

Inflation and Interest Rates with Endogenous Market Segmentation

Aubhik Khan

The Ohio State University

Julia K. Thomas*

The Ohio State University

December 2011

ABSTRACT

We examine a monetary economy wherein endogenous asset market segmentation permits the extent of household participation in open market operations to vary smoothly with changes in aggregate conditions. While we impose no stickiness at the microeconomic level in either prices or portfolio adjustment, we find that our flexible asset market segmentation can deliver gradual adjustment in the aggregate price level following a monetary shock and thus persistent non-neutralities. In our model economy, households incur fixed transactions costs when exchanging bonds and money and, as a result, carry money balances in excess of current spending to limit the frequency of such trades. As only a fraction of households choose to actively trade bonds and money at any given time, asset markets are endogenously segmented.

Because our households can alter the timing of their trading activities, the extent of market segmentation varies over time in response to real and nominal shocks, as does the distribution of real balances across households. We show that this added flexibility can substantially reinforce the sluggishness in aggregate price adjustment following a monetary shock, relative to models with exogenously segmented asset markets, and it can transform dramatic, transitory changes in real and nominal interest rates into more moderate and persistent liquidity effects. We also show that, following an endowment shock, changes in households' portfolio adjustment timing generate persistence in inflation and interest rates that is otherwise absent. Finally, we show that these changes can reshape aggregate quantity dynamics in a version of the model with production, generating hump-shaped responses in employment and output following a monotone shock to productivity.

*We thank Michael Dotsey, Chris Edmond, Robert King, Igor Livshits, Sylvain Leduc, Edward Nelson, Thomas Sargent and Pierre-Olivier Weill for useful comments and suggestions. We also thank seminar participants at Boston University, The Federal Reserve Banks of Kansas City, New York and St. Louis, NYU Stern and the University of British Columbia, as well as conference participants at the 2005 Society for Economic Dynamics meeting, 2006 Canadian Macro Study Group, 2007 Philadelphia Workshop on Monetary and Macroeconomics, Cowles Foundation 75th Anniversary Macroeconomics Conference on Lumpy Adjustment, 2008 WUSTL Conference on Money, Credit and Policy, and 2008 Texas Monetary Conference. Thomas gratefully acknowledges research support from the Alfred P. Sloan Foundation and the National Science Foundation (grant #0318163).

1 Introduction

There is a wealth of empirical research documenting not only the co-movement of real and nominal series at higher frequencies, but what is widely accepted as evidence of persistent responses in real variables following nominal disturbances. We study such co-movements using a monetary model wherein households face fixed costs of transferring wealth between interest-bearing assets and money. As a result of these transactions costs, participation in asset markets is endogenously segmented. Households infrequently access their interest income, and they carry money balances in excess of current spending to help finance their spending over coming periods. As is well known, market segmentation implies that open market operations can have real effects, because they directly involve only a subset of households. This paper establishes that, when market segmentation is endogenized in a setting where households hold inventories of liquid, low-yield assets, changes in the fractions of households participating in asset markets can add considerable persistence to movements in both nominal and real variables.

As in the many studies in monetary economics that precede us, several empirical relationships involving money, interest rates and prices motivate our work. First, short-term real interest rates are negatively correlated with expected inflation. Barr and Campbell provide direct evidence for this using U.K. data involving inflation-indexed bonds. Second, VAR studies consistently find evidence of liquidity effects; expansionary open market operations are associated with persistent reductions in short-term nominal interest rates. See, for example, Leeper, Sims, and Zha (1996) and Christiano, Eichenbaum, and Evans (1999). Third, the general price level appears to adjust slowly to nominal shocks. This finding is widely supported by the VAR literature as, for example, in the studies of Leeper, Sims, and Zha (1996), Christiano, Eichenbaum, and Evans (1999), and Uhlig (2005). Moreover, King and Watson (1996) show that, at business cycle frequencies, the price level is positively correlated with lagged real output. Additional evidence for the slow adjustment of the price level, discussed in Alvarez, Atkeson and Edmond (2009), is provided by the pattern of short-term movements seen between the ratio of money to consumption and velocity. The correlation between the ratio of money (M2) to consumption (PCE) and the corresponding measure of velocity is -0.89 for HP-filtered monthly data.

The most common theoretical approach used to organize such empirical evidence is the study of models where nominal prices are sticky at the firm level, and firms must produce to satisfy the demand arising at their given prices. These models have grown in sophistication with the inclusion of various aggregate frictions and stochastic driving processes, and are now widely used by central banks for monetary policy analysis.¹ However, there are still several open issues regarding the predictive viability of the often-termed *sticky-price* framework, beyond the well-known debate over the frequency of microeconomic price adjustment.² Two of these issues involve the model's predictions for interest rates; a third is the concern that money itself has no role in the model.

Because sticky price models assume a representative household, they have at their core a consumption Euler equation leading them to predict real interest rates tied to the growth rate of

¹Two of the most prominent examples are the models of Christiano, Eichenbaum and Evans (2005) and Smets and Wouters (2003).

²Chari, Kehoe and McGrattan (2000) initiated this debate with their observation that the degree of firm-level price stickiness imposed in the model fully determines the persistence in its aggregate responses to monetary shocks. Examining data on goods and service prices from the Bureau of Labor Statistics, Bils and Klenow (2004) find the typical duration of a price is only 4.3 or 5.5 months, depending on whether sale prices are excluded. Nakamura and Steinsson (2008) pursue a more detailed BLS data set including observations from the CPI Research Database and find evidence of far greater micro-level stickiness, documenting price durations on the order of 8 to 12 months.

aggregate consumption. Canzoneri, Cumby and Diba (2007) show that this prediction is refuted in the data not simply in terms of average levels, but more critically in the directions of change. Using U.S. consumption data, Canzoneri et al. retrieve the interest rate series implied by the model's Euler equation under a series of leading preference specifications (including five commonly used variants of habit persistence). They find that the model-implied interest rate series is consistently negatively correlated with observed U.S. interest rates.

Dotsey and King (2005) raise a related concern involving interest rate predictions. They resolve a long-standing issue for the sticky-price literature by developing an (S,s) model of nominal price setting that is consistent with empirical estimates of the persistence in inflation movements. In the process, however, they discover that the model predicts a *rise* in short-term nominal interest rates following a persistent positive shock to money growth rates. In other words, they show that the state-of-the-art micro-founded menu cost model contradicts the liquidity effect commonly seen as the first channel through which monetary policy is transmitted, which was described by Milton Friedman (1968) as follows: "*The initial impact of increasing the quantity of money at a faster rate than it has been increasing is to make interest rates lower for a time than they would otherwise have been.*"

Whether or not we choose to view the absent liquidity effect as a critical shortcoming of the paradigm, there is a more substantive disagreement between Friedman's view of the real effects of monetary policy and that underlying sticky-price models. Friedman viewed changes in the velocity of money as central in determining the mechanics of movements in output, employment and prices following an expansionary open market operation. Consider his testimony to the House of Commons Select Committee in 1979; "... *the initial effect of a change in monetary growth is an offsetting movement in velocity, followed by changes in the growth of spending initially manifested in output and employment, and only later in inflation.*" (Friedman, 1980).³ This view of the monetary transmission mechanism is firmly rejected by sticky-price models; the velocity of money, and indeed real balances themselves, are almost entirely irrelevant to such models' predictions.

We explore an alternative class of monetary model wherein the link between aggregate consumption growth and real interest rates is broken, a liquidity effect arises naturally, and real balances carry an essential role in facilitating goods market transactions. Once termed *limited participation* models, more recent examples of the framework are known as *segmented asset markets* models. We contribute to the segmented asset markets literature what the menu cost model contributes to the sticky-price literature. We endogenize the microeconomic stickiness at the heart of the model's non-neutralities - in this case, the infrequency of household portfolio adjustments. Our work builds on an important literature that studies monetary policy in models with exogenously segmented markets.⁴ As in the work of Grossman and Weiss (1983), Rotemberg (1984), and Alvarez, Atkeson and Edmond (2009), households in our model economy only periodically access the market for interest-bearing bonds (broadly interpreted as markets for relatively high-yield, illiquid assets). As a result, households carry inventories of money (interpreted as relatively low yield, liquid assets) to help finance their spending over several periods of future spending.⁵

In our model, as in the model of Alvarez, Atkeson and Edmond (2009), household spending rates (ratios of the value of current consumption to money holdings) are lowest among households that

³We thank Ed Nelson for bringing this to our attention.

⁴See Alvarez, Lucas and Weber (2001) and the references therein.

⁵A second branch of the exogenously segmented markets literature assumes a fixed subset of households is always participating in asset markets, while a second subset is permanently excluded. Recent examples include Occhino (2004, 2008) and Williamson (2008).

have recently transferred wealth held as bonds into money, and they rise with the time since such a transfer has occurred. In such an environment, changes in monetary policy alter the distribution of money holding across households with different spending rates, and this, in turn, has implications for interest rates and inflation. More specifically, an unanticipated expansionary open market operation increases the relative money holdings of households currently trading bonds for money. Because such households anticipate the longest wait before their next such trade, and thus have the lowest spending rates across all households, this leads to a fall in the aggregate spending rate, or aggregate velocity. As a result, nominal spending rises by less than the money supply, delaying the rise in the aggregate price level. Moreover, because real interest rates are determined not by aggregate consumption growth, but instead by the consumption of participating households at adjacent dates, a liquidity effect emerges directly from the fact that households trading in the bond markets at the date of the shock see the greatest rise in their consumption.

Our model is distinguished from other monetary models with segmented asset markets by the fact that the fraction (and composition) of households directly involved in an open market operation is endogenously determined. Our households individually choose the timing of their asset market participation by choosing when to pay fixed transactions costs required for such exchanges. In this respect, we build on the work of Alvarez, Atkeson and Kehoe (2002).⁶ Here, as there, households face idiosyncratic risk that influences when they adjust their portfolios of bonds and money. In contrast to the Alvarez, Atkeson and Kehoe (2002) setting, however, our households do not typically exhaust their money balances with their current consumption expenditure. Thus, our model delivers a nontrivial distribution of money that evolves across periods, and it permits the changes in aggregate velocity that give rise to sluggish adjustment in the aggregate price level.

Relative to other inventory-theoretic models of money, our model is unique in allowing the extent of market segmentation to vary smoothly over time with changes in the economy's state. In our model, the inclusion of idiosyncratic differences across households implies that the fraction of households choosing to participate in the asset markets remains nontrivial over time, as does the distribution of money.⁷ This has important implications for the propagation of nominal disturbances. Following an open market operation, endogenous changes in the timing of households' active participation in asset markets can gradualize aggregate price adjustment relative to that in a model with exogenous market segmentation. Furthermore, the resulting changes in household participation rates can transform sharp, temporary changes in real and nominal interest rates into moderate, persistent liquidity effects.⁸

Our findings here suggest that the segmented asset markets paradigm may be a promising alternative environment in which to explore the interaction of real and nominal variables, alongside their implications for monetary policy. Money has a prominent role in these models; indeed, changes in velocity are at the heart of the short-term non-neutrality they exhibit. The endogenous segmentation model of Alvarez, Atkeson and Kehoe (2002) succeeds in generating persistent liquidity effects and in reproducing the negative relation between real interest rates and anticipated

⁶Our work is also related to Chatterjee and Corbae (1991), who study an economy where some households choose to pay a fixed cost to trade bonds.

⁷Chiu (2007) has developed what is to our knowledge the only other monetary model with endogenously segmented markets; however, the absence of any source of heterogeneity there implies a fragile segmentation that tends either to collapse or remain unchanged in response to shocks.

⁸One exogenous segmentation model that also succeeds in this respect is the model of Williamson (2008) wherein households are permanently divided into groups with and without access to the asset markets. There, the assumption that households with (without) such access prefer to trade among themselves in the goods markets delivers a second type of segmentation that can lead to persistent liquidity effects.

inflation, while the Alvarez, Atkeson and Edmond (2009) inventory-theoretic model of money with exogenous segmentation separately delivers sluggish adjustment of the price level, and hence persistent inflation responses to nominal shocks. Drawing on elements from each of these models, we develop a model that simultaneously succeeds with regard to both sets of regularities. Moreover, as we have mentioned above, we find that endogenizing the microeconomic stickiness responsible for the model's non-neutralities can actually improve its performance.

