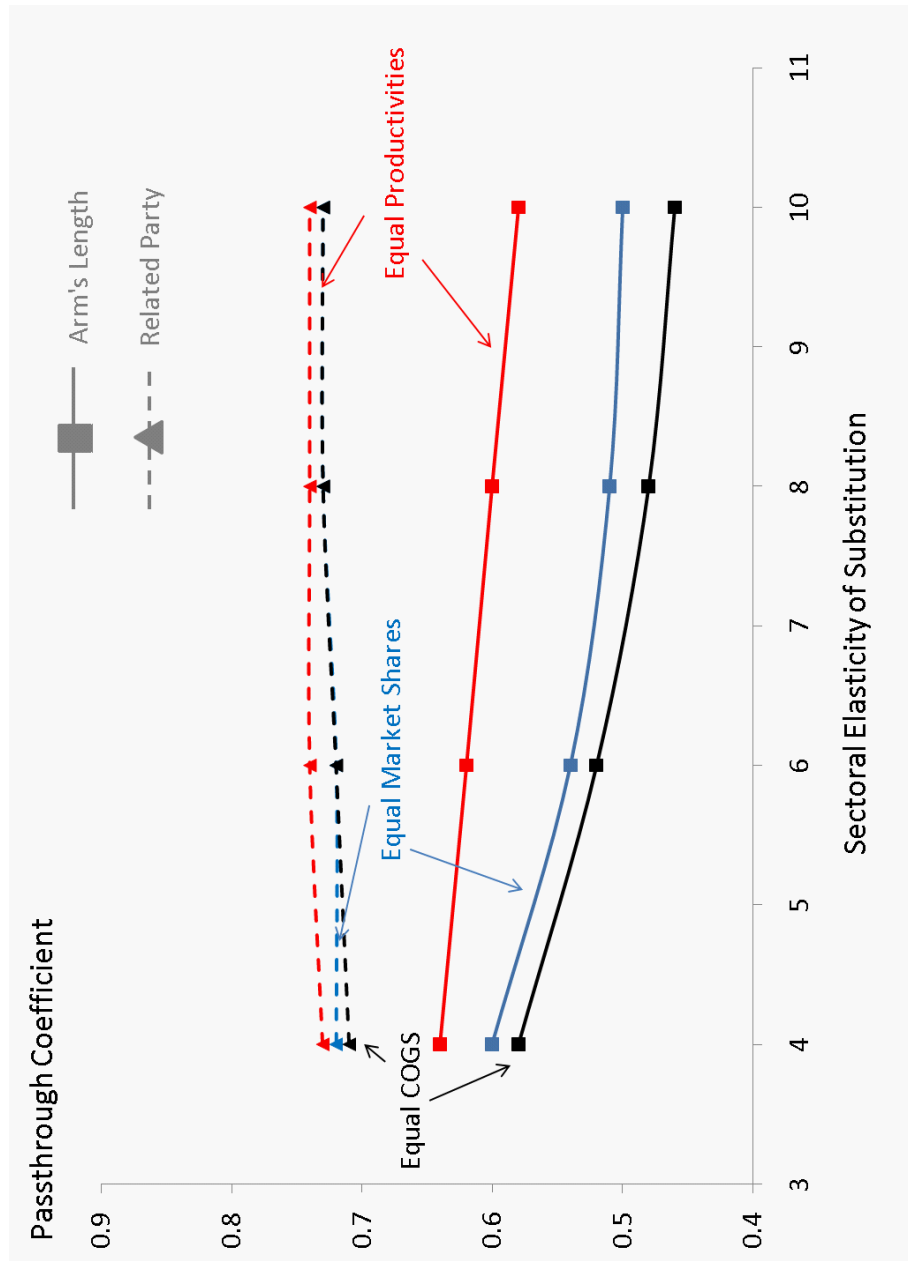


Figure 7: Simulated Arm's Length and Intrafirm Passthrough

Notes: Results from simulation detailed in Section 4 and Appendix B. All three simulated configurations generate incomplete passthrough for the arm's length firms and higher intrafirm passthrough.




Figure 8: Simulated Arm's Length and Intrafirm Synchronization

Notes: Results from simulation detailed in Section 4 and Appendix B. All three simulated configurations generate synchronization ratios greater than 1, indicating above-random levels of price chagne synchronization. The majority of the cases involve larger synchronization ratios for the sectors with two arm's length firms than for the hybrid sectors with one related party.

Appendix A: Additional Calculations and Proofs

This appendix gives details for several of the calculations made in the text.

