BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2009-11

# Information Flows and Aggregate Persistence

by Oleksiy Kryvtsov

www.bank-banque-canada.ca

Bank of Canada Working Paper 2009-11

April 2009

# Information Flows and Aggregate Persistence

by

**Oleksiy Kryvtsov**

Canadian Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
okryvtsov@bankofcanada.ca



ISSN 1701-9397

# Acknowledgements

This paper is a heavily revised Chapter 3 of my Ph.D. dissertation at the University of Minnesota. I am grateful to V.V. Chari, Larry Jones, Pete Klenow, and Ellen McGrattan for their invaluable guidance. I thank Francisco Covas, Mike Golosov, Chris Hellwig, Hashmat Khan, Guido Lorenzoni, Bob Lucas, Cesaire Meh, Rhys Mendes, Ricardo Reis, Malik Shukayev, Henry Siu, Yaz Terajima, Alexander Ueberfeldt, and numerous seminar participants for their comments and suggestions. Amanda Armstrong, Tom Carter, and Stephen Doxey provided excellent assistance.

# Abstract

Models with imperfect information that generate persistent monetary nonneutrality predominantly rely on assumptions leading to substantial heterogeneity of information across price-setters. This paper develops a quantitative general equilibrium model in which the degree of heterogeneity of information is determined endogenously. In the model, firms use two technologies to acquire information: costly updating to full information and costless learning from publicly observed market signals. Price changes of firms that update information infrequently are synchronized with market signals. This leads to an externality whereby less frequent updating increases the information conveyed by prices and quantities. When the model is calibrated to moments from a panel of BLS commodity sectors, it is found that the private value of costly updating to full information is close to zero, market signals are informative, and the real effects of monetary shocks are small.



*JEL classification: D83, E31, E32*
*Bank classification: Business fluctuations and cycles; Inflation and prices; Transmission of monetary policy*



# Résumé

Les modèles où l'information est imparfaite et où la monnaie n'est pas neutre en longue période reposent essentiellement sur des hypothèses qui introduisent une forte hétérogénéité dans l'information dont disposent les décideurs de prix. L'auteur construit un modèle d'équilibre général quantitatif dans lequel le degré d'hétérogénéité de l'information est déterminé de manière endogène. Les entreprises du modèle acquièrent de l'information par deux voies : au moyen d'une mise à jour qui leur donne accès à une information complète mais qui comporte un coût, et à la faveur d'une connaissance gratuite, déduite de signaux de marché de notoriété publique. Les entreprises qui actualisent peu fréquemment l'information modifient leurs prix au rythme des signaux du marché. Ce comportement donne lieu à une externalité du fait que la sporadicité des mises à jour accroît la quantité d'information véhiculée par les prix et les volumes. En étalonnant son modèle en fonction de moments calculés à partir des données du Bureau of Labor Statistics sur les prix d'une gamme étendue de produits, l'auteur constate que la mise à jour non gratuite est d'une valeur quasi nulle pour l'entreprise, que les signaux du marché apportent de l'information et que les chocs monétaires ont peu d'effets réels.



*Classification JEL : D83, E31, E32*
*Classification de la Banque : Cycles et fluctuations économiques; Inflation et prix; Transmission de la politique monétaire*

# 1 Introduction

What is the nature of short-run monetary nonneutrality and how important is it for business cycles? These questions have motivated business cycle researchers for decades. In seminal work, Lucas (1972) laid out foundations for the nonneutrality of money under rational expectations and imperfect information about fundamental disturbances. In Lucas' model, information becomes public knowledge soon after the monetary shock, which thus has strictly short-lived real effects. After Lucas (1972), the search for sources of persistent imperfect information that could generate some degree of monetary nonneutrality closer to that observed in the data[1] since become a priority for business cycle scholars.

The most recent generation of monetary business cycle models based on imperfect information has been successful in generating substantial monetary nonneutrality and inflation persistence. The key feature accounting for the long-lasting real effects of monetary policy in these frameworks is dramatic heterogeneity of information across price-setters, stemming from assumptions that leave information sets non-coincident across firms. Prices set by such firms are determined by the marginal cost of production and forecasts of competitors' prices.[2] Following a money shock, the speed at which a firm adjusts its price thus depends on forecasts as to the speed at which competitors' prices will responds. When information is very dispersed, modest price adjustments are expected among competitors, leading to slow price responses and, hence, long-lasting real effects.[3]

Heterogeneity of information in the existing models of adjustment under imperfect information is introduced by imposing restrictions on the way information is obtained. These restrictions can be divided in two types. The first type limits the frequency with which uninformed agents update their information sets.[4] It is typically assumed that updated information sets contain the full history of events. The frequency of

[1] Romer and Romer (1989, 2004), Bernanke and Mihov (1998), Christiano, Eichenbaum, and Evans (2000), provide some empirical evidence on the size of real effects of monetary shocks.

[2] An important, and standard, assumption here is complementarity in pricing, i.e., an increase in the average price level implies that each firm has an incentive to raise its own price. The degree of complementarity depends on goods' demand elasticity, returns to scale in production, and the elasticity of labor supply.

[3] See Hellwig (2008) for a review of incomplete information theories.

[4] Mankiw and Reis (2002, 2006), Bonomo and Carvalho (2004), Reis (2006).

updating is either exogenous, as in Mankiw and Reis (2002), or endogenous due to a fixed cost, as in Reis (2006). As long as the frequency of information updating is not too high (around once a year), the heterogeneity of information between informed and uninformed price-setters is significant, and sizeable monetary nonneutrality results.

Under the second type of restriction on information acquisition, firms update their information continuously by observing noisy signals of their price targets.[5] The noise is specific to each firm and is typically interpreted either as fluctuation in some information-processing technology or as a constrant on the firm's information-processing capacity, in the spirit of Sims (2003). With noise large enough relative to the signal, the firm's ability to accurately predict competitors' prices is limited, retarding the adjustment of prices following a monetary shock. Hence, heterogeneity of information, hardwired into models of imperfect information via either of these assumptions, ensures that monetary shocks will have considerable real effects.

This paper extends the existing framework to a setting in which the degree of heterogeneity in information is determined *endogenously.* This is achieved by assuming that price-setting firms have access to two technologies for acquiring information: public and private. Public technology is available every period at no cost and allows firms to infer state history from observing a market signal, namely the demand for the firm's output at a chosen price. This signal is public as it is perfectly correlated across firms in the same production sector; and it is endogenous since it reflects the price decisions of firms in the sector. Private technology is available at fixed cost and provides explicit knowledge of the full state history. For simplicity, use of the public technology will be called "learning", while the private technology is dubbed "updating".

In the model, firms trade off the use of efficient but costly updating for less costly but less efficient learning. More informative market signals improve the firm's forecasting ability and thus weakens the incentive to update the information set. In general equilibrium, there is an externality whereby the frequency of updating affects the amount of information conveyed by market prices and quantities. In the presence of fixed cost, updating is infrequent, which implies that there is a number of firms

[5] Woodford (2002, 2008), Hellwig (2002), Mackowiak and Wiederholdt (2009), Gorodnichenko (2008), Lorenzoni (2009).

that have not updated for some time. Information sets of such firms tend to overlap cosiderably since they mostly contain common observations of public signals. Price changes by less informed firms are thus synchronized more closely with public signals and with other such firms. Greater synchronization, in turn, implies that market signals convey more information about competitors' prices. In equilibrium, how informative market signals are affects private value of updating to full information thus determining the degree of heterogeneity of information across price-setters.

The model is calibrated to a cross-section of 111 BLS commodity sectors and is used to ask: what is the private value of information updating, how much information is conveyed by market signals, and what are the implications for the size of monetary nonneutrality? The main finding is that public market signals are very informative so that the private value of updating is close to zero. The incentive to economize on fixed costs by updating infrequently is reinforced by information externality: firms that do not update, improve the information content of market signals, which further weakens the incentives to update. When market signals are highly informative, firms can quickly disentangle nominal and real disturbances, which results in transient monetary nonneutrality. In the benchmark model, allowing for endogenous heterogeneity of information decreases persistence of real effects of monetary shocks from more than one and a half years to less than three quarters.

Price-setting models with endogenous information choice may face a problem of multiple equilibria, a point raised by Hellwig and Veldkamp (2009). They show that multiplicity stems from strong incentives for firms to synchronize the timing of their information updating. To circumvent this problem, it is assumed in the model that the fixed cost of updating is stochastic and drawn from a continuous dirstribution. This leads to staggering of updating decisions and implies a continuous relationship between distribution of firms over information sets and precision of the public signal, alleviating multiplicity of equilibria.

Endogenous determination of the number of informed agents and the precision of the public signal is akin to the setting in Grossman and Stiglitz (1980). As in their asset trading model, here market signals convey information from more to less informed agents, and the informativeness of the signals depends on the equilibrium number of informed agents. This number, in turn, depends on the cost of information, the extent of uncertainty, the strength of the market signal, and the amount of information gathered by the updating agents. In contrast to Grossman and Stiglitz (1980), this paper uncovers the dynamic aspect of the interaction between the number of firms that pay for information and the informativeness of the market signal. Firms that choose not to pay for information base their decisions on a long sequence of public signals and hence behave alike. This implies, in contrast to Grossman and Stiglitz (1980), that information about the state history conveyed by market variables *increases* with the number of agents not paying for information.

In Gorodnichenko (2008) observation of price as an endogenous free public signal leads to an externality whereby firms delay costly updating of information and instead learn costlessly from those firms that choose to update. Delays in price changes due to this information externality lead to inflation persistence. This paper differs from Gorodnichenko (2008) in two respects: firms' profits depend on the price level, and firms that do not update information are still allowed to change prices. The observed price level thus conveys information about prices set by less informed firms. Since these firms tend to behave alike, the public signal is informative. As a result, the payoffs from updating decrease, thus increasing the number of prices set by less informed firms and reinforcing the informativeness of the price level. Hence, in contrast to Gorodnichenko's case, the information externality in this paper reduces aggregate persistence.

The paper proceeds as follows: Section 2 develops the model and defines an equilibrium. Section 3 explains how the model is aggregated and solved. Section 4 presents and discusses the results of model simulations. Section 5 concludes.

# 2 Model

Consider a Lucas-type economy, which consists of an infinite number of structurally identical island economies. Each island is populated by a unit measure of final good producers and a unit measure of intermediate good producers. There is no trade or communication across islands. Time is indexed by $t = 1, 2, ....$ There are two sources of uncertainty on each island: the growth rate of the money stock, $\mu_t$, and the disturbance of the demand for intermediate goods, $\phi_t$.

It is assumed that innovations to the money growth process are identical across islands, whereas innovations to each demand shifter are i.i.d. across islands. Hence, money growth and the disturbance in demand for goods are interpreted as aggregate and island specific shocks to the aggregate economy.[6] In this setup, dynamics stemming from fluctuations in the money growth are identical across islands, whereas disturbances to island demand have no effect on the dynamics on other islands. Hence, to characterize aggregate dynamics, it is sufficient to characterize fluctuations on an island conditional on the aggregate shock.

Let $z_t = \{\mu_t, \phi_t\}$ denote realization of uncertainty on an island in period $t$, and $z^t = \{z_0, z_1, ..., z_t\}$ denote the island's history of states through period $t$.[7] The rest of this section lays out an island economy's setup and defines an equilibrium.

## 2.1 Final good producers on an island

Final good on each island is produced by a unit measure of competitive producers using intermediate goods as inputs. A Dixit-Stiglitz production function exhibits constant elasticity of substitution over a variety of differentiated input goods. Let $Y_t$ denote the quantity of the final good produced on an island, and $y_t(\tau)$ be the quantity of input goods purchased from an intermediate good producer $\tau$. Let $P_t$ and $P_t(\tau)$ respectively denote the prices of final and intermediate goods. Final good producers solve the following problem:

$$\max_{\{y_t(\tau)\}} \ P_t Y_t - \phi_t^{-1} \int P_t(\tau) y_t(\tau) d\tau$$

subject to

$$Y_t = \left[ \int y_t(\tau)^{\frac{\theta-1}{\theta}} d\tau \right]^{\frac{\theta}{\theta-1}} . \tag{1}$$

Here $\phi_t$ denotes a random shift in the demand for all intermediate input on the island. We assume that $\ln \phi_t$ follows a zero-mean AR(1) process $\ln \phi_{t+1} = \rho_\phi \ln \phi_t + \varepsilon_{\phi t+1}$,

[6]This island economy can also be set up as a multi-sector economy - with islands represented by sectors and consumption expenditure shares constant across sectors. An island paradigm is convenient for emphasizing that markets are physically separated.