Through a series of examples, we show that changes in the fractions of households choosing to exchange bonds and money can substantially reinforce the effects of market segmentation, so long as the difference between households' mean and maximum time between trades is not too large. As this statement suggests, the results obtained depend upon the underlying distribution of transactions costs, as this determines the average distribution of money holdings. When transactions costs are large, so that the mean time between asset trades is long, households tend to hold relatively large inventories of money. If, in addition, the maximum time between trades significantly exceeds the mean time, households on average return to the asset markets with substantial remaining balances. In such circumstances, we find that the persistence in inflation implied by the exogenous segmentation model can be eliminated, as households not participating in the asset markets sharply raise their spending rates following an open market operation. However, the transactions cost distributions yielding such results deliver an implausibly low frequency of household portfolio adjustments relative to the available microeconomic evidence.

When we assume a transactions cost distribution implying consistency with both average aggregate velocity and the frequency of households' asset market trades, the mean and maximum times between trades are similar, and the effects of market segmentation are strengthened by endogenous changes in household participation rates following a transitory money growth shock. In response to persistent money growth shocks, such changes lead to more sluggish price adjustment and more persistent liquidity effects in both nominal and real interest rates. When we examine real shocks with monetary policy following a Taylor rule, our economy generates persistence in the responses of inflation and interest rates altogether absent under time-invariant market segmentation. Moreover, in a version of our model with endogenous production, a persistent technology shock leads to non-monotone responses in employment and output, and these are again traced to changes in the timing of households' portfolio adjustments.

Beyond the findings noted above, our paper offers an independent theoretical contribution in formally establishing how the results of Alvarez, Atkeson and Kehoe (2002) may be extended to a model where there are persistent differences across households. Here, such extension is necessary, because cash-in-advance constraints do not always bind so that households carry inventories of money, thereby transmitting the effects of temporary idiosyncratic differences across periods. By assuming a full set of state-contingent nominal bonds that allow risk-sharing across households, we ensure that these differences across households are persistent, but not permanent. Because households must pay fixed transactions costs to access their bond holdings, the presence of state-contingent bonds in our economy does not lead to full insurance; households that are ex-ante identical diverge over time as idiosyncratic realizations of shocks drive differences in their money and bond holdings. Nonetheless, we prove that, whenever a heterogeneous group of households enters the bond market at the same time, all previous differences between them are eliminated. As a result, our model economy exhibits limited memory. Exploiting this property, we are able to apply the numerical approach to solving generalized (S,s) models developed by King and Thomas (2006) in a setting where the consumption and savings decisions of heterogeneous risk-averse households

are directly influenced by nonconvex costs. While this approach has been applied previously in solving models where risk-neutral production units face idiosyncratic fixed costs of adjusting their prices or factors of production (as in Dotsey, King and Wolman (1999) and Thomas (2002)), this is to our knowledge the first application involving heterogeneity among households.

2 Model

We begin by considering an endowment economy. (A version of the model with production is described in section 4.3.) Below, we provide an overview of the model. Thereafter, we proceed to a more formal description of households' problems, followed by the description of a financial intermediary that sells households claims contingent on both aggregate and individual states. Next, we show that there is an equivalent, but more tractable, representation of households' lifetime optimization problems, given their ability to purchase such individual-state-contingent bonds alongside the fact that they are ex-ante identical. Proofs of all lemmas are provided in the appendix.

2.1 Overview

The model economy has three sets of agents: a unit measure of ex-ante identical households, a perfectly competitive financial intermediary, and a monetary authority. Each infinitely-lived household values consumption in every date of life, with period utility $u(c)$, and it discounts future utility with the constant discount factor β , where $\beta \in (0, 1)$. In each period, households receive a common endowment, y . This endowment varies exogenously over time, as does the growth rate of the aggregate money supply, μ . Defining the date t realization of aggregate shocks as $s_t = (y_t, \mu_t)$, we denote the history of aggregate shocks by $s^t = (s_1, \dots, s_t)$, and the initial-period probability density over aggregate histories by $g(s^t)$.

Households have two means of saving. First, they have access to a complete set of state-contingent nominal bonds. These are purchased from a financial intermediary described below, and are maintained in interest-bearing accounts that we will refer to as households' *brokerage accounts*, following the language of Alvarez, Atkeson and Edmond (2009). Next, they also save using money, which they maintain in their *bank accounts* and use to conduct trades in the goods market.⁹ Households have the opportunity to transfer assets between their two accounts at the start of each period; this occurs after the realization of all current shocks, but prior to any trading in the goods market. As such, it is expositionally convenient to refer to each period as consisting of two subperiods that we will term *transfer-time* and *shopping-time*, although nothing in the environment necessitates this approach.

There are three inter-related frictions leading households to maintain money in their bank accounts. First, as in a standard cash-in-advance environment, households cannot consume their own endowments. Each household consists of a worker and a shopper, and the worker must trade the household endowment for money while the shopper is purchasing consumption goods. As a result, the household receives the nominal value of its endowment, $P(s^t)y(s^t)$, only at the end of the period after current goods trade has ceased.¹⁰ We assume that these end-of-period nominal

⁹When allowed to store money in their brokerage accounts, households never do so given positive nominal interest rates paid on bonds. Thus, we simplify the model's exposition here by assuming that money is held only in bank accounts and verify that nominal rates remain positive throughout our results.

¹⁰While this worker-shopper arrangement may appear stark in an endowment economy, it is less so if one envisions that each household's endowment is one of a unit measure of differentiated inputs that enter a consumption aggregator

receipts are deposited across their two accounts, with fraction λ paid into bank accounts and the remainder into brokerage accounts. Second, as all trades in the goods market are conducted with money, each household's consumption purchases are constrained by the bank account balance it holds when shopping-time begins.

Note that, absent other frictions, each household would, in every period, simply shift from its brokerage account into its bank account exactly the money needed to finance current consumption expenditure not covered by the bank account paycheck from the previous period. There is, however, a third friction that prevents this, leading households to deliberately carry money across periods; this is the assumption that they must pay fixed costs each time they transfer assets between their two accounts. Given these fixed costs, households maintain stocks of money to limit the frequency of their transfers, and they follow generalized (S,s) rules in managing their bank accounts.

Transfer costs are fixed in that they are independent of the size of the transfer; however, they vary over time and across households. Here, we subsume the idiosyncratic features that distinguish households directly in their fixed costs by assuming that each household draws its own current transfer cost, ξ , from a time-invariant distribution $H(\xi)$ at the start of each period. Because this cost draw influences a household's decision of whether to undertake any transfer, and hence its current consumption and money savings, each household is distinguished by its history of such draws, $\xi^t = (\xi_1, \dots, \xi_t)$, with associated density $h(\xi^t) = h(\xi_1) \cdots h(\xi_t)$. As will be seen below, households are able to insure themselves in their brokerage accounts through the purchase of nominal bonds contingent on both aggregate and individual exogenous states.

2.2 Households

At the start of any period, given date-event history (s^t, ξ^t) , a household's brokerage account assets include nominal bonds, $B(s^t, \xi^t)$, purchased in the previous period at price $q(s^t, \xi^t)$, as well as the fraction of its income from the previous period that is deposited there, $(1-\lambda)P(s^{t-1})y(s^{t-1})$. The remainder, the paycheck, $\lambda P(s^{t-1})y(s^{t-1})$, is deposited into the household's bank account and supplements its money savings there from the previous period, $A(s^{t-1}, \xi^{t-1})$. Given this start of period portfolio and its current fixed cost, the household begins the period by determining whether or not to transfer assets across its two accounts. Denoting the household's start-of-period bank balance by $M(s^{t-1}, \xi^{t-1})$, where

$$M(s^{t-1}, \xi^{t-1}) \equiv A(s^{t-1}, \xi^{t-1}) + \lambda P(s^{t-1})y(s^{t-1}), \quad (1)$$

the relevant features of this choice are summarized in the chart below.

	brokerage account withdrawal	shopping-time bank balance
$z(s^t, \xi^t) = 1$	$x(s^t, \xi^t) + P(s^t)\xi_t$	$M(s^{t-1}, \xi^{t-1}) + x(s^t, \xi^t)$
$z(s^t, \xi^t) = 0$	0	$M(s^{t-1}, \xi^{t-1})$

An *active* household is indicated by $z(s^t, \xi^t) = 1$. In this case, the household selects a nonzero nominal transfer $x(s^t, \xi^t)$ from its brokerage account into its bank account and has $M(s^{t-1}, \xi^{t-1}) + x(s^t, \xi^t)$ available in its bank account at the start of the current shopping subperiod. Here, the household's current fixed cost applies, so $P(s^t)\xi_t$ is deducted from its nominal brokerage wealth. Alternatively, the household may choose to undertake no such transfer, setting $z(s^t, \xi^t) = 0$ and

with identical weights to produce the single good consumed by all.

remaining *inactive*. In that case, it enters into the shopping subperiod with no change to its start-of-period bank and brokerage account balances.

We assume that the period utility function, $u(c)$, is strictly increasing, strictly concave and twice-continuously differentiable, and that $\lim_{c \rightarrow 0} u'(c) = \infty$. Each household chooses its state-contingent plan for the timing and size of its account transfers ($z(s^t, \xi^t)$ and $x(s^t, \xi^t)$), and its bond purchases, money savings and consumption ($B(s^t, s_{t+1}, \xi^t, \xi_{t+1})$, $A(s^t, \xi^t)$ and $c(s^t, \xi^t)$), to maximize its expected discounted lifetime utility,

$$\sum_{t=1}^{\infty} \beta^{t-1} \int_{s^t} \int_{\xi^t} u(c(s^t, \xi^t)) h(\xi^t) g(s^t) d\xi^t ds^t, \quad (2)$$

subject to the sequence of constraints in (3) - (6).

$$B(s^t, \xi^t) + (1 - \lambda) P(s^{t-1}) y(s^{t-1}) \geq \left[x(s^t, \xi^t) + P(s^t) \xi_t \right] z(s^t, \xi^t) \quad (3)$$

$$+ \int_{s_{t+1}} \int_{\xi_{t+1}} q(s^t, s_{t+1}, \xi_{t+1}) B(s^t, s_{t+1}, \xi^t, \xi_{t+1}) ds_{t+1} d\xi_{t+1}$$

$$M(s^{t-1}, \xi^{t-1}) + x(s^t, \xi^t) z(s^t, \xi^t) \geq P(s^t) c(s^t, \xi^t) + A(s^t, \xi^t) \quad (4)$$

$$A(s^t, \xi^t) + \lambda P(s^t) y(s^t) \geq M(s^t, \xi^t) \quad (5)$$

$$A(s^t, \xi^t) \geq 0 \quad (6)$$

Equation 3 is the household's brokerage account budget constraint associated with history (s^t, ξ^t) , and requires that expenditures on new bonds together with any transfer to the bank account and associated fixed cost not exceed current brokerage account wealth. Next, the bank account budget constraint in equation 4 requires that the household's money balances entering the shopping subperiod cover its current consumption expenditure and any money savings for next period.¹¹ Money balances for next period, in (5), are these savings together with the bank paycheck received after completion of current goods trade. Equation 6 prevents the household from ending current trade with a negative bank balance; thus, taken together with the restriction in (4), it requires that cash be used for consumption purchases. Finally, in addition to this sequence of constraints, we also impose a standard limit condition on household debt:

$$\lim_{t \rightarrow \infty} \int_{s^t} \int_{\xi^t} q(s^t, \xi^t) B(s^t, \xi^t) ds^t d\xi^t \geq 0. \quad (7)$$

Following the approach of Alvarez, Atkeson and Kehoe (2002), we find it convenient to model risk-sharing by assuming a perfectly competitive financial intermediary that purchases government bonds with payoffs contingent on the aggregate shock and, in turn, sells to households bonds with payoffs contingent on both the aggregate and individual shocks. In particular, given aggregate history s^t , the intermediary purchases government-issued contingent claims $B(s^t, s_{t+1})$ at price

¹¹In those periods when a household is active, it has a single unified budget constraint, $B(s^t, \xi^t) + P(s^{t-1}) y(s^{t-1}) + A(s^{t-1}, \xi^{t-1}) \geq P(s^t) [\xi_t + c(s^t, \xi^t)] + A(s^t, \xi^t) + \int_{s_{t+1}} \int_{\xi_{t+1}} q(s^t, s_{t+1}, \xi_{t+1}) B(s^t, s_{t+1}, \xi^t, \xi_{t+1}) ds_{t+1} d\xi_{t+1}$.

