Menu Costs, Aggregate Fluctuations and Large Shocks*

Peter Karadi[†] Adam Reiff[‡]

February 15, 2014

PRELIMINARY AND INCOMPLETE, PLEASE DO NOT DISTRIBUTE

Abstract

In menu cost models, real effects of aggregate nominal shocks are sensitive to unobserved characteristics of price setting. The standard way to calibrate key pricing parameters is to match the cross-sectional distribution of price changes. We argue that this unconditional distribution contains insufficient information for a clean identification. In particular, it is consistent with parameterizations with contradictory aggregate implications: one generating sizeable real effects, the other effective money neutrality. We argue, instead, that price change observations *conditional* on aggregate shocks can be sufficiently informative. To show this, we utilize new micro level evidence on price responses to large value-added tax shocks. We present a new menu cost model with idiosyncratic shocks and show that our model successfully predicts the magnitude of the observed price responses, outperforming alternative pricing models. The new model generates *small* real effects of monetary policy shocks implying that imposing the menu cost assumption alone fails to explain robust time-series evidence finding the opposite.

JEL codes: E31, E52

Keywords: Monetary Policy Transmission, State-Dependent Pricing, Price change distribution, Value-Added Tax

*The authors would like to thank Virgiliu Midrigan for constant encouragement and guidance and Peter Benczur, Michael Dotsey, Mark Gertler, Patrick Kehoe, John Leahy, Ellen McGrattan, Attila Ratfai, Harald Uhlig, and seminar participants at the Banque de France, Boston Fed, BIS, Central European University, European Central Bank, EEA 2011, Federal Reserve Board, Minneapolis Fed, New York University, University of Lausanne, Norges Bank for insightful comments. All remaining errors are ours. The views expressed are those of the authors and do not necessarily reflect the views of the Central Bank of Hungary, the European Central Bank or the Eurosystem.

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

[‡]reiffa@mnbb.hu, Magyar Nemzeti Bank (The Central Bank of Hungary)

1 Introduction

Small, fixed menu costs of price adjustment provide a classic explanation for staggered price setting (Mankiw (1985), Caplin and Spulber (1987)). Recent literature debates the extent to which aggregate nominal shocks influence real activity in models with menu costs. Golosov and Lucas (2007) find essentially money neutrality, while Midrigan (2011) finds large real effects in line with predictions of time-dependent pricing models (Calvo (1983)). They both reach their conclusions by calibrating unobserved price setting parameters, like the magnitude of the menu costs, to match key moments of the cross-sectional price change distribution. As Midrigan (2011) matches not only the frequency and the size, but also the dispersion of price changes, the current literature tends to accept his prediction of *large* real effects of nominal shocks in this class of models.

In the first part of this paper we revisit the result of Midrigan (2011). We consider quantitative menu cost models in which firms face idiosyncratic shocks (Golosov and Lucas (2007)). This helps matching the large observed magnitude of price changes that aggregate fluctuations can not explain. We also allow firms to produce multiple products as in Midrigan (2011), and reprice all of them for a single menu cost. As the price adjustment of the second product is 'free' in this case, the assumption allows us to generate frequent small price changes that we observe in the data. We claim that the unconditional cross-sectional distribution of price changes in such models does not identify well the extent of money non-neutrality. In particular, we propose a new idiosyncratic shock distribution, which makes the model match the cross-sectional distribution of price changes at least as well as Midrigan (2011), but has the opposite implication on aggregate price flexibility. In fact, aggregate price flexibility in our baseline model is almost as high as in Golosov and Lucas (2007).

What explains the insufficient information content of the unconditional price change distribution to identify the aggregate real effects of money shocks in menu cost models? Apparently, all hinges on the key role of the menu costs and the difficulty of inferring its correct magnitude. First, the degree of aggregate price flexibility in these models depends crucially on the strength of endogenous Caplin and Spulber (1987)-type selection of price changes, as emphasized by Golosov and Lucas (2007). Selection implies that price adjustments are non-random: firms with larger desired price changes are more likely to adjust. If the selection effect is high, those firms will respond to an aggregate shock that are prepared to change their prices substantially. This implies disproportionately large price level response and, thereby, small real effects of monetary policy shocks. But the extent of selection can not be directly measured, so it needs to be deduced from a model capturing key pricing characteristics well.

Second, the strength of the selection is closely linked to the size of the menu costs. If menu costs are small, only very small desired price changes get discouraged, implying small selection. With high menu costs, in contrast, large desired price changes will be heavily over-represented among the actual price changes, generating large selection effects. Conditional on

a fixed fraction of price changes,¹ *lower* menu costs *increase* the real effects of monetary policy shocks by reducing selection. Therefore, making the correct assessment of the size of menu costs is essential in order to obtain realistic estimates for the real effects of monetary policy shocks.

Third, menu costs are not directly observable. The standard way to calibrate them is to match key moments of the observed unconditional price change distribution, like the frequency and magnitude of price changes. These moments, however, also depend on the (also unobserved) distribution of idiosyncratic shocks. Assuming Gaussian shocks, Golosov and Lucas (2007) needed relatively large menu costs 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 most desired price changes are small, 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. Matching the high dispersion of the observed price change distribution provided further support for the model of Midrigan (2011).

And finally, small modifications in the distribution of idiosyncratic shocks may lead to sizable changes in the calibrated menu costs. To show this, we introduce a novel idiosyncratic shock distribution by mixing two Gaussian distributions, which is a possible outcome of a simple model with stochastic volatility. This model is a relatively minor departure from the Poisson distribution used by Midrigan (2011) (see also Danziger (1999) and Gertler and Leahy (2008)), which is a Gaussian shock arriving with a certain Poisson rate.² We calibrate the model to match key pricing moments. The idiosyncratic shock distribution favored by the data is close to the Poisson distribution and it is able to capture the low frequency, large average size and wide dispersion of price changes at least as well as Midrigan (2011). The larger dispersion of small price changes, however, is sufficient to require a substantially larger calibrated menu cost. This implies large selection, and small real effects of nominal shocks very different from Midrigan (2011) and not unlike that in the model of Golosov and Lucas (2007). We conclude that the unconditional price change distribution does not provide sufficient information to separately identify the menu costs and the idiosyncratic shock distribution, so it can not help us distinguish between prominent menu cost models with opposite implications on the real effects of monetary policy shocks.