[7]Throughout, island index is omitted.

with i.i.d. normal innovations $\varepsilon_{\phi t} \sim N(0, \sigma_\phi^2)$.

First-order conditions yield the demand function for intermediate goods:

$$y_t(\tau) = Y_t \phi_t^\theta \left( \frac{P_t}{P_t(\tau)} \right)^\theta , \qquad \forall \tau \in [0,1] \ . \tag{2}$$

Zero-profit condition implies that the island price index is

$$P_t = \phi_t^{-1} \left[ \int P_t(\tau)^{1-\theta} d\tau \right]^{\frac{1}{1-\theta}} \ . \tag{3}$$

As shown in Section 3, both nominal and real shocks have very similar effects on intermediate firms' outputs and prices. Sorting out their impact on the island market constitutes the signal extraction problem faced by the intermediate good producers.[8]

## 2.2 Intermediate good producers

On each island, there is a unit measure of ex-ante identical enterpreneurs who own firms producing perishable intermediate goods. Each enterpreneur is endowed with a unit of labor per period and with a production technology that converts labor input into firm-specific variety of intermediate good. For simplicity, it is assumed that an enterpreneur's labor is productive only for his own firm. Then it can be written that the firm produces $y(\tau)$ units of intemediate good at a cost of $\psi y(\tau)^{1+\xi} / (1+\xi)$ units of the final good, where $\xi$ represents the degree of decreasing returns to the labor input.[9] Markets for differentiated goods are monopolistically competitive, so each firm sets the price for its good. Asset markets are incomplete so that enterpreneurs cannot share any risks. Entepreneurs use their total income (possibly including transfers from the government) to finance their expenditures on consumption of final goods. Information that the enterpreneur obtains in the final goods market is not revealed to the firm.[10] Our assumptions about the structure of the labor and asset

[8]Alternatively, the island shock can represent the shift in production technology: $Y_t = \phi_t \left[ \int y_t(\tau)^{\frac{\theta-1}{\theta}} d\tau \right]^{\frac{\theta}{\theta-1}}$. In demand equation (2) $\phi_t^\theta$ is then replaced with $\phi_t^{\theta-1}$ and price equation (3) is unchanged. Results are unaffected.

[9]For example, technology $y = l^a$, where $l$ is the labor input, gives $\xi = 1/\alpha - 1$.

[10]For example, an enterpreneur can consist of the worker, who earns income, and the shopper, who spends it, and the worker and the shopper do not communicate.

markets prevent firms from inferring information conveyed by labor and asset prices. These assumptions can be relaxed but would require more uncertainty to ensure that information remains imperfect.

In contrast to final good producers, the intermediate good producers do not observe the full history of events on an island. They acquire information via two technologies, one that is more efficient and costly, "updating", and another that is less costly but is also less efficient, "learning". Updating is available at fixed cost, $\zeta_t(\tau)$, and provides observation of the full history through to the beginning of period $t$, i.e., $z^{t-1}$. The fixed cost, denominated in units of the final good, is assumed to be i.i.d. across enterpreneurs and over time, drawn from a distribution with differentiable c.d.f. $G(\cdot)$ satisfying [11]

$$\begin{aligned} &G(0) = 0 \ , \\ &G'(\zeta) > 0, \quad \zeta \in (0,1) \ , \\ &G(1) = \zeta_{\max} < \infty \ . \end{aligned}$$

Learning is costless: agents make inferences about the history $z^t$ from the sequence of public signals, $s^t = \{s_0, s_1, ..., s_t\}$, where $s_t$ is the signal in period $t$. It is assumed that in period $t$ these producers observe their own prices and quantities, $P_t(\tau)$ and $y_t(\tau)$, but not island prices and quantities. From (2), this is equivalent to observing the signal $s_t = Y_t \phi_t^{\theta} P_t^{\theta}$. Thus the signal is both public, as it is observed by all firms on the island, and endogenous, since it reflects price and quantity decisions made by other firms on the island.

At the beginning of period $t$, the information set of a producer that updated to full information $\tau$ periods ago is $I_{\tau,t} = \{z^{t-\tau}, s^t\}$. Information that is due to the market signal, $s^t$, is public as it is observed by all firms on the island, whereas information due to updating, $z^{t-\tau}$, is (partially) private as it is observed only by those firms that pay the fixed cost. The resulting hierarchical information structure on an island is thus very tractable: information sets of less informed firms are finer subsets of those for more informed firms, i.e., $I_{\tau+1,t} \subset I_{\tau,t}$, for $\tau = 1, 2, ....$ This implies that less informed firms mostly rely on information conveyed by public signals. The

[11] Dotsey, King, Wolman (1999) employ similar specification for the fixed cost of changing prices.

extent of heterogeneity of information is characterized by the distribution of firms across information sets $I_{\tau,t}$: firms are less heterogeneous in terms of information when a smaller fraction of firms opt to update to full information. The key feature of the model is that heterogeneity of information is endogenous in that firms' decisions regarding updating determine the distribution of firms across information sets, $I_{\tau,t}$.

The timing of events within a period is as follows. At the beginning of period $t$, producers form their prior distributions about the state history $z^t$. Next, the fixed costs of updating are realized and firms decide whether or not to update. Those that update pay fixed cost and directly observe the history through to last period, $z^{t-1}$. After that, the current state $z_t$ is realized, goods markets open, intermediate good firms observe current signal $s_t$, set current period prices $P_t(\tau)$, produce and trade in the final good market, collecting profit $\Pi_t(\tau) \equiv \frac{P_t(\tau)y_t(\tau)}{P_t} - \psi \frac{y_t(\tau)^{1+\xi}}{1+\xi}$. After goods markets are closed, firms form their posterior distributions about $z^t$ based on the signal history $s^t$.[12]

Firms whose last updating to full information occured in the same period have identical information sets. Combined with the assumption that the cost of updating is i.i.d., this implies that such firms also choose the same prices and quantities. Hence, the distribution of firms can be indexed by the number of periods since last updating: $\tau = 1, 2, ...$, where $\tau = 1$ denotes producers that updated at the beginning of the current period.

The type $\tau$ firm solves the following dynamic problem :

$$V_{\tau,t} = \max_{P_{\tau,t}, y_{\tau,t}} E_{\tau,t}\left[\frac{P_{\tau,t}y_{\tau,t}}{P_t} - \psi\frac{y_{\tau,t}^{1+\xi}}{1+\xi} + \frac{\beta}{\zeta_{\max}}\int_0^{\zeta_{\max}} \widetilde{V}_{\tau+1,t+1}(\zeta)\, dG(\zeta)\right], \quad (4)$$

$$\widetilde{V}_{\tau+1,t+1}(\zeta) = \max\left\{E_{\tau+1,t+1}(V_{1,t+1}) - \zeta, V_{\tau+1,t+1}\right\}, \quad (5)$$

subject to (2) and Bayesian laws of motion for distributions over individual state histories.[13] The initial distribution in period 0 is assumed to be equal to the uncon-

[12] It is assumed that after observing current signal, it takes time for firms to form their posterior distributions, so that price decisions are based on prior distributions of the state history. This assumption also separates pricing and information decisions and makes the problem more tractable without a loss of generality.

[13] Characterization these laws of motion is delayed till Section 3, in which for the linearized version of the model the Bayesian updating implies that the laws of motion for conditional expectations of the state are given by the Kalman filter equations.

ditional distribution of $z_0$.

In (4) and (5) $\beta \in (0,1)$ is a discount factor, and $E_{\tau,t}(\cdot) \equiv E(\cdot|I_{\tau,t})$ denotes the expectation based on information set $I_{\tau,t}$. Value function $\widetilde{V}_{\tau,t}(\zeta)$ $(V_{\tau,t})$ represents future profit streams to firm $\tau$ before (after) the updating decision in period $t$. The i.i.d. assumption on fixed costs of updating implies that $V_{\tau,t}$ does not depend on the realization of the fixed cost. After plugging (5) into (4) and integrating over $\zeta$ the objective (4)-(5) reduces to

$$
\begin{aligned}
V_{\tau,t} \quad &= \max_{P_{\tau,t},y_{\tau,t},\alpha_{\tau,t}} \quad E_{\tau,t}\left[\frac{P_{\tau,t}y_{\tau,t}}{P_t} - \psi\frac{y_{\tau,t}^{1+\xi}}{1+\xi}\right. \\
&\left. +\beta\left(\alpha_{\tau+1,t+1}V_{1,t+1} - \int_0^{G^{-1}(\alpha_{\tau+1,t+1})} dG(\zeta) + (1-\alpha_{\tau+1,t+1})V_{\tau+1,t+1}\right)\right], \qquad (6)
\end{aligned}
$$

where $G^{-1}(\cdot)$ denotes the inverse function of the c.d.f. $G(\cdot)$, $\alpha_{\tau+1,t+1} = G^{-1}(E_{\tau+1,t+1}(V_{1,t+1}) - V_{\tau+1,t+1})$ is firm $\tau$'s probability of updating in period $t+1$, and the integral in the last bracket represents fixed costs that the firm is expected to pay in period $t+1$.

Updating probabilities $\alpha_{\tau,t}$ imply that the distribution of households across types is given by

$$\nu_{\tau,t} = (1-\alpha_{\tau,t})\nu_{\tau-1,t-1}, \qquad \tau = 2,3,... \qquad (7)$$

$$\nu_{1,t} = 1 - \sum_{\tau=2}^{\infty}(1-\alpha_{\tau,t})\nu_{\tau-1,t-1} \ , \qquad (8)$$

where $\nu_{\tau,t}$ denotes the fraction of type $\tau$ households at the end of period $t$.

First-order conditions for problem (6) yield

$$P_{\tau,t}^{1+\xi\theta} = \frac{\theta}{\theta-1}\psi\left\{Y_t\phi_t^\theta P_t^\theta\right\}^\xi \frac{1}{E_{\tau,t}\left[P_t^{-1}\right]} \ , \qquad (9)$$

$$V_{\tau,t} = E_{\tau,t}V_{1,t} - G(\alpha_{\tau,t}) \ . \qquad (10)$$

According to the pricing equation (9), intermediate good prices in period $t$ depend on island demand $Y_t\phi_t^\theta P_t^\theta$ and the firm's expectations on the island price level $P_t$. Equation (10) is a cutoff rule saying that for a type $\tau$ firm there exists a value of fixed costs, $G(\alpha_{\tau,t})$, that makes it indifferent between updating and not updating to full information. It is shown in Appendix A that equations (10) imply a closed-form system for updating probabilities $\alpha_{\tau,t}$, given profits $\Pi_{\tau,t}$ :

$$\begin{aligned} G\left(\alpha_{\tau,t}\right) &= E_{\tau,t}\left\{\Pi_{1,t}-\Pi_{\tau,t}\right. \\ &\quad +\beta\left[\left(1-\alpha_{\tau+1,t+1}\right)G\left(\alpha_{\tau+1,t+1}\right)-\left(1-\alpha_{1,t+1}\right)G\left(\alpha_{1,t+1}\right)\right] \\ &\quad \left. +\beta\left[\int_0^{G^{-1}\left(\alpha_{\tau+1,t+1}\right)} dG\left(\zeta\right)-\int_0^{G^{-1}\left(\alpha_{1,t+1}\right)} dG\left(\zeta\right)\right]\right\}, \quad \tau=1,2,... \end{aligned} \tag{11}$$

## 2.3 Government

It is assumed that the government controls aggregate nominal demand, which is proportionately distributed among islands, so that an island's nominal demand, $M_t$, is given by:

$$M_t = P_t Y_t \ ,$$

where $P_t$, $Y_t$ are final good price and output on the island. Let $\mu_t$ denote the rate of growth of the nominal demand, $\mu_t = \frac{M_t}{M_{t-1}}$, and assume that $\ln \mu_t$ follows an AR(1) process

$$\ln \mu_{t+1} = (1-\rho_\mu)\ln \mu + \rho_\mu \ln \mu_t + \varepsilon_{\mu t+1}$$

with mean $\ln \mu$ and i.i.d. errors $\varepsilon_{\mu t} \sim N(0, \sigma_\mu^2)$.