$q(s^t, s_{t+1})$, and it sells them across households as claims contingent on individual transfer costs, ξ_{t+1} . Note that, as households' cost draws are not autocorrelated, the price of any such claim is $q(s^t, s_{t+1}, \xi_{t+1})$, independent of the individual history ξ^t .

For each (s^t, s_{t+1}) , the intermediary selects its aggregate bond purchases, $B(s^t, s_{t+1})$, and individual bond sales, $B(s^t, s_{t+1}, \xi^t, \xi_{t+1})$, to solve

$$\max_{\xi_{t+1}} \int \int_{\xi^t} q(s^t, s_{t+1}, \xi_{t+1}) B(s^t, s_{t+1}, \xi^t, \xi_{t+1}) h(\xi^t) d\xi^t d\xi_{t+1} - q(s^t, s_{t+1}) B(s^t, s_{t+1}) \quad (8)$$

subject to:

$$B(s^t, s_{t+1}) \geq \int \int_{\xi_{t+1}} B(s^t, s_{t+1}, \xi^t, \xi_{t+1}) h(\xi^t) h(\xi_{t+1}) d\xi^t d\xi_{t+1}. \quad (9)$$

The constraint in (9) requires that, for any (s^t, s_{t+1}) , the intermediary must purchase sufficient aggregate bonds to cover all individual bonds held against it for that aggregate history. Given s_{t+1} occurs, fraction $h(\xi_{t+1})$ of the households with history ξ^t to whom it sells such bonds will realize that state and demand payment. As shown in Lemma 1 below, the financial intermediary's zero profit condition immediately implies that the price of any individual bond associated with (s_{t+1}, ξ_{t+1}) is simply the product of the price of the relevant aggregate bond and the probability of an individual household drawing the transfer cost ξ_{t+1} .

Lemma 1. *The equilibrium price of state-contingent bonds issued by the financial intermediary, $q(s^t, s_{t+1}, \xi_{t+1})$, is given by $q(s^t, s_{t+1}, \xi_{t+1}) = q(s^t, s_{t+1}) h(\xi_{t+1})$.*

By assuming an initial period 0 throughout which households are identical, we allow them the opportunity to trade in individual-state-contingent bonds at a time when they have the same wealth and face the same probability distribution over all future individual histories. In this initial period, the government has some outstanding debt, \bar{B} , that is evenly distributed across households' brokerage accounts, and it repays this debt entirely by issuing new bonds. Households receive no endowment, draw no transfer costs and do not value consumption in this initial period. Rather, they simply purchase state-contingent bonds for period 1 subject to the common initial period brokerage budget constraint:

$$\bar{B} \geq \int \int_{s_1} B(s_1, \xi_1) q(s_1) h(\xi_1) d\xi_1 ds_1.$$

Following the proof of Lemma 1, section B of the appendix shows that the period 0 budget constraint above can be combined with the sequence of constraints in (3) to yield the following lifetime budget constraint common to all households.

$$\bar{B} \geq \sum_{t=1}^{\infty} \int \int q(s^t) h(\xi^t) \left(z(s^t, \xi^t) \left[x(s^t, \xi^t) + P(s^t) \xi_t \right] - P(s^{t-1}) (1 - \lambda) y(s^{t-1}) \right) d\xi^t ds^t, \quad (10)$$

where $q(s^t) \equiv q(s_1) \cdot q(s_1, s_2) \cdots q(s^{t-1}, s_t)$.

Finally, we assume that the monetary authority is subject to the sequence of constraints,

$$B(s^t) - \int_{s_{t+1}} q(s^t, s_{t+1}) B(s^t, s_{t+1}) ds_{t+1} = \bar{M}(s^t) - \bar{M}(s^{t-1}), \quad (11)$$

requiring that its current bonds be covered by a combination of new bond sales and the printing of new money. This sequence of constraints, alongside equilibrium in the money market, immediately implies that households' aggregate expenditures on new bonds in any period is exactly the difference between the aggregate of their current bonds and the change in the aggregate money supply:

$$\bar{M}(s^t) - \bar{M}(s^{t-1}) = B(s^t) - \int \int \int q(s^t, s_{t+1}) h(\xi_{t+1}) B(s^t, s_{t+1}, \xi^t, \xi_{t+1}) h(\xi^t) d\xi^t d\xi_{t+1} ds_{t+1}. \quad (12)$$

2.3 A risk sharing arrangement

Three aspects of the environment described above may be exploited to simplify our solution for competitive equilibrium: (i) fixed transfer costs are independently and identically distributed across households and time, (ii) households have perpetual access to a complete set of state-contingent claims in their brokerage accounts, and (iii) they are able to purchase these state-contingent bonds during an initial period in which they are perfectly identical. In this section, we show how these assumptions permit a more convenient representation of households' problems. In particular, exploiting the common lifetime budget constraint in (10) above, we will move from the household problem stated in section 2.2 to construct the equivalent problem of an extended family that manages all households' bonds in a joint brokerage account, and whose period-by-period decisions regarding bond purchases and account transfers implement the state-contingent lifetime plan selected by every household. In doing so, we transform our somewhat intractable initial problem into something to which we can apply the King and Thomas (2006) approach for solving aggregate economies involving heterogeneity arising due to (S,s) policies at the individual level.

Money as the individual state variable: A complete set of state-contingent claims in the brokerage account allows individuals to insure their bond holdings against idiosyncratic risk; these shocks only affect their bank accounts. Alternatively, an individual's money balance fully captures the cumulative effect of his history of idiosyncratic shocks. In Lemma 2, we prove that, prior to their current transfer cost draws, all relevant differences across households are fully summarized by their start-of-period money balances as they enter into any period.

Lemma 2. *Given $M(s^{t-1}, \xi^{t-1})$, the decisions $c(s^t, \xi^t)$, $A(s^t, \xi^t)$, $x(s^t, \xi^t)$ and $z(s^t, \xi^t)$ are independent of the history ξ^{t-1} .*

Since each ξ comes from an i.i.d. distribution, a household's draw in any given period does not predict its future draws, and thus directly affects only its asset transfer decision in that one period. This certainly affects current shopping-time money balances, and hence consumption. However, its only future effect is in determining the money balances with which the household will enter the subsequent period, given the household's ability to insure itself in its brokerage account by purchasing bonds contingent on both aggregate and individual shocks. In proving this result, we show that the solution to the original household problem from section 2.2, given the lifetime constraint in (10), is identical to the solution of an alternative problem where households pool risk period-by-period by each committing to pay the economywide average of the total transfers and associated fixed costs incurred across all active households in every period, irrespective of the timing and size of their own portfolio adjustments. It is immediate from this that households' bond holdings may be modeled as independent of their individual histories ξ^t . Thus, within every

period, the distinguishing features affecting any household's decisions can be summarized entirely by its start-of-period bank balance, $M(s^{t-1}, \xi^{t-1})$, and its current transfer cost, ξ_t .

Households as members of time-since-active groups: Our next lemma establishes that, within any period, all households that undertake an account transfer will select both a common consumption and a common end-of-period bank balance; hence they begin the subsequent period with the same bank (and brokerage) account balances.

Lemma 3. *For any (s^t, ξ^t) in which $z(s^t, \xi^t) = 1$, $c(s^t, \xi^t)$, $A(s^t, \xi^t)$ and $M(s^t, \xi^t)$ are independent of ξ^t .*

To understand this result, recall that household brokerage and bank accounts are joined in periods when they choose to adjust their portfolios, and all are identical when they make their state-contingent plans in date 0. Given this, in selecting their consumption for such periods, households equate their appropriately discounted marginal utility of consumption to the multiplier on the lifetime brokerage budget constraint from (10), which is common to all households. Next, in selecting what portion of their shopping-time bank balances to retain after consumption (hence their next-period balances), households equate the marginal utility of their current consumption to the expected return on a dollar saved for the next period weighted by their expected discounted marginal utility of next-period consumption. Given common inflation expectations and the common current consumption of active households, this implies that such households also share in common the same expected consumption for next period. Thus, all currently active households exit this period and enter the next period with common money holdings.

Note that Lemmas 2 - 3 combine to imply that, within any period, households that undertake balance transfers all enter shopping-time with the same bank balance, make the same shopping-time decisions, and then enter the next period as effectively identical. Moreover, of this group of currently active households, those households that do not undertake an account transfer again in the next period will continue to be indistinguishable from one another as they enter shopping-time, and hence will enter the subsequent period with common bank (and brokerage) account balances, and so forth. In other words, any household that was last active at some date t is effectively identical to any other household last active at that same date. This is useful in our numerical approach to solving for competitive equilibrium, since it allows us to move from identifying individual households by their current money holdings to instead identifying each household as a member of a particular time-since-active group, with all members of any one such group sharing in common the same start-of-period money balances.

Given the above results, we may track the distribution of households over time through two vectors, one indicating the measures of households entering the period in each time-since-active group, $[\theta_{j,t}]$, $j = 1, 2, \dots$, and the other storing the balances with which members of each of these current groups exited shopping-time in the previous period, $[A_{j,t}]$. From the latter, the current start-of-period balances held by members of each group are retrieved as $M_{jt} = A_{j,t} + \lambda P_{t-1} y_{t-1}$, where P_{t-1} represents the previous period's price level, and y_{t-1} the common endowment of the previous period. Households within any given start-of-period group j that do not pay their fixed costs move together into the current shopping subperiod with their starting balances M_{jt} . Across all start-of-period groups, those households that do pay to undertake a bank transfer will enter the current shopping subperiod in time-since-active group 0 with common shopping-time balances, $M_{0,t}$, which we refer to as the current *target* money balances.

Threshold transfer rules: Finally, we establish that households follow threshold policies in determining whether or not to transfer assets between their brokerage and bank accounts. Specifically, given its start-of-period money balances, each household has some maximum fixed cost that it is willing to pay to undertake an account transfer and adjust its balances to the current target.

Lemma 4. For any (s^t, ξ^{t-1}) , $\mathcal{Z} = \{\xi_t \mid z(s^t, \xi^t) = 1\}$ is a convex set bounded below by 0.

As our preceding results imply that all members of any given start-of-period group j are effectively identical prior to the draws of their current transfer costs, this last result allows convenient determination of the fractions of each such group undertaking account transfers, and thus the shopping-time distribution of households. Define the threshold cost ξ_{jt}^T as that fixed cost that leaves any household in time-since-active group j indifferent to an account transfer at date t . Households in the group drawing costs at or below ξ_{jt}^T pay to adjust their portfolios, while other members of the group do not. Thus, within each group j , the fraction of its members shifting assets to reach the current target bank balance is given by $\alpha_{jt} \equiv H(\xi_{jt}^T)$. Each such active household undertakes a transfer $x_{jt} = M_{0,t} - M_{j,t}$, and the total transfer cost paid across all members of the group are $\theta_j \int_0^{H^{-1}(\alpha_j)} \xi h(\xi) d\xi$.