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,

¹Lower menu costs *increase* price flexibility through increasing 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 (e.g. by reducing the standard deviation of the idiosyncratic shocks).

²Note also that the Gaussian distribution used by Golosov and Lucas (2007) is also a special case of our mixed normal distribution, when the probability of having one of the volatilities is 0.

we show that the price change distribution *conditional* on large aggregate shocks would provide valuable extra information about both the realistic size of the menu costs and the distribution of desired price changes, and thereby of the idiosyncratic shocks. This is intuitive. First, the 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 hits, the distribution of desired price changes, and also of the underlying idiosyncratic shocks are revealed.

We contrast pricing predictions to large, permanent shocks of our calibrated model to those of the alternatives and then we confront them with the data. For large aggregate shocks, all menu cost models predict highly *flexible* and *asymmetric* price responses, and highly *non-Gaussian distribution* of price changes when the shocks occur. Asymmetric aggregate price responses after large and symmetric positive and negative shocks emerge due to positive trend inflation, as in Ball and Mankiw (1994).³ The quantitative differences between their predictions can be contrasted with the data.

Our empirical 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 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 in affected prices (44% from 14%). Our model predicts this increase almost perfectly. We also observe a substantial asymmetry between the short-run inflation pass-through of the symmetric (+5% and -5%) positive and negative shocks (around 100% vs 30%). Our model matches almost perfectly this asymmetry. Finally, we observe highly dispersed price changes in the months of the tax shocks, with a lot of price changes close to the

³Time-dependent models (e.g. Calvo (1983)) imply symmetry even with trend inflation. The reason of asymmetry in menu cost models under trend inflation is that after a negative aggregate shock, firms can spare their menu costs by keeping their prices constant and allowing the trend inflation to reduce their relative prices for free; while inflation does not help them with positive shocks. Thereby, we can observe more price increases after a large positive shock than price decreases after an equally large negative shock, leading to asymmetry.

⁴Differently from the U.S. practice, where prices are posted net of sales tax. It is, however, similar to most other European countries.

⁵The changes took effect at different dates, they were preannounced and widely publicized to both stores and consumers. These features of the tax changes make us believe that alternative price rigidity explanations based on information frictions – for example Mackowiak and Wiederholt (2009), Mankiw and Reis (2002) and Woodford (2003a). – or uncertainty on the cost shocks – search models like Cabral and Fishman (2012) and Yang and Ye (2008). – can play only a minor role in explaining the observed outcomes.

⁶Our Hungarian sample is not special. The micro moments we hit are close to those in a sample of barcode data from a large US supermarket chain used, for example, by Midrigan (2011).

size of the tax change, but also a lot of large price changes. This is also matched remarkably well by our model.

Our model quantitatively outperforms alternative menu cost models. In particular, the frequency effects, the asymmetry, and the dispersion of price changes as a response to large shocks are all underestimated by the Golosov and Lucas (2007) model with Gaussian idiosyncratic shocks, and overestimated by the application of the Poisson distribution of Midrigan (2011). What explains the success of our model? The magnitude of the predicted flexibility, asymmetry and price dispersion are sensitive to the 1) distribution of idiosyncratic shocks and 2) the implied menu costs. In terms of flexibility, our fat-tailed idiosyncratic shocks imply that most desired price changes are small, so menu costs can be smaller to match the observed low steady-state frequency of price changes than in the Gaussian Golosov and Lucas (2007) model. When a large shock hits, smaller menu costs induce more firms to change their prices. This stronger extensive margin effect explains the excess price flexibility of our model relative to the Gaussian model. The model using the distribution of Midrigan (2011), however, assumes idiosyncratic shocks that imply too many small desired price changes, and thus require too small menu costs. The large shock induces too many firms to adjust, inconsistently with our evidence. In terms of asymmetry, the fat-tailed idiosyncratic shock distribution also influences its magnitude. Responses to aggregate shocks are asymmetric under trend inflation, because expected inflation raises the willingness of firms' to increase their prices. Fat tails further increase the proportion of price increases relative to price decreases. The predicted asymmetry is minimal in the model with Gaussian idiosyncratic shocks, but excessive with Poisson shocks. In contrast, in our baseline mixed normal model, they lead to asymmetry in line with the data, suggesting that this model describes the underlying idiosyncratic shock distribution better than the alternatives. Furthermore, the distribution of observed price changes when the large aggregate shocks hit provides further direct support for our underlying distributional assumption.

Our results also have important consequences for the real effects of *regular size* monetary policy shocks. For shocks of these size, the extensive margin becomes ineffective, implying small frequency variation and close to symmetric responses to positive and negative shocks. In terms of the real effects, what becomes important in this range is the endogenous selection of price changes, as we emphasized above. We find that our generalized idiosyncratic shock distribution, while it is close to the Poisson distribution of Midrigan (2011) and able to match the dispersion of prices at least as well, leads to small real effects closer to Golosov and Lucas (2007). The reason is that our distribution increases the calibrated menu cost sufficiently to lead to the reemergence of a strong selection effect, which substantially mitigates the real effects of monetary policy shocks.

Related literature Our paper is part of the literature aiming at exploring the macroeconomic predictions of realistic menu cost models (see Golosov and Lucas (2007), Klenow and Kryvtsov (2008), Midrigan (2011), Nakamura and Steinsson (2008), Klenow and Malin (2010),

Costain and Nakov (2011), Alvarez, Lippi, and Paciello (2011)). Most of this literature is using data collected during calm periods with low aggregate volatility. The lack of sufficient aggregate variation hinders empirical testing and stark comparison of the performance of alternative models. Our main empirical contribution is to use a unique natural experiment generated by large aggregate VAT shocks that generates distinct predictions across alternative models. In this sense, our paper is closer to papers that use pricing observations during high trend inflation periods (Golosov and Lucas (2007), Gagnon (2009), Alvarez, Gonzalez-Rozada, Neumeyer, and Beraja (2011)) to provide supporting evidence on the state-dependent menu cost assumptions. The focus on learning about responses to one-time large shocks and contrasting alternative menu cost models differentiate our paper from these papers.