By assumption, islands' nominal demand growth rates are all equal to $\mu_t$ :

$$\mu_t = \pi_t \frac{Y_t}{Y_{t-1}} \ . \tag{12}$$

## 2.4 Equilibrium

An equilibrium consists of sequences of firms' probability distributions over their respective information sets, prices $\{P_t, P_{\tau,t}\}_{t=0}^{\infty}$, allocations $\{Y_t, y_{\tau,t}\}_{t=0}^{\infty}$, value functions $\{V_{\tau,t}\}_{t=0}^{\infty}$, and updating weights and probabilities $\{\nu_{\tau,t}, \alpha_{\tau,t}\}_{t=0}^{\infty}$ such that, given the initial probability distributions for all cohorts in period 0, for all state histories,

(1) given island prices $\{P_t\}_{t=0}^{\infty}$, allocations and individual prices $\{P_{\tau,t}\}_{t=0}^{\infty}$ solve the optimization problem of the final and intermediate good producers;

(2) value functions $\{V_{\tau,t}\}_{t=0}^{\infty}$ satisfy equations (6);

(3) updating weights and probabilities $\{\nu_{\tau,t}, \alpha_{\tau,t}\}_{t=0}^{\infty}$ satisfy cutoff equations (10) and laws of motion (7)-(8);

(4) firms' probability distributions are updated according to the Bayes law;

(5) final good producers collect zero profits;

(6) all markets clear, island indexes $P_t$ and $Y_t$ are defined by (1), (3).

It is shown by Hellwig and Veldkamp (2009) that price-setting problems with discrete information choice based on public signals face multiple equilibria. Although analyzing the uniqueness of equilibrium does not appear to be possible in this model, I conjecture that equilibrium is indeed unique. The uniqueness follows from the specification that the stochastic fixed cost of updating is drawn from a continuous distribution, which smoothes the otherwise discrete choice to update. In the next section it is shown that this leads to staggered information decisions and smooth dynamics for the distribution of firms as a function of signal precision.[14]

# 3 Aggregation and Solution

This section provides a solution of the log-linearized version of the model introduced in Section 2. Log-linearization is used for two reasons. First, in the linear model, Bayesian updating implies that the laws of motion for conditional expectations of the state are given by the Kalman filter equations. Second, linearization allows for aggregation of the individual decision rules. Both of these features make the model tractable and suitable for numerical analysis.

## 3.1 Aggregation

Prices and quantities in the model are log-linearized around deterministic trend $\mu^t$ such that all prices grow with the rate of nominal demand growth, $\mu$, and all real

[14]In the numeric simulations of the model, I experimented with various initial conditions and degrees of computation accuracy, and found no other equilibria.

variables are constant :

$$
\begin{aligned}
P_\tau &= P = \mu^t \ , \\
y_\tau &= Y = \left(\frac{\theta - 1}{\theta\psi}\right)^{1/\xi} \ , \\
\Pi_\tau &= \left(\frac{\theta}{\theta - 1}\psi\right)^{-\frac{1}{\xi}} \frac{1 + \xi\theta}{(1 + \xi)\,\theta} \ .
\end{aligned}
$$

Pricing equation (9) in log-linearized form is

$$
\hat{P}_{\tau,t} = \frac{\xi}{1 + \xi\theta}\left(\hat{Y}_t + \theta\hat{\phi}_t\right) + \hat{P}_t + \frac{1}{1 + \xi\theta}\left[E_{\tau,t}\left[\hat{P}_t\right] - \hat{P}_t\right] \ , \tag{13}
$$

where hats denote log-linear deviations from trend. The first two terms in equation (13) are standard for price adjustment under perfect information: the price of a differentiated good relative to the island price level depends on the island output demand, as summarized by the first term. The constant $\frac{\xi}{1+\xi\theta}$ is smaller then 1, reflecting the dampening of price changes in response to changes in marginal costs due to strategic price complimentarity with the pricing decisions of other producers on the island. The last term characterizes the effect of imperfect information on the firm's price: to the extent that forecasts of the island's price level lag the actual level, firms dampen their price changes. Hence imperfect information is the only source of propagation of nominal shocks in this model, and propagation is inversely related to firms' forecasting accuracy.[15]

Substituting $\hat{P}_{\tau,t}$ into log-linearized equation for island price (3) yields equation for island price level :

$$
\hat{P}_t = -\hat{\phi}_t + \xi\hat{Y}_t + \sum_\tau \nu_\tau E_{\tau,t}\left[\hat{P}_t\right] + \sum_\tau P_\tau d\nu_{\tau,t} \ , \tag{14}
$$

where $d\nu_{\tau,t}$ denotes first-difference deviations of $\nu_{\tau,t}$ from the steady state. The last term represents the effect of changes in the distribution of firms on island price. Since $P_\tau$ is equal for all $\tau$, $\sum_\tau P_\tau d\nu_{\tau,t} = 0$, so in the linear model we only need to solve for a stationary distribution of firms, $\{\nu_{\tau,t}\}$.

The following proposition provides a closed-form system for the sequence of up-

[15] Section 4 discusses the extension with sticky prices as another potential source of propagation.

dating probabilities $\{\alpha_\tau\}$ that can be used to find the stationary distribution of firms $\{\nu_\tau\}$.

**Proposition 1.** *Let $MSE_\tau\left(\hat{P}_t\right)$ be the mean squared error of island price level for type $\tau$ producer in period $t$, $MSE_\tau\left(\hat{P}_t\right) = E\left[\hat{P}_t - E_{\tau,t}\hat{P}_t\right]^2$. Then up to a second order of magnitude,*

$$EE_{\tau,t}\left[\Pi_{1,t} - \Pi_{\tau,t}\right] = \Pi\frac{(1+\xi)\,\theta\,(\theta-1)}{(1+\xi\theta)^2}\left[MSE_\tau\left(\hat{P}_t\right) - MSE_1\left(\hat{P}_t\right)\right], \tag{15}$$

*and mean updating probabilities up to a first order of magnitude are given by*

$$\begin{aligned} G\left(\alpha_\tau\right) &= EE_{\tau,t}\left[\Pi_{1,t} - \Pi_{\tau,t}\right] \\ &+\beta\left[\left(1-\alpha_{\tau+1}\right)G\left(\alpha_{\tau+1}\right) - \left(1-\alpha_1\right)G\left(\alpha_1\right)\right] \\ &+\beta\left[\int_0^{G^{-1}(\alpha_{\tau+1})} dG\left(\zeta\right) - \int_0^{G^{-1}(\alpha_1)} dG\left(\zeta\right)\right]\Bigg\}, \qquad \tau = 1,2,... \end{aligned} \tag{16}$$

Proof. See Appendix B. ■

Proposition 1 is key to understanding the interaction between learning and updating highlighted in this paper. Equation (15) relates forecasting accuracy to firms' payoffs: the average expected gain in current profits after updating is proportional to the improvement in the mean squared error of island price level, up to a second order. Equation (16), in turn, links expected profit gains to the probability of updating: smaller predicted profit gains imply smaller updating probabilities.

There are two caveats stemming from Proposition 1. First, better forecasting accuracy decreases the incentive to update information sets. For example, more informative market signals decrease the number of firms willing to pay fixed costs to update their information. This is simply an implication of substitutatibility between two technologies for acquiring information. Second, endogenous updating determines the distribution of firms $\{\nu_\tau\}$ over information sets $I_{\tau,t}$, with less frequent updating implying greater weights on more outdated information sets. Such information sets mostly contain public information, so corresponding prices are more synchronized. Since market signals are determined by the distribution of prices, they convey more information when prices are less dispersed. Hence, the interaction between learning from market signals and the endogenous frequency of updating implies an *informa-*

*tion externality*: information conveyed by market signals determines the frequency of updating information by price-setting firms, which in turn affects the information contained in the market signals. Next section studies the implications of this information externatity for monetary nonneutrality.

The linearized equilibrium system is closed by the aggregate demand equation (12) in log-linearized form:

$$\hat{P}_t - \hat{P}_{t-1} = \widehat{\mu}_t - \hat{Y}_t + \hat{Y}_{t-1} \; . \qquad (17)$$

The resulting system has four linear equations (14), (15), (16) and (17) cast in terms of current period price and output levels, their average expectations and fundamental disturbances.

## 3.2 Solution

In the problem above firms draw inferences about the entire history of past shock realizations. Townsend (1983) showed that this type of problem does not have a finite-dimensional recursive representation and in general cannot be solved analytically. To solve this system numerically, I use a method of undetermined coefficients for a truncated state space.[16] The state history is approximated by a truncated vector

$$Z_t = \left[ z_t^{'}, z_{t-1}^{'}, ..., z_{t-T}^{'} \right]' \; .$$

This expanded state vector evolves according to

$$Z_t = AZ_{t-1} + B\varepsilon_t$$

where $\varepsilon_t = [\varepsilon_{\mu t}, \varepsilon_{\phi t}]'$, $A = \begin{bmatrix} \rho & 0_{2,2(T-1)} \\ I_{2(T-1)} & 0_{2(T-1),2} \end{bmatrix}$, $B = \begin{bmatrix} I_2 \\ 0_{2(T-1),2} \end{bmatrix}$, $\rho = \begin{bmatrix} \rho_\mu & 0 \\ 0 & \rho_\phi \end{bmatrix}$, and $I_N$ is an $N \times N$ matrix with ones on the diagonal and zeros otherwise. Each firm on an island, before setting its price, observes the signal $s_t = \hat{Y}_t + \theta\hat{\phi}_t + \theta\hat{P}_t$. Assume

[16] See Hellwig (2002), Lorenzoni (2009).

that the signal $s_t$ is a linear function of the state history:

$$s_t = HZ_t \ , \tag{18}$$

where $H$ is a $1 \times 2T$ matrix of unknown coefficients. After observing the signal and setting prices, firms update their expectations according to Kalman filter equations[17]:

$$E_{\tau+1,t+1}Z_{t+1} = AE_{\tau,t}Z_t + AK_\tau H[Z_t - E_{\tau,t}Z_t], \quad \tau = 1, ..., T-1 \ , \tag{19}$$

$$E_{1,t}Z_{t+1} = AZ_t \ , \tag{20}$$

where $K_\tau$ is the matrix of Kalman gains, and $E_{\tau,t}Z_t$ denotes expected mean of the state vector $Z_t$ conditional of prior distribution for the firm with information set $I_{\tau,t} = \{z^{t-\tau}, s^t\}$. Write this conditional expectation as a linear function of the state vector:

$$E_{\tau,t}Z_t = \Psi_\tau Z_{t-1} \ , \tag{21}$$

where $\Psi_\tau$ is a $2T \times 2T$ matrix of unknowns. Equations (14)-(20) yield a system of equations for unknown matrices $H$, $\{\Psi_\tau\}$ and average adjustment probabilities $\{\alpha_\tau\}$ that together characterize a linear equilibrium. Appendix C provides details of the numeric solution method.

# 4 Simulation of equilibrium dynamics

This section presents the results of model solution and simulation. It first describes the main mechanism behind the dissemination of information in the economy, namely the interaction between updating and learning. Then it discusses the implications of this interaction for equilibrium dynamics and for the monetary transmission mechanism.

[17]Hamilton (1994), Chapter 13.

## 4.1 Two channels of information acquisition: updating and learning

Information about the fundamental disturbances in the economy disseminates via two distinct channels: every period a fraction of agents updates to full information, and the remaining firms learn from observing island demand. To illustrate these channels, Figure 1 shows impulse responses of firms' expectations of the (detrended) quantity of money at the time of their price decisions. The top panel gives the responses in the model with learning. The most informed cohort 1 fully updates at the beginning of the second period, so its expectations jump to the actual level of money stock. The behavior of the higher cohorts is similar: expectations imperfectly track the increase in the quantity of money as they learn from the market signals, until they eventually update their information sets, at which point their expectations also jump to the actual level of money. In the model without learning (bottom panel), cohorts that do not update information, entirely miss the shock, and instead forecast that the quantity of money will stay at the pre-shock level. Hence, observing market signals allows firms to imperfectly infer the underlying shocks at times when paying for full information is too costly.