A family problem: Collecting the results above, and assuming that aggregate shocks are Markov, we may re-express the lifetime plans formulated by individual households as the solution to the recursive problem of an extended family that manages the joint brokerage account of all households and acts to maximize the equally-weighted sum of their utilities. In each period, given the starting distribution of households summarized by $\{\theta_j, A_j\}$ and the current price level P , the family selects the fractions of households from each time-since-active group to receive account transfers, α_j , (and hence the distribution of households over time-since-active groups at the start of next period, θ'_j), the shopping-time bank balance of each active household, M_0 , achieved by transfers from the family brokerage account, as well as the consumption and money savings associated with members of each shopping-time group, c_j and A'_{j+1} respectively, to solve the problem in (13) - (19) below. In solving this problem, the family takes as given the current endogenous aggregate state $K = [\{\bar{\theta}_j, \bar{A}_j\}, P_{-1}y_{-1}, \bar{M}_{-1}]$, and it assumes the future endogenous state will be determined by a mapping F that it also takes as given; $K' = F(K, s)$. In equilibrium, K' is consistent with the family's decisions.

$$V(\{\theta_j, A_j\}; K, s) = \max \sum_{j=1}^{\infty} \theta_j [\alpha_j u(c_0) + (1 - \alpha_j) u(c_j)] + \beta \int_{s'} V(\{\theta'_j, A'_j\}; K', s') g(s, s') ds' \quad (13)$$

subject to:

$$\sum_{j=1}^{\infty} \alpha_j \theta_j [M_0 - M_j] + P \sum_{j=1}^{\infty} \theta_j \left[\int_0^{H^{-1}(\alpha_j)} \xi h(\xi) d\xi \right] \leq \bar{M} - \bar{M}_{-1} + (1 - \lambda) P_{-1}y_{-1} \quad (14)$$

$$M_j = [A_j + \lambda P_{-1}y_{-1}], \text{ for } j > 0 \quad (15)$$

$$M_j \geq P c_j + A'_{j+1}, \text{ for } j \geq 0 \quad (16)$$

$$A'_{j+1} \geq 0, \text{ for } j \geq 0 \quad (17)$$

$$\sum_{j=1}^{\infty} \alpha_j \theta_j \geq \theta'_1 \quad (18)$$

$$\theta_j (1 - \alpha_j) \geq \theta'_{j+1}, \text{ for } j > 0 \quad (19)$$

Recall from equation 12 that money market clearing in each period requires that the aggregate of households' current bonds less their expenditures on new bonds must equal the change in the aggregate money supply. By imposing this equilibrium condition, we may use equation 14 to represent the family's budget constraint requiring that its joint brokerage assets cover all current transfers to active households and associated fixed costs, as well as all bond purchases for the next period. Next, equation 15 identifies the start-of-period money balances associated with each time-since-active group j , and (16)-(17) represent the bank account budget and cash-in-advance constraints that apply to members of each shopping-time group. Finally, equations 18 - 19 describe the evolution of households across groups over time. In (18), the total active households (shopping in group 0) in the current period is the population-weighted sum of the fractions of households made active from each start-of-period group, and these households move together to begin the next period in time-since-active group 1. In (19), households in any given time-since-active group j that are inactive in the current period will move into the next period as members of group $j + 1$.

2.4 Solution and parameter values

Recall that we imposed money-market clearing in formulating the family's problem above. As such, we can retrieve equilibrium allocations as the solution to (13) - (19) by appending to that problem the goods market clearing condition needed to determine the equilibrium price level taken as given by the family:

$$y = c_0 \sum_{j=1}^{\infty} \theta_j \alpha_j + \sum_{j=1}^{\infty} \theta_j (1 - \alpha_j) c_j + \sum_{j=1}^{\infty} \theta_j \left[\int_0^{H^{-1}(\alpha_j)} x h(x) dx \right]. \quad (20)$$

Equation 20 simply states that, within each period, the current aggregate endowment must satisfy total consumption demand across all active and inactive households together with the economywide fixed costs associated with account transfers.

In the results to follow, we abstract from trend growth in endowments, and we assume that money supply is increased at rate μ^* in the economy's steady-state. Thus, the steady-state is associated with inflation at rate μ^* and a stationary distribution of households over real balances described by $[\boldsymbol{\theta}^*, \mathbf{a}^*]$, where $\boldsymbol{\theta}^* = \{\theta_j^*\}$ and $\mathbf{a}^* = \{a_j^*\}$, with $a_j \equiv \frac{A_j}{P-1}$. As any given household travels outward across time-since-active groups, it finds its actual real balances for shopping time, $(a_j^* + \lambda y^*) / (1 + \mu^*)$, falling further and further below target shopping balances; thus, the maximum fixed cost it is willing to pay to undertake an account transfer rises. Given a finite upper support on the distribution of fixed transfer costs, this implies that no household will delay activity beyond some finite maximum number of periods, which we denote by J . Thus, the two vectors describing the distribution of households are each of finite length J . In solving the steady-state of our economy, we isolate J as that group j by which α_j is chosen to be 1.

Having arrived at the time-since-active representation described above, we can follow King and Thomas (2006) in applying linear methods to solve for our economy's aggregate dynamics local to the deterministic steady-state once we have addressed two remaining details. First, as the linear solution does not allow for a changing number of time-since-active groups, we must restrict J to be time-invariant. Thus, we assume that $\alpha_{J,t} = 1$ for all t , and we then verify that $\alpha_{j,t} \in (0, 1)$, for $j = 1, \dots, J - 1$, is selected throughout our simulations. Second, we assume that, in every date t , all households that enter shopping in time-since-active group $J - 1$ completely exhaust their money balances; $a_{J,t} = 0$. Given that any such household will undertake an account transfer with certainty

at the start of the next period, this assumption is consistent with optimizing behavior so long as we verify that nominal interest rates are always positive.¹²

In parameterizing our model, we set the length of a period to one quarter, and we choose the steady-state inflation rate μ^* to imply an average annual inflation at 3 percent. Period utility is iso-elastic, $u(c) = \frac{c^{1-\sigma}-1}{1-\sigma}$, with $\sigma = 2$, and we select the subjective discount factor β to imply an average annual real interest rate of 3 percent. The steady-state aggregate endowment is normalized to 1, and the fraction of the endowment paid to household bank accounts (which may be interpreted as household wages) is $\lambda = 0.6$, corresponding to labor's share of output. Holding these parameters fixed, we will consider several alternative assumptions regarding the distribution of the fixed costs that cause market segmentation in our model, as we discuss below.

3 Examples

We begin to explore our model's dynamics in this section through a series of examples considering the response to a rise in the money growth rate that, once observed, is known to be perfectly transitory. Throughout this section, we abstract from shocks to the endowment to study the effects of a monetary shock in isolation, and to isolate those aspects caused by the endogenous changes in the degree of market segmentation that distinguish our model. We consider each of three cases distinguished only by the distribution of fixed transfer costs, beginning with a baseline case where this distribution is uniform on the interval 0 to B . There, we set the upper support at $B = 0.25$ to imply that the maximum time that any household remains inactive is $J = 6$ quarters. For individual households, the result is a 4.82 quarter average duration between account transfers. In the aggregate, this calibration results in a steady-state velocity of 1.9, which corresponds to the U.S. average between 1995Q1 and 2005Q1.¹³ In our second example, we raise the maximum transfer cost to imply an aggregate velocity matching the U.S. postwar average, at 1.5. Retaining the assumption that transfer costs are distributed uniformly, this implies a mean household inactivity duration of 7 quarters and a substantially longer maximum inactivity spell, at 10 quarters. This large difference between a household's average expected period of inactivity versus the maximum such spell will be seen to have important qualitative implications for the model's aggregate dynamics. Thus, in our third example, we will move to consider a more flexible cost distribution under which aggregate velocity again averages 1.5, but mean and maximum durations are close at 9.55 and 10 quarters, respectively.

3.1 Steady-state

Before examining its responses to shocks, it is useful to begin with a discussion of household portfolio adjustment timing in our model's steady-state. We first consider how each of our three examples relates to the available micro-evidence provided by Vissing-Jørgensen (2002).

¹²Given positive nominal rates, if $a_{j,t} > 0$ ever were to occur, the family could have improved its welfare by reducing the target balances given to active households at date $t - (J - 1)$ and increasing its bond purchases at that date to finance increased transfers to a subsequent group of active households for whom the non-negativity constraint would eventually bind.

¹³For comparability, we follow Alvarez, Atkeson and Edmond (2003) in our measures of money and velocity. As in their paper, money is broadly defined as the sum of currency, checkable deposits, and time and savings deposits. They show that the opportunity cost of these assets, relative to short-term Treasury securities, is substantial and, as a whole, not very different from that of M1. Next, velocity is computed as the ratio of nominal personal consumption expenditures to money.

Using the Consumer Expenditure Survey, Vissing-Jørgensen computes that the fraction of households that actively bought or sold risky assets (stocks, bonds, mutual funds and other such securities), between one year and the next ranges from 0.29 to 0.53 as a function of financial wealth.¹⁴ For a direct comparison with each version of our quarterly model, we compute the steady-state unconditional probability that a household will undertake active trade within one year as $\sum_{j=1}^J (\theta_j - \theta_{j+4})$, with $\theta_{j+4} = 0$ for $j > J - 4$.¹⁵ We find that the fraction of households actively trading in an average year is 0.78 in our baseline example, which is quite high relative to the Vissing-Jørgensen data. This may be explained in part by the fact that the transfer costs in this example are calibrated to match aggregate velocity over the decade ending at 2005, when transactions costs were presumably lower than in her 1982-1996 sample period. When we instead calibrate to match aggregate velocity over the postwar period in our second example (with higher transactions costs), the fraction of households trading annually falls to 0.55, slightly above the empirical range. Our most successful example with regard to this evidence is the third, where high transactions costs are drawn from a distribution implying the same postwar aggregate velocity, but longer expected episodes of inactivity. There, the model predicts an average annual fraction of households conducting trades well within the empirical range, at 0.42.

We cannot compare our examples' mean inactivity durations to that implied by the Vissing-Jørgensen data without making some assumption about the shape of the empirical hazard. If one assumes that the probability of an active trade is constant from quarter to quarter in the data, then the range reported above implies a mean duration of household inactivity ranging from 7.5 to 13.8 quarters. Recall that the mean duration of inactivity in our baseline example is only 4.8 quarters, while that in our second example involving high transactions costs is 7 quarters. This again suggests that the frequencies of active trades implied by these two versions of our model are, if anything, high relative to the data. However, our third example with both high maximum and mean inactivity spells exhibits an average duration within the range implied by the data, at 9.55 quarters. Thus, we will study this third case as we move to examine our model's dynamic results in section 4.

We confine our remaining discussion of the model's steady-state to that arising under our baseline parameters, as the qualitative aspects that we will emphasize hold across all of our examples. Here, with both the aggregate endowment and the money growth rate fixed at their mean values, six groups of households enter into each period, with these groups corresponding to the number of quarters that have elapsed since members' last account transfer. As any individual household moves through these groups over time, its real money balances available for shopping fall further and further below the target value, 2.936, given both inflation and its expenditures subsequent to its last time active. To correct this widening distance between actual and target real balances, the household becomes increasingly willing to incur a fixed transfer cost. This implies that the threshold cost separating active households from inactive ones rises with households' time-since-active.

¹⁴The CEX interviews about 4500 households each quarter, and each household is interviewed five times, with financial information gathered in the final interview only. Vissing-Jørgensen (2002) limits her sample to 6770 households that held risky assets both at the time of the fifth interview and one year earlier. She finds that the probabilities of buying or selling risky assets do not significantly change when the sample, spanning 1982 - 1996, is split into subsamples according to interview dates.

¹⁵For example, in any date t of our model's steady-state, there are θ_1 households entering the period in time-since-active group 1. After one year, at the start of period $t + 5$, $\theta_1 - \theta_5$ of that original group have undertaken at least one trade. Thus, the fraction of them that have traded within a year is $\frac{\theta_1 - \theta_5}{\theta_1}$. The overall fraction trading within one year is the population-weighted sum of these fractions across each starting group, $j = 1, \dots, J$.

Thus, as transfer costs are drawn from a common distribution, the fraction of households currently active in Table 1 rises across start-of-period groups.

