Menu Costs, Aggregate Fluctuations, and Large Shocks[*]

Peter Karadi[†]   Adam Reiff[‡]

Sept, 2014

Abstract

How do frictions in price setting influence monetary non-neutrality? We revisit this classic question in a quantitative menu cost model with multi-product firms that face idiosyncratic shocks with unsynchronized stochastic volatility. The model matches the unconditional distribution of price changes and successfully predicts new evidence on pricing responses to large value-added tax shocks. In particular, it captures both the exploding fraction of price changes and the shape of their conditional distribution, outperforming alternative models. The model generates near money neutrality even to small nominal shocks.

JEL codes: E31, E52

Keywords: Monetary Policy Transmission, State-Dependent Pricing, Price Change Distribution, Value-Added Tax

[†]Corresponding author, peter.karadi@ecb.int, European Central Bank, Neue Mainzer Strasse 66, Frankfurt am Main, 60311, Germany and CEPR

[‡]reiffa@mnb.hu, Magyar Nemzeti Bank (The Central Bank of Hungary), Szabadsag ter 8-9, Budapest, Hungary

1 Introduction

How do frictions in price setting influence the impact of nominal shocks on economic activity? We revisit this classic question in a model where small, fixed menu costs of price adjustment generate staggered price setting (Mankiw (1985), Caplin and Spulber (1987)). Our quantitative model has three key features motivated by relevant micro-level pricing facts. First, firms face idiosyncratic shocks (Golosov and Lucas (2007)). This helps to match the large observed magnitude of price changes that aggregate fluctuations are unable to explain. Second, we assume that each firm produces a bundle of multiple products (Lach and Tsiddon (2007), Midrigan (2011), Alvarez and Lippi (2014)), and can reprice the whole bundle for a single menu cost. As the price change of each product except the first one is 'free', this assumption generates frequent small price changes that we also observe in the data. And third, as a novel feature, we assume that the idiosyncratic shocks exhibit stochastic volatility that is unsynchronized across firms.[1] In particular, we assume that shocks are realizations of a mixture of two Gaussian distributions with either a small or a large volatility. Two key idiosyncratic shock distributions used in the literature: the Gaussian (Golosov and Lucas (2007)) and Poisson distributions (Midrigan (2011))[2] are special cases of our distribution.

   We calibrate key parameters of the model to match the low frequency, large average size and wide dispersion of price changes, and find that the model generates small and temporary real effects after a monetary policy shock. This result is particularly striking in the context of the literature: our model matches the cross-sectional distribution of price changes at least as well as Midrigan (2011), but has the opposite implication on aggregate price flexibility. We find near money neutrality like Golosov and Lucas (2007), while Midrigan (2011) predicts large real effects in line with predictions of time-dependent pricing models (Calvo (1983)).

   What explains our near-neutrality result? It is the consequence of the interaction of two key factors. We find that both (i) the magnitude of the menu costs and (ii) the mass of firms on the verge of price adjustment are sizeable. Conditional on a fixed frequency of price changes,[3] both factors mitigate monetary non-neutrality by strengthening the endogenous Caplin and Spulber (1987)-type selection of price changes. High menu costs discourage small price changes and make firms with large desired price changes more likely to adjust. As a result, when an

---

[1]Vavra (2014), differently from us, assumes time-varying stochastic volatility that is synchronized across firms, so if the idiosyncratic volatility is high for one firm, it is also high for the others. He finds that money non-neutrality is significantly mitigated during periods of high idiosyncratic volatility, but stays high during periods of low volatility. In contrast, we find that using the empirically preferred distribution with unsynchronized stochastic volatility leads to money near-neutrality even for the average level of volatility.

[2]The Poisson distribution in Midrigan (2011) is a Gaussian shock arriving with a certain Poisson rate. See also Danziger (1999) and Gertler and Leahy (2008).

[3]Higher menu costs *reduce* price flexibility by reducing the fraction of price changes. As this fraction is directly observable, however, and remains unaffected by small aggregate shocks, it is reasonable to consider the exercise when the effect of lower menu costs on this fraction is neutralized by reparametrization of the model (e.g. by increasing the standard deviation of the idiosyncratic shocks).

aggregate shocks hits, those firms respond that are prepared to change their prices substantially. This effect is further reinforced if the measure of the adjusting firms is also high, as in our case. The two effects imply a disproportionally large price level response and small real effects of monetary policy shocks.

But what makes us infer large menu cost and a large fraction of firms that are ready to adjust? Both factors are the direct consequence of the assumed shape of our idiosyncratic shock distribution. Assuming Gaussian shocks, Golosov and Lucas (2007) needed menu costs even larger to match the large average size and low frequency of price changes. Midrigan (2011), in contrast, showed that with fat-tailed idiosyncratic shocks the same moments can be matched by small menu costs. The reason is that with his Poisson distribution most desired price changes are zero (no Poisson arrival), so small menu costs are sufficient to match the observed low price-change frequency; but some desired price changes are large, so conditional on a change, the average price change can still be large. Our mixed normal distribution is a generalization of these two distributions, and we let the data speak. The shape favored by the data is closer to the Poisson distribution than to the Gaussian, but it does not have a mass point at zero any more. The wider dispersion of small idiosyncratic shocks explains why our model requires calibrated menu costs that are higher than those found in Midrigan (2011). The significantly lower kurtosis of our idiosyncratic shock distribution, furthermore, implies a significantly higher mass of firms on the verge of adjusting their prices. The two factors reinforce each other, and imply large selection, and near money neutrality.

What does this example teach us? We have matched key moments of the unconditional price change distribution, and found that the conclusions on money neutrality are highly sensitive to the assumptions on the idiosyncratic shock distribution. Our particular distribution with stochastic volatility has implied small real effects, but the Poisson distribution of Midrigan (2011) implies the opposite, and alternative distributional assumptions might generate even different conclusions. Thus, without direct information on the menu costs or the distribution of the idiosyncratic shocks, the unconditional cross-sectional price change distribution does not identify the extent of money non-neutrality.

In the second part of the paper, we propose a solution to this identification problem. We argue that pricing responses to large nominal shocks provide sufficient information to clearly differentiate between models predicting various extent of aggregate price flexibility. In particular, we show that the price change distribution *conditional* on large aggregate shocks provides valuable extra information about both the realistic size of the menu costs and the distribution of desired price changes, and thereby that of the idiosyncratic shocks. This is intuitive. First, the directly observable fraction of adjusting firms after large shocks provide essential information about the realistic magnitude of menu costs. Second, if many firms adjust simultaneously after a large shock, the distribution of desired price changes, and, thereby, the closely related distribution of the underlying idiosyncratic shocks are revealed. We use these insights to provide empirical evidence supporting our model.

Our evidence comes from observations on pricing responses to episodes of large value-added tax (VAT) changes, which provide clear examples of measurable aggregate cost shocks. The episodes happened in Hungary, where posted prices include VAT,[4] so stores need to reset their prices to respond to these changes. Between 2004 and 2006, the Hungarian government changed the VAT-rates on several occasions. It increased the lower preferential rate from 12 to 15%, and then from 15 to 20%, while it decreased the standard rate from 25 to 20%.[5]

To study the price effects of these large tax shocks, we use micro-level price data underlying the consumer price index.[6] As an immediate response to the tax changes, we observe a major jump in the fraction of price changes of affected prices (from 13% to 60% after a 5% shock). Furthermore, we observe highly dispersed price changes in the months of the tax shocks, with an excess mass around the size of the tax change, but also frequent large price changes (the kurtosis increasing to over 8 from around 4).

We test the predictions of our baseline model and compare them to the predictions of alternative models by simulating pricing responses to large, permanent, aggregate shocks of the same size as the observed shock. By confronting the predictions with the data, we find that our baseline model with mixed normal idiosyncratic shock distribution predicts both the frequency of changes and the shape of the price change distribution at the months of the tax changes remarkably well. Furthermore, our model quantitatively outperforms alternative menu cost models. In particular, both the frequency effects, and the dispersion of price changes as a response to large shocks are underestimated by the Golosov and Lucas (2007) model with Gaussian idiosyncratic shocks, and overestimated by the model of Midrigan (2011) using the Poisson distribution. These results give strong support to our underlying assumptions about the distribution of idiosyncratic shocks and the implied magnitude of the menu costs. Our correct prediction of the fraction of adjusting firms after the shocks indicates that our calibrated menu costs are of the right magnitude. Furthermore, as the majority of prices responds to the large aggregate shocks, the observed price change distribution gets close to the usually unobserved desired price change distribution revealing essential information about the underlying idiosyncratic distribution. The fact that our baseline model gets close to matching the observed shape of the price change distribution when the shock hits provides direct evidence on the validity of our underlying mixed normal idiosyncratic shock distribution assumption.

Related literature Our paper is a contribution to the strand of literature that analyzes the extent of monetary non-neutrality in realistic menu cost models[7]. To calibrate key pricing

---

[4]As in most other European countries, but differently from the U.S., where prices are posted net of sales tax.

[5]The changes took effect at different dates, they were preannounced and widely publicized to both stores and consumers.

[6]Our Hungarian micro price sample is standard. The pricing moments of the homogeneous processed food sector we use are very close to those in a sample of barcode data of a large US supermarket chain used by Midrigan (2011).

[7]Recent key examples are Golosov and Lucas (2007), Midrigan (2011), Alvarez, Bihan, and Lippi (2014)

parameters, most of the literature uses moments of the price change distribution observed during calm periods with low aggregate volatility. We argue that these moments can imply a wide range of monetary non-neutralities,[8] and, therefore we use additional information from a series of natural experiments generated by large aggregate VAT shocks to support the validity of our model and calibrated parameter values. In this sense, our paper is closer to papers that use pricing observations during rare high trend inflation periods (Golosov and Lucas (2007), Gagnon (2009), Alvarez, Gonzalez-Rozada, Neumeyer, and Beraja (2011)). These papers provide supporting evidence on the state-dependent menu cost assumptions. Our aim is to go a step further and provide evidence on some of the key underlying assumptions of these menu cost models, like the shape of the idiosyncratic shock distribution and the size of the menu costs, that determine the extent of their predicted monetary non-neutrality.

Our work is also related to Elsby and Michaels (2014), who also find that the shape of the idiosyncratic shock distribution is key in terms of monetary neutrality in menu cost models. They find that continuous idiosyncratic shock distributions (like the normal or our mixed normal) imply monetary near-neutrality as the menu cost approaches 0; but discontinuous distributions with an atom (like the Poisson distribution) can imply money non-neutrality not dissimilar from the Calvo (1983) model even as the menu cost approaches 0. We argue that for realistic menu costs away from 0, the distinction between continuous and discontinuous distributions are not so clear cut. In particular, we provide a numerical example to show (see Figure 3) that the Poisson distribution is not a knife-edge case, and a whole range of continuous idiosyncratic shock distributions around it with positive, but still fairly small menu costs calibrated to match the steady state price change frequencies generate money non-neutrality similar to Calvo (1983). The reason is that these distributions require higher menu costs to match the observed price change frequencies and the costs will counteract the impact of the more dispersed idiosyncratic shock distributions. So ultimately choosing between relevant idiosyncratic shock distributions is an empirical question. Our paper takes on the empirical challenge and use large aggregate shocks to figure out which of these distributions are more realistic in pricing models, and find that the empirically relevant distributions imply aggregate near-neutrality.

Price responses to value-added tax shocks have been documented in the literature (see Gagnon (2009) and the references therein). These papers document significant increases in the fraction of price changes in the months of the tax changes in line with the predictions of state-dependent pricing models. In a recent example, Gagnon, López-Salido, and Vincent (2012)

---

[8]This result is not inconsistent with the conclusion of Alvarez, Bihan, and Lippi (2014). They also find that the extent of monetary non-neutrality is unidentified by the kurtosis and the frequency of price changes – sufficient statistics in a large class of models – if the cost distribution has fat-tails with infrequent and large jumps, as in our case. We provide an example (see Figure 3) when models calibrated to have the same frequency and kurtosis can generate aggregate outcomes ranging from monetary non-neutrality similar to Calvo (1983) to near-neutrality of Golosov and Lucas (2007) under various idiosyncratic shock distributions. The identification of the degree of monetary non-neutrality requires further information about the shape of the cost shock distribution or the size of the menu cost, both of which are revealed after a large aggregate shock presented in our paper.

are using Mexican value-added tax increases to argue for the presence of strong extensive margin and selection effects. The evidence presented in our paper contributes to this evidence by analyzing responses to positive *and* negative VAT change for a large set of products, within a short time span. We document that the aggregate responses to these shocks are asymmetric, in line with predictions of menu cost models with trend inflation (Ball and Mankiw (1994)). The responses are asymmetric because after a VAT decrease, firms can save on their menu costs by keeping their nominal prices constant and allowing the trend inflation to reduce their relative prices for free. Differently from the above mentioned predominantly descriptive papers, furthermore, we use the new evidence to test the quantitative predictions of our structural model. Our model has no problem in quantitatively matching also the observed asymmetry of the inflation responses.

The Hungarian episode provides new evidence that alternative pricing models should strive to explain. Standard pricing models with *information frictions* (see e.g. Mackowiak and Wiederholt (2009), Mankiw and Reis (2002), Woodford (2003a)), for example, assume that it is the costs of collecting information on shocks that limit firms' optimal response to aggregate shocks. Value-added tax shocks, however, are large and measurable shocks that are widely-publicized months before their implementation. These information costs, thereby, should be minimal, and can be expected to be outweighed by the potential losses caused by suboptimal price setting. Still we could observe some 40% of the affected firms failing to respond to a positive shock, suggesting that information frictions alone are insufficient to explain observed price rigidities. An alternative strand of the literature uses *search frictions* instead to explain price stickiness (Stiglitz (1984), Cabral and Fishman (2012)). Their common key assumption is that the marginal consumers are uninformed about the cost shocks faced by price setting firms. To prevent them from initiating a search, firms optimally refrain from responding to each cost changes. This assumption, however, is arguably strong for VAT-shocks, which were also known by most of the consumers. Still some affected firms did not change their prices, suggesting that the above explanation can neither be the sole reason for infrequent price adjustments. The observed asymmetry between the impact effect of positive and negative VAT shocks, which we explain by the interaction of trend inflation and menu costs, also challenges some alternative price rigidity explanations. Muted price-decreases after the reduction of the VAT rates were considered highly unfair by the general public, as can be witnessed from the newspaper articles of the time. But this fairness consideration – emphasized by Rotemberg (2005) and Rotemberg (2011) – were clearly not sufficient to dissuade firms from following asymmetric pricing strategies.

Structure: Section 2 presents our model. In Section 3, we describe our the data, calibrate our model, show its predictions on money neutrality and explain the intuition behind our results. In Section 4, we present the pricing impacts of large value-added tax changes and use them to test our model and contrast its predictions to alternative menu cost models. Section 5 presents robustness results and Section 6 concludes.

2 The Model

We have a general equilibrium monetary model with a representative consumer, a continuum of firms producing differentiated goods and a government setting monetary policy and tax rates. The consumer supplies labor, consumes and saves. Firms use labor to produce a bundle of multiple products and face idiosyncratic shocks influencing their optimal prices. To change any of their preset prices, firms need to pay a fixed menu cost. Firms are assumed to satisfy all demand at their posted price. Government sets an exogenous process for the money supply, and the level of value-added tax rates.[9]

2.1 Consumers

The representative consumer consumes a Dixit-Stiglitz aggregate (C_t) of a basket of multiple goods g $(g = 1, 2, \ldots, G)$ purchased from firms i $(i \in [0, 1])$, and supplies labor L_t to maximize the expected present value of her utility:

$$\max_{C_t(i,g), L_t, B_t} E_0 \sum_{t=0}^{\infty} \beta^t \left(\log C_t - \frac{\mu}{1+\psi} L_t^{1+\psi} \right), \tag{1}$$

where β is the discount factor, μ is the disutility of labor, and ψ is the inverse Frisch-elasticity of labor supply. The aggregate consumption is $C_t = \left(\int C_t(i)^{(\theta-1)/\theta} di \right)^{\theta/(\theta-1)}$, where $C_t(i) = \left(\frac{1}{G} \sum_{g=1}^{G} [A_t(i,g) C_t(i,g)]^{(\gamma-1)/\gamma} \right)^{\gamma/(\gamma-1)}$. C_t is a CES-aggregate of individual good consumptions $C_t(i,g)$ (with across-firm and across-good elasticity parameters θ and γ, respectively). The measure of firms i is normalized to 1, and the number of goods per firm is G. $A_t(i,g)$ reflects the quality of the good, with higher quality goods providing larger marginal utility of consumption. These good- and producer-specific idiosyncratic shocks influence optimal prices generating dispersion that help us to match the empirical price change size distribution.