The interaction between updating and learning affects the transmission mechanism in this paper, and in particular, has implications for the extent of real persistence after monetary shocks. To build intuition behind these implications, consider a stylized example involving the output response following a +1% i.i.d. innovation in money growth.[18] Equations (14) and (17) imply that the output response is given by

$$\begin{aligned} Y_0 &= \frac{1}{\xi + 1} \\ Y_j &= \frac{\left(1 - \sum_{l=1}^{j} \nu_l\right) FE_j}{\xi + \left(1 - \sum_{l=1}^{j} \nu_l\right) FE_j}, \quad j = 1, 2, ... \end{aligned}$$

where subscript $j$ denotes the number of periods (say, quarters) after the impulse, and $FE_j = 1 - \sum_{l=j+1}^{\infty} \frac{\nu_l}{1-\sum_{l=1}^{j} \nu_l} \frac{E_{l,t+j-1} P_{t+j}}{P_{t+j}}$ is the average forecast error of the island price level in $j^{th}$ quarter after the shock among firms that have not updated their

[18] Recall that conditional on a given money shock, the island-specific and aggregate reponses are identical.

information sets. The error is expressed as a fraction of the island price level in quarter $j$. Assume, for simplicity, both a constant rate of updating $\alpha_j = \alpha$, and a constant forecast error $FE_j = x$. Define the half-life of the output response as the time $J$ that it takes for output impulse response to decrease by half relative to its response in quarter 0. Then the fraction of updating can be written as a decreasing function $\alpha(J)$ of the half-life of the output response:

$$\alpha(J) = 1 - \left( \frac{\xi x^{-1}}{2\xi + 1} \right)^{1/J} . \tag{22}$$

Figure 2 compares $\alpha(J)$ for two cases: $x = 1$ (no learning), and $x = 1/2$ (firms can predict half of the fluctuations in price level). Consider the economy with no learning and with updating once a year (point A). The corresponding half-life of the output response is around 7 quarters. With learning, ceteris paribus, forecast errors decrease by $1 - x$, i.e. by 50% (point B). From (15) and (16), we know that smaller forecast errors weaken the incentive to updating so the frequency of updating decreases and, ceteris paribus, the persistence of the output response increases somewhat (point C). A reduced-form (22), however, ignores the effect of the decrease in the frequency of updating on the dynamics of output and the price level, and thus also on the informativeness of the market signal. In Section 4 it is shown that less frequent updating, through greater synchronization of price changes, tends to improve information conveyed by market signals, leading to more accurate price decisions and reduced persistence of the output response (point D).

Before beginning a quantitative investigation into the implications of learning and updating for aggregate persistence, the model is calibrated to some salient facts concerning U.S. inflation.

## 4.2 Calibration

The benchmark model is calibrated to match key characteristics of quarterly U.S. time series. Calibrated parameters are given in Table 1. The discount factor $\beta$ is $0.97^{1/4}$ implying that the interest rate is 3%. The scale parameter $\psi$ is set to $\frac{\theta}{\theta-1}$, normalizing steady state output to 1. The degree of decreasing returns to labor, $\xi$, is 0.18 implying returns to scale of 0.85. Since labor is the only productive input in the model, this value represents a compromise between more conventional constant returns and a labor share of 0.7. The elasticity of substitution between intermediate goods, $\theta$, is 5, which is consistent with values in the IO literature - between 2 and 5 - and in the macro literature - between 5 and 10.[19] To calibrate aggregate and island demand shocks, I use the BLS data for consumer prices in 111 commodity categories in the U.S. over the period 1978:1-1997:4. The sample accounts for 49% of 1997 consumption expenditures.[20] The standard deviation and serial correlation of the island demand shock process are chosen to match the weighted average standard deviation and serial correlation of (detrended) inflation across the goods categories, respectively 1.97% and -0.05.

Parameters of the aggregate demand (money growth) process, are picked to match standard deviation and serial correlation of CPI inflation excluding food and energy for the period 1957:1-2008:4, respectively 0.68% and 0.81.[21] This method for calibrating the aggregate demand is preferred to estimating it directly from the data on money supply for two reasons. First, in the model, the money growth shock is meant to capture a variety of nominal disturbances impacting the aggregate rate of inflation (e.g., shocks to the velocity of money). Second, due to differences in information filtering between the models with and without learning, assuming a common money shock process for both cases would result in different moments of aggregate inflation. Hence, for each model, the shock parameters are calibrated to match the same target moments for inflation.

Table 1 shows that calibrated money growth shocks in the model without learning are much less persistent than in the benchmark model. In the model without learning, low covariance of price levels across firms reinforces slow adjustment of prices due to

[19] The degree of strategic complementarity in pricing, given by $\frac{\xi}{1+\xi\theta}$, equals to 0.095, which is close to the range of 0.10 to 0.15 given by Woodford (2003). Changing $\xi$ and $\theta$ to better accord with Woodford's estimates does not affect main results.

[20] By restricting the sample is to a shorter time period, 1990:1-1997:4, the number of sectors can be expanded to 159, covering 93% of expenditures. Since the focus in the paper is on time series moments, I prefer a narrower but longer sample.

[21] Inflation moments from CPI less food and energy are between moments obtained using the GDP deflator and total CPI. For GDP deflator standard deviation and serial correlation are 0.59% and 0.86, and they are respectively 0.78% and 0.67 for total CPI. Hence, out of three most commonly used aggregate price indeces, CPI less food and energy is the most representative of U.S. inflation in the postwar data.

outdated information. This implies strong persistence of inflation. In fact, models with infrequent information updating tend to predict *too much* inflation peristence.[22] To offset this feature, calibrated money growth shocks in the model without learning have almost zero serial correlation. In the benchmark model, access to public signals implies that inflation is more volatile but less persistent. Hence, in order to match the same calibration targets, the serial correlation of money shocks in the benchmark model is greater but the innovations to money shocks are on average smaller, relative to the model without learning.

Finally, we assume that the distribution of the fixed cost of updating is uniform. This impiles that we only have to calibrate one parameter: the maximal fixed cost, $\zeta_{\max}$. This parameter is chosen so that, without learning, information is updated once a year.[23] This frequency of updating is consistent with empirical studies by Carroll (2003), Mankiw, Reis, and Wolfers (2003), and Khan and Zhu (2006). The implied average fixed cost of updating is 1.5% of the firm's quarterly revenue, in line with micro evidence in Zbaracki, Ritson, Levy, Dutta, and Bergen (2004). They document that the managerial cost of price adjustment, including the costs of information gathering and analysis, totals about 1% of quarterly revenue.

## 4.3 Aggregate implications of imperfect information

This subsection explores the implications of the interaction of costly updating and costless learning for the degree of persistence of output and inflation. I consider the dynamics of aggregate output and inflation after an unexpected increase in nominal expenditures, and focus particularly on the degree of persistence of impulse responses. Chari, Kehoe, and McGrattan (2000) argue that the empirically plausible half-life of the output impulse to a monetary impulse is around 10 quarters. This number is consistent with Christiano, Eichenbaum, and Evans (2000), who conduct VAR analysis of the U.S. postwar economy.

Figure 3 illustrates the responses of aggregate output and the price levels following a +1 standard deviation innovation in money growth in both the benchmark model

[22] See, e.g., Mankiw and Reis (2002), Reis (2006).

[23] For the model with learning, results are not very sensitive to calibration of the fixed cost distribution.

and the model without learning.[24] The top panel shows the responses of aggregate output. In both model economies, output response is hump-shaped, peaking after 1 quarter. The magnitude of the response in the model without learning is about twice as large as in the benchmark model, due to a larger size of the shock. The biggest difference between the two responses is their persistence. In the model without learning, the time after shock that it takes for output response to fall to half its maximuml, i.e., its "half-life", is 6.7 quarters. Since calibrated money growth shocks in the model without learning are not persistent, this half-life should be interepreted as a lower bound on persistence in models where firms acquire information only by updating. Nevertheless, even this lower bound is more than twice as large as the half-life in the benchmark model, which is comes in around 2.6 quarters.[25]

The bottom panel of Figure 3 shows the reponses of price levels in two models. A money growth impulse implies a permanent increase in the aggregate price level in each model. The increase is larger in the model without learning. In line with output responses, it takes only half a year after the shock for the price level to halve the distance to its new level in the benchmark model, whereas it takes 3 times longer in the model without learning. Thus the model with learning implies much faster adjustment than the model without learning.

It is important to note that the model with learning, despite featuring less aggregate persistence, is validated empirically. Table 2 shows that it matches several auxiliary moments from the BLS data as well as the model without learning. Specifically, the correlation of money growth with aggregate and sector-specific inflation rates is 0.39 and 0.14 in the model, fairly close to 0.30 and 0.28 in the data. Both models match closely the across-time volatility of the interquartile range: 0.26% and 0.28% in the benchmark and in the no-learning model respectively (0.29% in the data). Both models exaggerate somewhat the dispersion of inflation rates across sectors: the weighted interquartile range of inflation rates for an average quarter in the benchmark model is 2.52 % points, as compared against 1.76 % points in the BLS

[24] Thus the shock is representative *for each model*. In the benchmark model, the shock is smaller but more persistent.

[25] In the model without learning, the half-life of 10 quarters can be obtained by increasing the average fixed cost of updating to 4.5% of revenue. The average time between updating then increases from 1 to 1.5 years. In the model with learning, however, half-life is virtually unaffected by increases in the fixed cost.

data.

The model is also consistent with some regularities concerning inflation expectations that have been highlighted in the empirical literature. Mankiw, Reis, and Wolfers (2003) find that the interquartile range of expectations of annual aggregate inflation for 2003 is between 1.5% and 2.5% for economists, and 0% to 5% for households. The benchmark model places the disagreement in inflation expectations at 1.11%, whereas in the model without learning the disagreement is smaller, 0.52%. Forecast errors for annual inflation in expectations surveys ranged between 1.07% to 1.29% for the period from September 1982 to March 2002. Both models fall short of this range, with the benchmark model performing slightly better, placing forecast errors at 0.35%.[26]

Finally, in the model with learning, forecast errors conditional on the monetary shock converge at similar rates across firms due to the synchronizing effect of market signals. In contrast, in the model without learning, forecast errors of updating firms converge faster than those of uninformed firms. Moreover, dispersion of price expectations is virtually unaffected by the shock in the benchmark model, whereas it increases in the model without learning. Both predictions are documented in Coibon and Gorodnichenko (2008), who study responses to structural shocks in empirical expectations data.

## 4.4 The role of updating and learning in aggregate persistence

I start by characterizing the extent of updating and learning in the two model economies, then study implications for aggregate persistence.

Figure 4 provides measures of firms' learning efficiency and updating frequency as a function of the number of quarters since the last updated. The top panel provides forecast errors for the island-specific price level. In the model without learning, forecast errors steadily increase with time since last updating, rising from 1.7% points

[26] In the data, expected inflation rates and forecast errors are reported at an annual horizon. To convert from quarterly to annual frequency in the model, quarterly rates are multiplied by $1 + \rho + \rho^2 + \rho^3$, where $\rho$ is the corresponding serial correlation; namely, 0.81 for the aggregate quarterly rate of inflation, 0.21 (-0.05) for the average island-specific forecast error in the benchmark (no learning) model.

for the most informed cohort to more than 5% for cohorts that have not updated for more than 2 years. The average forecast error is 2.5% points. In the benchmark model, forecast errors are virtually flat: 1.3% points for cohort 1 and 1.4% points for all cohorts that have not updated for more than a year. This is a consequence of learning from a sequence of commonly observed market signals: it takes about a year to infer from market signals the effect of fundamental shocks on aggregate price level. Hence, in the model with learning, agents' beliefs are "synchronized" in the sense that an agent who has not updated for many years on average forecasts inflation just as well as the agent who updated one year ago.

According to the cutoff equation (16), forecast errors for island-specific inflation determine the average probability of updating conditional on the time past since last updating (i.e., hazard rates). In particular, hazard rates are increasing for the model economy without learning, and they are flat for the benchmark model.[27] Even more remarkable is that updating in the benchmark model is extremely infrequent, namely once every 25 years, in contrast to the model without learning, which features annual average frequency of updating. In the benchmark model, less frequent updating synchronizes information sets across firms, leading to smaller forecast errors. As a result, incentives to pay for updating are very weak, reinforcing the low frequency of updating.[28] Hence, very low frequency of updating information sets is due to an externality whereby the number of firms updating information affects the informativeness of market signals.

At the aggregate level, high forecasting accuracy by price-setting firms implies that their price changes are better aligned with fluctuations in nominal demand, reducing the potential for monetary nonneutrality. Figure 5 plots a general equilibrium version of Figure 1 – that is, it gives the half-life of the output response to a +1 standard deviation monetary impulse as a function of equilibrium frequency of updating. Specifically, for each of two models, with and without learning, I conduct a number of simulations, each time changing the cost of updating, to obtain a range

[27] It is insightful to parallel flat hazard rates of information updating in this paper with flat hazard rates of *price* adjustment in micro data, reported in Klenow and Kryvtsov (2008). If information updating is important under infrequent price adjustment, the mechanism in this paper can aid in explaining flat or declining hazard rates in the micro data.

[28] This implication of the model is reminiscent of Grossman and Stiglitz (1980), where costly information updating and homogeneity of beliefs can lead to thin markets for acquiring information.

of average frequencies of updating. In the model without learning, the half-life of output responses quickly increases as the average frequency of updating decreases. In contrast, for model economies with learning, the half-life increases only modestly, from 0.5 quarters for average frequency 99% to 2.6 quarters for the frequency 1%. The conclusion from this quantitative exercise is that, in the model with learning, the greatest degree of monetary nonneutrality is achieved when only a few firms find it optimal to update to full information, so that most firms opt instead to acquire information from market signals.