TABLE 1: Determination of steady-state shopping-time distribution

time-since-active group	0	1	2	3	4	5	6
start-of-period populations		0.208	0.205	0.196	0.174	0.136	0.082
<i>fraction currently active</i>		<i>0.011</i>	<i>0.045</i>	<i>0.113</i>	<i>0.218</i>	<i>0.397</i>	<i>1.000</i>
shopping-time real balances	2.936	2.510	2.095	1.691	1.301	0.929	n/a
shopping-time populations	0.208	0.205	0.196	0.174	0.136	0.082	0

In figure 1A, we plot the steady-state distribution of households across groups as they enter shopping-time from the final row of table 1. Corresponding to the rising fractions of active households shown above, the dashed curve reflecting the measures of households in each shopping-time group monotonically declines across groups. The solid curve in the figure illustrates the ratios of real consumption expenditure relative to real balances, individual velocities, associated with the members of each shopping-time group. Because households are aware that they must use their current balances to finance consumption not only in the current period but also throughout subsequent periods of inactivity, individual spending rates rise across groups in response to a declining expected duration of future inactivity. Currently active households, those households in group 0, face the longest potential time before their next balance transfer, and thus have the lowest individual velocities. By contrast, households currently shopping in group 5 will undertake a transfer with certainty at the start of the next period; thus, individual velocity is 1 for members of this last group.

Two aspects distinguishing our endogenous segmentation model will be relevant in its responses to shocks below. First, on average, a household's probability of becoming active monotonically rises with the time since its last active date, as seen above. Second, these probabilities change over time as shocks influence the value households place on adjusting their bank balances. To isolate the importance of these two elements below, we will at times contrast the responses in our economy to those in a corresponding economy that has neither. In that otherwise identical *fixed duration* model, the timing of any household's next account transfer is certain and is not allowed to change with the economy's state. For comparability with our endogenously segmented economy, where households' mean duration of inactivity is 4.8 quarters, households in the corresponding fixed duration model undertake transfers exactly once every 5 quarters.

Figure 1B displays the steady-state of the fixed duration model. There, households enter every period evenly distributed across 5 time-since-active groups. Throughout groups 1 through 4, fraction 0 of each group's members are allowed to undertake account transfers, while fraction 1 of the members of group 5 are automatically made active. Thus, 20 percent of households enter into shopping in each time-since-active group 0 through 4, and this shopping-time distribution remains fixed over time. As in our model with endogenously timed household portfolio adjustments, here too individual velocities monotonically rise with time-since-active and hit 1 in the final shopping group. However, given its lesser maximum duration of inactivity (5 quarters here versus 6 in the endogenous segmentation model), households in the fixed duration economy exhibit somewhat higher spending rates throughout the distribution relative to those in panel A.

3.2 Money injection: a baseline example

Throughout the remainder of this section, we study the endogenous segmentation economy’s dynamics using a series of examples designed to illustrate its underlying mechanics. In each case, we examine the effects of an unanticipated one period rise in the money growth rate.

Fixed duration model. For reference, we begin in figure 2 with an examination of the aggregate response in the fixed duration model, where the fractions of active households across groups are fixed and dictated by $\alpha_{FD} = [0\ 0\ 0\ 0\ 1]$.¹⁶ As seen in the top panel, the aggregate price-level rises only halfway at the date of the money supply shock, with the remaining price adjustment staggered across several subsequent periods. This inflation episode continues until those households who were active at the shock date have traveled through all time-since-active groups and are once again active, at the start of date 6.

The aggregate price-level adjusts gradually in this exogenously segmented markets economy for precisely the reasons explained by Alvarez, Atkeson and Edmond (2009). Open market operations that inject money into the brokerage accounts must be absorbed by active households.¹⁷ However, as they will be unable to access their brokerage accounts again for 5 periods, these households retain large inventories of money relative to their current consumption spending. As noted above, their spending rate is the lowest among all households in the economy. Consequently, total nominal spending does not rise in proportion to the money supply, and a rise in the share of money held by active households leads to a rise in aggregate real balances. Equivalently, in this endowment model, velocity falls.

Formally, in a fixed duration model with J time-since-active groups, aggregate velocity may be expressed as the sum of two terms, one associated with the common velocity of currently active households and one associated with the velocities of inactive households across their respective groups:

$$V_t = \frac{1}{J} \frac{M_{0t}}{\bar{M}_t} v_{0t} + \sum_{j=1}^{J-1} \frac{1}{J} \frac{M_{jt}}{\bar{M}_t} v_{jt}. \quad (21)$$

From this equation, it is clear that the rise in relative money holdings of active households must reduce aggregate velocity, so long as individual velocities do not rise much in response to the shock. As seen in the bottom panel of figure 2, in our fixed duration example, half of the money injection is absorbed by an initial fall in aggregate velocity. As households that were active at the time of the shock travel through time-since-active groups in subsequent periods, their spending rate rises, pulling aggregate velocity back up. During this episode nominal spending rises faster than the money supply, the price level grows above trend, and aggregate real balances return to their long-run level.

Turning to the response in interest rates shown in the middle panel of figure 2, note that the money injection causes a large, but purely transitory, liquidity effect. In economies with segmented markets, real rates are determined by the marginal utilities of consumption among active households in adjacent periods, given that only these households can transform interest-bearing assets into consumption. In the fixed duration model, only those households that are active at the date of

¹⁶Figures in this and the subsequent section reflect the effects of a temporary 0.1 percentage point rise in the money growth rate. Given that our model is solved linearly, we have re-scaled all responses to correspond to a 1 percentage point rise for readability.

¹⁷No household unable to shift assets from its brokerage account into its bank account will accept the additional money, given the rate-of-return dominance implied by positive nominal interest rates.

the shock experience a rise in their lifetime wealth. As a result, their consumption rises, while the consumption of households active in subsequent dates remains unchanged, thus explaining the large but temporary fall in the real interest rate.

Endogenous segmentation model. Figure 3 displays the aggregate response to the same temporary shock in the baseline case of our model. The endogenous segmentation economy exhibits somewhat sharper initial price adjustment, associated with a smaller fall in aggregate velocity, and it has a more protracted response in inflation. Although the average time between a household's account transfers is 4.8 periods in our model's steady-state, its high inflation episode following the purely transitory money shock lasts 8 periods. The initial decline in interest rates is substantially smaller than were those in figure 2, at about one-tenth of the size of the money growth shock. However, in contrast to the immediate correction seen under fixed duration, the real interest rate here remains persistently low for 6 quarters. These differences in amplitude and propagation arise from the two elements distinguishing our model, the nontrivial rising hazard reflecting the fractions of households undertaking bank transfers from each time-since-active group, and the movement in this hazard in response to an aggregate shock. The first of these elements is central to our model's larger initial rise in inflation, while the second is entirely responsible for its substantially different real interest rate response.

Similar to (21) above, aggregate velocity in our model is determined by a weighted sum of the individual velocities of active and inactive households, with weights determined by the measures of households in each time-since-active group and their relative individual money holdings:

$$V_t = \left(\sum_{j=1}^J \alpha_{j,t} \theta_{j,t} \right) \frac{M_{0,t}}{\bar{M}_t} v_{0,t} + \sum_{j=1}^{J-1} (1 - \alpha_{j,t}) \theta_{j,t} \frac{M_{j,t}}{\bar{M}_t} v_{j,t} + \frac{P_t}{\bar{M}_t} \sum_{j=1}^J \theta_{j,t} \int_0^{H^{-1}(\alpha_j)} x h(x) dx. \quad (22)$$

The final term reflects the proportion of the aggregate money stock used in paying transfer costs, and was absent in (21). However, as this term is quantitatively unimportant both on average and following the shock, it cannot explain our economy's lesser decline in aggregate velocity relative to the fixed duration model. The first-order difference lies in the second term, the weighted velocities of inactive households.

In the fixed duration model, every household spends all of its money between any one balance transfer and the next, because this timing is certain. By contrast, the average household in our economy typically has some left-over money in its bank account when it undertakes its next transfer, because this timing is uncertain. Given their ability to alter this expected left-over money, our inactive households are more flexible in responding to the money growth shock.¹⁸ In response to the rise in anticipated inflation, their spending rates, v_{jt} , rise between 0.3 and 0.5 percent with the money injection, roughly twice as much as in the fixed duration model. As a result, our economy experiences a lesser decline in the second (and largest) term determining aggregate velocity at the date of the shock due to its nontrivial hazard. This is mitigated to some extent by changes in the hazard, as discussed below.

¹⁸Consumption choices among each inactive group of households, $j = 1, \dots, J - 2$, satisfy:

$$u'(c_{jt}) = \beta E_t \left[\frac{P_t}{P_{t+1}} \left((1 - \alpha_{j+1,t+1}) u'(c_{j+1,t+1}) + \alpha_{j+1,t+1} u'(c_{0,t+1}) \right) \right].$$

As active households have higher consumption than inactive households (and $c_{0,t+1}$ rises with the money injection in our economy), the positive and increased probability of becoming active in the next period compounds the effect of anticipated continued high inflation in discouraging money savings.

Because the money injection implies an inflationary episode that will reduce inactive households' real balances, it increases the value of actively converting bonds held in the brokerage account into money. Thus, a greater than usual measure of households become active. However, this rise in the number of active households has only limited impact in reducing aggregate velocity, since it implies that in equilibrium each active household receives a lesser share of the total money injection. As a result, the weight $\frac{M_{0,t}}{M_t}$ is smaller in (22) than it is in (21), which in turn implies a lesser initial rise in the consumption of active households in our economy. The smaller rise in each active household's money holdings also implies that their velocity falls by less than in the fixed duration model (0.4 versus 1.6 percent).

While endogenous market segmentation reduces the initial real effect of a monetary shock, it also propagates it through changes in the timing of households' transfer activities, which are summarized in panel A of figure 4. Following a substantial initial rise, the overall measure of active households falls below its steady-state value for a number of periods, despite persistently high activity rates across groups, α_{jt} . This is because large initial rises in these rates shift the household distribution to imply higher than usual money balances for the mean household in subsequent dates, thereby reducing its incentive to transfer funds from the brokerage account. Alternatively, households' initial rises in activity rates thereafter imply unusually high membership in low-numbered time-since-active groups where, given the rising hazard, activity rates are lowest.

Those persistent changes in the distribution of households are responsible for the persistent real effects in our economy. In dates following the shock, although money growth has returned to normal, the measure of active households is sufficiently below average that each such household receives an above-average transfer of real balances in equilibrium. Thus, the rise in the consumption of active households in our economy is not purely transitory as it was in the fixed duration model. Rather, as seen in panel B of figure 4, it returns to steady state gradually as the distribution resettles. This explains why the real interest rate has a much smaller initial decline in our economy, and why it remains persistently low.

Figure 5 verifies the importance of changes in our economy's endogenous timing of household transfer activities by displaying the aggregate response in an otherwise identical model where such changes are not permitted. In this *time-dependent* activity model, a nontrivially rising hazard governs the timing of household account transfers; in fact, it is precisely that from our economy's steady-state in table 1. Here, however, this hazard is held fixed throughout time. From the comparisons in panels A and C, it is clear that changes in group-specific activity rates serve to reduce aggregate velocity in our model economy, yielding more gradual price adjustment, as was argued above. Next, the time-dependent model's interest rate responses in panel B confirm that our economy's persistent liquidity effects in real rates arise entirely from changes in the hazard, rather than its average shape. Absent these changes, the interest rate decline is completely transitory just as in the fixed duration model.

3.3 Money injection: high transfer costs examples

In the preceding example, the fixed transfer costs causing our economy's market segmentation were selected to yield average aggregate velocity at 1.9, and implied a 4.8 quarter average duration of inactivity among households. Here, we examine our model's response to the same temporary money growth shock in an example with high transfer costs implying aggregate velocity matching the U.S. postwar average, at 1.5, and a mean inactivity duration of 7 quarters. In this case, the maximum inactivity spell facing a household is substantially longer, at 10 quarters, and only

about 14 percent of households are active in the average period (versus 21 percent above). Thus, households are spread across far more groups and, as a result, carry larger inventories of money on average.

Figure 6 is the high transfer cost counterpart to figure 3. Here, in contrast to the previous example, the aggregate price level actually rises by more than the money growth shock at its impact, given a rise in aggregate velocity, and the inflationary episode is entirely temporary. Moreover, the persistent decline in the real interest rate of figure 3 has also evaporated. These dramatic changes in the model's response may be traced to two features of the mechanics discussed above that become more pronounced when households face the possibility of very extended absence from their brokerage accounts: the rise in activity rates at the date of the shock and, more importantly, the rise in individual spending rates among inactive households.