Claim 1 *We wish to show:*

$$\Omega_{AL} = \frac{s_{AL}\varepsilon(\varepsilon-1)^2}{\varepsilon(\varepsilon-1) + (\rho-\eta)(\rho-1)s_{AL}(1-s_{AL})} c_{AL}.$$

With flexible prices, the arm's length firm's profits can be written as:

$$\pi_{AL}^A = \frac{1}{\varepsilon} c_{AL} p_{AL}.$$

Partially differentiating with respect to the optimal arm's length flexible price gives:

$$\begin{aligned} \frac{\partial \pi_{AL}^A}{\partial p_{AL}} &= -\frac{1}{\varepsilon^2} \frac{\partial \varepsilon}{\partial p_{AL}} c_{AL} p_{AL} + \frac{1}{\varepsilon} \frac{\partial c_{AL}}{\partial p_{AL}} p_{AL} + \frac{1}{\varepsilon} c_{AL} \\ &= -\frac{1}{\varepsilon^2} \frac{\partial \varepsilon}{\partial s_{AL}} \frac{\partial s}{\partial p_{AL}} c_{AL} p_{AL} + c_{AL} \left(\frac{1-\varepsilon}{\varepsilon} \right) \\ &= -\frac{1}{\varepsilon^2} c_{AL} [\varepsilon(\varepsilon-1) + (\rho-\eta)(\rho-1)s_{AL}(1-s_{AL})] \\ &= -c_{AL} \left(\frac{\partial p_{AL}}{\partial m_{AL}} \right)^{-1}. \end{aligned}$$

This implies that we can write: $\frac{\partial \pi_{AL}^A}{\partial m_{AL}} = \frac{\partial \pi_{AL}^A}{\partial p_{AL}} \frac{\partial p_{AL}}{\partial m_{AL}} = -c_{AL}$. Differentiating again, we get:

$$\begin{aligned} \frac{\partial^2 \pi_{AL}^A}{\partial m_{AL}^2} &= -\frac{\partial c_{AL}}{\partial p_{AL}} \frac{\partial p_{AL}}{\partial m_{AL}} \\ &= \frac{c_{AL}}{p_{AL}} \frac{\varepsilon^3}{\varepsilon(\varepsilon-1) + (\eta-\rho)(1-\rho)s_{AL}(1-s_{AL})}. \end{aligned}$$

Substituting into the form: $\frac{1}{2} \frac{\partial^2 \pi_{AL}^A}{\partial m_{AL}^2} (dm_{AL})^2 = \frac{1}{2} \Omega_{AL} \hat{m}^2$ demonstrates the claim.

Claim 2 *We wish to show:*

$$\Omega_{RP} = s_{RP} \varepsilon_{RP} c_{RP}.$$

As above, we start with the flexible price expression for related party profits: $\pi_{RP}^A = \frac{1}{\eta} p c = \frac{1}{\eta} p^{1-\eta}$. Partially differentiating with respect to the distributor's unit input cost gives: $\frac{\partial \pi_{RP}^A}{\partial x} = \frac{1-\eta}{\eta} p^{-\eta} \frac{\partial p}{\partial x} = -c$ because $\frac{\partial p}{\partial x} = \frac{\eta}{\eta-1}$. Using:

$$\begin{aligned} \frac{\partial x}{\partial p_{RP}} &= \left[\gamma_{AL}^\rho p_{AL}^{1-\rho} + (1-\gamma_{AL})^\rho p_{RP}^{1-\rho} \right]^{\frac{\rho}{1-\rho}} (\gamma_{RP})^\rho p_{RP}^{-\rho} \\ &= c_{RP}/c, \end{aligned}$$

we can write: $\frac{\partial \pi_{RP}^A}{\partial m_{RP}} = \frac{\partial \pi_{RP}^A}{\partial x} \frac{\partial x}{\partial p} \frac{\partial p}{\partial m_{RP}} = -c_{RP}$, because $p_{RP} = m_{RP}$ and, hence, $\frac{dp_{RP}}{dm_{RP}} = 1$. The remaining steps follow those in Claim 1.

Claim 3 We define $\eta / (2\eta - 1) = \overline{s_{AL}}$ and wish to show that:

$$s_{AL} < \overline{s_{AL}} \implies \Omega_{RP} > \Omega_{AL}.$$

We write:

$$\begin{aligned} \Omega_{RP} &= s_{RP} \varepsilon_{RPC} x \\ &= (1 - s_{AL}) (\eta - s_{AL} (\eta - \rho)) c x \\ &= (\eta (1 - 2s_{AL} + s_{AL}^2) + \rho s_{AL} (1 - s_{AL})) c x, \end{aligned}$$

and

$$\begin{aligned} \Omega_{AL} &= (\varepsilon_{AL} - 1) \alpha s_{AL} c x \\ &= (\eta s_{AL}^2 + \rho s_{AL} (1 - s_{AL}) - s_{AL}) \alpha c x. \end{aligned}$$

In this form, it is easy to see:

$$(\alpha \Omega_{RP} - \Omega_{AL}) / \alpha c x = \eta (1 - 2s_{AL}) + s_{AL}.$$

Factoring out the arm's length share, we see that:

$$s_{AL} < \eta / (2\eta - 1) \implies \alpha \Omega_{RP} > \Omega_{AL},$$

and since $\alpha < 1$, this implies $\Omega_{RP} > \Omega_{AL}$.