This conclusion warrants three caveats. First, when updating is very infrequent, information filtering problem is no longer exacerbated by the need to forecast competitors' price changes, as those are highly synchronized with the signal. In this case, persistence depends only on the relative volatility and persistence of the shocks, given the signal that is a linear combination of the shocks.[29] In the data, the inflation rate for a typical sector is more volatile but much less persistent than the aggregate inflation rate. This implies that in the calibrated model it should be easy to disentangle shocks to island-specific inflation from monetary shocks.

Second, the frequency with which the signal is observed determines how quickly the shocks can be identified. In the benchmark model, despite very infrequent sampling of market signals (only once a quarter) the underlying shocks are identified within roughly 3 quarters. With more frequent signal observations, the model would predict even less monetary nonneutrality.

Finally, in the model, market signals, although very informative, do not reveal full information. This implies that the private value of updating is always positive, so firms that draw very small fixed cost choose to update. This setting with a range of fixed costs of updating information thus avoids the "fundamental conflict between the efficiency with which markets spread information and the incentives to acquire information" highlighted by Grossman and Stiglitz (1980). In addition, by ruling out an equilibrium with prices fully revealing information, this setting prevents multiple equilibria of the sort discussed in Hellwig and Veldkamp (2009).

[29] This also means that the exact source of the signal is almost irrelevant for persistence. For example, a signal $s_t = a\hat{Y}_t + b\hat{\phi}_t + c\hat{P}_t$, where $a, b, c$ are real numbers, would result in almost the same degree of persistence as in the benchmark model.

I conclude that the benchmark model provides a reasonable upper bound - less than 3 quarters - on the persistence with which agents respond to monetary shocks in models with imperfect information. The upper bound is given by the time that it takes uninformed firms to infer underlying nominal and real disturbances from market signals. Information externalities of the sort described in this paper weaken private incentives to update and thus enhance the informativeness of market prices and quantities. An important corollary is that business cycles driven by purely nominal disturbances lack the persistence that is typical in the data.

### 4.5 Robustness

Here I establish robustness of my main results with respect to some assumptions about the precision of the market signals, the shape of the fixed cost distribution, and price duration.

First, I relax the assumption that market signals are accurately observed. Instead, I now assume that the market signal $s_t$ is observed with a measurement error $\varepsilon_{st}$, where $\varepsilon_{st}$ is a zero-mean i.i.d. draw from a normal distribution. For my solution method to work, I assume that measurement errors are common for all firms on an island, e.g., inaccurate information about market prices and quantities is published in an island newspaper. This extension then does not increase the heterogeneity of information sets among firms on an island, but rather reduces the informativeness of public signals vis-à-vis fundamental disturbances. I find that in order to increase the half-life to 6.8 quarters, as in the model without learning, the standard deviation of measurement errors has to be 7 times the standard deviation of money growth, or 10 times the standard deviation of the island-specific demand shock. Hence, measurement errors would need to be implausibly large if they are to ensure considerable monetary nonneutrality in the benchmark model.

Second, I experiment with the shape of the continuous distribution from which firms draw their fixed costs of information updating. The uniform distribution assumed in the benchmark model assigns equal weights to each realization of the fixed cost from 0 to $\zeta_{\max}$. This distribution was disturbed in both directions, i.e., the model was simulated given the distribution that puts more (less) weight on extreme fixed cost realizations.[30] It was found that, for both models – i.e., with and without learning – the size of monetary nonneutrality is not sensitive to the shape of the fixed cost distribution, conditional on the average fixed cost remaining. This, to even larger extent, concerns the benchmark model, because the level of the probability of updating is small regardless of the underlying fixed cost distribution.

Finally, I relax the assumption that prices freely adjust every quarter. Price duration as found in the micro data ranges from 4 to 11 months.[31] Now I allow intermediate good producing firms to adjust their prices only every other quarter. In particular, half adjusts in even quarters and the other half in odd quarters. The implied duration of fixed prices is 6 months, in line with empirical evidence. To maintain tractability, I also assume that a firm cannot update its information set unless it also adjusts its price. This assumption does not crucially affect the average incentive to update: although the incentives for a given firm double because the firm is presented with an opportunity to update less frequently, only half of all firms are eligible for updating at any quarter. It is found that main results are not affected. Even though the persistence of the output response following a monetary shocks increases due to sticky prices, the half-life remains less than a year in the model with learning, and lasts a bit more than two years in the model without learning. In this sense, sticky prices do not induce any endogenous persistence in real responses, but rather delay the response in the price level. Given that empirically plausible price stickiness is relatively modest, adding sticky prices to the benchmark model thus does not help in confronting the persistence problem.

# 5 Conclusion

The biggest challenge for monetary models is to generate persistent business cycles. Recent research claims that the persistence problem can be solved using assumptions

[30]For instance, two extreme cases of fixed cost distributions are: bi-modal distribution density, which splits the unit weight between very small and very large costs, implying Calvo-like adjustment, and uni-modal density that puts most of the weight on one value of fixed costs, implying purely state-dependent adjustment.

[31]For example, in U.S. data on consumer prices, this range is 4 to 7 months in Klenow and Kryvtsov (2008), and 8 to 11 months in Nakamura and Steinsson (2008).

that lead to dramatic but exogenous dispersion of information about fundamental shocks. This research predominantly abstracts from the role of prices and quantities in conveying information. In this paper, firms are free to observe and learn from market signals, and heterogeneity of information across price-setters is instead an endogenous outcome of the decision to eschew these signals in favour of updating. It is found that this decision creates an externality by shifting the weight in the price distribution towards those prices that have been informed by market signals. As market signals coordinate firms' price decisions, the costly updating of information becomes less frequent and the accuracy of forecasting competitors' prices improves. Better forecasting accuracy in turn leads to short-lived real effects for monetary shocks. I conclude that imperfect information alone cannot explain the substantial degree of monetary nonneutrality found in the data.

# Appendix A. Value functions and adjustment probabilities

Denote the optimal payoff of the $\tau$ cohort in period $t$ by $\Pi_{\tau,t}$,

$$\Pi_{\tau,t} = \frac{P_{\tau,t} y_{\tau,t}}{P_t} - \psi y_{\tau,t}^{1+\xi} / (1+\xi) \ ,$$

where $P_{\tau,t}$ and $y_{\tau,t}$ are respectively the prices and output chosen by cohort $\tau$ in period $t$.

Then equation (6), after using equation (10) for $E_{\tau,t}V_{1,t}$, becomes

$$V_{\tau,t} = E_{\tau,t} \left\{ \Pi_{\tau,t} + \beta V_{\tau+1,t+1} + \beta \left[ \alpha_{\tau+1,t+1} G\left(\alpha_{\tau+1,t+1}\right) - \int_0^{G^{-1}(\alpha_{\tau+1,t+1})} dG\left(\zeta\right) \right] \right\} \ .$$

Denote $\Gamma\left(\alpha_{\tau+1,t+1}\right) \equiv \int_0^{G^{-1}(\alpha_{\tau+1,t+1})} dG\left(\zeta\right)$. Taking expectations $E_{\tau,t}$ of equation above for $\tau = 1$ and subtracting the equation for $\tau$, we obtain

$$\begin{aligned} E_{\tau,t}V_{1,t} - V_{\tau,t} \quad = \quad & E_{\tau,t} \left\{ \Pi_{1,t} - \Pi_{\tau,t} \right. \\ & + \beta \left[ V_{1,t+1} - V_{\tau+1,t+1} - E_{1,t+1}V_{1,t+1} + E_{\tau+1,t+1}V_{1,t+1} \right] \\ & + \beta \left[ \alpha_{1,t+1} G\left(\alpha_{1,t+1}\right) - \int_0^{G^{-1}(\alpha_{1,t+1})} dG\left(\zeta\right) \right] \\ & \left. - \beta \left[ \alpha_{\tau+1,t+1} G\left(\alpha_{\tau+1,t+1}\right) - \int_0^{G^{-1}(\alpha_{\tau+1,t+1})} dG\left(\zeta\right) \right] \right\} \ , \end{aligned}$$

where by the law of iterated expectations, $E_{\tau,t}E_{\tau+1,t+1}V_{1,t+1} = E_{\tau,t}E_{1,t+1}V_{1,t+1}$. Notice that the law of iterated expectations applies because the information structure in the model implies that information sets of less informed firms are finer subsets of those of more informed firms.

Using (10) once again yields (11):

$$\begin{aligned} G\left(\alpha_{\tau,t}\right) \quad = \quad & E_{\tau,t} \left\{ \Pi_{1,t} - \Pi_{\tau,t} \right. \\ & + \beta \left[ \left(1 - \alpha_{\tau+1,t+1}\right) G\left(\alpha_{\tau+1,t+1}\right) - \left(1 - \alpha_{1,t+1}\right) G\left(\alpha_{1,t+1}\right) \right] \\ & \left. + \beta \left[ \int_0^{G^{-1}(\alpha_{\tau+1,t+1})} dG\left(\zeta\right) - \int_0^{G^{-1}(\alpha_{1,t+1})} dG\left(\zeta\right) \right] \right\} \ . \end{aligned}$$

For the case, when the fixed cost of updating is uniformly distributed, this equation becomes:

$$\alpha_{\tau,t} = \zeta_{\max}^{-1}\left[E_{\tau,t}\Pi_{1,t} - \Pi_{\tau,t}\right] + \frac{\beta}{2}\left[(\alpha_{1,t+1}-1)^2 - (\alpha_{\tau+1,t+1}-1)^2\right], \quad \tau = 1, 2, ...$$

# Appendix B. Proof of Proposition 1

Cohort $\tau$'s profits are

$$\begin{aligned}\Pi_{\tau,t} &= \left(\frac{\theta}{\theta-1}\psi\right)^{\frac{1-\theta}{1+\xi\theta}}\left(Y_t\phi_t^{\theta}P_t^{\theta}\right)^{\frac{1+\xi}{1+\theta\xi}}P_t^{-\frac{\theta(1+\xi)}{1+\xi\theta}}\left[\frac{1}{P_tE_{\tau,t}\left[P_t^{-1}\right]}\right]^{\frac{1-\theta}{1+\xi\theta}} \\ &\quad -\frac{\psi}{1+\xi}\left(\frac{\theta}{\theta-1}\psi\right)^{\frac{-\theta(1+\xi)}{1+\xi\theta}}\left(Y_t\phi_t^{\theta}P_t^{\theta}\right)^{\frac{1+\xi}{1+\xi\theta}}P_t^{\frac{-\theta(1+\xi)}{1+\xi\theta}}\left[\frac{1}{P_tE_{\tau,t}\left[P_t^{-1}\right]}\right]^{\frac{-\theta(1+\xi)}{1+\xi\theta}}.\end{aligned}$$

Up to a second order of magnitude these payoffs are

$$\begin{aligned}\Pi_{\tau,t} &= \Pi\Bigg[1+\frac{1+\xi}{1+\xi\theta}\left(\hat{Y}_t+\theta\hat{\phi}_t+\theta\hat{P}_t\right) \\ &\quad +\frac{1}{2}\left(\frac{1+\xi}{1+\xi\theta}\right)^2\left(\hat{Y}_t^2+\theta^2\hat{\phi}_t^2+\theta^2\hat{P}_t^2+2\theta\hat{Y}_t\hat{\phi}_t+2\theta\hat{Y}_t\hat{P}_t+2\theta^2\hat{P}_t\hat{\phi}_t\right) \\ &\quad +\frac{1}{2}\left(\frac{1+\xi}{1+\xi\theta}\right)^2\left(\hat{Y}_t^2+\theta^2\hat{\phi}_t^2+\theta^2\hat{P}_t^2+2\theta\hat{Y}_t\hat{\phi}_t+2\theta\hat{Y}_t\hat{P}_t+2\theta^2\hat{P}_t\hat{\phi}_t\right) \\ &\quad +\frac{1}{2}\left(\frac{\theta(1+\xi)}{1+\xi\theta}\right)^2\hat{P}_t^2-\frac{1+\xi}{1+\xi\theta}\left(\hat{Y}_t+\theta\hat{\phi}_t+\theta\hat{P}_t\right)\frac{\theta(1+\xi)}{1+\xi\theta}\hat{P}_t \\ &\quad -\frac{1}{2}\frac{(1+\xi)\,\theta\,(\theta-1)}{(1+\xi\theta)^2}\left\{\left[\hat{P}_t-E_{\tau,t}\hat{P}_t\right]^2+E_{\tau,t-1}\left[\hat{P}_t-E_{\tau,t}\hat{P}_t\right]^2\right\}\Bigg],\end{aligned}$$

so that

$$\begin{aligned}\Pi_{1,t}-\Pi_{\tau,t} &= \Pi\frac{1}{2}\frac{(1+\xi)\,\theta\,(\theta-1)}{(1+\xi\theta)^2}\Bigg[\left[\hat{P}_t-E_{\tau,t}\hat{P}_t\right]^2+E_{\tau,t-1}\left[\hat{P}_t-E_{\tau,t}\hat{P}_t\right]^2 \\ &\quad -\left[\hat{P}_t-E_{1,t}\hat{P}_t\right]^2-E_{1,t-1}\left[\hat{P}_t-E_{1,t}\hat{P}_t\right]^2\Bigg].\end{aligned}$$