The consumer's budget constraint for each time period t is given by

$$\int \sum_{g=1}^{G} P_t(i,g) C_t(i,g) di + B_{t+1} = R_t B_t + W_t L_t + \tilde{\Pi}_t + T_t, \tag{2}$$

where $P_t(i,g)$ is the nominal gross price of firm i for product g, B_t is a nominal bond with gross return R_t, W_t is nominal wage, $\tilde{\Pi}_t$ is nominal profits, and T_t is a lump-sum transfer.

The aggregate price level in this economy is $P_t = \left(\int P_t(i)^{1-\theta} di \right)^{\frac{1}{1-\theta}}$, where the producer-level aggregate price $P_t(i) = \left(\frac{1}{G} \sum_{g=1}^{G} [P_t(i,g)/A_t(i,g)]^{1-\gamma} \right)^{1/(1-\gamma)}$. This implies that aggregate expenditure is given by $P_t C_t$. The representative consumer's demand for each individual good g of producer i can be expressed as

$$C_t(i,g) = A_t(i,g)^{\gamma-1} \left(\frac{P_t(i,g)}{P_t(i)} \right)^{-\gamma} \left(\frac{P_t(i)}{P_t} \right)^{-\theta} C_t. \tag{3}$$

---

[9]In our framework with no intermediate production, value-added tax is equivalent to sales tax.

The Euler-equation implies that $1/R_t = \beta P_{t-1} C_{t-1}/(P_t C_t)$. The labor supply equation is given by $\mu L_t^\psi C_t = W_t/P_t$.

2.2 The government and the central bank

We assume that money supply growth rate follows an autoregressive process with a drift

$$\log(M_t/M_{t-1}) = g_{Mt} = \mu_M + \rho_M g_{Mt-1} + \varepsilon_{Mt} \tag{4}$$

The growth rate of money supply introduces inflation to the model (with 0 technology growth, inflation rate in the non-stochastic steady state equals $\pi = \mu_M/(1 - \rho_M)$). The extra money supply $M_t - M_{t-1}$ is redistributed to the consumer in a lump-sum way.

We assume that the value-added tax rate follows an exogenous unit root process $\tau_t = \tau_{t-1} + \varepsilon_{\tau t}$. The net revenue of the firm is the fraction $1/(1 + \tau_t)$ of gross revenues, and the fraction $\tau_t/(1+\tau_t)$ is paid for the government. These tax revenues of are also redistributed in a lump-sum way. Without loss of generality, we assume balanced budgets: $M_t - M_{t-1} + (\tau_t/(1+\tau_t))P_t C_t = T_t$.

2.3 Firms

Each firm i is assumed to sell its products (i, g) $(g = 1, 2, \ldots, G)$ in a monopolistically competitive market; post gross nominal prices $P_t(i, g)$ and satisfy all demand at these prices. Firms can adjust any of their prices $(P_t(i, 1), P_t(i, 2), \ldots, P_t(i, G))$ for a single fixed menu cost $\phi P_t C_t$. Firms maximize the expected discounted present value of their profits $E_0 \sum_{t=0}^\infty \frac{1}{\prod_{q=0}^t R_q} \tilde{\Pi}_t(i)$, where the periodic profit level (net of menu costs) is the difference between nominal revenues and production costs: $\tilde{\Pi}_t(i) = \sum_{g=1}^G [(1/(1 + \tau_t))P_t(i, g)Y_t(i, g) - W_t L_t(i, g)]$.

We assume that firms use a constant returns to scale technology with labor as the only factor. The production function is $Y_t(i, g) = L_t(i, g)/A_t(i, g)$ with higher quality products requiring extra labor input.[10] We assume the log of the good quality follows a random walk:

$$ln A_t(i, g) = ln A_{t-1}(i, g) + \varepsilon_t(i, g), \tag{5}$$

where innovations $\varepsilon_t(i, g)$ are mean-zero i.i.d. random variables with variance σ_A^2. We assume that good quality innovations $\varepsilon_t(i, g)$ are uncorrelated across firms, but within firms they are correlated across goods g with the correlation parameter ρ_ε. Furthermore, we assume a that $\varepsilon_t(i, g)$ is drawn from a mean-zero normal distribution with stochastic volatility: with probability p it has a (small) variance $\lambda^2\sigma^2$, and with probability $1 - p$ it has a (larger) variance σ^2. We denote the proportionality between the standard deviations by $0 \leq \lambda \leq 1$.

$$\varepsilon_t(i, g) = \begin{cases} N(0, \lambda^2\sigma^2) & \text{with probability } p \\ N(0, \sigma^2) & \text{with probability } 1 - p \end{cases}$$

---

[10]Similarly to Midrigan (2011), this assumption helps us to reduce the dimensionality of the problem.

To ensure that the variance of $\varepsilon_t(i,g)$ is indeed σ_A^2, we set $\sigma^2 = \sigma_A^2/(p\lambda^2 + 1 - p)$. The parameters p and λ give us degrees of freedom to influence the shape of the distribution. The Gaussian idiosyncratic shock innovations of Golosov and Lucas (2007) and the Poisson shock innovations of Midrigan (2011) are special cases (with $\lambda = 1$ and $\lambda = 0$, respectively).

The production function implies an individual and aggregate labor demand $L_t(i,g) = A_t(i,g) \cdot Y_t(i,g)$ and $L_t = \int \sum_{g=1}^{G} L_t(i,g) di$. We obtain a stationary period profit function (net of menu costs) by substituting labor demand and the households demand (equation (3)) into the periodic profit function $\tilde{\Pi}_t(i)$, using the equilibrium condition $Y_t(i,g) = C_t(i,g)$ and the definition of the aggregate price index, and normalizing the resulting nominal profit function with the nominal GDP ($P_t Y_t$):

$$\bar{\Pi}_t(i) = \sum_{g=1}^{G} \left[\frac{1}{1+\tau_t} p_t(i,g)^{1-\gamma} - w_t p_t(i,g)^{-\gamma} \right] \left(\frac{1}{G} \sum_{g=1}^{G} p_t(i,g)^{1-\gamma} \right)^{(\gamma-\theta)/(1-\gamma)}. \tag{6}$$

The variable $p_t(i,g) = \frac{P_t(i,g)}{A_t(i,g)P_t}$ is the (good quality-adjusted) relative price of firm i of product g, and $w_t = W_t/P_t$ is the real wage. As it is apparent from equation (6), we can write the normalized profit function of the firm as $\bar{\Pi}_t(\mathbf{p}_t(i), w_t, \tau_t)$.

The firms' decision on whether to adjust its prices depends on its last period relative prices and the current idiosyncratic shocks it faces. Our setup allows us to collapse these two idiosyncratic variables into a single state variable $\mu_{t-1}(i,g) = \frac{P_{t-1}(i,g)}{A_t(i,g)P_{t-1}} = p_{t-1}(i,g)\frac{A_{t-1}}{A_t}$ for each firm i and good g. It is its last period relative price deflated by the current good specific idiosyncratic shock. We assume no aggregate uncertainty about the aggregate variables,[11] so our aggregate state in each period is the paths of the value-added tax rate (τ_t), the level of the money supply M_t and the cross-sectional distribution of adjusted relative price vectors Γ_t known at time t, which we denote by $\Omega_t = \{\tau_i, M_i, \Gamma_i\}_{i=1}^{\infty}$.

In each period, after observing the idiosyncratic good quality innovations, firms decide to change nominal prices or keep them constant. If they change any of their prices, they pay the menu costs, which enable them to set their relative price vector optimally. If they decide not to change prices, their relative price vector gets equal to the one inherited from the previous period deflated by this period's equilibrium inflation rate. The normalized value of the firm if it chooses not to change its prices is

$$V^{NC}\left(\mu_{t-1}(i), \Omega_t\right) = \bar{\Pi}\left(\frac{\mu_{t-1}(i)}{1+\pi_t}, w_t, \tau_t\right) + \beta E_t V\left(\frac{\{\mu_{t-1}(i,g)e^{\varepsilon_{t+1}(i,g)}\}_{g=1}^{G}}{1+\pi_t}, \Omega_{t+1}\right), \tag{7}$$

where $\mu_{t-1}(i)$ is the vector of idiosyncratic states – the relative price vector inherited from the previous period and adjusted with this period's good quality shock – at the time of the price

---

[11]This is different from Golosov and Lucas (2007) and Midrigan (2011) who calibrate their models taking aggregate volatility into consideration. The simplification has no qualitative impact on the results: their conclusions on aggregate monetary non-neutrality stay unaffected in a model without aggregate uncertainty, as we show later.

change decision, and $e^{\varepsilon_{t+1}(i,g)} = A_{t+1}(i,g)/A_t(i,g)$ is the next period's productivity innovation for each product g.[12]

If the firm chooses to change its prices, it chooses the new relative price vector $\mathbf{p}_t^*(i)$ optimally, and it pays menu costs, so its normalized value is

$$V^C(\Omega_t) = \max_{\mathbf{p}_t^*(i)} \left\{ \Pi(\mathbf{p}_t^*(i), w_t, \tau_t) - \phi + \beta E_t V\left(\{p_t^*(i,g)e^{\varepsilon_{t+1}(i,g)}\}_{g=1}^G, \Omega_{t+1}\right) \right\}. \tag{8}$$

Finally, the value of the firm is determined by the upper envelope of V^{NC} and V^C.[13]

$$V(\mu_{t-1}(i), \Omega_t) = \max_{\{C,NC\}} \left[V^{NC}(\mu_{t-1}(i), \Omega_t), V^C(\Omega_t) \right]. \tag{9}$$

2.4   The equilibrium

We consider rational expectations equilibria without aggregate uncertainty. The equilibrium is a set of policy rules $\{C_t(i,g), L_t, B_t\}$, price setting rules $\{P_t(i,g)\}$, prices $\{P_t, R_t, W_t\}$, policy shocks $\{\tau_t, M_t\}$ and adjusted relative price distributions $\{\Gamma_t\}$ for each t, such that

1. The representative consumer's policy functions $\{C_t(i,g), L_t, B_t\}$ maximize her utility function (1) given her budget constraint (2),

2. The firms' nominal price setting rule $\{P_t(i,g)\}$ maximizes their normalized value functions (9), (7), (8) and they determine production $\{Y_t(i,g)\}$ and labor demand $\{L_t(i,g)\}$ to satisfy demand. The firms form correct beliefs about the random process of the idiosyncratic shocks $\{A_t(i,g)\}$.

3. Money supply equals aggregate demand $M_t = P_t C_t$ in each period t. The seignorage revenue with the tax revenues are redistributed in a lump-sum way.

4. Goods markets $C_t(i,g) = Y_t(i,g)$, bond market $B_t = 0$ and labor market $\int \sum_{g=1}^G L_t(i,g)di = L_t$ clear in each period t.

5. The adjusted relative price distribution $\{\Gamma_t\}$ develops consistently with the idiosyncratic shock distribution and the price setting rules.

We solve for this equilibrium numerically with standard global solution methods. Our numerical solution algorithm for the steady state and for the transition path after an unexpected persistent money growth and pre-announced permanent tax shocks are detailed in the Appendix.

---

[12]The relative price that was inherited from this period if there was no price change, $\mu_{t-1}(i,g)/(1+\pi_t)$, needs to be adjusted with the next period's productivity innovations $e^{\varepsilon_{t+1}(i,g)} = A_{t+1}(i,g)/A_t(i,g)$, to obtain the relative price vector at the time of next period's price change decision. The expectation is taken over the future idiosyncratic state variables, conditional on their current values.

[13]The value function of the time-dependent pricing model of Calvo (1983) is

$$V(\mu_{t-1}(i), \Omega_t) = (1-\kappa)V^{NC}(\mu_{t-1}(i), \Omega_t) + \kappa V^C(\Omega_t),$$

where the menu cost is set to zero $\phi = 0$ and κ is now an exogenous probability of price change.

3   Calibration

We calibrate unobserved parameters of our model to match key moments of the steady state price change distribution.

3.1   The data

We use a data set of store-level price quotes in Hungary between December 2001 and December 2006 underpinning the consumer price index.[14] We focus on the processed food sector, for several reasons. First, this is the largest homogeneous sector in our sample, with a CPI-weight of 16.1%. Second, the moments of the price change distribution in this sector are essentially identical to the US retail store data used by Midrigan (2011) facilitating comparison. And third, in this sector, the composition of the subgroups facing the various value-added tax rates are very similar.[15] This permits us to disregard composition effects, and contributes to a clean comparison of the price effects of various tax changes.[16]

We observe the prices of 128 different products in our dataset,[17] and each product is observed in 123 stores on average each month.[18] We sales-filter our data to focus on regular prices, because – as Kehoe and Midrigan (2010) argues – they determine the aggregate behavior of the economy. To do this, we first, exclude price changes that are flagged as sales in the data. Then, we also drop any remaining price changes that are (1) at least 10 %, (2) and are completely reversed within 1 month.[19] The CPI micro-data we use records actual posted prices (not unit value indices as most barcode datasets), so it is less susceptible to the kind of measurement error analyzed in Eichenbaum, Jaimovich, Rebelo, and Smith (2014).[20] We mitigate the impact of potential measurement error further by dropping price change observations that are larger than

---

[14]For a detailed description of the data set, see Gabriel and Reiff (2010).

[15]An example of similar products facing different tax rates is 'cookies' facing the lower rate and 'chocolate-chip cookies' facing the higher. In the two subgroups, the steady-state yearly inflation rates are 4.3% vs. 4.1%; the steady-state monthly frequency of price changes are 11.5% vs. 12.9%; the average absolute size of price changes are 9.7% vs. 10.8%; the kurtosis are 3.98 vs. 3.96; and the interquartile range of price changes are 8.1% and 8.3%, respectively.

[16]We provide further evidence on the lack of the composition bias in the Appendix, by studying the inflation effects of yet another 5%-point VAT-increase in July 2009, which affected predominantly the products facing a VAT-decrease earlier.

[17]98 affected by the tax increase and 30 by the tax decrease

[18]The number of item replacements and substitutions are small.

[19]Midrigan (2011) introduces endogenous sales directly into his model, but reports moments of regular price changes. By disregarding sales, we make our model directly comparable also to the paper of Golosov and Lucas (2007) that does not model sales either.

[20]The processed food products we use are not considered 'problematic' by Eichenbaum, Jaimovich, Rebelo, and Smith (2014). The only processed food they consider problematic (cigarettes) are excluded from our sample, because they were subject to frequent large excise tax changes.

100% in absolute value.[21]

Our aim is to calculate key moments of a representative product. To do this, we first calculate each moment at the product level, and then aggregate them using the expenditure-based CPI-weights. This way, we avoid pooling potentially heterogeneous product level price change distributions.[22] The moments presented in the first column of Table 1 are sample averages excluding observations during tax-change months.

Table 1: Moments of regular price change distribution

Used in	Data		Models		
calibration	Baseline	Midr. (2011)	Mixed	Poisson	Normal
Frequency	12.6%	11.6%	12.6%	12.6%	12.6%
Size	9.9%	11%	9.9%	9.9%	9.9%
Kurtosis	3.98	4.02	3.98	3.98	*1.97*
Interquartile range	8.13%	8%	8.13%	*9.55%*	*6.3%*
Inflation	4.23%	0%	4.23%	4.23%	4.23%
Other moments					
1st decile of abs size distr	2.8%	3%	2.8%	2.9%	3.6%
1st quartile of abs size distr	4.6%	5%	4.9%	4.4%	6.8%
Median of abs size distr	7.7%	9%	7.5%	6.6%	10.0%
3rd quartile of abs size distr	12.7%	13%	13.0%	13.9%	13.1%
9th decile of abs size distr	19.4%	21%	20.8%	21.6%	15.7%
Fraction < 0.25 of mean	9.8%	8%	10.3%	9.9%	6.5%
Fraction < 0.5 of mean	27%	25%	26.8%	37.0%	13.4%

The second column of Table 1 confirms that our data is standard: our moments of the distribution of regular price changes mimics those observed in a major US grocery store chain Dominick's used by Midrigan (2011). In line with his findings, regular prices change infrequently: around once every 8 months. When they do, the price changes are sizeable: they are close to 10%. Furthermore, one can observe sizable dispersion of price changes: a lot of price changes as small (10% of the price changes are less than 3% in absolute value), but some price change are large (10% is larger than 19%). Furthermore, the difference between the first and the third quartile is more than 8% and the kurtosis of the price change distribution is around 4 that is

---

[21]Alvarez, Bihan, and Lippi (2014) suggest evidence that such trimming in CPI data brings measured moments closer to those observed in online data presumably less susceptible to measurement error.