Price responses to value-added tax shocks have been documented in the literature (see Gagnon (2009) and the references therein). These papers, similarly to those utilizing high inflation environments, document significant increases in the fraction of price changes in the months of the tax changes in support of state-dependent pricing assumptions. In a recent example, Gagnon, López-Salido, and Vincent (2012) 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 products⁷, within a short time span.⁸ Differently from these predominantly descriptive papers, furthermore, we use the new evidence to test the predictions of our model and contrast quantitative predictions of structural menu cost models.

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 60% of the affected firms failing to respond, 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 and asymmetric price responses (Stiglitz (1984), Cabral and Fishman (2012), Yang and Ye (2008)). 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

⁷The negative change in the sample of Jonker, Folkertsma, and Blijenberg (2004) in the Netherlands influences only 2 products

⁸In the sample of Baudry, Le Bihan, Sevestre, and Tarrieu (2007) for France, there were 5 years between the VAT increase (2%-points) and the VAT decrease (1% points).

prices, suggesting that the above explanation can neither be the sole reason of infrequent price adjustments.⁹

Structure: In section 2, we present the model and explain our method of calibration. In section 3, we present the Hungarian VAT experiment, our data and show the predictive ability of our baseline model. We contrast it to the alternative menu cost models of Golosov and Lucas (2007) and Midrigan (2011). In section 4, we run counterfactual experiments with our calibrated model. In particular, we show how the predictions of our model with mixed normal shock compares to the Midrigan (2011), Golosov and Lucas (2007) and the Calvo (1983) models for different monetary policy shock sizes and inflation rates. Section 5 concludes.

2 The Model

To study aggregate fluctuations in menu cost models, we use a heterogeneous firm multi-product menu cost model with fat-tailed idiosyncratic shocks, as in Midrigan (2011).¹⁰ We introduce nominal shocks into the model as exogenous permanent innovations to the money supply. We also introduce exogenously determined value-added tax rates.¹¹ As we show in the Appendix, nominal shocks and value-added tax shocks lead to equivalent inflation effects in our model. For simplicity, we assume away aggregate uncertainty, which means that consumers and firms always expect money supply to grow by its exogenously given rate, and expect tax rates to stay constant.¹²

⁹Our evidence on inflation asymmetry is not in line with some alternative frameworks either. Muted pricing response to VAT decreases 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. The convex adjustment cost model of price adjustment (Rotemberg (1982)) would also be inconsistent with our evidence. This model would imply asymmetric inflation response, but with the opposite sign: in an inflationary environment, the additional response of firms to positive shocks would be *smaller* than to a negative one, because the marginal cost of price increase (additional to the trend-inflation induced optimal positive change) would be even higher. Finally, the asymmetric inflation effect due to the asymmetric shape of the profit function (Devereux and Siu (2007), Ellingsen, Friberg, and Hassler (2006)) is in fact incorporated in our model, but quantitative results suggests that this type of asymmetry is negligible relative to the asymmetry caused by positive trend inflation.

¹⁰Besides fat-tailed shocks and multi-product firms, Midrigan (2011) has also introduced endogenous sales into his model. We disregard from sales in this paper by sales-filtering the data, making the model more directly comparable to the Golosov and Lucas (2007) that does not model sales either. Gertler and Leahy (2008) also used fat-tailed idiosyncratic shocks, with real rigidities, which we do not have. Gertler and Leahy (2008) argued that without real rigidities, fat-tailed idiosyncratic shocks alone do not lead to realistic real effects of monetary shocks. However, they do not match the shape (in particular, the kurtosis) of the price change size distribution, which can be one reason of their different conclusion from Midrigan (2011).

¹¹In our framework with no intermediate production, value-added tax is equivalent to sales tax.

¹²This assumption simplifies the analysis, but has no qualitative implications on the results.

2.1 Consumers

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

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

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 (with across-firm and across-good elasticity parameters θ and γ , respectively) of individual good consumptions $C_t(i, g)$. The measure of individual firms i is normalized to 1, and the number of individual goods is G . We assume that the across-firm elasticity of substitution parameter θ is larger than the across-good elasticity of substitution parameter γ , reflecting the fact that different goods purchased from the same firms are not as close substitutes as the same goods purchased from different firms.

Similarly to Midrigan (2011), we call $A_t(i, g)$ as the quality of the good. This term enters the consumption aggregate in a way that higher quality goods provide larger marginal utility of consumption. These good- and store-specific idiosyncratic shocks provide a mechanism through which the optimal price of firms keeps changing, giving an incentive to price changes. In essence, these idiosyncratic quality shocks have similar role as the idiosyncratic productivity shocks in Golosov and Lucas (2007): they help us to match the empirical price change size distribution.¹³

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

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

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

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

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 $\mu L_t^\psi C_t = W_t / P_t$.

¹³As we discuss later, higher quality products also have higher marginal cost of production. Similarly to Midrigan (2011), this assumption helps us to reduce the dimensionality of the problem.

2.2 The government and the central bank

The central bank is assumed to follow a constant money supply growth rule: $\log(M_t/M_{t-1}) = g_M$. We also assume that in equilibrium, nominal spending must be equal to the money supply: $M_t = P_t C_t$, which means that the nominal income $P_t Y_t$ is going to grow at a constant rate $g_{PY} = g_M$ as well.¹⁴ The extra money supply $M_t - M_{t-1}$ in the economy is redistributed in a lump-sum way.