Noting that

$$E\left\{\left[\hat{P}_t - E_{\tau,t}\hat{P}_t\right]^2 + E_{\tau,t-1}\left[\hat{P}_t - E_{\tau,t}\hat{P}_t\right]^2\right\} = 2MSE_\tau\left(\hat{P}_t\right) ,$$

we obtain:

$$EE_{\tau,t-1}\left[\Pi_{1,t} - \Pi_{\tau,t}\right] = \Pi\frac{(1+\xi)\,\theta\,(\theta-1)}{(1+\xi\theta)^2}\left[MSE_\tau\left(\hat{P}_t\right) - MSE_1\left(\hat{P}_t\right)\right] .$$

Taking unconditional mean of both sides of (11) we obtain

$$\begin{aligned}
EG\left(\alpha_{\tau,t}\right) &= EE_{\tau,t}\left\{\Pi_{1,t} - \Pi_{\tau,t}\right. \\
&\quad +\beta\left[\left(1-\alpha_{\tau+1,t+1}\right)G\left(\alpha_{\tau+1,t+1}\right) - \left(1-\alpha_{1,t+1}\right)G\left(\alpha_{1,t+1}\right)\right] \\
&\quad \left.+\beta\left[\int_0^{G^{-1}(\alpha_{\tau+1,t+1})} dG(\zeta) - \int_0^{G^{-1}(\alpha_{1,t+1})} dG(\zeta)\right]\right\}, \qquad \tau = 1,2,...
\end{aligned}$$

Up to first order of magnitude this becomes

$$\begin{aligned}
G\left(\alpha_\tau\right) &= EE_{\tau,t}\left[\Pi_{1,t} - \Pi_{\tau,t}\right] \\
&\quad +\beta\left[\left(1-\alpha_{\tau+1}\right)G\left(\alpha_{\tau+1}\right) - \left(1-\alpha_1\right)G\left(\alpha_1\right)\right] \\
&\quad +\beta\left[\int_0^{G^{-1}(\alpha_{\tau+1})} dG(\zeta) - \int_0^{G^{-1}(\alpha_1)} dG(\zeta)\right]\Bigg\}, \qquad \tau = 1,2,...
\end{aligned}$$

# Appendix C. Numeric solution method

Kalman gains matrix $K_\tau$ and the mean squared error matrix $\Sigma_{\tau+1|\tau}$ satisfy

$$\begin{aligned}
&K_\tau = \Sigma'_{\tau|\tau-1}H'(H\Sigma'_{\tau|\tau-1}H')^{-1} , \\
&\Sigma_{\tau+1|\tau} = A[\Sigma_{\tau|\tau-1} - \Sigma'_{\tau|\tau-1}H'(H\Sigma'_{\tau|\tau-1}H')^{-1}H\Sigma_{\tau|\tau-1}]A' + Q , \\
&vec(\Sigma_{1|0}) = [I - (A\otimes A)]^{-1}\cdot vec(Q) ,
\end{aligned}$$

and $Q = B\begin{bmatrix}\sigma_\mu^2 & 0\\ 0 & \sigma_\phi^2\end{bmatrix}B'$.

Let

$$E_{\tau,t}Z_t = \Psi_\tau Z_{t-1} .$$

Since the state variable is truncated, we use approximation:

$$Z_{t-1} = UZ_t \ ,$$

where $U = \begin{bmatrix} 0_{2T-2,2} & I_{2T-2} \\ 0_{2,2} & 0_{2,2T-2} \end{bmatrix}$.

Then the system (19)-(20) yields the equations that can be solved iteratively for matrices $\Psi_\tau$ :

$$\begin{aligned} \Psi_{\tau+1} &= A\Psi_\tau U + AK_\tau H \left[ I - \Psi_\tau U \right], \quad \tau \geq 1 \ , \\ \Psi_1 &= A \ . \end{aligned}$$

Let $\hat{Y}_t = H_Y Z_t$. The equilibrium equation (14)

$$-\left(1+\xi\right)\hat{Y}_t = -\hat{\mu}_t - \hat{\mu}_{t-1} - ... - \hat{\phi}_t + \sum_{\tau \geq 1} \nu_\tau E_{\tau,t} \left[ -\hat{Y}_t + \hat{\mu}_t + \hat{\mu}_{t-1} + ... \right]$$

in turn, yields equation for $H_Y$ :

$$H_Y \left\{ \left(1+\xi\right) I - \Psi U \right\} = \left[1,0,1,0,...\right] \left(I - \Psi U\right) + \left[0,1,0,0,...\right] \ ,$$

where $\Psi = \sum_{\tau \geq 1} \nu_\tau \Psi_\tau$.

# References


Bernanke, B. S., and I. Mihov (1998): "Measuring Monetary Policy," *Quarterly Journal of Economics*, 113, 869–902.

Bonomo, M., and C. Carvalho (2004): "Endogenous Time-Dependent Rules and Inflation Inertia," *Journal of Money, Credit and Banking*, 36(6).

Carroll, C. (2003): "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics*, 118, 269–298.

Chari, V., P. Kehoe, and E. McGrattan (2000): "Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?," *Econometrica*, 68(5), 1151–1179.

Christiano, L. J., M. Eichenbaum, and C. Evans (2000): "Monetary Policy Shocks: What Have We Learned and to What End?," *Handbook of Macroeconomics, eds. Taylor, Woodford.*

Coibon, O., and Y. Gorodnichenko (2008): "What can survey forecasts tell us about informational rigidities?," *unpublished paper.*

Gorodnichenko, Y. (2008): "Endogenous information, menu costs and inflation persistence," *unpublished paper.*

Grossman, S. J., and J. E. Stiglitz (1980): "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70(3), 393–408.

Hamilton, J. D. (1994): *Time Series Analysis.* Princeton University Press, Princeton.

Hellwig, C. (2002): "Public Announcements, Adjustment Delays adn the Business Cycle," *UCLA Economics Online Papers 208.*

——— (2008): "Monetary Business Cycle Models: Imperfect Information," in *The New Palgrave Dictionary of Economics and the Law, 2nd edition*, ed. by Durlauf, and Blume. Palgrave-Macmillan, London.

HELLWIG, C., AND L. VELDKAMP (2009): "Knowing What Others Know: Coordination Motives in Information Acquisition," *Review of Economic Studies*, 76, 223–251.

KHAN, H., AND Z. ZHU (2006): "Estimates of the Sticky-Information Phillips Curve for the United States," *Journal of Money, Credit, and Banking*, 38(1), 195–207.

KLENOW, P. J., AND O. KRYVTSOV (2008): "State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?," *Quarterly Journal of Econoimcs*, 123(3), 863–904.

LORENZONI, G. (2009): "A Theory of Demand Shocks," *forthcoming in American Economic Review.*

LUCAS, R. E. J. (1972): "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4, 103–124.

MACKOWIAK, B., AND M. WIEDERHOLDT (2009): "Optimal Sticky Prices under Rational Inattention," *forthcoming in American Economic Review.*

MANKIW, G. N., AND R. REIS (2002): "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics*, 117(4), 1295–1328.

——— (2006): "Pervasive Stickiness," *American Economic Review*, 96(2), 164–169.

MANKIW, G. N., R. REIS, AND J. WOLFERS (2003): "Disagreement about Inflation Expectations," *NBER Macroeconomics Annual 2003.*

NAKAMURA, E., AND J. STEINSSON (2008): "Five Facts About Prices: A Reevaluation of Menu Cost Models," *Forthcoming, Quarterly Journal of Economics.*

REIS, R. (2006): "Inattentive Producers," *Review of Economic Studies*, 73(3), 793–821.

ROMER, C. D., AND D. H. ROMER (1989): "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz," *NBER Macroeconomics Annual*, 4, 121–170.

——— (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *unpublished paper*.

Sims, C. A. (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50(3).

Townsend, R. M. (1983): "Forecasting the Forecasts of Others," *The Journal of Political Economy*, 91(4), 546–588.

Woodford, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton.

——— (2008): "Information-Constrained State-Dependent Pricing," *unpublished paper*.

Zbaracki, M., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004): "Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets," *Review of Economics and Statistics*, 86(2), 514–533.

## Table 1. Model parameters

| Parameter | | Benchmark | No learning |
|---|---|---|---|
| Discount factor | $\beta$ | $0.97^{1/4}$ | |
| Demand elasticity for intermediate goods | $\theta$ | 5 | |
| Returns to scale in intermediate good production | $(1+\xi)^{-1}$ | 0.85 | |
| Average fixed cost of updating, % revenue | $E(\zeta)$ | 1.5 | |
| Serial correlation of money growth | $\rho_\mu$ | 0.365 | 0.064 |
| St. dev. of innovations to money growth, % | $\sigma_\mu$ | 0.98 | 2.12 |
| Serial correlation of island demand | $\rho_\varphi$ | 0.245 | 0.57 |
| St. dev. of innovations to island demand, % | $\sigma_\varphi$ | 1.49 | 1.98 |

Table 2. Moments in calibrated models with and without learning

| Moment | Data | Benchmark | No learning |
|---|---|---|---|
| | | *Used in calibration* | |
| Ser. corr. of aggregate inflation | 0.81 | 0.81 | 0.81 |
| St. dev. of aggregate inflation, % | 0.68 | 0.68 | 0.68 |
| Ser. corr. of island inflation | -0.05 | -0.05 | -0.05 |
| St. dev. of island inflation, % | 1.97 | 1.97 | 1.97 |
| | | *Additional moments:* | |
| Corr. of agg inflation with money growth | 0.30 | 0.39 | 0.46 |
| Corr. of island inflation with money growth | 0.28 | 0.14 | 0.17 |
| Dispersion of island inflation rates, % points | 1.76 | 2.52 | 2.45 |
| St. dev. of dispersion of island inflation, % | 0.29 | 0.26 | 0.28 |
| Dispersion of expectations of aggregate annual inflation rate, % points | 1.5 to 2.5 | 1.11 | 0.52 |
| Forecast errors of aggregate annual inflation, % points | 1.07 to 1.29 | 0.35 | 0.28 |
| Frequency of updating, % | | 1.0 | 17.5 |
| Duration between updating, quarters | | 100 | 4.0 |
| Aggregate output persistence, quarters | | 2.6 | 6.7 |

Note: Sectoral inflation moments are from the BLS data for inflation in 111 commodity categories in the U.S. in 1978:1-1997:4. Aggregate inflation moments are for CPI less food and energy in 1957:1-2008:4. Inflation expectations moments are from Mankiw, Reis and Wolfers (2003). Dispersion in a quarter is measured by the weighted interquartile range. Forecast errors are root mean squared errors averaged across cohorts and time. Aggregate output persistence is measured by half-life of aggregate output response to +1 st. dev. of money growth impulse.