[22]An alternative would be to calculate moments of a pooled distribution of price changes that are standardized at the product-store level (see Klenow and Kryvtsov (2008) and Midrigan (2011), Alvarez, Bihan, and Lippi (2014)). As standardization at the product level is naturally achieved by our method, we do not expect this choice to influence our results.

significantly larger than the kurtosis of the normal distribution (3).[23] One notable difference between the Hungarian and the US dataset is the level of the inflation rate. They are set to 0 in Midrigan (2011), while it is around a, still low, 4.2% in Hungary. Later we show that while inflation rate influences the level of price flexibility in our model, our qualitative results are insensitive to the level of observed inflation rate.

3.2 Calibration

We use standard values for some basic parameters. We set the monthly discount factor to $\beta = 0.96^{1/12}$, implying a 4% yearly real rate. We set the the inverse Frisch-elasticity of the labor supply (ψ) to zero, implying a perfect partial wage-elasticity of labor supply. Under this assumption, nominal wages move in lockstep with the money supply, similarly to the models of Golosov and Lucas (2007) and Midrigan (2011). This condition also generates full long term inflation pass-through of the value-added tax shocks that we also observe in the data. We set the money growth ($\mu_M/(1 - \rho_M)$) equal to the yearly observed trend inflation rate, 4.2%, but also show predictions for different inflation rates (in particular for 2%, 0%) in Section 5. The persistence of the money shock (ρ_M) is calibrated to be 0.61 and its standard deviation to 0.18% as in Midrigan (2011).

We set the elasticity of substitution between products of different retail firms (θ) to be 5.[24] In line with Midrigan (2011), we assume each firm produces 2 goods ($G = 2$), and we set the elasticity of substitution for goods within a firm to be close to 1, in line with the observation that they are not close substitutes. We set the persistence of the idiosyncratic technology shock ρ_A to unity, the value chosen by Gertler and Leahy (2008) and Midrigan (2011).[25]We assume that the within-firm correlation between the idiosyncratic shocks are 60% that is also in line with the calibration of Midrigan (2011).[26]

We calibrate the menu cost and the parameters determining the shape of the idiosyncratic shock distribution to match key moments of the steady state price change distribution. In our baseline model, we calibrate the menu cost parameter ϕ and the standard deviation of the idiosyncratic productivity shocks σ_A to match the frequency and the average size of absolute price changes. We calibrate the probability of low variance shock p to match the kurtosis of the price change distribution, and the variance-proportionality parameter λ to match the interquartile range of the absolute price change distribution. The four moments exactly identify our four free parameters.

---

[23]**?**) finds that a kurtosis of 4 is a reasonable kurtosis estimate in the US for a general set of products, including services.

[24]This is an intermediate value between those used by Midrigan (2011) (3) and Golosov and Lucas (2007) (7). The value influences the estimates of menu costs and the standard deviation of idiosyncratic shocks, but it has no significant influence on our aggregate conclusions.

[25]We show in Section 5, that our conclusions are unchanged with a persistence parameter of 0.7.

[26]In Section 5, we show that our results still hold when this correlation is set to 0.

The alternative models are special cases of our baseline calibration, with less free parameters. We calibrate these by exactly matching moments that are sufficient for their identification. For the Poisson distribution of Midrigan (2011), we set the λ parameter to 0 and match the frequency, size and kurtosis of the price change distribution. For the Gaussian distribution of Golosov and Lucas (2007), we set the λ parameter to 1 and the p parameter to 0 and match the frequency and the size of the price change distribution with the menu cost and the standard deviation parameters.

Table 2 presents the calibrated parameters of our baseline model (mixed) and the alternatives.

Table 2: Calibration

Parameters	Mixed	Poisson	Normal
ϕ	2.4%	1.6%	5.0%
σ_A	4.3%	4.4%	3.8%
p	91.2%	90.6%	0
λ	8.8%	0	1

As Table 2 shows, our baseline parametrization implies a menu cost of 2.4% of steady-state revenues, paid with a 12.5%/2 probability (in a two-product framework, half of the price changes are free). So the overall cost of regular price adjustment is around 0.15% of steady-state revenues. We consider it of a reasonable magnitude. Levy et al. (1997), for example, estimate the costs related to price adjustment in supermarkets to be 0.7% of revenues, which is higher than our estimates, but their measure includes the costs related to sales that comprise most of the observed price changes in supermarkets. Similarly to previous quantitative menu cost models with idiosyncratic shocks, the model needs volatile idiosyncratic shocks to hit the large average absolute size of the price changes. The calibrated standard deviation is 4.3% (per month). The data favors a mixed normal idiosyncratic shock distribution where the low volatility shock is arriving with a high 91.2% probability and with a standard deviation that is only 8.8% of the high volatility shock.

Assuming a Poisson distribution as Midrigan (2011) or a Gaussian distribution as Golosov and Lucas (2007) have notable effects on the calibrated menu costs, but minor effects on the calibrated standard deviations.[27] The Poisson distribution implies a probability of 0 idiosyncratic shock that is very close to our baseline (90.6% vs. 91.2%), but disallowing any variation in the low volatility shock results in a menu cost estimate that is 33% lower than our baseline. The reason is that smaller menu costs are sufficient to keep the frequency of steady state price changes on target. Not so with the Gaussian distribution. There, the underestimated fraction

---

[27]Note that while these unconditional standard deviations are rather similar in the different menu cost models, the conditional standard deviation of a non-zero shock in the Poisson case is much larger, $\sigma_A/\sqrt{1-p} = 14\%$. Similarly, the productivity shock innovation has a standard deviation of $\sigma = \sigma_A/\sqrt{(p\lambda^2 + 1 - p)} = 14\%$ in the large standard-deviation case in the mixed normal model.

of small price changes requires menu costs that are more than twice as high as our baseline to keep the steady state frequency at bay.

The last three columns of Table 1 shows the ability of our baseline model to match moments of the steady state price change distribution. The table confirms that we match all the targeted moments, and with this we are also very close to match the dispersion measured by the quartiles and the first and last decile of the absolute price change distribution. Also, we come very close to matching the proportion of small price changes, here measured as the fraction of price changes smaller than the quarter and the half of the mean absolute price change, respectively. In this, we are doing better than the alternatives, that should not be surprising given we have extra parameters in matching key moments.

Using the Poisson distribution, the steady state price change distribution somewhat over-estimates the observed interquartile range (9.35% instead of the 8.1%). The reason is that it marginally overestimates the number of small price changes as confirmed by the slightly under-estimated first decile (2.4% instead of the 2.8% in the data) and the somewhat overestimated fraction of changes that are smaller than half of the mean change (34% instead of the 27% in the data). The differences between the data and the moments using the Poisson distribution are small. Matching the observed price dispersion is much less successful with the normal distribution as in Golosov and Lucas (2007), even though here, differently from the original model, we have small price changes as a result of our two-product assumption. The model underestimates both the kurtosis (1.9 instead of 3.98 in the data) and the interquartile range of the absolute price change sizes (6.3% instead of 8.1%), and underestimates the fraction of small price changes (13.4% of less than half of the mean, instead of 27% in the data).

All this can be seen on Figure 1 that shows a histogram of price changes in the data and those implied by the models. It confirms that our baseline model does a very good job at matching the distribution. From this figure, its shape is very close to that of the Poisson distribution, and they both match the observed price changes much better than the implied price changes of the model using the normal distribution.

3.3 Macro implications

As we have shown in the previous section, our baseline model is able to match the frequency, the average size and the dispersion of steady state price changes well, even marginally better than the Poisson distribution used by Midrigan (2011). In this section, we show that its macro implications are very different from the model of Midrigan (2011), and actually closer to that of Golosov and Lucas (2007). In particular, our model implies *small* real effects of standard monetary policy shocks.

To show this, we take our calibrated models, and hit them with an unexpected persistent money growth shock. To maintain comparability with Midrigan (2011), we set the inflation rate

Figure 1: Steady state distribution of price changes
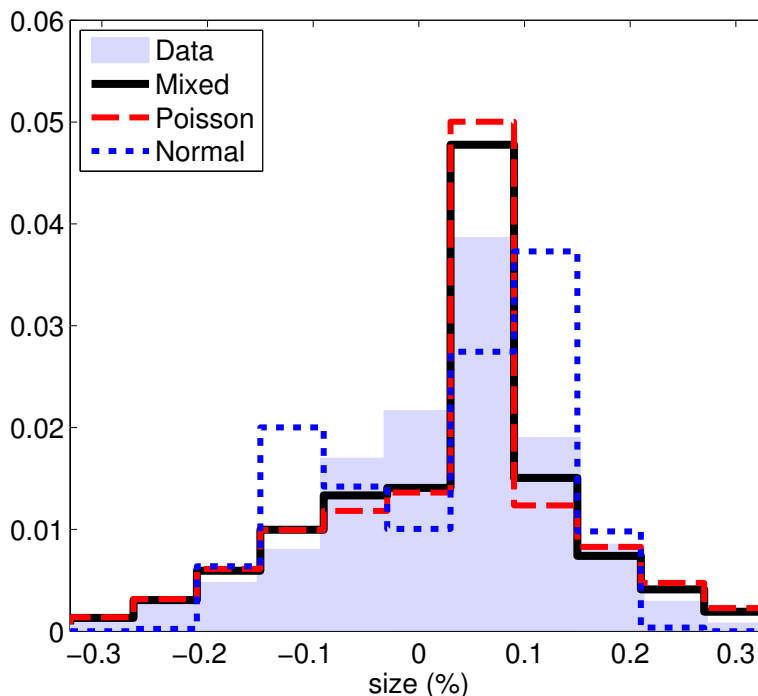
The figure plots the steady state price change distributions in the data and in different menu cost models. The baseline (mixed) and the Poisson model are similarly good at matching the observed distribution (see also Table 1), while the normal model is less successful in matching the dispersion of the price changes.

to 0.[28] Later, we show how the inflation rate influences the implications of the model. Figure 2 presents the impulse responses.

We measure the pass-through as the proportion of the price level effect of the shock over the cumulated change of the money supply, or formally as

$$\gamma_t = \frac{\sum_{i=0}^{t} \pi_i - \bar{\pi}}{\sum_{i=0}^{t} \Delta m_i}. \tag{10}$$

The real effect is measured as the output effect of the shock. The figure confirms that our baseline model implies small and temporary real effects (the cumulative real effects are the multiple of 1.59 of the monetary shock). It is close to the implications of the model using the Gaussian distribution as in Golosov and Lucas (2007) (0.65). Its implications are very far from those of the model with Poisson distribution (7.60) that generates real effects close to that of

---

[28]Note that setting the inflation rate to zero for calibrations made at positive inflation rates is not equivalent to recalibrating the models with 0 inflation. The steady state moments of the models change. The frequency, for example, of our baseline model is 11%; while it is 8.9% with Poisson and 12.4% with Gaussian shocks. For maintaining comparability to the Calvo model, we reset its frequency to the 8.9% of the Poisson case. In the robustness section, we show that recalibrating the models to zero inflation does not influence our main conclusions.

Figure 2: Impulse responses to a monetary shock
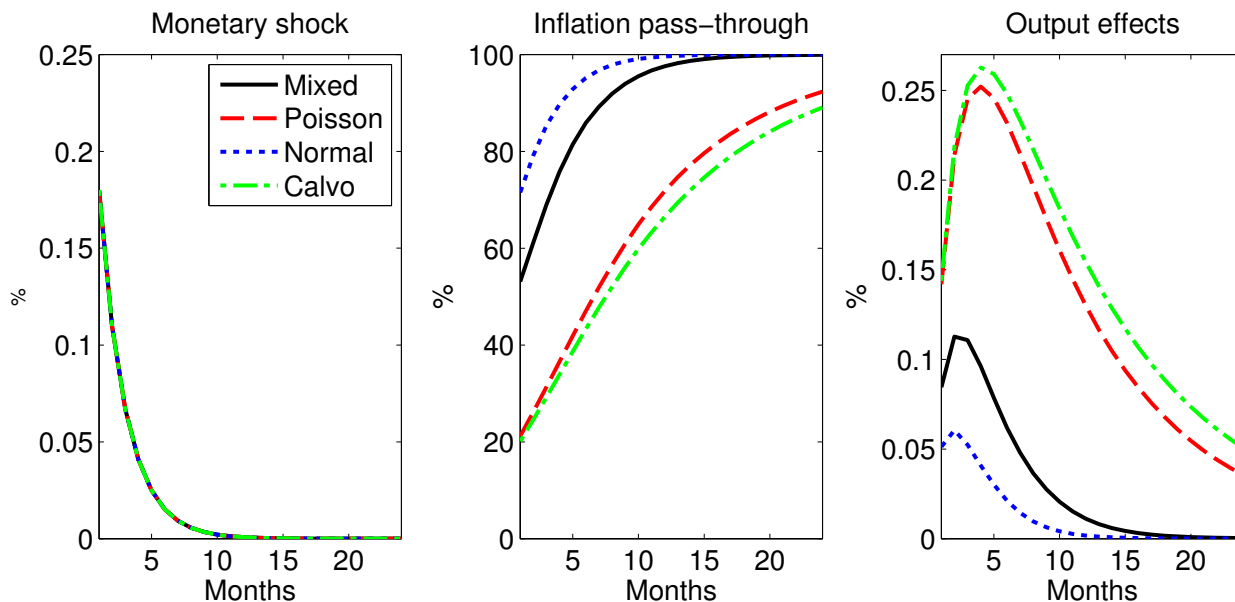
The figure plots the impulse responses of the baseline multi-product model (mixed) and alternatives to a one standard deviation persistent money growth shock. The figure shows that the baseline model implies small and temporary real effects (the cumulative real effects are the multiple of 1.59 of the monetary shock). They are substantially lower than the Calvo (1983) (8.93) and the Poisson model of Midrigan (2011) (7.60), and only somewhat larger than the model with the Normal distribution of Golosov and Lucas (2007) (0.65).

the Calvo (1983) model (8.93). These results imply that the macro-implications of these models are highly sensitive to the shape of the idiosyncratic shock distribution.

What drives these results? The response of the aggregate price level to the monetary policy shock can be decomposed into three key margins of adjustment. The intensive margin characterizes adjustment in the size of the price changes, the extensive margin characterizes the change in the fraction of adjusting firms, and the selection effect measures variation in the composition of the adjusters. The composition of adjusters might matter, because, as Caplin and Spulber (1987) and Golosov and Lucas (2007) argue, if the new adjusters are ready to change their prices by a wide margin, they can make the price level highly flexible. We are about to show that it is the magnitude of this selection effect that is highly sensitive to the idiosyncratic shock distribution.

To derive the decomposition formally, it is instructive to derive the inflation rate as a function of the desired price changes (the adjustments that firms would make if menu costs disappeared for a period x), their hazard function $h(x)$ and distribution $f(x)$ (see e.g. Caballero and Engel

(2007), Midrigan (2011) and Costain and Nakov (2011)).[29] The inflation rate equals

$$\pi = \int x h(x) f(x) dx, \tag{11}$$

and the impact of an aggregate shock in period t can be decomposed into an intensive, extensive and selection margins, as

$$\pi_t - \pi = \underbrace{\overline{h(x)} \Delta \overline{x}}_{\text{intensive}} + \underbrace{\Delta \overline{h(x)} \overline{x} + \Delta \overline{h(x)} \Delta \overline{x}}_{\text{extensive}} + \underbrace{\Delta \int_x x \left(h(x) - \overline{h(x)} \right) f(x) dx}_{\text{selection}}, \tag{12}$$

where an average of variable y denoted by $\overline{y} = \int x f(x) dx$, and Δy denotes difference from the steady state $y_t - y$.