We assume that all goods face the same value-added tax rate τ_t that is exogenously determined by the government. This means that a fraction $1/(1 + \tau_t)$ of gross revenues will be the firms' net revenue, and the fraction $\tau_t/(1 + \tau_t)$ will be paid for the government as a value-added tax. These tax revenues 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 The firms

Each firm i is assumed to sell its products (i, g) ($g = 1, 2, \dots, G$) in a monopolistically competitive market; post gross nominal prices $P_t(i, g)$ and satisfy all demand given these prices. However, if firms choose to change their gross nominal prices, they must pay a fixed menu cost $\phi P_t C_t$. As in Midrigan (2011), the payment of this fixed menu cost enables the firm to change all of its prices $(P_t(i, 1), P_t(i, 2), \dots, P_t(i, G))$.

The firms' problem is to 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: $\tilde{\Pi}_t(i) = \sum_{g=1}^G [(1/(1 + \tau_t))P_t(i, g)Y_t(i, g) - W_t L_t(i, g)]$.

For the production process, we assume that firms use a constant returns to scale technology with a single production factor of labor. The production function is $Y_t(i, g) = L_t(i, g)/A_t(i, g)$, reflecting the fact that higher quality products require more labor input.¹⁵ As the idiosyncratic good quality process does not have a trend, the growth rate of money supply and nominal expenditures ($g_{PY} = g_M$) will determine the trend inflation ($\pi = g_{PY} = g_M$).

We specify the log of the good quality process as a random walk process:

$$\ln A_t(i, g) = \ln A_{t-1}(i, g) + \varepsilon_t(i, g), \quad (4)$$

where innovations $\varepsilon_t(i, g)$ are mean-zero i.i.d. random variables with variance σ_A^2 . We assume that within firms, good quality innovations $\varepsilon_t(i, g)$ are correlated across goods g , with the correlation parameter ρ_ε , but they are uncorrelated across firms. Further, similarly to Gertler and Leahy (2008) and Midrigan (2011), we assume a fat-tailed, leptokurtic distribution for the innovations, with which we match the excess kurtosis of the observed price change distribution. To

¹⁴From the Euler-equation, this implies that the gross nominal interest rate $R_t = e^{g_{PY}}/\beta$ will also be constant over time.

¹⁵This technical assumption is taken from Midrigan (2011), and reduces both the state space and the computational burden.

generate this extra kurtosis we assume that $\varepsilon_t(i, g)$ follows a mixed normal distribution: with probability p , it is drawn from a mean-zero normal distribution with (small) variance σ_1^2 , and with probability $1 - p$ it is drawn from another mean-zero normal distribution with (larger) variance $\sigma_2^2 > \sigma_1^2$.

$$\varepsilon_t(i, g) = \begin{cases} N(0, \sigma_1^2) & \text{with probability } p \\ N(0, \sigma_2^2) & \text{with probability } 1 - p \end{cases}$$

We denote the proportionality between the standard deviations by $0 \leq \lambda \leq 1$: $\sigma_1 = \lambda \sigma_2$. To ensure that the variance of $\varepsilon_t(i, g)$ is indeed σ_A^2 , we set $\sigma_2^2 = \sigma_A^2 / (p\lambda^2 + 1 - p)$, and $\sigma_1^2 = \sigma_A^2 \lambda^2 / (p\lambda^2 + 1 - p)$.

With this specification, a higher parameter p will increase the kurtosis of the good quality shock innovations. Parameter λ will change the shape of the distribution in a way that will enable us to match the interquartile range of the empirical price change size distribution. Note that by setting $p = 0$, we have the model with Gaussian idiosyncratic shock innovations, as in Golosov and Lucas (2007); and by setting $\lambda = 0$, the idiosyncratic shock innovations will follow a Midrigan (2011)-type Poisson-distribution.

The production function implies an individual and aggregate labor demand $L_t(i, g) = A_t(i, g)Y_t(i, g)$ and $L_t = \int \sum_{g=1}^G L_t(i, g) di$. Substituting this individual 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 smoothly growing nominal GDP ($P_t Y_t$), we obtain a stationary period profit function (net of menu costs) as

$$\bar{\Pi}_t(i) = \sum_{g=1}^G \left[\frac{1}{1 + \tau_t} p_t(i, g)^{1-\gamma} - w_t p_t(i, g)^{-\gamma} \right] \left(\frac{1}{G} \sum_{g=1}^G p_t(i, g)^{1-\gamma} \right)^{(\gamma-\theta)/(1-\gamma)}. \quad (5)$$

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 store i on product g , and $w_t = W_t/P_t$ is the real wage. Importantly, this normalized profit function does not depend on the level of good quality shocks, but only on the vector of relative prices of store i . This means that the only idiosyncratic state variable of the firms is the vector of good-specific relative prices $\{p_t(i, g)\}_{g=1}^G$ or in short, $\mathbf{p}_t(i)$. The aggregate state variables are w_t (real wage), τ_t (value-added tax rate), π_t (inflation rate) and Γ_t (the cross-sectional distribution of relative price vectors), which we denote by $\Omega_t = (w_t, \tau_t, \pi_t, \Gamma_t)$. As it is apparent from equation (5), we can write the normalized profit function of the firm as $\bar{\Pi}_t(\mathbf{p}_t(i), w_t, \tau_t)$.