Appendix B: Model Solution and Simulation

This appendix gives details of the projection method used to find an approximate solution to the model in Section 2 and to generate simulated data.¹⁵ Application of these methods to a model of adjustment costs follows Midrigan (2006, forthcoming) and their use for solving a dynamic game follows Miranda and Vedenov (2001). Miranda and Fackler (2002) provides an accompanying MATLAB toolbox (CompEcon) that was used extensively.

I approximate each of the two expected value functions (10) with a linear combination of orthogonal (Chebyshev) basis polynomials:

$$V_j(\Theta^t) \approx \sum_{i_1=1}^{N_1} \sum_{i_2=1}^{N_2} \sum_{i_3=1}^{N_3} \sum_{i_4=1}^{N_4} b_{i_1 i_2 i_3 i_4} \psi_{i_1}(p_1^{t-1}) \psi_{i_2}(p_2^{t-1}) \psi_{i_3}(m_1) \psi_{i_4}(m_2) \quad (\text{B1})$$

where ψ_{i_j} is an i_j th degree Chebyshev polynomial and is a function of the j th state variable. The collocation method requires the approximation (B1) to hold exactly at specific points called collocation nodes: $\{p_1^{t-1}(i_{n_1}), p_2^{t-1}(i_{n_2}), m_1(i_{n_3}), m_2(i_{n_4})\}$ for $i_{n_k} = 1 \dots N_k$ and $k = 1 \dots 4$. Since there are two value functions to estimate (one for each firm), this reduces the problem to solving a system of $2N_1 N_2 N_3 N_4$ equations in $2N_1 N_2 N_3 N_4$ unknown coefficients, $b_{i_1 i_2 i_3 i_4}$.

The algorithm starts with a guess for the coefficients on the Chebyshev basis polynomials and the optimal policies for each firm at each collocation node. Since the approximated function is an expected (rather than realized) value function, this policy is the profit maximizing price, conditional on an adjustment cost sufficiently low to warrant a price change. This potential price (together with the distribution function $G_j(\phi)$) implicitly defines the probability of price adjustment.

Given the initial set of collocation coefficients and taking the guess for the other firm's optimal policy as given, I use a modified Newton routine to solve simultaneously for each firm's optimal price, conditional on adjustment, at each collocation node. The first order condition (FOC) has a term reflecting profits given an adjustment price as well as the expected continuation value given this price. In order to approximate this latter term, I discretize the joint distribution of cost (exchange rate) shocks and integrate using Gaussian quadrature. After each Newton step, I calculate the probability of adjustment, ξ_j , implied by the optimal adjustment price because this probability enters the competitor's own optimization problem (9). This process continues until the FOC of both firms is sufficiently close to zero and the probability of adjustment does not change with additional iterations.

Finally, a combination of function iteration with dampening and Newton's method with backstepping is used to determine the next set of Chebyshev polynomial coefficients to consider. With this new set of collocation coefficients, a new set of equilibrium policies is found. The process is repeated until the changes in the basis coefficients and optimal policies in each iteration, as well as

¹⁵I thank Uli Doraszelski for his very helpful advice on the numerical methods detailed in this section.

the differences between the right-hand side and left-hand side of the expected value function (10) at the collocation nodes, are extremely small.

The accuracy of the approximations can be gauged by calculating the difference between the left- and right-hand sides of the firm's expected value functions at a set of nodes denser than the collocation nodes. For some of the parameter configurations tested, these errors are larger than would be desirable, at average respective levels of about $1e-4$ and $5e-4$ and maximum levels of about $7e-4$ and $4e-3$ for the related party and arm's length firms, when expressed as a share of the expected value functions. This lack of precision, in addition to the two-firm structure, precludes treatment of the simulation as a true calibration exercise. The consistency of the comparative statics and qualitative results across approximations with varying numbers of collocation nodes, however, suggests this level of accuracy is sufficient to demonstrate the key points in this paper.¹⁶

The above procedure generates a solution for a given set of parameter values. To consider other parameter values, I start with the solution to a close by problem (in the sense that the parameter values are close) and use simple continuation methods. There are no guarantees these will work, however, and I often had to try varying multiple parameters, including the number of collocation nodes itself, in order to move around the parameter space.¹⁷ Once a solution to the above system of equations has been approximated, I simulate the cost shocks and generate simulated pricing responses from the firms.

There is no way to guarantee a suitable starting guess for policies from new locations in the parameter space (after random cost shocks), so the algorithm occasionally does not converge. In such cases, I simply draw a different shock value and try again. These instances account for far less than one percent of all simulated good-periods. With simulated cost and pricing data, I generate measures for key statistics such as the unconditional duration (or stickiness) of prices, the synchronization of price changes, and the pass-through of cost shocks.

¹⁶Given the very similar results for varying numbers of nodes, most results in the figures reflect faster simulations with less nodes than that used to measure the size of approximation errors.

¹⁷See Chapter 5 of Judd (1988) for a discussion of simple continuation methods.