The intensive margin is the product of the average frequency and the change in the average desired price change: in a Calvo-model with fixed frequency and random selection, this would be the only component of the pass-through. The extensive margin is defined here as the aggregate effects caused by the changes in the price change frequency. It is the sum of two products: the product of the frequency increase and the average desired price change and the product of the frequency change and the average desired price change. For small shocks that do not influence the aggregate frequency, this term is negligible (but it can play an important role for large shocks with large frequency effect as we show later). The third factor is the selection effect, coming from the fact that 'new' price changers are going to have higher than average desired price changes. The third term expresses this by measuring the increased correlation between the desired price change and the adjustment hazard after a shock.

To obtain a single price-flexibility measure from impulse responses, we calculate an average shock pass-through. For this, we first calculate a 'marginal' pass-through (γ_t^m) for each period t, measured as the inflation effect in period t relative to the fraction of cumulative money shocks yet to be passed through:

$$\gamma_t^m = \frac{\pi_t - \bar{\pi}}{\left(\sum_{i=0}^t \Delta m_i - \sum_{i=0}^{t-1} (\pi_i - \bar{\pi}) \right)}. \tag{13}$$

We weight these marginal pass-throughs based on their relative size ($\omega_t = (\pi_t - \bar{\pi})/(\Delta m/(1-\rho))$) to arrive at a measure of the level of aggregate price flexibility:

$$\bar{\gamma} = \sum_{t=1}^T \omega_t \gamma_t^m. \tag{14}$$

Note that in case of full price flexibility, this weighted average marginal pass-through is one, and in the Calvo model, it equals to the periodic marginal pass-through, that is constant over time, and for marginal aggregate shocks equals to the price change frequency.

---

[29]The desired price changes are functions of the individual state variable p_{-1} (the last period's quality-adjusted relative price) and aggregate states that we suppress for notational convenience.

The impulse responses shown on Figure 2 imply a price flexibility measure ($\bar{\gamma}$) of 40% in our baseline mixed normal model, and 12%, 62% and 11% in the Poisson, Gaussian and Calvo models respectively. We find that the difference across models is driven by the difference in the selection effect. For the small money growth shocks, extensive margin effect is negligible. The intensive margin effect, in turn is very close in each models, the differences coming only from the variation in the frequency of price changes at 0 inflation rates (they are 15.5% in our baseline and 17.5% in the Gaussian models, but only 11% with Poisson and Calvo). The selection effects, however, are very different: contributing to the measured pass-through by 24%-points in our baseline mixed normal case and 44.5% in the Gaussian case, while they are only 1% in the Poisson case. What explains this difference?

We illustrate the sensitivity of the selection effect on the shape of the idiosyncratic shock distribution by using a random menu cost model (Dotsey, King, and Wolman (1999)). This is less realistic than our baseline, but substantially more tractable. Instead of the multi-product assumption, we assume here that each firm produces a single product $G = 1$, but with a certain probability q it can change its price for free ($\phi_i = \phi$ with probability $1 - q$ and 0 with probability q). To mimic our baseline two-product assumption, we calibrate the probability of free menu cost to $q = 6.29\%$ that is the half of the observed frequency. This way, half of the observed price changes are going to be 'free', just as in our multi-product setting. The main advantage of this simplification is that we lose a state variable (the second relative price), relieving us from a large fraction of the computational burden. The constant probability of free price adjustment, furthermore, is a valid assumption in an environment with small shocks, like the one standard deviation money growth shock, which does not influence the price-change frequency as argued also by Midrigan (2011).[30] The key difference from the multi-product setting is the correlation between the idiosyncratic shocks of the price changes with and without menu costs. In this model, the free adjusters are independent of the costly adjusters, while in the multi-product setting their idiosyncratic shocks are correlated with a coefficient $\rho_\varepsilon = 60\%$. As a result, the steady state price change distribution of this model is less realistic than our baseline model. The key results, however, go through: the model with the mixed normal distribution has real effects close to that using the normal distribution, while the model with the Poisson distribution is very close to the Calvo (1983) model (for details, see Section 5).

The key parameter of our mixed normal distribution is the relative standard deviation of the small variance distribution relative to the large variance distribution (λ). For $\lambda = 0$, we get the Poisson distribution and for $\lambda = 1$ we get the Gaussian distribution. The parameter best matching our observed steady state price change distribution in this random menu cost model is

---

[30]It limits the model's applicability for large shocks with influence on the price-change frequency, however, because the assumption unrealistically fixes the absolute number of 'free' adjustments, instead of allowing it to increase proportionally to the price change frequency as a response to the large shock, as in the multi-product setting.

18%.[31]. For each λ value, we recalibrate the model to match the frequency, size and kurtosis of the steady state price change distribution,[32] and simulate one standard deviation money growth shocks to calculate the price flexibility measure ($\bar{\gamma}$) – as a proxy for the speed of pass-through –, its components and the real effects of monetary policy shocks to our distributional assumptions. Figure 3 shows our results.

Figure 3: Idiosyncratic distribution and real effects
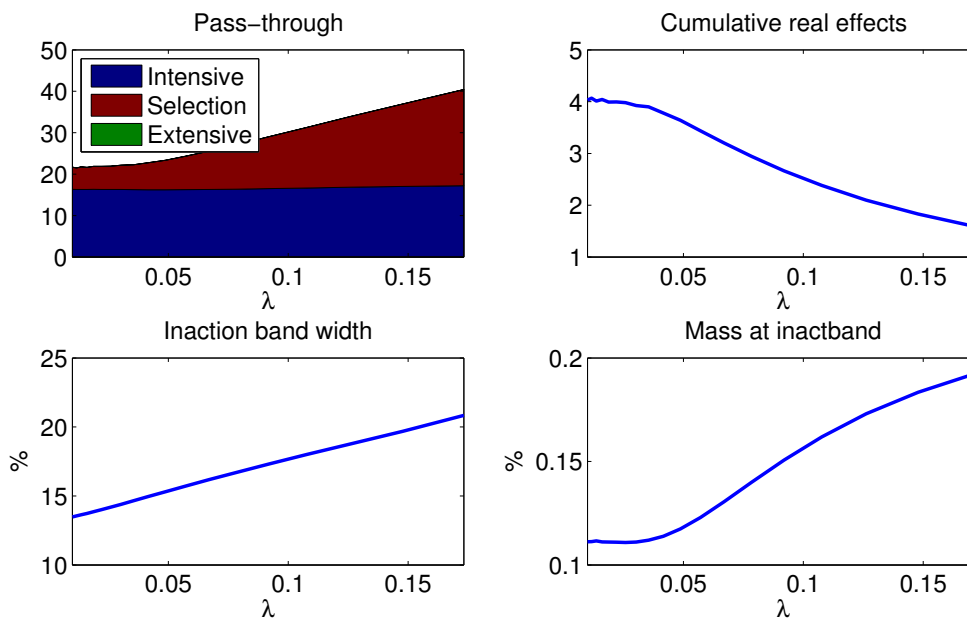
The figure shows the implied price level flexibility and real effects as a function of the shape of the idiosyncratic shock distribution. The parameter λ measures the relative variability of the small volatility shock relative to the large volatility shock in the mixed normal distribution. $\lambda = 0$ is the Poisson, $\lambda = 1$ is the Gaussian case. The figure only plots λ values where there exist parameters combinations matching the observed frequency, size and kurtosis of the steady state price change distribution. This was impossible for $\lambda > 0.2$ (note that the Gaussian distribution does not match the kurtosis). The figure confirms that small deviations from the Poisson distribution quickly leads to substantially higher price level flexibility and low real effects. The source of the increase is the selection effect.

The first panel of the figure shows our price level flexibility measure, the average marginal pass-through for various idiosyncratic shock distributions. It confirms that the flexibility is very sensitive to small deviations from the Poisson distribution, and the main driving force of the flexibility increase is the selection effect. The second panel on the figure shows the cumulative real effects of the baseline shock relative to the cumulative shock effect. The figure shows that

---

[31]Somewhat higher than 8.8% in our baseline case, but still much closer to the Poisson case (0) than to the Normal (1)

[32]We do this with positive inflation as in our baseline case, switching inflation off for the shock experiments. The results are not sensitive to doing this exercise with calibration to 0 inflation.

the cumulative real effects for very small λ values are high (the same measure in the Calvo (1983) model is close to 6), but it quickly decreases to much smaller levels. The exercise is interesting in light of recent results in the literature. First it shows that in the presence of infrequent and large idiosyncratic jumps (brought about by the large variance component of our mixed normal distribution), the same frequency, size and kurtosis can imply a very wide range of potential monetary non-neutralities, and, as concluded also by Alvarez, Bihan, and Lippi (2014), further information is required to identify the realistic extent of price flexibility. Second, the example shows that there is no discontinuity in the extent of real effects as the idiosyncratic distribution approaches the discontinuous Poisson distribution (as $\lambda \to 0$). This suggests that the Poisson distribution is not a knife-edge case as Elsby and Michaels (2014) implicitly imply and for realistic menu costs away from 0, there is a continuum of continuous distributions around the Poisson that generate real effects not dissimilar from Calvo (1983). The third and the fourth panels of the figure show the width of the inaction band and the mass of firms at the inaction thresholds, both increasing significantly with the λ parameter. We turn to their role in influencing the selection effect right now.

For this, it is instructive to look at an even simpler, but closely related model, the standard single-product version of the menu cost model ($G = 1$, $\phi_i = \phi$, $\bar{\pi} = 0$). The key channels in this model are very similar to the more realistic versions, but the intuition is much cleaner. The impact of a marginal monetary policy shock ($\Delta m \to 0$) can be expressed as

$$\frac{\pi}{\Delta m} = \underbrace{\int_x f(x)h(x)dx}_{\text{intensive}} + \underbrace{\int_x xf(x)h'(x)dx}_{\text{selection}}.$$

The adjusting firms increase their desired price change by the unit of the monetary shock generating an intensive margin effect. The selection effect is determined by the changing composition of price adjusters ($h'(x)$).[33] In a single-product model with permanent idiosyncratic shocks, the hazard function takes a particularly simple Ss form: firms have a lower and an upper adjustment thresholds: they will keep their prices unchanged if their desired price change is within this inaction band, and they will adjust with probability 1 if it is outside. This implies that $h'(x)$ is exactly 1 at the adjustment thresholds and 0 otherwise. The selection effect, thereby equals the product of the desired price changes at these inaction thresholds (half the width of the inaction bands in the symmetric case as this), and the mass of products at the inaction thresholds.

Figure 4 shows the desired price change distributions and the inaction bands for a single-product model calibrated to match the same steady state moments as in the multi-product experiment with our baseline mixed normal distribution and the alternatives.

The shape of the idiosyncratic distribution has a strong influence on the mass of firms around the adjustment threshold, and through its influence on the calibrated menu cost, it also influences

---

[33]The extensive margin effect will be zero on the margin, because the measure of new price increasing firms and the measure of new non-price-decreasing firms cancel each other out, keeping the measure of price adjustments constant.

Figure 4: Desired price change distributions and inaction bands
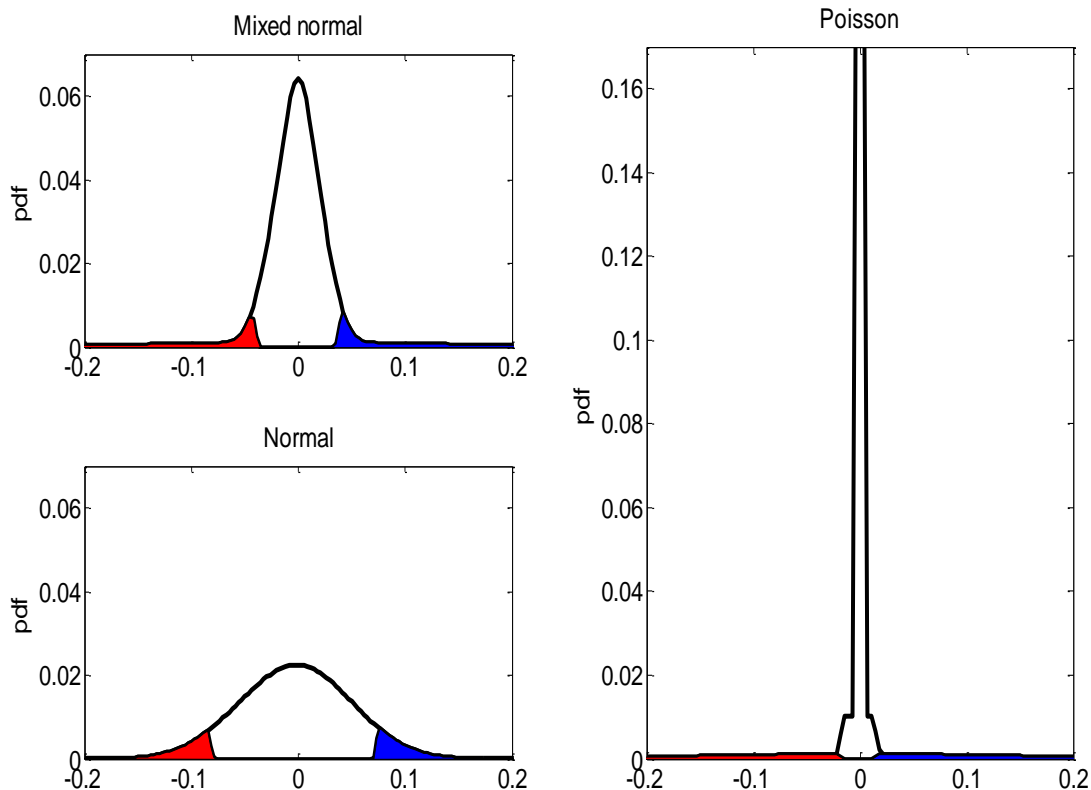
The figure plots steady state desired price change distributions and inaction bands obtained from calibrated single-product models. The shaded areas show the implied actual price change distribution with price decreases on the left and price increases on the right.

the width of the inaction band. Both the measure of firms at the inaction thresholds and the width of the inaction band are high in the model with Gaussian idiosyncratic shocks as in Golosov and Lucas (2007). This results in large selection effects. In contrast, the Poisson distribution of Midrigan (2011) implies both a low mass at the inaction thresholds and a narrow inaction band. As a result, the selection effect here is particularly low, close to the Calvo (1983) model. As our baseline model with mixed normal distribution deviates from the Poisson distribution of Midrigan (2011) (as λ increases), both the width of the inaction band and the mass of firms around the inaction threshold increase substantially (see panel 3 and 4 on Figure 3), reinforcing each other's effect. As a result, relatively small deviation from the Poisson distribution will have significant effect on the selection, increasing the inflation pass-through quickly close to the

Golosov and Lucas (2007) model.

4   Large shocks

In the previous section, we have shown that the shape of the idiosyncratic shock distribution and the resulting calibration of the menu cost determine the macroeconomic implications of menu cost models. We have argued that the empirically observable steady state price change distribution contains insufficient information for valid identification, by showing that our baseline model with a mixed normal distribution can match the steady state distribution as well as a Poisson distribution used by Midrigan (2011), but implies small real effects of monetary shocks contrary to his results. In this section we provide supporting evidence for our model and in particular our assumption of mixed normal idiosyncratic distribution. Our evidence is coming from pricing responses to large, aggregate nominal shocks, which we argue can generate sufficient information for identification of both the idiosyncratic shock distribution and of the magnitude of menu costs; and also provide evidence for the menu cost approach against popular alternatives.

Why are large aggregate shocks informative? First, the number of adjusting firms as a response to a shock of a known size contains a lot of information about the reasonable magnitude of menu costs. A too small menu cost would imply too many, while a too large menu cost would imply too few price changes. Consider Figure 5 that plots the effects of permanent shocks of various sizes in our calibrated baseline model and the alternatives. The figure confirms that large shocks make the price level more flexible (first panel, measured by the average pass-through $\bar{\gamma}$) and the cumulative real effects smaller (second panel) by increasing the fraction of adjusting firms (third panel). Importantly, the calibrated models have highly heterogeneous frequency responses to similar sized shocks partly because of their different calibrated menu costs, facilitating their comparison.