**Figure 1. Responses of actual and expected money stock to +1% i.i.d. money growth impulse**

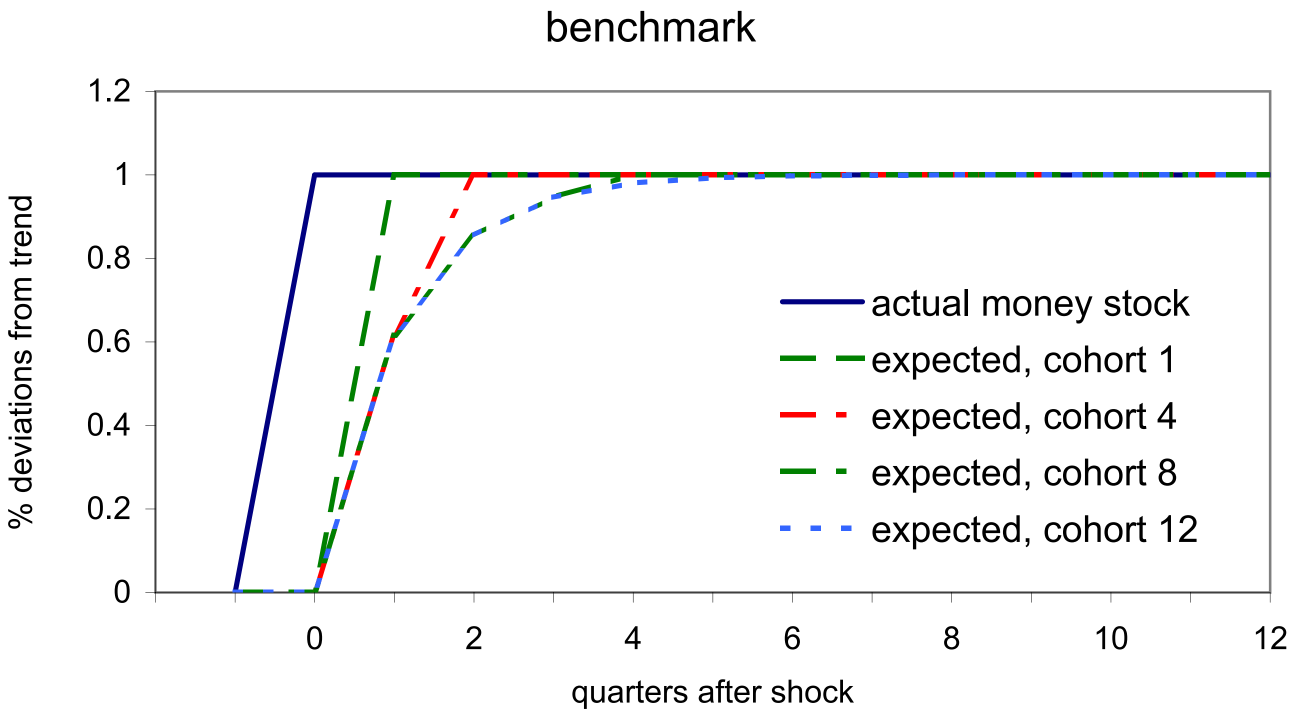


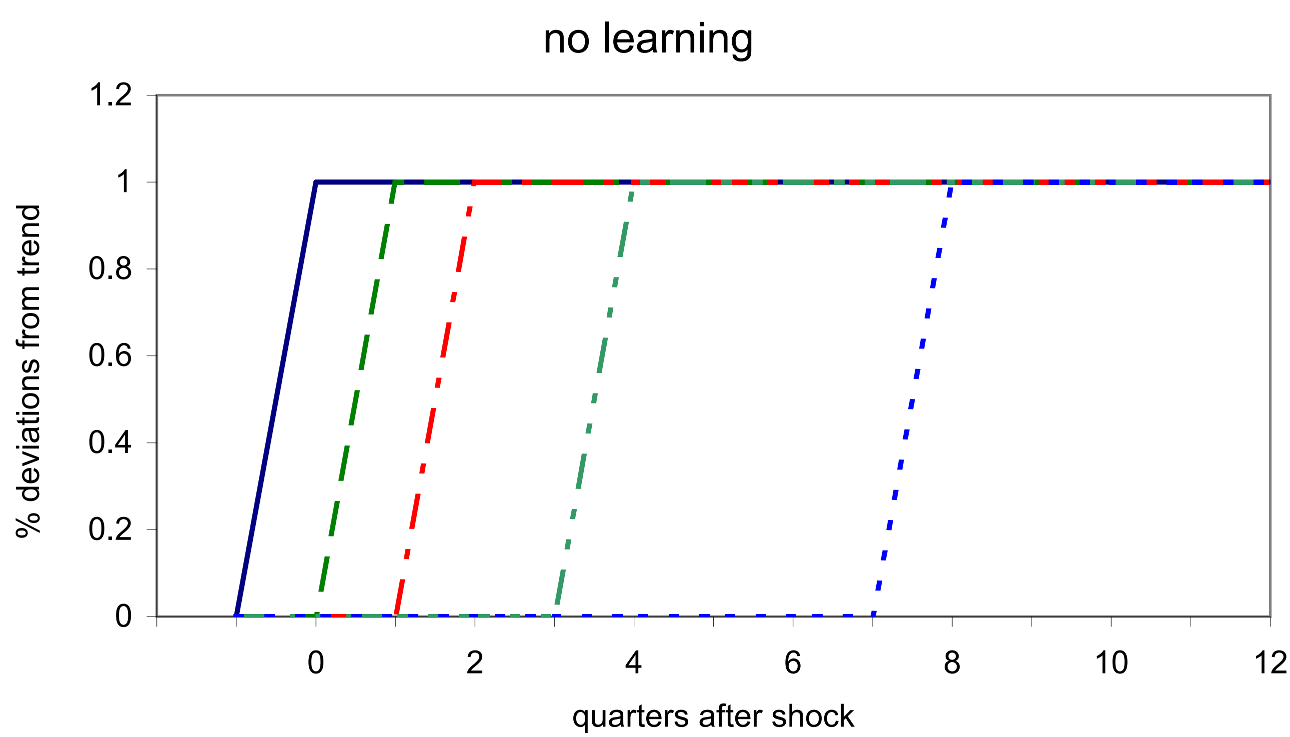

Figure2. Half-life of output responses to +1% i.i.d. money growth impulse, partial equilibrium