In each period, after observing the idiosyncratic good quality innovations, firms have to decide whether they would like to change nominal prices or not. Changing prices means that they must pay the (normalized) menu cost ϕ , which will enable them to set optimally their relative price vector. If they decide not to change prices, their relative price vector will be equal to the one inherited from the previous period, deflated by this period's inflation rate. The value of the firm if it chooses not to change its prices is

$$V^{NC}(\mathbf{p}_{t-1}(i), \Omega_t) = \bar{\Pi}\left(\frac{\mathbf{p}_{t-1}(i)}{1 + \pi_t}, w_t, \tau_t\right) + \beta E_t V\left(\frac{\{p_{t-1}(i, g)e^{\varepsilon_{t+1}(i, g)}\}_{g=1}^G}{1 + \pi_t}, \Omega_{t+1}\right), \quad (6)$$

where $\mathbf{p}_{t-1}(i)$ is the relative price vector – inherited from the previous period and adjusted with this period's good quality shock – at the time of the price 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 all product g .¹⁶

If the firm chooses to change its prices, it chooses the new relative price vector $\mathbf{p}_t^*(i)$ optimally, and it will have to pay a (normalized) menu cost ϕ , so its value will be

$$V^C(\Omega_t) = \max_{\mathbf{p}_t^*(i)} \left\{ \Pi(\mathbf{p}_t^*(i), w_t, \tau_t) - \phi + \beta E_t V\left(\{p_t^*(i, g)e^{\varepsilon_{t+1}(i, g)}\}_{g=1}^G, \Omega_{t+1}\right) \right\}. \quad (7)$$

Finally, the value of the firm will be determined by the upper envelope of V^{NC} and V^C :

$$V(\mathbf{p}_{t-1}(i), \Omega_t) = \max_{\{C, NC\}} [V^{NC}(\mathbf{p}_{t-1}(i), \Omega_t), V^C(\Omega_t)]. \quad (8)$$

2.4 The equilibrium

We consider a rational expectations equilibrium of the model. As there is no aggregate uncertainty in the model, firms know the equilibrium paths of aggregate state variables $\tau_t, \Gamma_t, w_t, \pi_t$. The equilibrium conditions are the following:

1. The representative consumer chooses $C_t(i, g), L_t, B_t$ to maximize her utility function (1) given her budget constraint (2), taking goods prices $\{P_t(i, g)\}$, interest rates R_t and the nominal wage W_t as given.
2. The firms are assumed to set nominal prices $P_t(i, g)$ to maximize their value function (8), (6), (7), given their current relative price vector $\mathbf{p}_{t-1}(i)$, and the values and future paths of aggregate state variables $\{\tau_t, \Gamma_t, w_t, \pi_t\}$. The firms also form correct beliefs about the random process of the idiosyncratic shocks $\{A_t(i, g)\}$.
3. The central bank increases money supply M_t with a constant growth rate g_M , with which it also keeps the nominal output growth constant. The seignorage revenue is redistributed in a lump-sum way.
4. The government sets value-added tax rates exogenously. It redistributes tax revenues in a lump-sum way.

¹⁶The relative price that was inherited from this period if there was no price change, $\mathbf{p}_{t-1}(i)/(1 + \pi_t)$, 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 and aggregate state variables, conditional on their current values.

5. Market clearing in all goods markets $C_t(i, g) = Y_t(i, g)$.
6. The net supply of nominal bonds is zero: $B_t = 0$.
7. Equilibrium in the labor market, implying that the nominal wage W_t equates aggregate labor demand and labor supply.

We solve for this equilibrium numerically using global solution techniques. Details of our numerical solution algorithm are in the Appendix. There we also describe our solution method for the transition path after unexpected monetary or tax shocks.

2.5 Model variants

We compare our baseline model (labeled “Mixed”) above with the model of Golosov and Lucas (2007), Midrigan (2011) and Calvo (1983). As discussed, the Golosov-Lucas model (labeled “Normal”) is nested in the baseline model when the probability of the small-variance idiosyncratic shock innovation is $p = 0$. The Midrigan-model (labeled “Poisson”) is also nested in the baseline model when we set the proportionality parameter between the variances at $\lambda = 0$.

In the Calvo-model (labeled “Calvo”), we set the menu cost to $\phi = 0$ and introduce an exogenous probability of price change, κ . The value function in this case is

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

which reflects that price change is not the firms’ decision any more, but it is an exogenous opportunity arriving with probability κ . We solve the Calvo-model with the same global solution methods that we use for the numerical solution of menu cost models.

3 Calibration

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

3.1 The data

To calculate moments of price change distribution, we use a data set of store-level price quotes in Hungary between December 2001 and December 2006 underpinning the consumer price index.¹⁷ We focus on processed food items, for several reasons. First, this is the largest 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, one aim of our paper is to compare the effects of positive and negative value-added tax changes on the prices. In this sector, the composition of the

¹⁷For a detailed description of the data set, see Gabriel and Reiff (2010).

subgroups facing different value-added tax rates are very similar.¹⁸ This permits us to disregard composition effects, and contributes to a clean measurement of the effects of positive and negative tax changes.¹⁹

We observe the prices of 128 different products in our dataset,²⁰ and each product is observed in 123 stores on average each month.²¹ 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.

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. The moments presented in the first row of Table 1 are sample averages disregarding observations during the tax-change months.

The second row of Table 1 confirms that our data is not special: our moments of the price change distribution 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 last quartile is more than 8% and the kurtosis of the price change distribution is around 4 that is significantly larger than the kurtosis of the normal distribution (3). One notable difference between the Hungarian and the US dataset is the level of the inflation rate. They are set to 0 in the US case, while it is around a still low 4.2% in Hungary. We are going to show that while inflation rate has substantial influence on the price flexibility of our model economy, our qualitative results are insensitive to the level of observed inflation of our sample.

3.2 Calibration

For some basic parameters, we use standard calibrations. We set the monthly discount factor $\beta = 0.96^{1/12}$, implying a 4% yearly real rate. We set the the inverse Frisch-elasticity of the

¹⁸An example of similar products facing different tax rates is ‘cookies’ facing the lower rate and ‘chocolate-chip cookies’ facing the higher.

¹⁹In 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. We provide further evidence on the lack of the composition bias 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.

²⁰98 affected by the tax increase and 30 by the tax decrease

²¹The number of item replacements and substitutions are small. Even in the rare case of item or store substitution, there is a variable in the data set advising us about this.