Second, the observed distribution of price changes after a large aggregate shock directly reveals the 'desired' price change distribution. This is not the case with small aggregate shocks, because then the menu cost prevents the realization of desired price changes, especially the small ones. As the distribution of desired price changes is in tight relationship with the idiosyncratic shock distribution, the observed distribution at the month of large aggregate shocks contains key information about the underlying idiosyncratic shock distribution.

The aggregate shocks we are utilizing are changes in value-added tax rates. They are suitable for our purposes, because they are easily measurable and transparent aggregate nominal cost shocks influencing the optimal gross prices of all affected firms. Similarly to other European countries, in Hungary, where our data is coming from, gross prices are posted, so price adjustment requires paying the menu costs – this is different from the standard practice in the US, where sales taxes are only added ex post at the counter.[34] The tax changes are the consequence of the

---

[34]This practice is also reinforced by a consumer protection law requiring that "consumers can not be forced to calculate prices in their head". 1997. CLV. Law on Consumer Protection and 7/2001. (III. 29.) Ministry of

Figure 5: Effects of various shock sizes
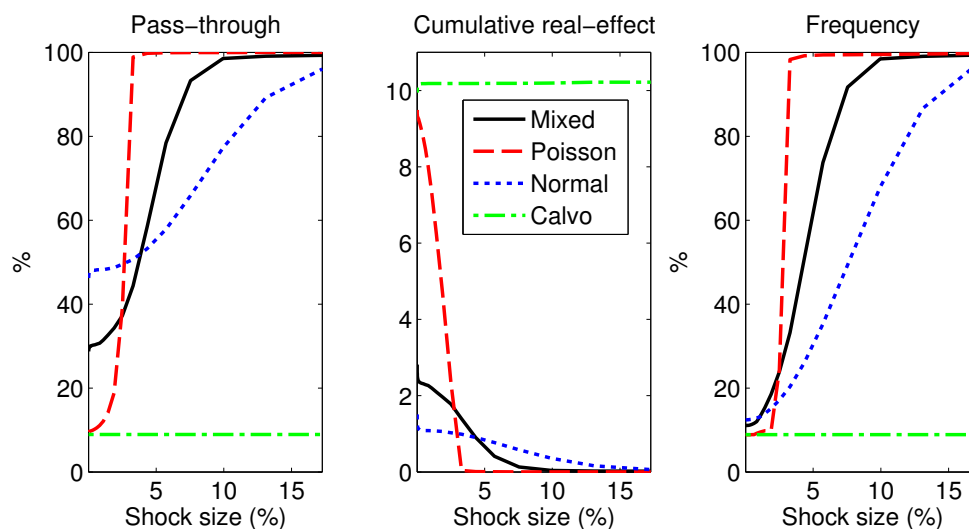
The figure plots the average inflation pass-through ($\bar{\gamma}$), the cumulative real effects and the frequency responses to permanent shocks of various sizes in different models for 0 inflation rate. Intuitively, price level becomes more flexible and real effects become smaller with larger shock sizes. The results show different sensitivity to shock sizes that is especially prevalent with the Poisson model: while for small shocks it has price rigidity (first panel) and real effects (second panel) as high as the Calvo (1983) model, for large shocks it generates price flexibility even higher, and real effects even lower than our baseline. The third panel on the figure also confirms that the models have heterogeneous frequency responses. As frequency is readily observable, it allows direct testing of the model assumptions.

Hungarian authorities sequentially closing the gap between two different value-added tax rates in 2004 and 2006 to simplify the tax code. In January 2004, they increased the preferential rate from 12% to 15%. Then, in January 2006, before the general elections, they decreased the standard rate from 25% to 20%, and in September, after the election, they increased the preferential rate again from 15% to 20%. These changes were announced 5, 6 and 3 months in advance, respectively, which we take into account in or model simulations. All of the tax changes can be considered permanent over any reasonable price-setting horizon, and were also taken as such by the public. [35] As we show in the Appendix, during this period the country has enjoyed a stable macroeconomic environment with stable growth, low inflation rates and stable exchange rates. Monetary policy, furthermore, refrained from responding to the tax shocks: the inflation targeting central bank had expressed in advance that it was "seeing through" the direct


Economy decree and its explanation.

[35]The VAT-decrease of January 2006 was ultimately reversed in 2009, when the government has increased the 20% rate to 25% to finance its budget shortfall following a financial crisis related drop in its revenues. In the Appendix we describe the pricing implications of this VAT change. We do not directly include this VAT change in our main analysis, because of the extreme cyclical downturn during its implementation.

effects of the tax shocks, because they only affected the price level with only temporary effects on the measured inflation rate.

4.1 Aggregate effects

Figure 6: Monthly inflation (processed food sector, Hungary)
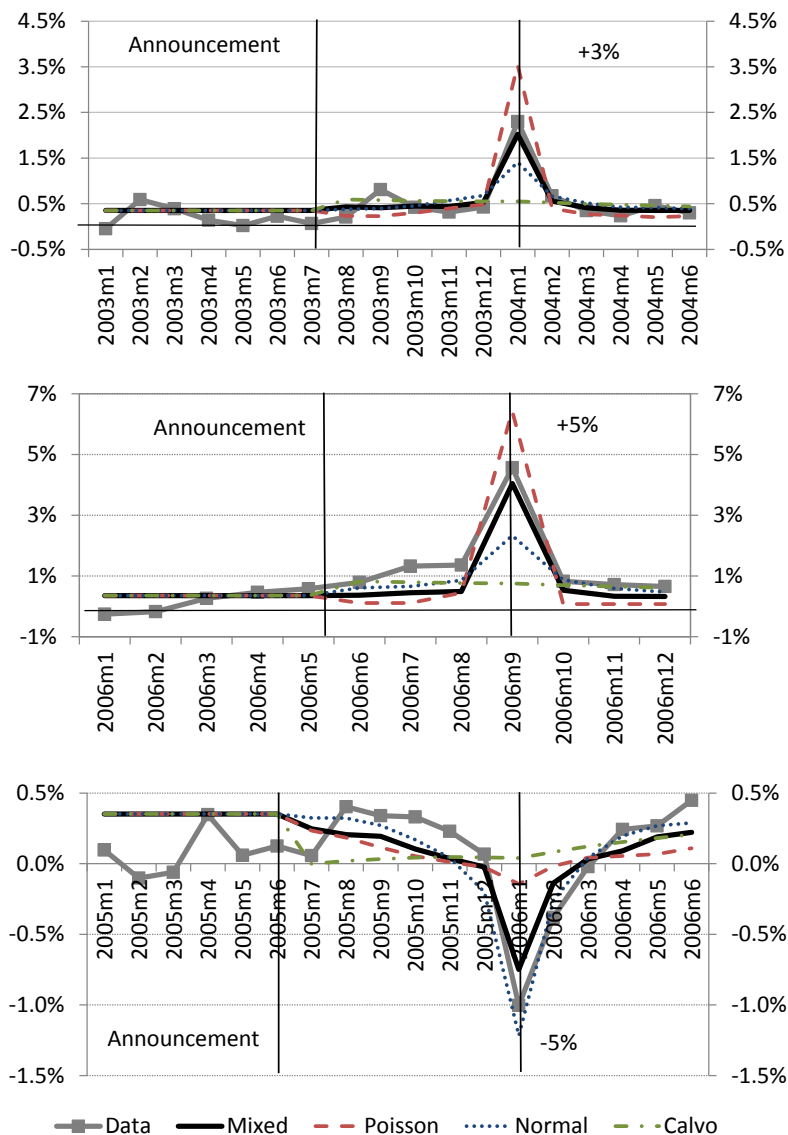
The figures plots the monthly inflation rate in the processed food sector in the data and in different menu cost models. The baseline model (mixed) is successful in matching the changes in the inflation rates. In contrast, the Poisson model overestimates, and the normal model underestimates the inflation effects of the shocks (see also Table 3). In the Calvo model, very differently from the data and the menu cost models the inflation effect peaks right after the announcement: the adjusting firms respond to the expected tax changes right away.

Figure 6 plots the dynamics of the monthly inflation rate in the affected groups of the processed food sector in the data, and the first column of Table 3 presents the measured inflation pass-through $(\pi_t - \bar{\pi})/\Delta\tau_t$ and the frequency of price changes of the affected firms at the months of the tax changes. Two observations stand out. First, there was a substantial adjustment on the extensive margin. In contrast to the 12.6% regular price change frequency, 52% and 62% of firms adjusted within one month of the positive 3% and 5%-point changes, respectively.[36] As a result, we observed large immediate inflation pass-through of the shocks (74% and 99%). Second, the effects of the positive and negative shocks were asymmetric: the frequency response to a -5% shock was only 27% (compared to the 62% to the similar sized positive shock), and its inflation effect was only 33% (compared to the 99% of the positive shock).

Both of these effects are in line with the predictions of menu cost models. First, they predict frequency increase, as the large shock makes the gains from price adjustment overcome the menu cost for a large fraction of the firms. Second, they predict asymmetry under trend inflation (see e.g. Ball and Mankiw (1994)), because, after a negative shock, firms can save their menu cost by refraining from adjustment and waiting for the trend inflation to adjust their relative prices for free. The quantitative predictions of the various menu cost models, however, depend crucially on their particular assumptions.

Table 3: Moments of price adjustment at the months of tax changes

Moment	Size	Data	Mixed	Poisson	Normal	Calvo
	+3%	74%	64%	127%	41%	7.0%
Inflation pass through	+5%	99%	88%	142%	49%	8.0%
	-5%	33%	27%	12%	39%	6.6%
	+3%	52%	32%	60%	18%	12.6%
Frequency	+5%	62%	55%	90%	25%	12.6%
	-5%	27%	19%	11%	17%	12.6%

We use these observations to test the predictions of our model. For this, we take our baseline model, calibrated to match the steady-state price change distribution, and hit it with permanent tax shocks that are announced before their implementation.[37] We do a similar experiment with

---

[36]Gagnon (2009) documents similarly large frequency effect after large nominal shocks in Mexico.

[37]Disregarding preannouncement would not qualitatively change our conclusions, but it would have sizable quantitative impact on the results. Its effect is particularly apparent in terms of the negative aggregate shocks: disregarding preannouncement would substantially increase the pass-through of the shock in the month of the tax decrease (to 48% in our baseline mixed normal case from the current 27%, while the impact of positive shocks would essentially be unchanged). What drives these results? It is predominantly related to the positive trend inflation that reduces the fraction of adjusting firms in the month of the negative shock. Firms foresee this, and as a result, they start adjusting to the shock already before it hits by and postponing costly price increases and reducing the magnitude of their price changes when they occur. In contrast, the preannouncement of a

alternative menu cost models with Poisson and Gaussian distributions and the Calvo (1983) model. The size of the shocks are calibrated to match the size of the tax changes, and the lags between the announcement and the tax changes are similarly calibrated to be in line with reality.

Figure 6 also presents the simulated inflation paths in our baseline mixed normal model and in the alternatives. The baseline model is successful in matching the changes in the inflation rates and also broadly consistent with its dynamics. In contrast, the Poisson model overestimates, and the normal model underestimates the inflation effects of the shocks (see also Table 3). In the Calvo model, very differently from the data and the menu cost models, the inflation effect peaks right after the announcement: the adjusting firms respond to the expected tax changes right away. The second column of Table 3 concurs that our baseline model predicts remarkably well the observed pricing effects of the tax shocks. First, we come very close in hitting the inflation pass-throughs observed during the tax changing months (the model predicts 67%, 91% and 28% versus the observed 74%, 99% and 33% rates for the 3%, 5% and -5% shocks). The model achieves this by mainly matching the large frequency increases well, except during the 3% shock, when it somewhat underestimates the frequency change. Second, the baseline model also does quite well in generating substantial asymmetry in the inflation pass-through by predicting three times larger pass-through for the positive shock than for the negative one (91% vs 28%), similarly to the data. We consider this all the more remarkable, as we do not have any free parameters to influence the behavior of the model when the shocks hit. The results provide strong support to the claim that the calibrated menu cost of our baseline model is of the right magnitude.

Alternative menu cost models are unable to match the data moments, as columns 3-5 of Table 3 and Figure 6 show. Using the Poisson distribution we substantially overestimate the frequency response and the pass-through of the positive shocks, and also overestimate the asymmetry between the effects of similar sized positive and negative shocks. The normal distribution, in contrast, underestimates the frequency and pass-through of the positive shocks, and underestimates the asymmetry. Needless to say, the Calvo model has no chance of matching the large frequency response and the pass-through because of its hard-wired assumption of constant price-change frequency.
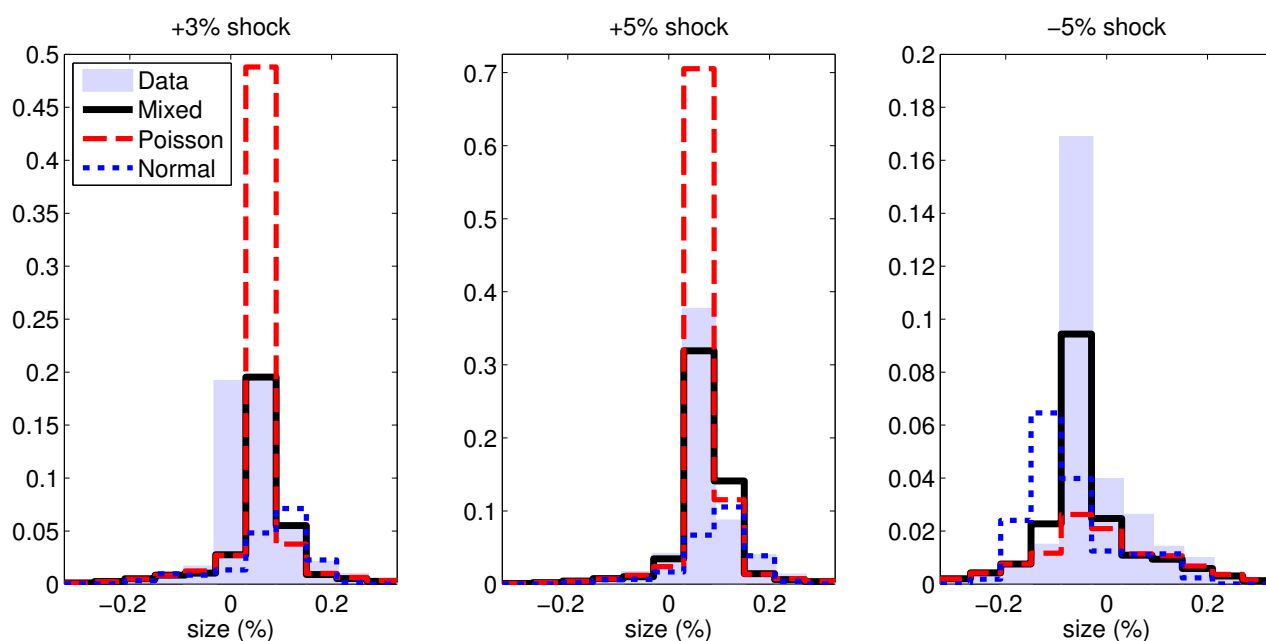
4.2   Distribution

As we argued earlier, the distribution of price changes in the months of tax changes reveal valuable information about the idiosyncratic shock distribution. The shaded areas on Figure 7

---

positive shock brings about less adjustment before the shock, because firms foresee that with high probability they are going to adjust in the month of the tax increase. Their reduced planning horizon, furthermore, might even limit their incentives to respond to idiosyncratic shocks (see Hobijn, Ravenna, and Tambalotti (2006)). Preannouncement in the Calvo model brings about a sizeable announcement effect (look at Figure 6) differently from the menu cost models, because the probability of adjustment in the month of the shock there is fixed at a low level, so firms start adjusting to it whenever they have the opportunity.

show the realized distribution of non-zero price changes at the months of the tax changes. The histograms are normalized to sum to the observed fraction of price changes during each particular tax changing months. The observed distributions show excess kurtosis, each with a sharp peak and fat-tails. The figures also show the predicted histograms calculated from our baseline model and the alternatives. The figures show that our baseline model is fairly successful in matching the distributions. The figures suggest, however, that the Poisson model overestimates, while the Gaussian model underestimates the kurtosis of the realized distributions.