With the money injection comes a permanent upward shift in the path of the aggregate price level. This is far more costly for inactive households in this example relative to the previous one, because these households are, at impact, holding much higher inventories of money in preparation for longer horizons of potential inactivity. This leads to a percent rise in activity rates (4.2 percent) similar to that in our previous example with much smaller transactions costs. However, high transactions cost draws keep most households inactive. In an effort to offset the fall in their real consumption spending, these households increase their spending rates.¹⁹ As a result, the percentage rises in v_{jt} over time-since-active groups 1 through 5 are roughly double those in our previous endogenous segmentation example, and these rises are also large in each of the higher-numbered groups new to this example, at around 0.75 percent. Moreover, even with increased activity rates, inactive households make up roughly 75 percent of all households at the date of the shock. Thus, as these households release substantially more money into the goods market than usual, total nominal spending actually increases by more than the money supply.

There is virtually no real interest rate response at all in this example. The high current price level, and the large change in activity rates (which have a lower steady state value in this example), are sufficient to imply that each active household receives no greater percentage rise in real money holdings than did its counterpart in our previous example. However, active households here have significantly lower spending rates on average (given potentially very long absences from the brokerage accounts). As a result, their consumption rises by less than one-third the rise seen in figure 4, remaining very close to that of households active in subsequent dates.

We have referred to our second example as one with an increased average duration of inactivity. However, what distinguishes this example is the substantial difference between the mean duration of inactivity (7 quarters) versus the maximum (10 quarters). This leads to additional precautionary accumulation of money and, on average, households return to the bond market with sizeable left-over money balances. Thus, inactive households at the date of the money injection have substantial flexibility in raising their current spending rates by reducing their expected future left-over balances. This allows the sharp initial rises in individual velocities central in the results above. Alternative examples where the maximum length of inactivity is similarly high, but the mean duration is close to it, more closely resemble the baseline example in the section above. One such example follows.

To obtain a high maximum inactivity duration example where the mean duration is similarly high, we abandon our assumption that transfer costs are distributed uniformly. In this third case, we

¹⁹As noted above, probabilistic timing of future activity is important, because it implies that, on average, most households return to their brokerage accounts with some money remaining. The money with which a typical household ordinarily returns to the brokerage account is higher in this example relative to the previous one, which allows greater rises in spending rates among inactive households.

assume a beta distribution parameterized by $[\alpha = 3, \beta = 1/3]$ and set the maximum transfer cost at $B = 0.50$. This results in an average aggregate velocity again at 1.5, a maximum inactivity duration of 10 quarters and a mean duration of 9.55 quarters. With this change in the cost distribution, our model's steady-state hazard describing average activity rates looks much like that of a fixed duration model, in that activity rates are near zero for all groups below J . As such, one might imagine that its dynamic response would resemble a $J = 10$ version of figure 2. However, because our households are able to change the timing of their transfers, this is not the case. In fact, figure 7 reveals that the response to the temporary money shock is instead quite similar to that in our baseline endogenous segmentation example. Again, adjustment in the aggregate price level is slow, resulting in a persistently high inflation episode and, unlike a fixed duration model, the real interest rate is persistently low. The one new feature here relative to both the fixed duration model and our baseline example is a persistent liquidity effect in nominal interest rates. From this, it is clear that the distribution of the costs responsible for market segmentation can have important effects on aggregate dynamics.

Recall that this third example also improves upon those above in its consistency with the microeconomic evidence regarding the frequency of active trades. Here, the model predicts that, on average, the fraction of households undertaking active trades within one year is 0.42. Unlike either of the preceding examples, this prediction lies inside the range estimated by Vissing-Jørgensen (2002), 0.29 – 0.53. Thus, we pursue this case of high mean and maximum inactivity duration as we examine our model's results throughout the remainder of the paper.

4 Results

The examples we have considered thus far are useful in illustrating the mechanics of our model, at least qualitatively. However, the analysis of a purely random increase in the money supply is not what most would view as reflective of inflation and interest rate dynamics in an actual economy. In this section, we present results for our model with beta-distributed transactions costs under more plausible assumptions about monetary policy.

First, we examine a persistent shock to the money growth rate. Next, we consider an environment where the monetary authority implements changes in the money supply towards stabilizing inflation in the face of persistent endowment shocks. Finally, in section 4.3, we consider a version of our model with production wherein the monetary authority follows a Taylor rule responsive to both deviations in inflation and output. As we examine the resulting dynamics in each of these cases, we will draw upon our analyses of the examples above for explanations.

4.1 Persistent money growth shock

We begin by considering the aggregate response to a persistent rise in the money growth rate, now assuming that the money growth rate follows mean-zero AR-1 process in logs with autocorrelation 0.57, as in Chari, Kehoe, and McGrattan (2002).²⁰ To see how endogenous changes in the extent of market segmentation influence this response, we contrast our state-dependent (SD) endogenous segmentation economy to the corresponding time-dependent (TD) exogenous

²⁰The persistence of the monetary measure used to calibrate our model is actually substantially higher, at 0.93 over the sample period 1954:1 to 2003:1. Since our results are not qualitatively changed, we use the Chari, Kehoe and McGrattan M1-based value for comparability.

segmentation model where activity rates are permanently fixed at the steady-state adjustment hazard implied by our model, so that such changes are not permitted.²¹

Figure 8 shows log deviations from the initial trend for the money stock and for the price levels of our model and the time-dependent model. The impact of the shock on the money supply is largely finished by period 7, while the price levels are more sluggish in their adjustment, with above-average inflation continuing for 3 or 4 additional periods. It is worth re-emphasizing that our model’s sluggish aggregate price adjustment is obtained without recourse to any imposed stickiness in micro-level portfolio adjustment. Indeed, the comparison with the exogenous segmentation (TD) model in figure 8 reveals that endogenous variation in the extent of segmentation actually slows down adjustment in the aggregate price level. Moreover, while both models overshoot the new price trend, this is less pronounced in our model.

In the time-dependent adjustment model, large wealth effects for households active in the early periods of the shock lead to sharper increases in prices. By eroding the real balances of inactive households, inflation redistributes consumption to active households. By contrast, in our model, rises in the numbers of active households reduce the increase in their individual money holdings, and thus the extent to which, in equilibrium, consumption must be redistributed. Compared to the transitory shock studied in figure 3, since the rise in the money growth rate is now persistent, some households delay their early return to the brokerage account by a period or two in hopes of lower transfer costs. As a result, the rise in the number of active households is initially smaller, and it persists for several periods, thereby protracting the distributional effects of the shock. Relative to the time-dependent model, a persistently smaller redistribution of real balances from inactive to active households in our model explains its lower rates of inflation in periods after the shock. Moreover, because the early periods with above-average numbers of active households are followed by 6 periods in which this number falls below steady-state, the episode with high real balances per active household is extended (as discussed above in section 3.2). This implies greater persistence in the increased consumption of active households in our model, and thus yields a persistent liquidity effect in both real and nominal interest rates (not shown).

4.2 Endowment shock under a Taylor rule

Here, we allow for shocks to the real endowment received by households and posit a monetary authority acting to stabilize inflation. More specifically, we assume that the common household endowment follows a persistent lognormal process,

$$\log(y_t) = \rho \log(y_{t-1}) + \varepsilon_t, \varepsilon \sim N(0, \sigma_\varepsilon^2),$$

with $\rho = 0.90$ and $\sigma_\varepsilon = 0.007$, and that the monetary authority follows a Taylor rule responding only to deviations in inflation:

$$i_t = i^* + 1.5[\pi_t - \pi^*].$$

Figure 9 shows our economy’s aggregate response to a persistent rise in the endowment, alongside the corresponding response in the time-dependent model where the hazard dictating households’ probabilities of becoming active remains fixed over time. Taken on its own, absent any response in

²¹The response in the corresponding fixed-duration model (with deterministic portfolio adjustment timing) is similar to that of the time-dependent model shown here. Recall from our discussion of figure 7 that the beta distribution from which transfer costs are drawn in our model leads to a steady-state hazard resembling the hazard of a fixed duration model.

the money growth rate, the rise in endowments would imply a fall in the inflation rate to increase real balances. As such, in both model economies, the Taylor rule dictates a fall in the nominal interest rate. Indeed, given the active policy rule we have assumed, where the nominal interest rate responds by more than inflation, the real interest rate must fall.

In the time-dependent model there is a sharp, unanticipated fall in prices at the initial date of the shock. Interest rates fall, and an increase in the real balances of active households finances subsequent purchases of the increase in output. Households active after the shock do not experience a rise in their consumption and do not require anything beyond the usual transfer of real balances from their brokerage accounts. Thus, inflation returns to its average value in the second period, as do interest rates.

By contrast, in our economy with endogenous market segmentation, there are persistent responses in inflation and interest rates, with a half-life of roughly 3 quarters. The fall in nominal interest rates gives households an incentive to hold more money. The resulting rise in activity rates implies that the increase in the money supply, relative to trend, is spread over more households than usual, thus lowering the rise in each individual withdrawal. Relative to the exogenous segmentation model, this reduces the rise in consumption among households active in date 1 of the shock. Thereafter, with fewer households remaining in other groups at this initial date, subsequent active populations are reduced, thereby raising their individual withdrawals. As a result, consumption among active households rises less sharply upon the shock's impact and is more evenly spread across subsequent active groups. Thus, the real interest rate initially falls by less than it would otherwise, and thereafter gradually returns to steady state. This generates persistence in the nominal interest rate that translates, through the Taylor rule, into persistently low inflation.

We stress that our economy's persistent response to the endowment shock cannot be attributed to the Taylor rule itself, noting again that there is no such persistence in the exogenous segmentation counterpart. Rather, it arises entirely from households' ability to change the timing of their portfolio adjustments in response to the economy's aggregate state. Viewed alternatively, it is a consequence of endogenous changes in the extent of market segmentation.

4.3 Real shock in an economy with production

We next consider the joint mechanics of market segmentation, output and employment in a version of our model with production. Households now value both consumption and leisure, and their utility function is given by $u(c, n) = \frac{(c - \eta n^\gamma)^{1 - \sigma}}{1 - \sigma}$, where c is consumption, n denotes hours of work, and we restrict $\gamma > 1$ and $\sigma > 0$. Because differences in money holdings imply differences in consumption, our introduction of variable labor supply in this section implies differing hours worked for households in different time-since-active groups. Let $n_{j,t-1}$ describe the hours of work in period $t - 1$ by a household that, at that time, had last been active j periods in the past, $j = 0, \dots, J - 1$.

We assume a perfectly competitive representative firm that hires labor from households at common real wage w_t and produces the single consumption good. The firm is owned equally by all households, there is no trade in shares, and all profits are returned to owners. Denoting aggregate employment by N_t , the firm's production is given by

$$Y_t = z_t N_t^\rho, \tag{23}$$

and profits are $\Pi_t = Y_t - w_t N_t$. The logarithm of total factor productivity, z_t , follows an autoregressive process with persistence parameter, ρ .

The timing of events within a period is as follows. At the beginning of each period, households observe the current aggregate state, draw their transfer costs, determine their portfolio adjustments, and then participate in the labor and goods markets. Fraction λ_N of labor income and λ_Π of profits are paid into the bank account. As before, income earned this period is available at the start of the next. Thus, there remains a cash-for-goods constraint on consumption purchases. While we use w_{t-1} to describe the real wage in period $t-1$, and Π_{t-1} to represent real profits, all income is paid in nominal units. Thus, at the beginning of period t , the real payment into the bank account of each household of time-since-active type j , $j = 1, \dots, J$, is $[\lambda_N(w_{t-1}n_{j-1,t-1}) + \lambda_\Pi\Pi_{t-1}]\frac{P_{t-1}}{P_t}$, and $[(1 - \lambda_N)(w_{t-1}n_{j-1,t-1}) + (1 - \lambda_\Pi)\Pi_{t-1}]\frac{P_{t-1}}{P_t}$ is deposited into the brokerage account.