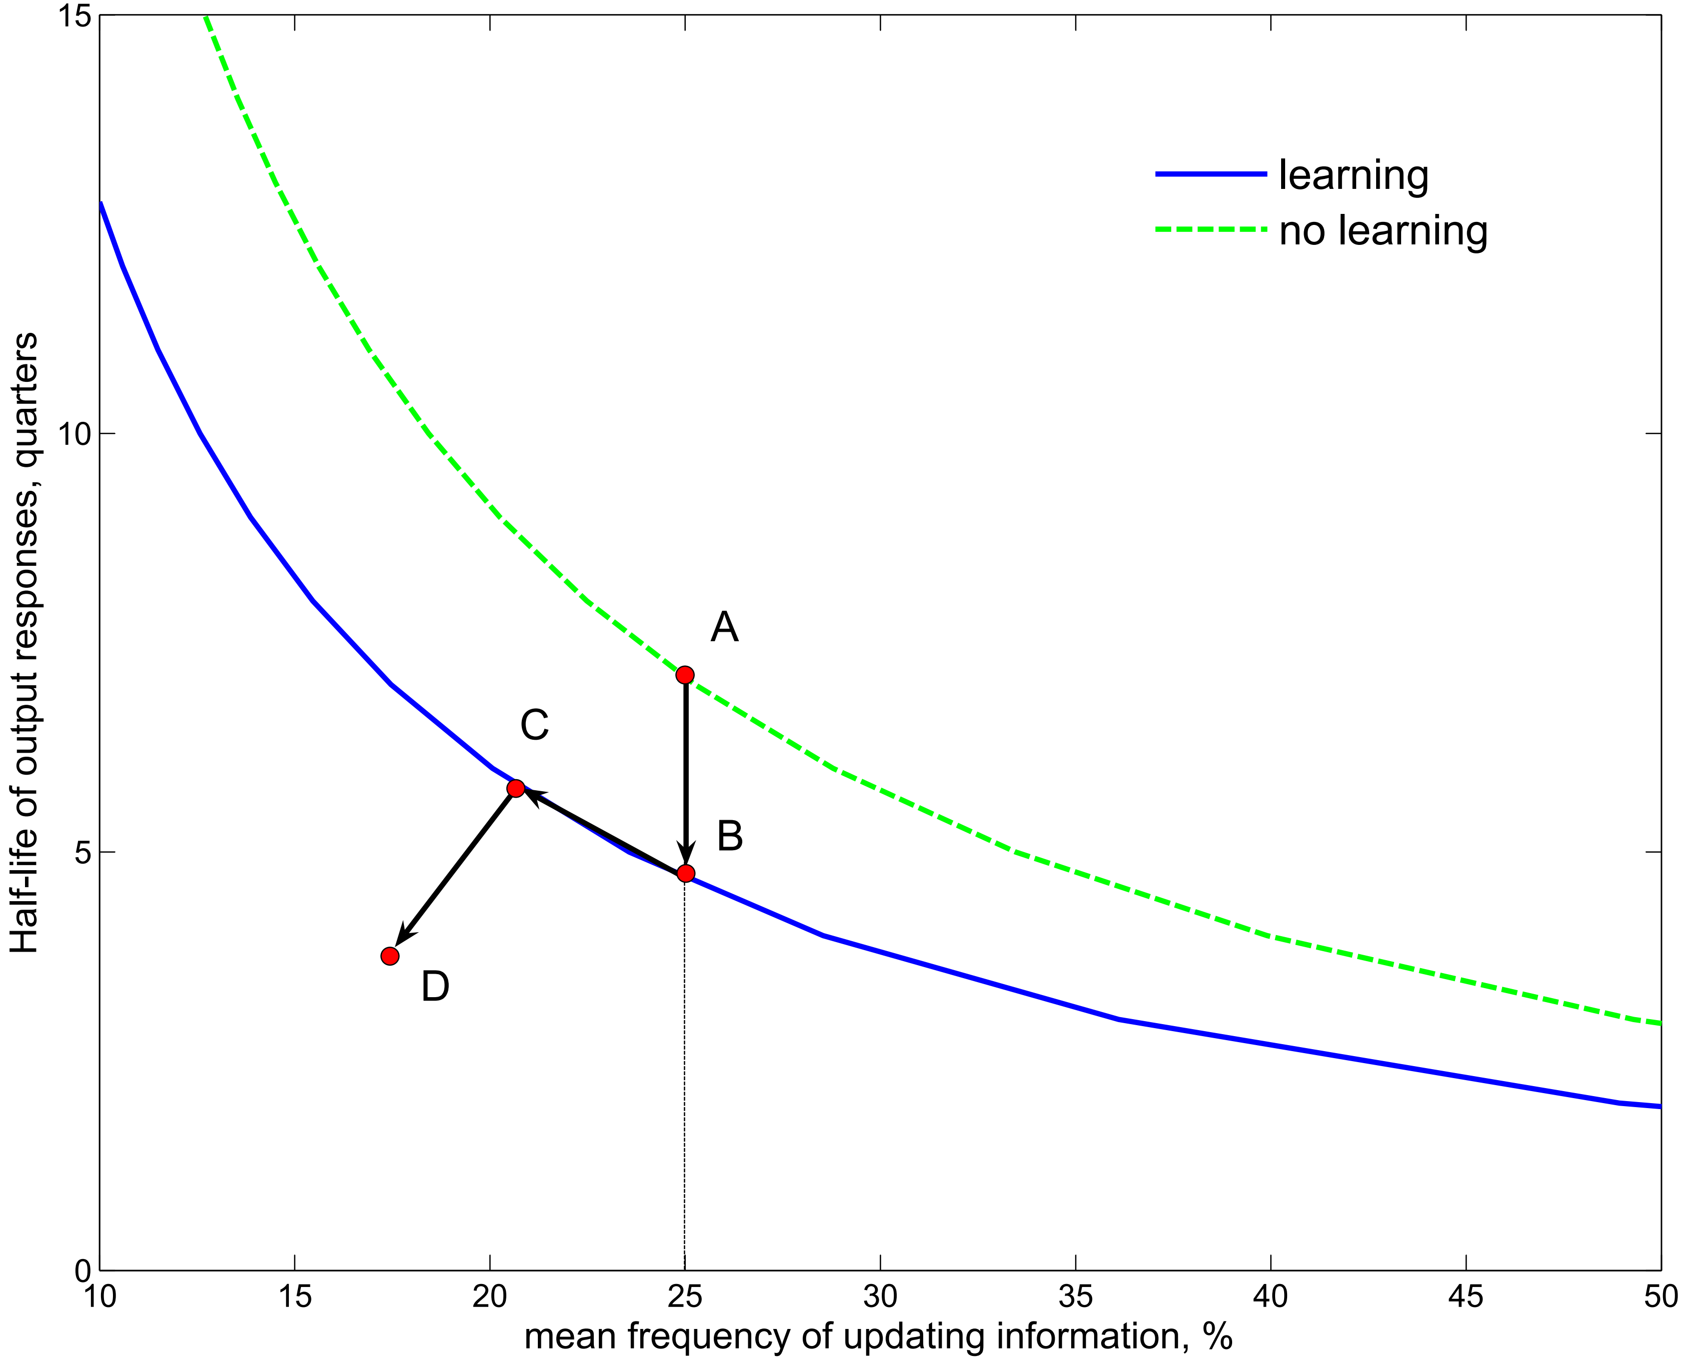

**Figure 3. Responses of aggregate output and price levels to +1 st.dev. of money growth impulse**

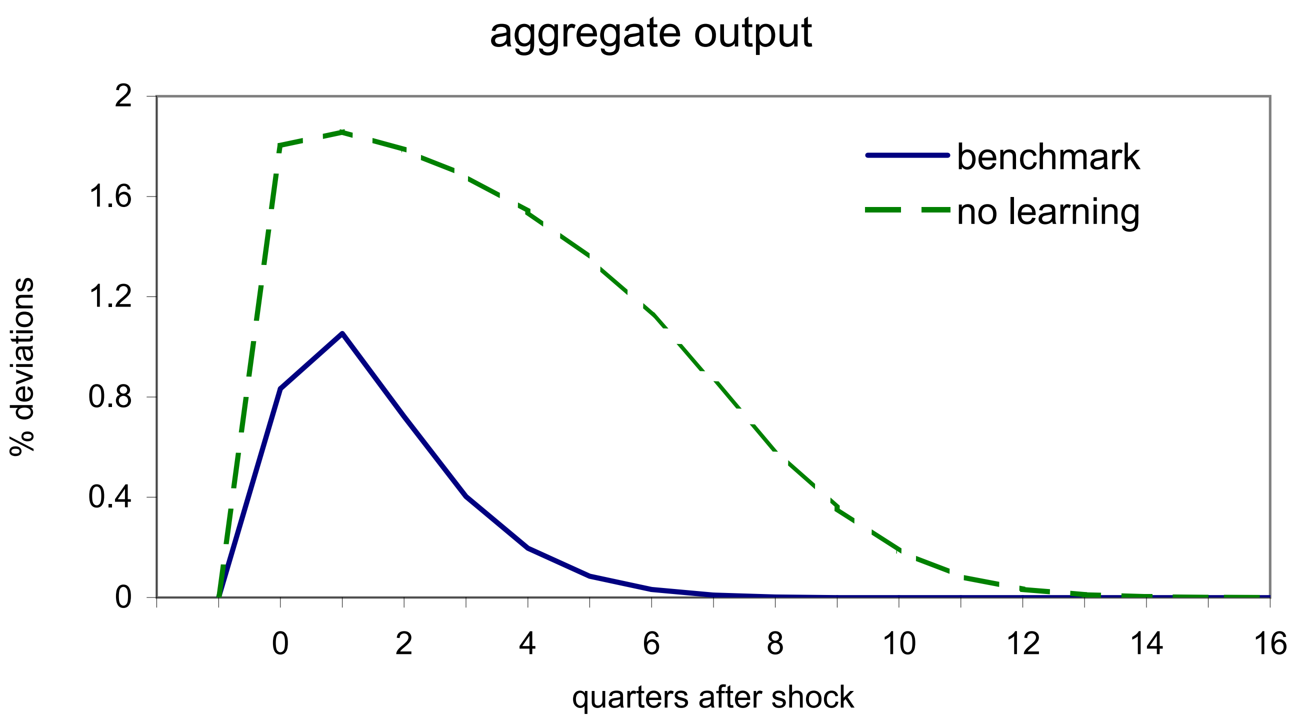


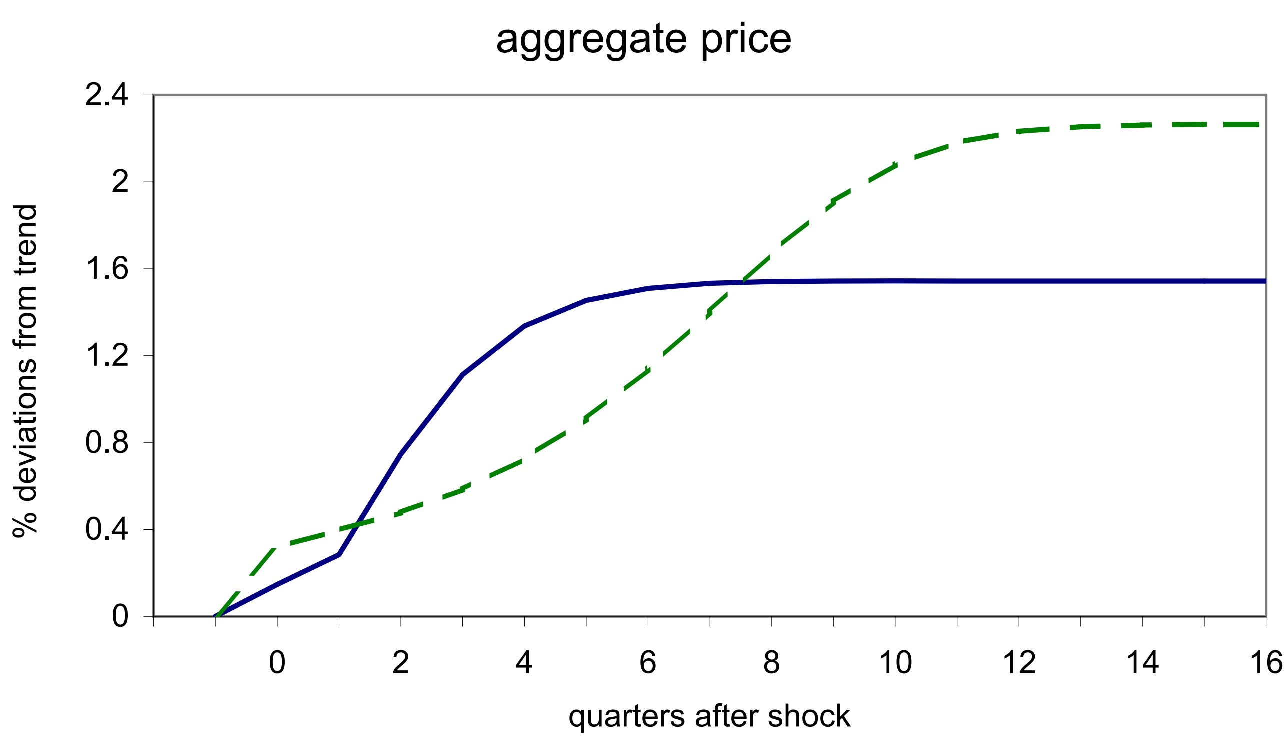

**Figure 4. Forecast errors and updating probabilities**

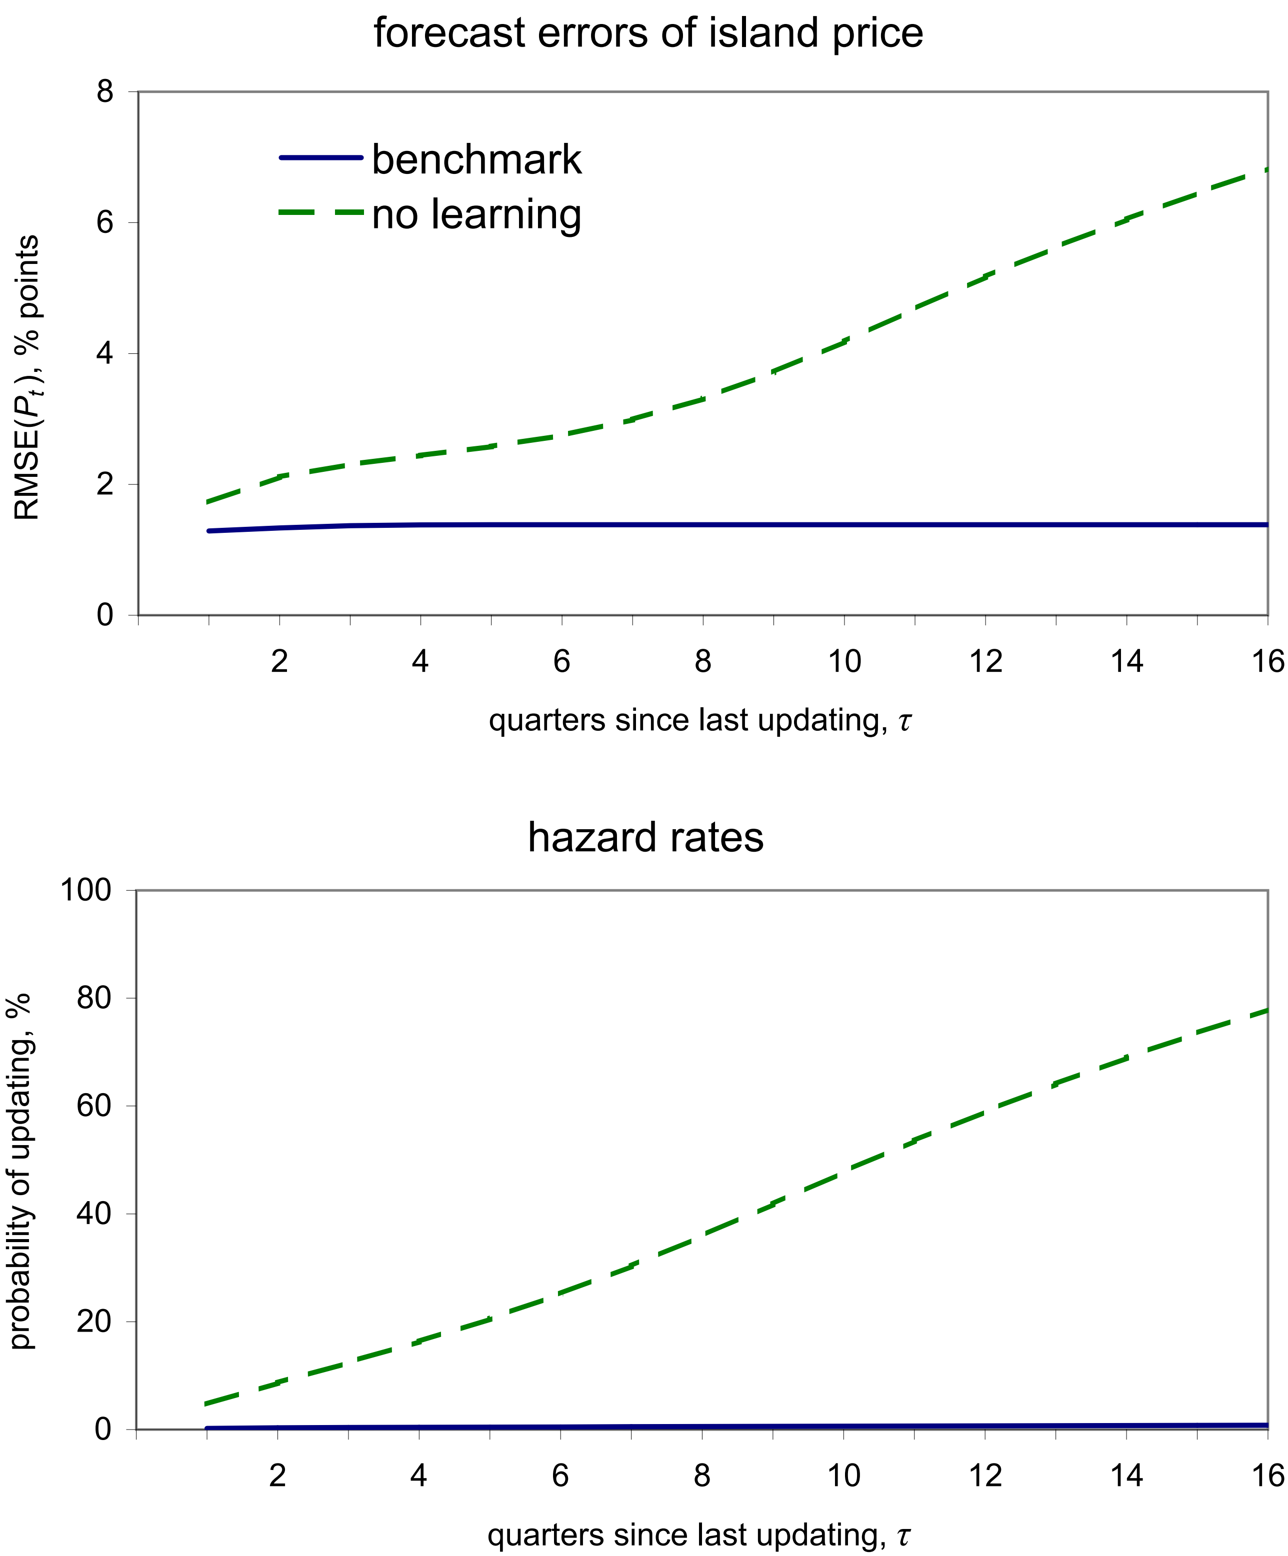

Figure 5. Half-life of output responses to +1 st. dev. impulse to money growth, general equilibrium