Table 1: Moments of regular price change distribution

Used in calibration	Data		Models		
	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	1.93
Interquartile range	8.13%	8%	8.13%	9.35%	6.3%
Inflation	4.23%	0%	4.23%	4.23%	4.23%
Other moments					
1st decile of abs size distr	2.8%	3%	2.7%	2.4%	3.1%
1st quartile of abs size distr	4.6%	5%	4.9%	3.9%	6.3%
Median of abs size distr	7.7%	9%	7.5%	6.1%	9.5%
3rd quartile of abs size distr	12.7%	13%	13%	13.3%	12.6%
9th decile of abs size distr	19.4%	21%	20.8%	20.1%	15.2%
Fraction < 0.25 of mean	9.8%	8%	10.5%	9.3%	6.5%
Fraction < 0.5 of mean	27%	25%	26.7%	34.8%	13.4%

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 menu cost models of Midrigan (2011) and Golosov and Lucas (2007). This condition also implies full long term inflation pass-through of value-added shocks that we also observe in the data. We set the money growth (g_M) equal to the yearly trend inflation rate, 4.2%; the persistence of the money shock 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 5.²² In line with Midrigan (2011), we set the elasticity of substitution for goods within a retail 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 Midrigan (2011) and Gertler and Leahy (2008).²³ We assume that the correlation between the idiosyncratic shocks are 60% that is also in line with the calibration of Midrigan (2011).

We calibrate the size of 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

²²This 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 the aggregate behavior of the model economy.

²³There 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)). We show that our results are robust to using alternative values.

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.

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 contains the calibrated parameters of our baseline model (mixed) and the alternatives.

Table 2: Calibration

Parameters	Mixed	Poisson	Normal
ϕ	2.4%	1.8%	5.1%
σ_A	4.3%	4.3%	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 for free). So the overall cost of price adjustment is around 0.15% of steady-state revenues, which is on the lower end of empirical estimates.²⁴ 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 calibration favors a mixed normal idiosyncratic shock distribution where the low volatility shock is arriving with a high 91% probability and with a standard deviation that is 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.²⁵ The Poisson distribution implies a probability of 0 idiosyncratic

²⁴Levy et al. (1997) estimate this cost to be 0.7% of revenues, while Zbaracki et al. (2004)'s estimate is 1.2%.

²⁵Note 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 $\sigma_2 = \sigma_A/\sqrt{(p/\lambda + 1 - p)} = 14\%$ in the large standard-deviation case in the mixed normal model.

shock that is very close to our baseline, but disallowing any variation of the low volatility shock results in a menu cost estimate that is 25% 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 of small shocks require menu costs that are more than twice as high as our baseline to keep steady state frequency at bay.

The third column of Table 1 shows the ability of our baseline model in matching 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 in matching 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 price changes. 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 overestimates 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 underestimated 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 when using 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 the multi-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 virtually indistinguishable from that of the Poisson distribution, but they both dominate the implied price changes of the normal distribution.

3.3 Macro implications

As we have shown in the previous section, our baseline model is able to match 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.

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 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 to 0.²⁶ Later, we are going to show how the inflation rate influences the implications of the model. Figure 2 presents the impulse responses.

The pass-through on the figure is measured as the proportion of the price level effect over the money supply effect of the shock that is expressed as

$$\gamma_t = \frac{\sum_i 0^t \pi_i - \bar{\pi}}{\sum_i 0^t \Delta m_i}. \quad (9)$$

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. It is close to the implications of the model using the Gaussian distribution as in Golosov and Lucas (2007). Its implications are very

²⁶Note 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 somewhat. In the Appendix, we show in a related model 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 real effects. They are substantially lower than the Calvo (1983) and the Poisson model of Midrigan (2011), and almost as small as the model with the Normal distribution of Golosov and Lucas (2007).

far from those of the model with Poisson distribution that generates real effects close to that of the Calvo (1983) model. These results imply that the macro-implications of these models are 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) argues, 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, we follow Costain and Nakov (2011). The inflation rate can be expressed as

$$\pi = \int x^*(p_{-1})\lambda(p_{-1})\psi(p_{-1})dp_{-1}, \quad (10)$$

where x^* is the desired nominal price change, λ is the hazard function of the price change (equal to 1 or 0), and ψ is the steady state distribution of products. All functions depend on the

individual state variable p_{-1} (the last period's productivity-adjusted relative price). Dependence on the aggregate states are suppressed for notational convenience.

The period t inflation effect of an aggregate shock is given by

$$\pi_t - \pi = \underbrace{\bar{\lambda}\Delta\bar{x}^*}_{\text{intensive}} + \underbrace{\Delta\bar{\lambda}\bar{x}^* + \Delta\bar{\lambda}\Delta\bar{x}^*}_{\text{extensive}} + \underbrace{\Delta \int_{p_{-1}} (x^* - \bar{x}^*) \lambda \psi}_{\text{selection}}, \quad (11)$$

where an average of variable y denoted by $\bar{y} = \int y(p_{-1})\psi(p_{-1})dp_{-1}$, 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. 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 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:

$$\gamma_t^m = \frac{\pi_t - \bar{\pi}}{(\sum_{i=0}^t \Delta m_i - \sum_{i=0}^{t-1} (\pi_i - \bar{\pi}))}. \quad (12)$$

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

$$\bar{\gamma} = \sum_{t=1}^T w_t \gamma_t^m. \quad (13)$$

The impulse responses shown on Figure 2 implies a price flexibility measure of 68% in our baseline mixed normal model, and 25%, 84% and 18% in the Poisson, Gaussian and Calvo models respectively. In our experiment we find that the difference between the models is driven by the difference in the selection effect. For the small money growth shocks of our experiments, extensive margin effect is negligible. The intensive margin effect, in turn is very close in each models, comparable to the pass-through of 17.6% of the Calvo model. The selection effects, however, are very different: contributing to the measured pass-through by 39% in the Mixed

Normal case and 55% in the Gaussian case, while they are only 2% in the Poisson case. What explains the difference?