We represent the current type- j household's real wage earnings from the previous period as e_{jt} in the problem that follows. The other wealth specific to a household is the money retained in its bank account at the end of the previous period. Let a_{jt} denote the real value of the money saved by a household currently of type j , as of the end of last period, which implies real balances of $\frac{P_{t-1}}{P_t}a_{jt}$ for the current period. A type- j household then has real balances m_{jt} in its bank account at the start of this period, where:

$$m_{jt} = \frac{P_{t-1}}{P_t}[a_{jt} + \lambda_N e_{jt} + \lambda_\Pi \Pi_{t-1}], \text{ for } j = 1, \dots, J. \quad (24)$$

As before, we retrieve competitive allocations by solving the recursive problem of an extended family. Relative to the family objective function (13), the only change is that $u(c, n)$ replaces the original period utility function $u(c)$. Noting that we now use $a_{j,t}$ to define the support of the distribution of real balances, we briefly describe the constraints involved here. Most are simply real-valued analogues of equations 14 - 19 from section 2.3 with endowment income replaced by labor and profit income. Beyond these, we have the aggregate production function in (23), alongside the new constraints associated with the introduction of variable labor supply. The first set identifies the labor income a household obtains from the current period:

$$w_t n_{jt} \geq e_{j+1,t+1} \text{ for } j = 0, \dots, J-1. \quad (25)$$

For convenience, we define a new variable χ to summarize total income entering the brokerage accounts each period. Specifically, let χ_{t+1} represent the end of period t value of real labor income and profits to be deposited into the family brokerage account at the start of period $t+1$;

$$(1 - \lambda_\Pi)\Pi_t + (1 - \lambda_N) \left[\sum_{j=1}^J \theta_{jt} \alpha_{jt} e_{1,t+1} + \sum_{j=1}^{J-1} \theta_{jt} (1 - \alpha_{jt}) e_{j+1,t+1} \right] \geq \chi_{t+1}. \quad (26)$$

The brokerage account constraint (14) then becomes

$$\frac{P_{t-1}}{P_t} \chi_t + \frac{P_{t-1}}{P_t} \mu_t \bar{a}_{t-1} \geq \sum_{j=1}^J \theta_{jt} \alpha_{jt} (m_{0t} - m_{jt}) + \sum_{j=1}^J \theta_{jt} \varphi(\alpha_{jt}), \quad (27)$$

where $\varphi(\alpha_{jt}) \equiv \int_0^{H^{-1}(\alpha_j)} \xi h(\xi) d\xi$ as before, m_{0t} is the shopping-time real balances allocated to each currently active household, and \bar{a}_{t-1} represents economywide real balances as of the end of date $t-1$.

Turning to the constraints involving households' bank accounts, we replace equations (15) and (16) with their real analogues in (24) and (28).

$$m_{j,t} \geq c_{jt} + a_{j+1,t+1} \text{ for } j = 0, \dots, J \quad (28)$$

For each shopping-time group $j = 0, \dots, J$, the cash-for-goods requirement is imposed by $a_{j+1,t+1} \geq 0$, the counterpart to (17).

The laws of motion for the distribution of households over time-since-active groups, (18) and (19), continue to apply, and (24) and (28) determine the associated group-specific real balances as a function of lagged real balances and labor incomes, $\{a_{jt}, e_{jt}\}_{j=1}^J$, lagged profits, Π_{t-1} , and inflation, $\pi_t = \frac{P_t}{P_{t-1}} - 1$. As before, the date- t distribution of real balances is fully described by $\{\theta_{jt}, m_{jt}\}_{j=1}^J$.

Given the one-period lag in the deposit of labor and profit income, the family's state vector is now $(\{\theta_{jt}, a_{jt}, e_{jt}\}_{j=1}^J, \Pi_{t-1}, \chi_t)$. The aggregate resource constraint in (20) still applies. Equilibrium in the labor market is imposed by:

$$N_t = \left(\sum_{j=1}^J \alpha_{jt} \theta_{jt} \right) n_{0t} + \sum_{j=1}^{J-1} (1 - \alpha_{jt}) \theta_{jt} n_{jt}. \quad (29)$$

Finally, we assume that monetary policy follows an active Taylor rule of the form:

$$i_t = i^* + 1.5[\pi_t - \pi^*] + 0.35[y_t - y^*]. \quad (30)$$

The production economy is parameterized as follows. We set the labor income share to $\nu = 0.6$ and the persistence of total factor productivity at $\rho = 0.9$. Next, in specifying period utility, we set $\sigma = 2$ and $\gamma = 2.5$; this implies a relatively low wage-elasticity of labor supply in comparison with the values often used in macroeconomic research. We set $\eta = 2.7$, so that aggregate hours worked are on average approximately $\frac{1}{3}$ of time. We maintain all other parameters from the endowment model studied above in this section, continuing to assume beta-distributed transactions costs. Finally, we set the steady-state level of total factor productivity at 0.33 so as to match our previous measure of aggregate velocity from the endowment economy without changing parameters of the transaction cost distribution.

Figures 10 and 11 illustrate our economy's response following a persistent shock to total factor productivity. As output increases with the productivity shock, the demand for real balances rises. Given the money supply, this reduces the price level relative to its trend. The monetary authority responds according to the Taylor rule, lowering the nominal interest rate. As the return on bonds falls, the fraction of households actively trading bonds for money rises.

As we have now seen in several instances above, the initial rise in the number of traders has dynamic implications. In periods following the shock, the number of active households falls. This, in turn, implies that active households in those later periods also carry real balances exceeding the level typically held by an active household. In response, they increase their consumption spending, which drives further increases in output and employment. As a result, we observe non-monotone responses in these series, particularly in employment.

Note that the humped-shaped quantity responses in figure 10 arise entirely from our households' ability to choose when they trade bonds for money, and to change the timing of these trades in response to changes in income, prices, wages, and interest rates. Quantity responses are monotone when we suppress such choices, as can be seen from the dashed lines depicting the responses from the corresponding time-dependent model where activity rates are fixed. In the time-dependent model, the full benefit of the productivity shock falls to those households that are active at the date

of the shock. A large, one-period drop in the nominal interest rate (in the middle panel of figure 11) induces these households to take on raised real balances sufficient to fully absorb the extra output arising from the shock among themselves. This, in turn, delivers a sharp, but transitory drop in inflation. In our model, by contrast, endogenous changes in activity rates spread the benefits of the shock across consecutive groups of active households. The rise in real balances per active household lasts beyond the first period, as the monetary authority responds to the persistent fall in inflation by keeping the nominal rate low. By reducing the rise in inequality, our model transforms the large, transitory movements in inflation and interest rates from the time-dependent model into more moderate and persistent declines.

4.4 Interest rates and aggregate consumption growth

In our introduction, we noted that sticky price models, and indeed all models assuming a representative household, predict a tight link between real interest rates and the growth rate of aggregate consumption, through the household consumption Euler equation. It is well understood that Euler-implied real interest rates obtained using aggregate consumption data are a poor proxy for actual market rates. However, Canzoneri, Cumby and Diba (2007) show that the discrepancy extends beyond a problem of mean spreads, or even volatilities. Using U.S. consumption data to retrieve the real rate series predicted by the Euler equation under a series of preference specifications both without and with habit persistence of varying forms, they show that the Euler-implied interest rate series is consistently negatively correlated with observed real rates.

In this final subsection of results, we show that our endogenous segmentation model with production is consistent with the empirical findings of Canzoneri et al. Specifically, we establish that, as in the U.S. data, our model's actual real interest rate series is negatively correlated with the series that would be imputed from its aggregate consumption growth rates using the assumption of a representative household. For simplicity, we implement their baseline exercise with constant relative risk aversion preferences, which is as follows. Let r^E represent the consumption Euler-based real interest rate, let C represent aggregate consumption, and let r^M be the market real rate. Given CRRA utility, the representative household's Euler equation is

$$C_t^{-\sigma} = \frac{1}{1 + \rho} E_t[(1 + r_{t+1}^e) C_{t+1}^{-\sigma}],$$

where ρ is the household's subjective discount rate. Taking logarithms, and imposing certainty equivalence, we infer an approximate ex-post real interest series via

$$r_{t+1}^E \simeq \sigma \gamma_{t+1}^C + \rho, \tag{31}$$

where γ^C is the gross growth rate of aggregate consumption between dates t and $t + 1$.²² From the equation above and quarterly U.S. consumption data between 1966 and 2004, Canzoneri et al. obtain the Euler-implied interest rate series $\{r_t^E\}$. Comparing this to the U.S. Federal Funds rate over the same period $\{r_t^M\}$, they find the correlation between the two series is -0.37 .

We conduct this same investigation using data from a 10,000 period simulation of our production economy model from section 4.3. Here, we include aggregate shocks to productivity,

$$\log(z_t) = 0.95 \log(z_{t-1}) + \varepsilon_t^z,$$

²²Canzoneri et al. (2007) report that comparing ex-ante rates obtained by adjusting nominal rates with one-quarter ahead inflation forecasts yields the same basic results as the ex-post rate comparison undertaken here.

and to the interest rate rule specified in equation 30, now appended by innovations ε_t^i . Each shock has mean-zero normal innovations; $\varepsilon_t^z \sim N(0, \sigma_{\varepsilon^z}^2)$ and $\varepsilon_t^i \sim N(0, \sigma_{\varepsilon^i}^2)$. To ensure that our model-based comparison is proximate to the empirical comparison described above, we select the volatilities of the two shocks so that our simulated aggregate consumption growth rates and market real interest rates have standard deviations closely matching those in the data.²³

Figure 12 presents a 75-period subsample of interest rate data drawn from our full-length simulation. Based on a quick inspection of the figure, it is easy to see that the Euler-derived interest rate implied by aggregate consumption growth in our model does not co-move positively with the actual real interest rate, which in our environment is determined by the marginal utilities of active (trading) households in adjacent dates. In fact, when we examine the results from the full 10,000 period simulation summarized in Table 2, we see that the model predicts a negative correlation between the two series surprisingly close to that observed in the U.S. data.

TABLE 2: Consumption growth and real interest rates

	$\sigma(\gamma^C)$	$\sigma(r^M)$	$\text{corr}(r^M, r^E)$
U.S. data	0.696	2.39*	-0.37*
simulated model data	0.680	2.21	-0.23

*From Table 1 of Canzoneri et al (2007).

As in any model with segmented asset markets, interest rates in our model are not determined by the consumption path of a representative household but, instead, by subsets of households actively trading in the market for government bonds. This implies a sharp separation between changes in aggregate consumption and the equilibrium real interest rate. As a result, we find that an empirically plausible version of the model reproduces the negative correlation between the market clearing real interest rate and the interest rate implied by aggregate consumption growth.

5 Concluding remarks

In the sections above, we have developed a monetary endowment economy where fixed transactions costs lead households to reallocate their wealth between money and interest-bearing securities infrequently. Given these costs, some fraction of households choose not to withdraw interest income in any given period; thus, on average, households carry inventories of money to finance their consumption spending over multiple periods. In this segmented asset markets environment, open market operations directly affect only a subset of all households. Moreover, the size of that subset varies over time as households respond to changes in economic conditions by varying the timing of their portfolio adjustments. As a result, changes in money growth rates are followed by gradual adjustment in the aggregate price level. Further, because they disrupt the distribution of real wealth across households without generating excessive inequality, such changes can have persistent effects on real interest rates.

Our approach to endogenizing market segmentation has emphasized idiosyncratic risk at the household level, through the assumption that transactions costs vary randomly across households. We have shown through a series of examples that the underlying distribution of these costs has

²³Market real rate volatility is matched to the standard deviation of the U.S. Federal Funds rate reported by Canzoneri et al. (2007). Aggregate consumption growth rate volatility is matched to the standard deviation of quarterly real PCE growth rates in the U.S. over 1954:Q1 through 2005:Q1. The resulting parameter values are $\sigma_{\varepsilon^z} = 0.004$ and $\sigma_{\varepsilon^i} = 0.0075$.

important implications for the resulting distributions of households over money holdings, and thus the real effects of monetary policy. Given this finding, we have restricted the distribution of transactions costs to ensure consistency with both macroeconomic data on aggregate velocity, and microeconomic data on the frequency with which households buy or sell securities.