Figure 7: Distribution of price changes at the months of tax changes

The figure plots the price change distributions in the data and in different menu cost models at the months of the tax changes. The baseline model (mixed) is successful in matching the frequency and the distribution of price changes. In contrast, the Poisson model overestimates the frequency and overestimates the kurtosis of the price change distribution, and the normal model underestimates the frequency and the kurtosis (see also Table 4).

The moments of the price change distributions measured during the tax change months confirm these observations. The first column of Table 4 shows that the kurtosis of the price change distribution averaged around 9 in the months of the tax changes, increasing substantially from 4 observed normally. The interquartile range of the absolute price change distribution during tax changes became tighter, averaging around 5.5% from the 8.2% observed normally. In terms of the average size of price changes, it is also interesting to note the significant decrease during the tax changing months (less than 9% vs. the 9.9% during normal times). All these observations are in line with the assumption of fat-tailed idiosyncratic shocks. The kurtosis increases, because a large fraction of the desired price change distribution with high kurtosis

gets revealed. The interquartile range becomes tighter for a similar reason: the range of the desired price change distribution is much lower, but normally most of the small price changes are not realized because of the menu costs. As the aggregate shock makes these price changes worthwhile, the observed interquartile range collapses. The lower average price change is the consequence of a lot of new below-average actual price changes. They are the result of the high relative mass of firms around the center of the fat-tailed desired price change distribution pushed over the inaction thresholds.

The second column of Table 4 shows that our baseline model is successful in predicting most of the moments. The kurtosis is fairly close to those observed in the data on average (even though it predicts significantly higher kurtosis for the 5% positive shock than for the 5% negative shock that we do not observe in the data). The predicted interquartile range also matches the observed values quite closely, as is the case with the average size.

In contrast, the model with the Poisson distribution (column 3) substantially overestimates kurtosis, and underestimates the interquartile range and the observed size, after positive shocks. These can all be a reflection of an idiosyncratic shock distribution with too high kurtosis. We observe the opposite with the Gaussian distribution (column 4): it underestimates the kurtosis and overestimates both the interquartile range and the average prices during tax changing months. This suggests that their idiosyncratic distribution does not have sufficient kurtosis.

Table 4: Moments of price dispersion at the months of tax changes

Moment	Size	Data	Mixed	Poisson	Normal
	+3%	10.6	9.4	16.5	4.2
Kurtosis	+5%	8.1	12.9	20.9	5.6
	-5%	9.2	6.0	3.5	3.5
	+3%	5.1%	4.7%	2.7%	6.5%
Interquartile range	+5%	5.9%	4.3%	2.7%	6.6%
	-5%	5.0%	5.7%	11.4%	6.6%
	+3%	6.5%	8.3%	6.8%	10.4%
Absolute size	+5%	9.0%	8.5%	7.8%	10.7%
	-5%	8.6%	8.4%	10.2%	10.3%

To conclude, we consider the pricing responses to the large and asymmetric pass-through of symmetric tax shocks in Hungary as a strong empirical evidence supporting menu cost pricing models. It is inconsistent with standard time-dependent pricing models (Calvo (1983)). Furthermore, the Hungarian tax experiment provides strong quantitative support to our baseline model with mixed normal idiosyncratic shocks over both the Poisson model similar to Midrigan (2011) and the model with Gaussian shocks like Golosov and Lucas (2007). It matches well the frequency and the inflation pass-through of these shocks as well as the observed distribution of

the price changes during tax changing months.

5   Robustness

In this section, we study various modifications of our baseline model in order to investigate the robustness of our conclusions. In each case, our main focus continues to be the extent of monetary non-neutrality after a benchmark monetary policy shock under our mixed normal idiosyncratic shock distribution and under the alternatives. Despite some quantitative differences, we find that our main conclusions hold up: the mixed normal idiosyncratic shock distribution generates monetary near-neutrality in all of these models.
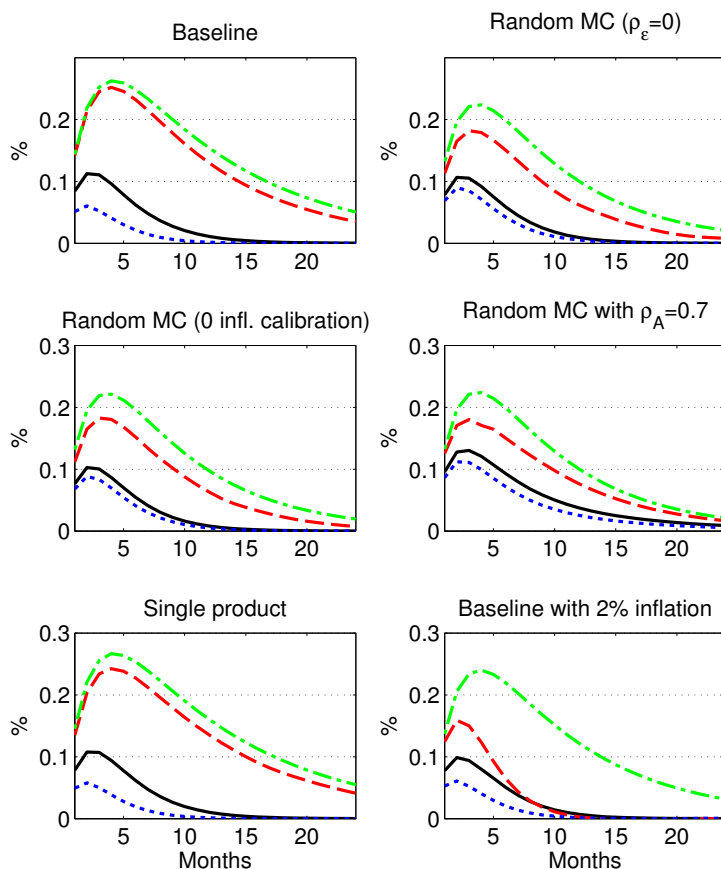
5.1   Random menu cost model

In Section 3.3, we illustrated the selection effect by using a single-product approximation of our baseline multi-product model, the "random menu cost model". In this model, we assumed that each single-product firm can change its price for free with a probability that equals the half of the steady-state price change frequency, and with the inverse probability it has to pay the fixed menu cost. This assumption mimics the baseline two-product economy because in both models, half of the actual price changes are for free, which – in line with empirical evidence – generates some very small price changes in the models. The main difference between the baseline model and this single-product approximation is that while in the baseline model, the idiosyncratic shocks of costly and free price changes are positively correlated (through the within-firm correlation of idiosyncratic shocks, ρ_ε), in the single-product approximation this correlation is $\rho_\varepsilon = 0$. Furthermore, this approximation should work relatively well only if the extensive margin is not important, as in the heterogenous menu cost model the frequency of free price changes is pinned down by assumption. Thus, for small aggregate shock the approximation is reasonable.[38]

The top right panel of Figure 8 shows the real effect of a small, one standard deviation, and persistent shock in the random menu cost model, for all three types of idiosyncratic shock distributions and the Calvo-model. The results are very similar to that in the baseline model (repeated here on the top left panel): the real effect in the mixed-normal model is small and close to that of the model with Gaussian shocks, while in the model with Poisson shocks the real effects are almost as large as in the Calvo-model. As the first two rows of Table 5 show, the cumulative real effect in the baseline and Gaussian models are 1.48 and 1.12 (expressed in multiples of the cumulative money shock), respectively; the same numbers were 1.59 and 0.65 in the baseline model. The cumulative real effect in the Poisson and Calvo-models are 4.06 and 5.78, respectively; these are similarly close to each other as the respective numbers in the baseline multiproduct model (7.60 and 8.93).

---

[38]Please see the Appendix for the calibrated parameters in the robustness exercises.

Figure 8: Impulse responses to a benchmark monetary shock

The figure plots the real effects of small (one standard deviation) persistent monetary shock. The top left panel is our benchmark, the top right panel shows the random menu cost model (see section 5.1). In the second row, the left panel shows the random menu cost model, re-calibrated under zero trend inflation (section 5.3), and the right panel shows the calibration if the persistence of the idiosyncratic shocks are reduced from $\rho_A = 1$ to $\rho_A = 0.7$. The bottom left panel is the simple menu cost model (section 5.2), and the bottom right panel is the baseline multi-product model with 2% per year trend inflation (section 5.4). The figure shows that our benchmark results are robust when the trend inflation is zero. If trend inflation increases to a standard 2% level (bottom right panel), all menu cost models imply small and temporary real effects.

5.2 Simple menu cost model

Relative to the random menu cost model of the previous subsection, we now make a further simplification to our model: while maintaining the single-product assumption, we assume that each firm has to pay the same menu cost whenever it changes its price. This gets us back to the simple menu cost model. This version is highly tractable, but has been criticized of being unable to generate small price changes.

The bottom left panel of Figure 8 shows the real effects of the same, small aggregate monetary policy shock. Again, results are similar to those at the baseline multi-product model and the

Table 5: Cumulative real effects

Parameters	Mixed	Normal	Poisson	Calvo
Baseline	1.59	0.65	7.60	8.93
Random MC	1.48	1.12	4.06	5.78
Random MC, 0 infl. calib	1.39	1.09	4.04	5.79
Random MC, $\rho_A = 0.7$	2.85	2.16	4.67	5.93
Simple-MC	1.52	0.60	7.85	9.34
Baseline + Infl	1.30	0.65	1.78	7.53

random menu cost model approximation: in terms of real effects, the model with mixed normal idiosyncratic innovations is similar to the model with Gaussian shocks, while the model with Poisson shocks resembles the Calvo-model.[39] The fifth row of Table 5 shows that the cumulative real effects are 1.52 and 0.60 times the cumulative money shocks in the mixed normal and Gaussian models, while they are 7.85 and 9.34 times the cumulative shocks in the model with Poisson shock and the Calvo-models, respectively. These numbers are remarkably close to the baseline model with multi-product firms.

5.3 Recalibrating to zero inflation

So far, we had calibrated the parameters matching the observed inflation rate in Hungarian data (4.2% per year), and then we "switched off" inflation. Because of this, the main moments of the alternative models (frequency and size of price change, kurtosis and interquartile range of the size distribution) were slightly different from each other. To show that these differences did not drive the results, we recalibrated the random menu cost version of the models to have exactly the same moments under zero inflation rate.

The right panel in the second row of Figure 8 shows the real effects of small aggregate shocks in these re-calibrated models. The results are essentially the same as in the random menu cost model without recalibration: the models with mixed normal and Gaussian shocks are again very similar to each other, and so are the model with Poisson shocks and the Calvo-model. The third row of Table 5 shows the cumulative real effects in the model variants: they are the multiple of 1.39, 1.09, 4.04 and 5.79 of the cumulative nominal shocks in the models with mixed normal, Gaussian, Poisson shock, and in the Calvo-model, respectively.

---

[39]Note that the real effect of the Calvo-model is smaller in this case than in the random menu cost model. The reason of this is the difference between the steady-state price change frequency. Both models were calibrated to positive trend inflation – matching our Hungarian data –, but then we switched off inflation. This means that the steady-state frequency of price change will go down in the menu cost models, and it will go down more in the single product model (as here there are no free price changes). For the sake of comparability, we adjusted the Calvo-parameters to reflect these changes in the frequencies of price changes: therefore in the simple menu cost model the Calvo-parameter is slightly smaller, leading to somewhat larger real effects.

5.4 Changing the persistence of the idiosyncratic shock

In our baseline model, we assumed that the idiosyncratic shocks have permanent impact on the quality of the product like in Gertler and Leahy (2008) and Midrigan (2011). This has allowed us to substantially reduce computational burden by eliminating one of the state variables. The right panel in the second row of Figure 8 shows the real effects of the random menu cost model if we reduce the persistence of the idiosyncratic shocks to $\rho_A = 0.7$ that is a reasonable midpoint of the parameters used in the literature.[40] The results show that this modification does not change our main conclusions. As the fourth row of Table 5 shows, the cumulative real effects are the multiple of 2.85 and 2.16 in the mixed normal and normal and 4.67 and 5.93 in the Poisson and the Calvo models, respectively.

5.5 Impact of trend inflation

Trend inflation influences the real effects of monetary policy shocks. The most obvious channel is its influence on the frequency of price changes: higher inflation makes firms adjust their prices more frequently, so they respond more also to monetary policy shocks. Somewhat less intuitively, trend inflation also influences the selection effect. The main reason is that it influences the shape of the steady state price change distribution, and thereby the mass of firms at their inaction thresholds. Its quantitative effect is strongly influenced by the shape of the idiosyncratic shock distribution, as we are about to show.

The bottom right panel of Figure 8 shows real effects of small aggregate money shocks in our baseline multi-product model when we introduce a standard 2% inflation rate per year. The figure shows that the models, particularly the model with the Poisson distribution, are sensitive to even such a low inflation rate: now all of the menu cost models predict small and temporary real effects with cumulative real effects of 1.30, 0.65, 1.78 times the cumulative monetary policy shock, much lower than that of the Calvo (1983) model (7.53) (see the fifth row of Table 5).

What explains the changing magnitude of the real effects? We can explore this by studying the random menu cost model calibrated for various λ values. Figure 9 shows the results for different inflation rates from 0% to the 4.2% in our sample. The figure shows that for very small values of λ (i.e. cases close to the Poisson distribution), even mild trend inflation leads to substantial increase in the flexibility of prices and hence a significant decrease in the real effects of monetary policy shocks. The key factor leading to these effects are the increased selection effect caused by the substantial increase in the mass of firms at the inaction thresholds.

The shape of the desired price change distribution of the Poisson case in the single product version of the model can explain the high mass of firms at the inaction band (see Figure 10). The steady state desired price change distribution is highly asymmetric, with lot of firms at the positive inaction threshold. These firms got to the inaction threshold just because of the trend

---

[40]There is a wide range of values used in the literature, like 0.45, 0.66, 0.678 and 0.95 (Golosov and Lucas (2007), Nakamura and Steinsson (2008), Klenow and Willis (2006) and Costain and Nakov (2011)).

Figure 9: Idiosyncratic distribution and real effects
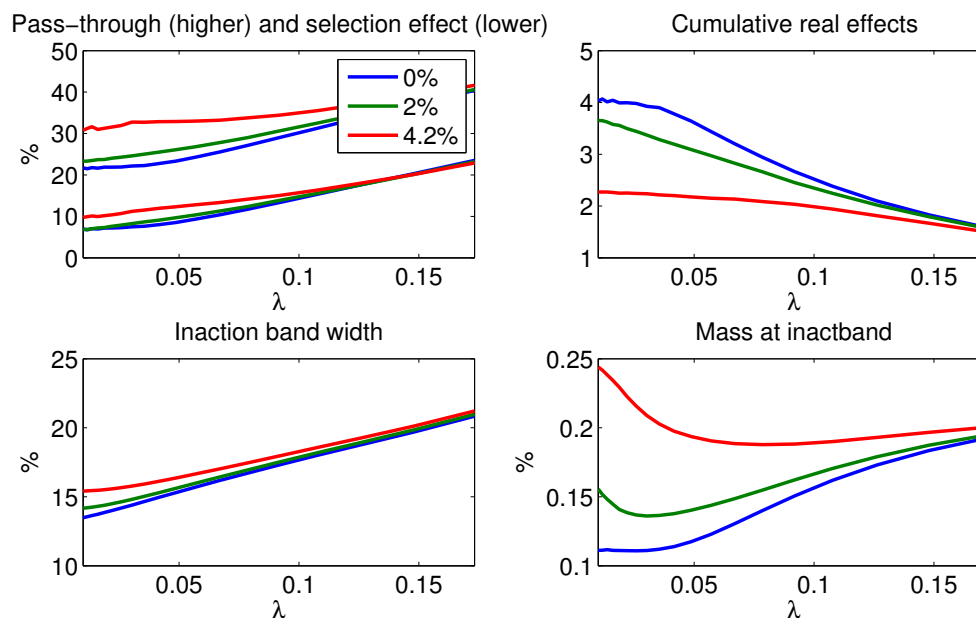
The figure shows the implied price level flexibility and selection effects, real effects, the width of the inaction band and the mass at inaction thresholds as functions of the shape of the idiosyncratic shock distribution and the trend inflation rate. The figure shows that for very small values of λ, even mild trend inflation leads to substantial increase in the flexibility of prices and small real effects of monetary policy shocks. The key factor leading to these effects are the substantial increase in the selection effect caused by the increase in the mass of firms at the inaction thresholds.

inflation, without idiosyncratic shocks; this explains the relatively high mass at the threshold.