We can illustrate the sensitivity of the selection effect on the shape of the idiosyncratic shock distribution by using a menu cost model that is less realistic than our baseline, but more tractable. Instead of the multi-product assumption, we assume instead that each firm produces a single product, but with a certain probability it can change its price for free. To mimic our baseline two-product assumption, we calibrate the probability of free menu cost to 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, making the computational burden much easier. 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).²⁷ 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. We detail this model and its calibration in the Appendix. The key parameter of our mixed normal distribution is the relative standard deviation of the small variance normal distribution relative to the large variance component (λ). For $\lambda = 0$, we have the Poisson distribution and for $\lambda = 1$ we have the Gaussian distribution. The parameter best matching our observed steady state price change distribution in this heterogeneous menu cost model is 18%.²⁸ For each λ value, we recalibrate the model to match the frequency, size and kurtosis of the steady state price change distribution,²⁹ and simulate one standard deviation money growth shocks to calculate the sensitivity of price level frequency, its components and the real effects of monetary policy shocks to our distributional assumptions. Figure 3 shows our results.

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

²⁷It limits the model's applicability for large shocks with influence on the price-change frequency, because the assumption unrealistically limits the endogenous increase of 'free' adjustment in the multi-product setting.

²⁸Somewhat higher than 8.8% in our baseline case, but still much closer to the Poisson case than to the Normal

²⁹We do this with positive inflation as in our baseline case, switching inflation off for the shock experiments. In the Appendix, we show that the results are not sensitive to doing this exercise with calibration of 0 inflation.

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 were 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 cumulative real effects for very small λ values are high (the same measure in the Calvo (1983) model is somewhat over 5), but it quickly decreases to insignificant levels. The third and the fourth panel shows two factors in this particular model: 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 model, the standard single-product version of the menu cost model. The key channels in this model are very similar to the more realistic versions, but the intuition is much cleaner. To support the visualization of the model, it is instructive to disentangle the actual price change distribution to the product of two objects: 1) the distribution of 'desired' price changes: the adjustments that firms would make if menu costs magically disappeared for a period and 2) the adjustment hazard that determines

the probability of adjustment for a particular price change. In a single-product model with permanent idiosyncratic shocks, this hazard function takes a particularly simple sS form: firms have a lower and an upper adjustment threshold and 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.

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.

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.

An aggregate nominal shock uniformly increases the desired price changes of all firms: it shifts the desired price change distribution to the right. A marginal shock will keep the fraction

of price adjusters constant: the measure of new price increasing firms and the measure of new non-price-decreasing firms cancel each other out. As a result, the extensive margin effect will be zero. The adjusting firms respond to the aggregate shock by increasing their prices more or decreasing it less generating the intensive margin effect. The selection effect is determined by the changing composition of price adjusters: the new price increasers (on the right of the inaction band) and the new price non-decreasers (on the left of the inaction band), and its magnitude depends on two factors: 1) the mass of firms at the inaction thresholds and 2) the width of the inaction bands. As both the new price increasers and the new non-price decreasers contribute to the aggregate effects, their combined mass influences the strength of the pass-through. The width of the inaction band determines the desired price changes of these new adjusters: the wider it is, the higher the price adjustment going to be contributing to the aggregate price flexibility.

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 realistic menu cost, it also influences the width of the inaction band. Both the measure of firms at the inaction thresholds and the width of the inaction band is 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 inaction threshold 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), 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). As a result, relatively small deviation from the Poisson distribution will have significant effect on the selection, increasing the inflation pass-through close to the Golosov and Lucas (2007) model.

3.4 Effects of trend inflation

Trend inflation, mainly through its effect on the steady state price change distribution has strong effect on the measure of pass-through for very low levels of λ , and the Poisson case ($\lambda = 0$) in particular. Figure 5 show the results. 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 mass of firms at the inaction thresholds.

The shape of the desired price change distribution of the Poisson case can explain the high mass of firms at the inaction band (see Figure 6). The steady state desired price change distribution is going to be highly asymmetric, with lot of firms at the positive inaction threshold. These firms got to the inaction threshold just because of the trend inflation, without idiosyncratic shocks explaining their high mass.

[TO BE COMPLETED]

Figure 5: Idiosyncratic distribution and real effects

The figure shows the implied price level flexibility, real effects, the width of the inaction band and the mass at inaction thresholds as a function 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 mass of firms at the inaction thresholds.

Figure 6: 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.

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 steady state price change distribution might contain 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 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 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 implies too many, while a too large menu cost implies too few price changes. 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 costs 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.³⁰ The tax changes are the consequence of the 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 lower rate from 12% to 15%. Then, in January 2006, before the general elections, they decreased the top rate from 25% to 20%, and in September, after the election, they increased the lower rate from 15% to 20%. These changes were announced 5, 6 and 3 months in advance, respectively, which we take into account in our 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.³¹ Over

³⁰This 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 Economy decree and its explanation.

³¹The VAT-decrease of January 2006 was ultimately reversed in 2009, when the government has increased the

the period, the country has enjoyed a stable macroeconomic environment with stable growth, low inflation rates and stable exchange rates. Monetary policy, furthermore, did not respond to the shocks: the inflation targeting central bank had expressed in advance that it was “seeing through” the direct 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

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.³² 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 costs 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
Inflation pass through	+3%	74%	67%	133%	42%	10.1%
	+5%	99%	91%	148%	39%	10.1%
	-5%	33%	28%	13%	48%	8.6%
Frequency	+3%	52%	33%	63%	18%	12.6%
	+5%	62%	55%	90%	25%	12.6%
	-5%	27%	20%	11%	17%	12.6%

We use these observations to test the predictions of our model. For this, we take our baseline

20% rate to 25% to finance its budget shortfall following a financial crisis related drop in its revenues. In section XX 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.

³²Gagnon (2009) documents similarly large frequency effect after large nominal shocks in Mexico.

Figure 7: 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.

model, calibrated to match the steady-state price change distribution, and hit it with a one-time permanent tax shock that is announced before its implementation. We do a similar experiment with alternative menu cost models with Poisson and Gaussian distributions and the Calvo (1983) model. The size of the shock is calibrated to match the size of the tax change, and the lags

between the announcement and the tax changes are similarly calibrated to be in line with reality.