These results suggest that depending on the underlying idiosyncratic shock distribution, the level of trend or target inflation can have significant influence on the effectiveness of monetary policy. Increasing the inflation target from 2% to over 4% increases the selection effect and makes monetary policy even less effective. This result is true in our baseline model – even though monetary policy is not very effective even at 0% inflation rates – but they are particularly important with the Poisson distribution. While the real effects there are as effective as in the Calvo (1983) model with 0 inflation, their effectiveness reduce substantially with a 2% inflation and become almost completely ineffective already with a 4% inflation (not shown), a somewhat surprising result.

5.6   Robustness to large shocks

So far we have shown that slight modifications of the baseline model – the random menu cost model and the simple menu cost model – leave our small shock results intact. In this subsection

Figure 10: Desired price change distribution and inaction bands with a positive trend inflation
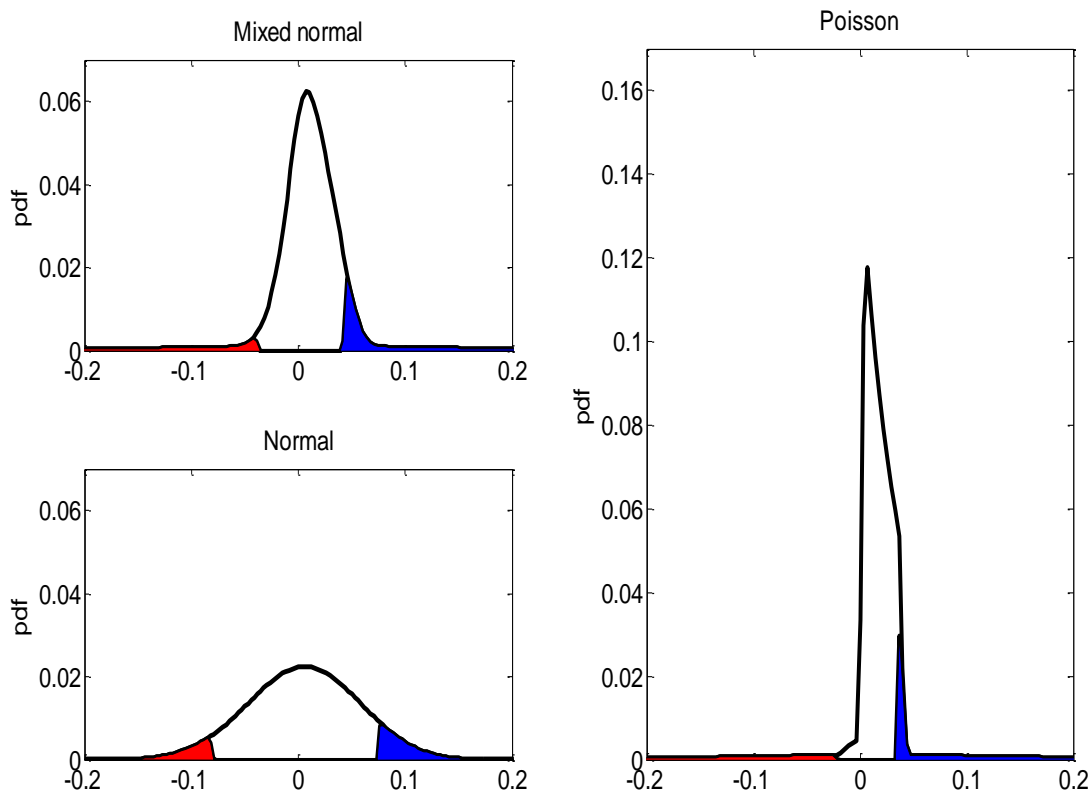
The figures plot the desired price change distributions and inaction bands for the baseline model and the alternatives. The shaded areas show the actual price changes.
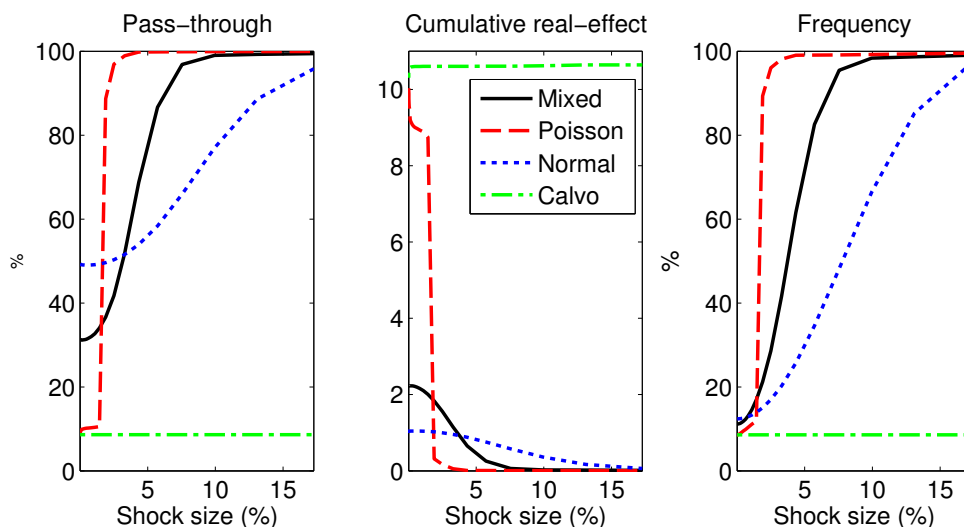
we revisit the models' responses to large shocks in the simple menu cost model.[41]

   Figure 5 in Section 4 plots the inflation pass-through, the cumulative real effects and the frequency responses to permanent shocks of various sizes in our baseline multi-product models. In Figure 11, we prepared the same figure for the simple menu cost model. The two figures are remarkably similar. First, in both cases the menu cost models become more flexible and the real effects get smaller as the shock size increases. Second, the model with Poisson shocks has

---

[41]As discussed before, the random menu cost model is not a good approximation of the baseline multi-product model for large shocks, as in this case the extensive margin becomes effective in the baseline model, but not in the random menu cost model (because the frequency of the 'free' price changes is fixed by assumption). So the only model variant in which we can reasonably expect similar results to the baseline model is the simple menu cost model.

aggregate price rigidity (similar to the Calvo case) for small shock sizes, but for large shocks it becomes the most flexible of all menu cost models in both cases. This is because of the strong extensive margin effect of the Poisson model in both the multi- and single-product versions. The relationship between aggregate price flexibility and shock sizes, and cumulative real effects and shock sizes are also remarkably similar for the menu cost models with mixed normal and Gaussian shocks.

Figure 11: Effects of various shock sizes

The figure plots the inflation pass-through, the cumulative real effects and the frequency responses to permanent shocks of various sizes in different single-product menu cost models for 0 inflation rate. This figure is very similar to the same figure in the baseline multi-product model in Section 4.

One interesting consequence of these results is that the simplest and most tractable menu cost model, in which firms produce a single product and pay a single menu cost after each nominal price changes, has very similar aggregate implications to the baseline menu cost model with two-product firms for both small and large shocks.

6 Conclusion

We have calibrated and tested a new quantitative pricing model. We have shown that our proposed menu cost model with an idiosyncratic shock distribution with unsynchronized stochastic volatility is highly successful in matching key moments of the steady state price change distribution. Furthermore it makes excellent out-of-sample predictions on the observed frequency increase and price change distribution following large value-added tax shocks in Hungary. We argued that existing alternative models would have hard time matching the facts.

Our baseline menu cost model implies high Caplin and Spulber (1987)-type selection-effect. So even though our model is highly successful in matching micro-pricing facts, the predicted low real effects of monetary policy shocks are in contrast to macroeconomic VAR evidence (Bernanke and Blinder (1992), and Christiano, Eichenbaum, and Evans (1999)). This suggests that alternative frictions, like real rigidities (Woodford (2003b), Gertler and Leahy (2008), Nakamura and Steinsson (2010)), wage rigidities (Erceg, Henderson, and Levin (2000)), or information frictions on aggregate shocks (VAT is a much easier and more well-publicized shock than monetary policy shocks that is the baseline shock in information friction models, see e.g. Lucas Jr (1972)) can be necessary to explain the observed effects of aggregate nominal shocks.[42] These additional rigidities, however, should also generate flexible and asymmetric responses to large aggregate shocks consistent with the new evidence presented in our paper. We leave this for further research.

References

Alvarez, F., H. L. Bihan, and F. Lippi (2014, May). Small and large price changes and the propagation of monetary shocks. Working Paper 20155, National Bureau of Economic Research.

Alvarez, F., M. Gonzalez-Rozada, A. Neumeyer, and M. Beraja (2011). From hyperination to stable prices: Argentina's evidence on menu cost models. Unpublished manuscript.

Alvarez, F. and F. Lippi (2014). Price setting with menu cost for multiproduct firms. *Econometrica 82*(1), 89–135.

Ball, L. and N. G. Mankiw (1994). Asymmetric price adjustment and economic fluctuations. *The Economic Journal 104*(423), 247–261.

Bernanke, B. S. and A. S. Blinder (1992). The federal funds rate and the channels of monetary transmission. *American Economic Review 82*(4), 901–21.

Bils, M. and Y. Chang (2000). Understanding how price responds to costs and production. *Carnegie-Rochester Conference Series on Public Policy 52*(0), 33 – 77.

---

[42]Bils and Chang (2000) also emphasizes differences in observed price responses to demand and cost shocks (see also Gagnon and López-Salido (2014)), and suggests potential explanations based on market structure.

Caballero, R. J. and E. M. Engel (2007). Price stickiness in ss models: New interpretations of old results. *Journal of Monetary Economics 54*, 100–121.

Cabral, L. and A. Fishman (2012). Business as usual: A consumer search theory of sticky prices and asymmetric price adjustment. *International Journal of Industrial Organization 30*(4), 371 – 376.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics 12*(3), 383 – 398.

Caplin, A. S. and D. F. Spulber (1987). Menu costs and the neutrality of money. *The Quarterly Journal of Economics 102*(4), 703–726.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999). Monetary policy shocks: What have we learned and to what end? Volume 1, Part A of *Handbook of Macroeconomics*, pp. 65 – 148. Elsevier.

Costain, J. and A. Nakov (2011). Distributional dynamics under smoothly state-dependent pricing. *Journal of Monetary Economics 58*(68), 646 – 665.

Danziger, L. (1999). A dynamic economy with costly price adjustments. *American Economic Review*, 878–901.

Dotsey, M., R. G. King, and A. L. Wolman (1999). State-dependent pricing and the general equilibrium dynamics of money and output. *The Quarterly Journal of Economics 114*(2), 655–690.

Eichenbaum, M., N. Jaimovich, S. Rebelo, and J. Smith (2014). How frequent are small price changes? *American Economic Journal: Macroeconomics 6*(2), 137–155.

Elsby, M. and R. Michaels (2014). Fixed adjustment costs and aggregate fluctuations. Unpublished manuscript.

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics 46*(2), 281 – 313.

Gabriel, P. and A. Reiff (2010). Price setting in Hungary: A store-level analysis. *Managerial and Decision Economics 31*(2-3), 161–176.

Gagnon, E. (2009). Price setting during low and high inflation: Evidence from Mexico. *The Quarterly Journal of Economics 124*(3), 1221–1263.

Gagnon, E. and D. López-Salido (2014). Small price responses to large demand shocks. Unpublished manuscript.

Gagnon, E., D. López-Salido, and N. Vincent (2012, June). *Individual Price Adjustment along the Extensive Margin*, pp. 235–281. University of Chicago Press.

Gertler, M. and J. Leahy (2008). A Phillips curve with an Ss foundation. *Journal of Political Economy 116*(3), 533–572.

Golosov, M. and R. E. Lucas (2007). Menu costs and Phillips curves. *Journal of Political Economy 115*(2), 171–199.

Hobijn, B., F. Ravenna, and A. Tambalotti (2006). Menu costs at work: restaurant prices and the introduction of the euro. *The Quarterly Journal of Economics 121*(3), 1103–1131.

Kehoe, P. J. and V. Midrigan (2010, September). Prices are sticky after all. NBER Working Papers 16364, National Bureau of Economic Research, Inc.

Klenow, P. J. and O. Kryvtsov (2008). State-dependent or time-dependent pricing: Does it matter for recent U.S. inflation? *The Quarterly Journal of Economics 123*(3), 863–904.

Klenow, P. J. and J. L. Willis (2006, March). Real rigidities and nominal price changes. Research Working Paper RWP 06-03, Federal Reserve Bank of Kansas City.

Lach, S. and D. Tsiddon (2007). Small price changes and menu costs. *Managerial and Decision Economics 28*(7), 649–656.

Levy, D., M. Bergen, S. Dutta, and R. Venable (1997). The magnitude of menu costs: Direct evidence from large U. S. supermarket chains. *The Quarterly Journal of Economics 112*(3), 791–825.

Lucas Jr, R. E. (1972). Expectations and the neutrality of money. *Journal of economic theory 4*(2), 103–124.

Mackowiak, B. and M. Wiederholt (2009). Optimal sticky prices under rational inattention. *The American Economic Review 99*(3), 769–803.

Mankiw, N. G. (1985). Small menu costs and large business cycles: A macroeconomic model of monopoly. *The Quarterly Journal of Economics 100*(2), 529–537.

Mankiw, N. G. and R. Reis (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve. *The Quarterly Journal of Economics 117*(4), 1295–1328.

Midrigan, V. (2011). Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica 79*(4), 1139–1180.

Nakamura, E. and J. Steinsson (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics 123*(4), 1415–1464.

Nakamura, E. and J. Steinsson (2010). Monetary non-neutrality in a multisector menu cost model. *The Quarterly Journal of Economics 125*(3), 961–1013.

Rotemberg, J. J. (2005). Customer anger at price increases, changes in the frequency of price adjustment and monetary policy. *Journal of Monetary Economics 52*(4), 829–852.

Rotemberg, J. J. (2011). Fair pricing. *Journal of the European Economic Association 9*(5), 952–981.

Stiglitz, J. E. (1984). Price rigidities and market structure. *The American Economic Review 74*(2), 350–355.

Vavra, J. (2014). Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *The Quarterly Journal of Economics 129*(1), 215–258.

Woodford, M. (2003a). Imperfect common knowledge and the effects of monetary policy. In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, pp. 25–58. Princeton University Press.

Woodford, M. (2003b). *Interest and prices: Foundations of a theory of monetary policy.* Princeton University Press.

7 Appendix

7.1 Equivalence of inflation effects of money and tax shocks

In this section, we show that for our baseline parametrization (similarly to those of Golosov and Lucas (2007) and Midrigan (2011)), permanent money shocks and value added tax shocks have equivalent effects on the inflation path - even though their effects on the output is different. The proof is using a guess and verify method, showing that assuming the same price level paths, the money supply and the tax rate have equivalent effects on the optimal price choices; indeed justifying the equal price effect assumption.

Nominal wage moves together with the money supply, under our assumptions on the labor supply equation with separable utility that is logarithmic in consumption and linear in labor ($\psi = 0$). In this case

$$W_t = \mu P_t Y_t = \mu M_t. \tag{15}$$

Substituting this into equation (6) about the periodic normalized profit, one can easily re-write the profit equation as

$$\bar{\Pi}_t(i) = \left(\frac{1}{1+\tau_t}\right) \sum_{g=1}^{G} \left[p_t(i,g)^{1-\gamma} - \frac{\mu M_t(1+\tau_t)}{P_t} p_t(i,g)^{-\gamma} \right] \left(\frac{1}{G} \sum_{g=1}^{G} p_t(i,g)^{1-\gamma}\right)^{(\gamma-\theta)/(1-\gamma)}. \tag{16}$$

Let's guess that the present and future path of the price level $\{P_t\}$ is the same for a permanent tax shock and a permanent money level shock. The optimal price choices of firms depend on their normalized value function, which is a present discounted value of their future profits. As the derivation shows, the tax rate influences the level of profits, but its influence on the optimal price choice is equivalent to that of the money supply.[43] As we also assume lump-sum redistribution of taxes, the variables will not influence the budget constraints either. It means that the assumption of equivalent price level development is indeed verified. So we are justified to use evidence gained from value added tax shocks to test the predictions of our model to large permanent money shocks.