The second column of Table 3 shows 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 predict 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 as high 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 show. Using the Poisson distribution we substantially overestimate the frequency response and the pass-through of the positive shocks, and also overestimates 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 8 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.

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 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

Figure 8: 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).

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 show 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 ob-

serve 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 there does not have sufficient kurtosis.

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

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

To conclude, the empirical evidence on the large and asymmetric pass-through of symmetric tax shocks in Hungary as a strong 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 evidence supporting 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 Conclusion

We have calibrated and tested a new quantitative pricing framework. We have shown that our proposed menu cost model with a generalized, mixed-normal idiosyncratic shock distribution is highly successful in matching key moments of the steady state price change distribution and makes excellent out-of-sample predictions on the observed frequency increase and asymmetry of large value-added tax shocks in Hungary, outperforming alternative models. We have shown that these effects are the consequence of the adjustment on the extensive margin.

Our baseline menu cost model that is highly successful in matching pricing facts also implies high Caplin and Spulber (1987)-type selection-effect. The predicted low real effects of monetary policy shocks are in contrast to macro-economic 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)), 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. Mackowiak and Wiederholt (2009), Mankiw and Reis (2002), Woodford (2003a)) can be necessary to explain the observed effects of aggregate nominal shocks.³³ 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., M. Gonzalez-Rozada, A. Neumeyer, and M. Beraja (2011). From hyperinflation to stable prices: Argentina's evidence on menu cost models.
- Alvarez, F. E., F. Lippi, and L. Paciello (2011). Optimal price setting with observation and menu costs. *The Quarterly Journal of Economics* 126(4), 1909–1960.
- Ball, L. and N. G. Mankiw (1994). Asymmetric price adjustment and economic fluctuations. *The Economic Journal* 104(423), 247–261.
- Baudry, L., H. Le Bihan, P. Sevestre, and S. Tarrieu (2007). What do thirteen million price records have to say about consumer price rigidity? *Oxford Bulletin of Economics and Statistics* 69(2), 139–183.
- Bernanke, B. S. and A. S. Blinder (1992, September). 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.
- 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.

³³Bils and Chang (2000) also emphasizes differences in observed price responses to demand and cost shocks, and suggests potential explanations based on market structure.

- 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.
- Devereux, M. B. and H. E. Siu (2007). State dependent pricing and business cycle asymmetries. *International Economic Review* 48(1), 281–310.
- Ellingsen, T., R. Friberg, and J. Hassler (2006). Menu costs and asymmetric price adjustment. Discussion Paper 5749, Centre for Economic Policy Research.
- 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., 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.
- Jonker, N., C. Folkertsma, and H. Blijenberg (2004). An empirical analysis of price setting behaviour in the Netherlands in the period 1998-2003 using micro data. Working Paper 413, European Central Bank.
- 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 B. A. Malin (2010). Microeconomic evidence on price-setting. Volume 3 of *Handbook of Monetary Economics*, pp. 231 – 284. Elsevier.
- 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.
- 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.

- 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.
- Rotemberg, J. J. (1982). Monopolistic price adjustment and aggregate output. *The Review of Economic Studies* 49(4), 517–531.
- 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.
- 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.
- Yang, H. and L. Ye (2008). Search with learning: Understanding asymmetric price adjustments. *The RAND Journal of Economics* 39(2), 547–564.
- Zbaracki, M. J., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004). Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *The Review of Economics and Statistics* 86(2), 514–533.

6 Appendix

6.1 Equivalent 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)), 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

$$W_t = \mu P_t Y_t = \mu M_t. \quad (14)$$

Substituting this into equation (5) about the periodic normalized profit, one can easily re-write the profit equation as

$$\bar{\Pi}_t(i) = \left(\frac{1}{1 + \tau_t} \right) \sum_{g=1}^G \left[p_t(i, g)^{1-\gamma} - \frac{\mu M_t (1 + \tau_t)}{P_t} p_t(i, g)^{-\gamma} \right] \left(\frac{1}{G} \sum_{g=1}^G p_t(i, g)^{1-\gamma} \right)^{(\gamma-\theta)/(1-\gamma)}. \quad (15)$$

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.³⁴ 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 compare the predictions of the model for large monetary shocks.

6.2 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 output, so under unchanged money supply, we will have full pass-through into the gross nominal prices in this case as well.³⁵

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

³⁴For 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.

³⁵We are also using the flexible price solution as starting values for the iterative procedures in our numerical solution method.

choose prices to maximize this, implying the following optimal relative price:

$$p_t^i, g^* = (1 + \tau_t) \frac{\gamma}{\gamma - 1} w_t, \quad (16)$$

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

$$w_t = \frac{\gamma - 1}{\gamma(1 + \tau_t)}. \quad (17)$$

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

$$Y_t = \frac{\gamma - 1}{\gamma\mu(1 + \tau_t)}. \quad (18)$$

The nominal price level can be obtained as

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

The expected growth rates are

$$E(g_Y) = -E(g_{1+\tau_t}), \quad E(\pi_t) = g_M + E(g_{1+\tau_t}) \quad (20)$$

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

6.3 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} .³⁶ 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:

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).

³⁶There is one other relevant aggregate variable: R_t (return on nominal bonds). We showed earlier that our constant growth assumption on nominal income and money stock implies that $R_t = \frac{e^{g_{PY}}}{\beta}$.

2. Given this guess and the steady-state inflation rate, we use a fine grid on relative prices and idiosyncratic shocks to solve for the optimal pricing policies of individual firms. We use value function iteration.
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 the resulting distribution is the same as the distribution one period earlier.
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 is exactly equal to 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 $\{\pi_1, \dots, \pi_T\}$. Our initial guess is the full immediate pass-through.
3. Given this guess, we calculate the resulting equilibrium path of the real wages: $\{w_1, \dots, w_T\}$. 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.
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.

6.4 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 5 indicates, 102 of the 128 products in our original processed food sample were hit by this new 5%-point tax increase.³⁷ 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 5: 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%

The third column of Table 5 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)

³⁷The 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.

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.³⁸

³⁸If 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.