7.2 The flexible price equilibrium

The algebraic solution for the flexible price equilibrium provides useful information about the long-term pass-through of the permanent tax- and money shocks. Money shocks, naturally, have no real effects under flexible prices, so we will have full pass-through to the price level. A permanent value-added tax shock, for our parametrization, will imply a unit drop in the real

---

[43]For that to be exactly true, we need to assume that menu costs are tax deductible, so their effective costs drop with higher value-added taxes together with the value functions.

output, so under unchanged money supply, we will have full pass-through into the gross nominal prices in this case as well. [44]

To gain some insight into why value-added tax shocks imply a unit drop in output, it is useful to look at the firms' static profit maximization problem. Under flexible prices, firms will choose prices to maximize this, implying the following optimal relative price:

$$p_t i, g^* = (1 + \tau_t)\frac{\gamma}{\gamma - 1}w_t, \tag{17}$$

where $w_t = W_t/P_t$ is the real wage. The equation shows that each firms want to increase their *relative* prices as a response to a tax increase. As all firms can not do this in equilibrium, real wages have to endogenously drop. It requires lower labor demand and output; and as household wage income will drop in parallel, the aggregate demand will adjust sufficiently to satisfy general equilibrium.

As all firms will choose the same productivity-adjusted relative price, the Dixit-Stiglitz-aggregate of these relative prices – that needs to be equal to 1 by definition – is also $\gamma w_t(1 + \tau_t)/(\gamma - 1) = 1$. We find that

$$w_t = \frac{\gamma - 1}{\gamma(1 + \tau_t)}. \tag{18}$$

From the labor market equation, we know that $w_t = W_t/P_t = \mu Y_t$, and any demand is going to be satisfied at this wage. The equilibrium output is, thus, given by

$$Y_t = \frac{\gamma - 1}{\gamma\mu(1 + \tau_t)}. \tag{19}$$

The nominal price level can be obtained as

$$P_t = \frac{M_t}{Y_t} = M_t\frac{\gamma\mu(1 + \tau_t)}{\gamma - 1}. \tag{20}$$

The expected growth rates are

$$E(g_Y) = -E(g_{1+\tau_t}), \quad E(\pi_t) = g_M + E(g_{1+\tau_t}) \tag{21}$$

This shows that a permanent increase in the tax will imply a full and immediate inflation pass-through under flexible prices.

7.3 Numerical solution algorithm

This subsection describes our numerical solution algorithm. It consists of two parts.

First, we solve for the steady-state aggregate variables π^{SS}, w^{SS} and Γ^{SS}. As we assumed no aggregate uncertainty, aggregate variables will converge to their steady-state values. The steady-state inflation rate is equal to the growth rate of money stock: $\pi^{SS} = g_M = g_{PY}$. Then we calculate the steady-state real wages (w^{SS}) and the distribution of firms over their idiosyncratic state variables (Γ^{SS}) with the following iterative procedure:

---

[44]We are also using the flexible price solution as starting values for the iterative procedures in our numerical solution method.

1. We start with a guess for w^{SS}, w_0. Initially, this guess is equal to the flexible-price steady-state of w, that we can calculate analytically (see previous subsection).

2. Given this guess and the steady-state inflation rate, we use a fine grid on relative prices[45] and idiosyncratic shocks[46] to solve for the optimal pricing policies of individual firms. We use value function iteration with quadratic and linear interpolation.

3. With the resulting policy functions, we calculate the steady-state distribution of firms over their idiosyncratic state variables. For this, we use the same set of grids as for the value function iteration. We again do this numerically: starting from a uniform distribution, we calculate the resulting distribution after idiosyncratic shocks hit, and also after firms re-price. Then again calculate the resulting distributions after a new set of idiosyncratic shocks and new re-pricing. We do this until convergence.

4. We calculate the (Dixit-Stiglitz) average relative price in the resulting steady-state distribution. If this is smaller (larger) than 1, then we increase (decrease) our initial guess (w_0) of the real wages.

5. We repeat these steps until the average relative price in the calculated steady-state distribution equals 1.

In the second part of our numerical algorithm, we calculate equilibrium paths of aggregate variables after an unexpected shock at $t = 0$ to the money supply, assuming that initially all aggregate variables were in their steady states. We calculate the equilibrium paths of π (inflation), w (real wages) and Γ (distribution of firms over their idiosyncratic state variables) with the following shooting algorithm:

1. We assume that aggregate variables will reach their steady-state in a finite (large) number of periods, T.

2. We start with a guess for the equilibrium inflation path $\{\pi_1, \ldots, \pi_T\}$. Our initial guess is the full immediate pass-through.

3. Given this guess, we calculate the resulting equilibrium path of the real wages: $\{w_1, \ldots, w_T\}$. As $w_t = \mu Y_t$, we do this by calculating the equilibrium real GDP path, which we know from the equilibrium inflation path (and the constant nominal growth assumption).

4. Given the inflation and real wage paths, we calculate the path of value and policy functions. We do this by backward iteration from T, where the economy and the value functions are assumed to converge to a steady state.

---

[45]We have 2,000 gridpoints when idiosyncratic shock innovations are permanent (i.e. when the normalized relative price is the only idiosyncratic state variable).

[46]In the temporary idiosyncratic shock case (when the idiosyncratic shock is also a state variable) we have 101 gridpoints for the quality shocks, and 500 for the relative prices.

5. Starting from period 1, and using the steady-state distribution of firms over their idiosyncratic state variables as initial distribution, we use the sequence of policy functions (together with the idiosyncratic shock processes) to calculate the resulting path of Γ, the distribution of firms over their idiosyncratic state variables.

6. From the resulting sequence of distributions, we calculate the resulting inflation path, and compare it with our initial guess. If the two are different, we update our guess to the linear combination of our previous guess and the resulting inflation path.

7. We do these iterations until the resulting inflation path is the same as our initial guess.

7.4 Calibrated parameters of the robustness exercises

This section presents the calibrated parameters of the robustness exercises. The parameters were calibrated to match the observed frequency of price changes and size of absolute price changes (in all), kurtosis of price changes (Poisson and mixed normal) and interquartile range of the absolute price change distribution (mixed normal[47]). The table shows that there are sizable differences between the parameters across different models, but this does not lead to differences between the models' aggregate conclusions as Section 5 shows.

Table 6: Calibrated parameters of the robustness exercises

	Simple MC			Random MC		
	Mixed	**Poisson**	**Normal**	**Mixed**	**Poisson**	**Normal**
ϕ	0.8%	0.6%	2.2%	5.8%	4.5%	6.7%
σ_A	4.2%	4.4%	3.5%	4.3%	4.7%	4.0%
p	90.8%	90.5%	0	95.3%	90.3%	0
λ	7.6%	0	1	18.2%	0	1
	Random MC, 0 infl. calib.			Random MC, $\rho_A = 0.7$		
	Mixed	**Poisson**	**Normal**	**Mixed**	**Poisson**	**Normal**
ϕ	5.6%	4.6%	6.7%	6.1%	5.5%	6.7%
σ_A	4.5%	5.2%	4.0%	6.5%	5.5%	6.8%
p	95.9%	90.3%	0	96.2%	86.6%	0
λ	17.4%	0	1	31.5%	0	1

7.5 Macroenvironment in Hungary

In this section, we describe the macroenvironment in Hungary around the value-added tax changes of 2004-2006.

---

[47]In the random menu cost versions of the model, the interdecile range was matched.

Figure 12: The macro-environment in Hungary
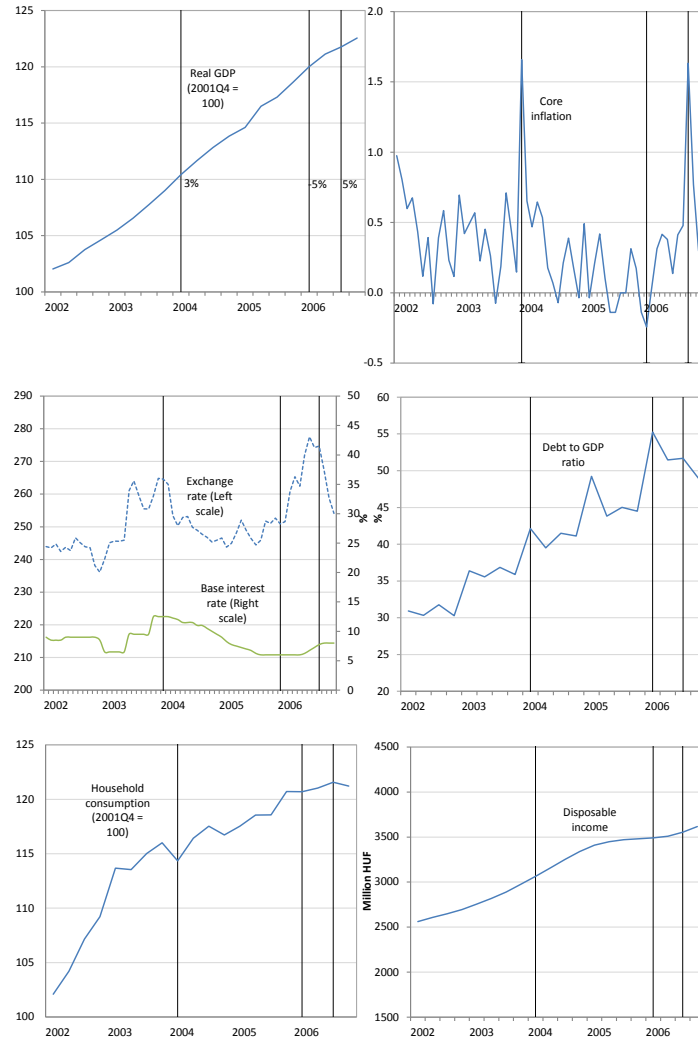
The figure plots development of key indicators in Hungary around the value-added tax changes. The 12% rate was increased to 15% in January 2004 (first horizontal bar), the 15% rate was increased to 20% in September 2006 (third horizontal bar) and the 25% rate was decreased to 20% in January 2006 (second horizontal bar). The indicators show a steady growth in GDP, retail consumption with small inflation rates (4%) and relatively stable exchange rates.

Between 2002 and 2006, Hungarian real GDP grew by an average of 4.2% per year. As panel A of Figure 12 shows, this growth rate was remarkably stable: the yearly growth rates fluctuated between 3.9% (2003 and 2006) and 4.8% (2004). Meanwhile, the core inflation (see panel B) averaged at a low 4.1%.[48]

This relatively quick real growth can partly be explained by the rapid growth rate of private debt: the approximately 30% private debt-to-GDP ratio of 2002 increased steadily to around 55% by 2006 (panel D).[49] In parallel, the budget deficit was also running high (between 6.4% of GDP in 2005 and 8.9% of GDP in 2003), and government debt increased from 52% of GDP in 2002 to around 62% of GDP in 2006. Under these circumstances, it is hardly surprising that household consumption (panel E) and the volume of retail trade also increased steadily during this period.[50]

The aim of the value-added tax increase of 2004 was to cut the high budget deficit (8.9% of GDP) of 2003. It was part of various tax-related measures, some tax increases, some tax decreases with small effects on the net personal disposable income (see panel F). In September 2006, the aim was similar: the government wanted to cut the budget deficit by implementing a series of tax increases. Most of the tax measures influencing personal income took effect in January 2007, while the VAT-increase (September) and some regulated price increases (e.g. heating gas, electricity) were implemented mid-year. The value-added tax decrease of January 2006, with only minor additional tax measures, came as a popularity increasing measure before the April 2006 general elections (which resulted in the government being re-elected).

The inflation targeting central bank explicitly communicated that it would not react directly to the value-added tax changes, as their direct effect would disappear from the inflation rate at the policy horizon. It added that it would monitor any second-round effects through changing expectations. Indeed, there are no immediate changes in the policy rate right after the value-added tax changes, and during this period it seems that the central bank was mostly responding to depreciating exchange rates (see panel C). The exchange rate was mostly stable during the period, temporary shocks to it were counteracted by the interest policy of the inflation targeting central bank.

One could claim that factors missing from our model had significant influence on the observed pass-through of the VAT-changes. First, we saw an approximately 10% depreciation of the local currency, the Hungarian Forint (relative to the Euro) in the summer of 2006, i.e. just before the 5%-point tax hike. However, this depreciation was only temporary and counteracted by interest rate hikes by the exchange-rate smoothing central bank. Furthermore, given the long time lag at which exchange rate movements pass through into processed food prices in Hungary, we can safely assume that the exchange rate had minor impact on the CPI developments. Second, there

---

[48]Panel A shows the level of GDP, assuming that it was 100 in 2001Q4. Panel B shows the m-o-m core inflation rates, which has an average of 0.345%, i.e. approximately 4.1% per year.

[49]Panel D shows the sum of household and corporate debt, as a fraction of GDP.

[50]Panel E shows the level of households' consumption expenditures, assuming that it was 100 in 2001Q4.

were a series of regulated electricity and gas price increases in August and September of 2006 (together with the VAT-increase). However, they affected only the consumer prices of gas and electricity and had no impact on producer prices, so this could only have had a minor impact on retail firms pricing behavior in September 2006. Finally, fiscal policy measures in parallel with the VAT-changes have not had a significant impact on the net disposable income (see panel F), so their effect on prices could only have been minimal. All in all, we argue that most of the observed movements in the inflation rates were due to the VAT-changes, and other factors can be safely disregarded.

7.6 A 5% tax increase in 2009

Although the product groups affected by the 2004 and 2006 VAT changes were similar, these groups were not identical. This may lead to composition bias in estimating moments. In this subsection we use evidence from a 5%-point VAT-increase in 2009 to evaluate the possible size of this bias in estimating the asymmetry in inflation effects in 2006. This extra evidence helps us, because now we have products that were hit by both a VAT-decrease and a VAT-increase, so we can directly compare their price responses. The difficulty, however is that during the 2009 increase, the economy was undergoing a serious recession that might have a substantial impact on the inflation pass-through. Controlling for the business cycle effects implies a comparable level and asymmetry of the pass-through as in our baseline experiment.

In July 2009, in an attempt to increase government revenues during the financial crisis, Hungarian authorities decided to increase the by now unified VAT-rate of 20% to 25%. As the second column of Table 7 indicates, 102 of the 128 products in our original processed food sample were hit by this new 5%-point tax increase.[51] Our estimate for the inflation pass-through of this tax change is 56.6%, which is relatively small. One possible explanation for this moderate pass-through is the ongoing large recession (6.8% fall in Hungarian real GDP in 2009; as opposed to the 3.9% real GDP growth in 2006).

Table 7: Inflation pass-through in 2009 and 2006

	Tax inc in 2009	Inc09-Dec06	Inc09-Inc06
Number of items	102	29	73
CPI-weight	10.75	3.37	7.38
Infl PT in 2009 (+5%)	56.6%	68.8%	51.1%
Infl PT in 2006 Jan (-5%)	–	32.6%	–
Infl PT in 2006 Sep (+5%)	–	–	88.0%

[51]The remaining 26 items (all of them basic food items) got into a newly created 18% VAT-category. In essence, this created again a multiple-rate VAT-system with rates 18% and 25%, but now a much smaller proportion of the consumption basket had the lower VAT-rate than before 2006.

The third column of Table 7 contains information about the inflation pass-through for those 29 processed food items that were hit by both the 5%-point tax decrease in January 2006 and the 5%-point tax decrease in July 2009. Even in this group (not subject to composition bias) we see substantial asymmetry, 32.6% vs 68.8%. This is despite the fact that the pass-through in 2009 was in general much smaller than in 2006: according to column 4, for the 73 items that were hit by both tax increases of September 2006 and July 2009, the respective pass-throughs were 51.1% and 88.0%. So the substantial asymmetry in column 3 is likely to be underestimated due to business cycle effects, thus our original estimate (33% vs 99%) seems reasonable.[52]

---

[52]If we assumed that the 2009 pass-through was proportionally lower for each single product (i.e. only 51.1/88.0=58.1% of the 2006 pass-through), then the asymmetry in the third column would be 32.6% vs 118.5%, not far from our original estimate.