In our economy, small changes in the number of households actively participating in open market operations lead to changes in the distribution of money holdings across households that differ markedly from those in monetary models where the degree of market segmentation is fixed. When shocks to money growth rates are transitory, these changes in distribution add persistence to inflation, and they transform sharp temporary movements in interest rates into more moderate and gradual responses. Under persistent money growth shocks, they imply more gradual price adjustment. When monetary policy is governed by an active Taylor rule, persistent real shocks cause persistent movements in inflation and interest rates only if households are allowed to respond with changes in the timing of their portfolio reallocations. Finally, in a version of the model with production, we observe hump-shaped responses in employment and output, alongside persistent declines in inflation and real and nominal interest rates following a persistent, monotone shock to technology. We establish that time-variation in the extent of market segmentation is essential to these aspects of our economy's response; all arise from changes in household activity rates prompted by the shock itself and the accompanying monetary policy response.

Arguably a benchmark criterion for any model attempting to explain how monetary policy affects the aggregate economy is that it should be consistent with the relation between short-term interest rates and real economic activity. We have shown above that the endogenously segmented markets model effortlessly delivers a liquidity effect. Moreover, it resolves a prevalent difficulty facing models of monetary economies, their prediction of a strong co-movement between aggregate consumption growth and short-term real interest rates. As the real interest rate is determined by subsets of households in a segmented markets environment, there is no close relation between it and aggregate consumption growth.

Our findings here suggest that monetary models with endogenous market segmentation may be useful laboratories in which to derive a more complete understanding of the relations between movements in real and nominal aggregate series observed in the data. An important next step in this line of research is a careful generalization of the model to include capital accumulation while maintaining the endogenous asset market segmentation. Such extension will signal the extent to which richer versions of the segmented markets framework model are likely to be useful in predicting the joint movement of real and nominal variables.

References

- [1] Alvarez, Fernando, Andrew Atkeson and Chris Edmond (2009) ‘Sluggish Responses of Prices and Inflation to Monetary Shocks in an Inventory Model of Money Demand’ *Quarterly Journal of Economics* **124**(3), pages 911-67.
- [2] Alvarez, Fernando, Andrew Atkeson, and Patrick Kehoe (2002) ‘Money, Interest Rates and Exchange Rates with Endogenously Segmented Markets’ *Journal of Political Economy* **110**(1), pages 73 - 112.
- [3] Alvarez, Fernando, Robert E. Lucas, Jr., and Warren Weber (2001) ‘Interest Rates and Inflation’ *American Economic Review* **91**(2), pages 219 - 225.
- [4] Bills, Mark and Peter J. Klenow (2004) ‘Some Evidence on the Importance of Sticky Prices,’ *Journal of Political Economy* **112**(5), pages 947-85.
- [5] Canzoneri, Matthew B., Robert E. Cumby and Behzad T. Diba (2007) ‘Euler equations and money market interest rates: A challenge for monetary policy models’ *Journal of Monetary Economics* **54**, pages 1863-81.
- [6] Chari, V.V., Patrick J. Kehoe and Ellen R. McGrattan (2000), ‘Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem?’ *Econometrica* **68**(5), pages 1151-1179.
- [7] Chatterjee, Satyajit and Dean Corbae (1992) ‘Endogenous Market Participation and the General Equilibrium Value of Money’ *Journal of Political Economy* **100**(5), pages 615 - 646.
- [8] Chiu, Jonathan (2007) ‘Endogenously Segmented Asset Market in an Inventory Theoretic Model of Money Demand’ Bank of Canada Working Paper 2007-46.
- [9] Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (1999) ‘Monetary Policy Shocks: What Have We Learned and to What End?’ chapter 14 in Michael Woodford and John B. Taylor (eds.) *Handbook of Macroeconomics* **IA**. Elsevier Science.
- [10] Dotsey, Michael, Robert G. King, and Alexander L. Wolman (1999) ‘State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output’ *Quarterly Journal of Economics* **114**, pages 655–690.
- [11] Friedman, Milton (1968) ‘The Role of Monetary Policy’ *The American Economic Review* **58**(1), pages 1-17.
- [12] Friedman, Milton (1980) *Memorandum to U.K. Treasury and Civil Service Committee regarding ‘Enquiry Into Monetary Policy,’* House of Commons, United Kingdom, 1980 July.
- [13] Grossman, Sanford, and Laurence Weiss (1983) ‘A Transactions-Based Model of the Monetary Mechanism’ *American Economic Review* **73**(5), pages 871-880.
- [14] King, Robert G. and Julia K. Thomas (2006) ‘Partial Adjustment without Apology’ *International Economic Review* **47**(3), pages 779-809.
- [15] King, Robert G. and Mark W. Watson (1996) ‘Money, Prices, Interest Rates and the Business Cycle’ *Review of Economics and Statistics* pages 35 - 53

- [16] Leeper, Eric M., Christopher A. Sims, and Tao Zha (1996) ‘What Does Monetary Policy Do?’ *Brookings Papers on Economic Activity* **1996**(2), pages 1 - 78.
- [17] Nakamura, Emi and Jón Steinsson (2008) ‘Five Facts About Prices: A Reevaluation of Menu Cost Models’ *Quarterly Journal of Economics* **123** (4), pages 1415-64.
- [18] Occhino, Filippo (2004) ‘Modeling the Response of Money and Interest Rates to Monetary Policy Shocks: A Segmented Markets Approach’ *Review of Economic Dynamics* **7**(1), pages 181-97.
- [19] Occhino, Filippo (2008) ‘Market Segmentation and the Response of the Real Interest Rate to Monetary Policy Shocks’ *Macroeconomic Dynamics*, **12** (5), pages 591-618.
- [20] Rotemberg, Julio J. (1984) ‘A Monetary Equilibrium Model with Transactions Costs’ *Journal of Political Economy* **92**(1), pages 40-58.
- [21] Thomas, Julia K. (2002) ‘Is Lumpy Investment Relevant for the Business Cycle?’ *Journal of Political Economy*, **110**(3), pages 508–534.
- [22] Uhlig, Harald (2005) ‘What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure’ *Journal of Monetary Economics* **52**(2), pages 381-419.
- [23] Vissing-Jørgensen, Annette (2002) ‘Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures’ NBER working paper 8884.
- [24] Williamson, Stephen D. (2008) ‘Monetary Policy and Distribution’ *Journal of Monetary Economics* **55**, pages 1038-53.

Appendix

A Risk sharing

As described in the main text, bonds issued by a competitive financial intermediary are contingent on both aggregate and idiosyncratic shocks. The supply of these bonds is created using purchases of government bonds and free entry into intermediation implies a zero-profit condition.

Our first lemma derives the equilibrium price of households' bonds. Given the aggregate history s^t , the intermediary's profit when next period's aggregate shock is s_{t+1} solves the problem in (8) - (9). There, recall that households identified by the history ξ^t purchase $B(s^t, s_{t+1}, \xi^t, \xi_{t+1})$ units of nominal bonds that pay one unit of currency next period if the aggregate shock is given by s_{t+1} and their idiosyncratic shock is ξ_{t+1} , and the measure of such households is $h(\xi^t)$.

Lemma 1 *The equilibrium price of state-contingent bonds issued by the financial intermediary, $q(s^t, s_{t+1}, \xi_{t+1})$, is given by $q(s^t, s_{t+1}, \xi_{t+1}) = q(s^t, s_{t+1}) h(\xi_{t+1})$.*

Proof. *Substitution of the right-hand side of (9) for $B(s^t, s_{t+1})$ in (8) gives*

$$\begin{aligned}
 & \int_{\xi^t} \int_{\xi_{t+1}} q(s^t, s_{t+1}, \xi_{t+1}) B(s^t, s_{t+1}, \xi^t, \xi_{t+1}) h(\xi^t) d\xi_{t+1} d\xi^t \\
 & - q(s^t, s_{t+1}) \int_{\xi^t} \int_{\xi_{t+1}} B(s^t, s_{t+1}, \xi^t, \xi_{t+1}) h(\xi^t) h(\xi_{t+1}) d\xi_{t+1} d\xi^t \\
 = & \int_{\xi^t} \int_{\xi_{t+1}} \left(q(s^t, s_{t+1}, \xi_{t+1}) - q(s^t, s_{t+1}) h(\xi_{t+1}) \right) h(\xi^t) B(s^t, s_{t+1}, \xi^t, \xi_{t+1}) d\xi_{t+1} d\xi^t \\
 = & 0.
 \end{aligned}$$

The last equality follows from the zero profit condition. ■

B Characterizing household behavior under risk-sharing

This section provides the main results we use in our numerical approach to solving the model. Lemma 2 proves that, given risk-sharing in the brokerage account, beginning of period money balances capture all relevant differences across households at the start of any period. Additional results in lemma 3 establish that whenever households access their brokerage account, there is uniformity of actions in that they choose the same consumption and bank account balance regardless of their past history of idiosyncratic shocks. The importance of this is that any differences across households will be limited by the number of periods since active transactions between the bank and brokerage account. Such differences disappear whenever households with different bank account balances access their brokerage account at the same time. Finally, lemma 4 proves that households' follow a threshold rule in determining the timing of their account transfers; specifically, it shows that households choose to become active whenever their current fixed transfer cost falls below some threshold value that is common to all households with the same beginning-of-period bank balance.

We begin by deriving a lifetime budget constraint associated with the brokerage account. Using lemma 1, the brokerage budget constraint for any household in the initial period is:

$$\bar{B} \geq \int_{s_1} \int_{\xi_1} q(s_1) h(\xi_1) B(s_1, \xi_1) d\xi_1 ds_1. \tag{32}$$

In period 1, the brokerage account constraint is:

$$B(s_1, \xi_1) \geq \int_{s_2} \int_{\xi_2} q(s_1, s_2) h(\xi_2) B(s_1, s_2, \xi_1, \xi_2) d\xi_2 ds_2 \\ - (1 - \lambda) P(s_0) y(s_0) + \left(x(s_1, \xi_1) + P(s_1) \xi_1 \right) z(s_1, \xi_1)$$

which, substituted into (32), gives:

$$\bar{B} \geq \int_{s_1} \int_{\xi_1} q(s_1) h(\xi_1) \left[\int_{s_2} \int_{\xi_2} q(s_1, s_2) h(\xi_2) B(s_1, s_2, \xi_1, \xi_2) d\xi_2 ds_2 - (1 - \lambda) P(s_0) y(s_0) \right. \\ \left. + \left(x(s_1, \xi_1) + P(s_1) \xi_1 \right) z(s_1, \xi_1) \right] d\xi_1 ds_1 \\ = \int_{s^2} \int_{\xi^2} q(s^2) h(\xi^2) B(s_1, s_2, \xi_1, \xi_2) d\xi^2 ds^2 - \int_{s_1} \int_{\xi_1} q(s_1) h(\xi_1) \left[(1 - \lambda) P(s_0) y(s_0) \right. \\ \left. - \left(x(s_1, \xi_1) + P(s_1) \xi_1 \right) z(s_1, \xi_1) \right] d\xi_1 ds_1,$$

where we use the notation $q(s^t) = q(s_1) \cdot q(s_1, s_2) \cdots q(s^{t-1}, s_t)$ and $h(\xi^t) = h(\xi_1) \cdots h(\xi_t)$. Repeated substitutions for $B(s^j, \xi^j)$ for $j = 2, \dots, t$ using (3) leads to the following equation.

$$\bar{B} \geq \int_{s^t} \int_{\xi^t} q(s^t) h(\xi^t) B(s^t, \xi^t) d\xi^t ds^t - \sum_{j=1}^{t-1} \int_{s^j} \int_{\xi^j} q(s^j) h(\xi^j) \left[(1 - \lambda) P(s^{j-1}) y(s^{j-1}) \right. \\ \left. - \left(x(s^j, \xi^j) + P(s^j) \xi_j \right) z(s^j, \xi^j) \right] d\xi^j ds^j.$$

Taking the limit of the above equation, given the No-Ponzi condition,

$$\lim_{t \rightarrow \infty} \int_{s^t} \int_{\xi^t} q(s^t) h(\xi^t) B(s^t, \xi^t) ds^t d\xi^t \geq 0,$$

we arrive at the following lifetime budget constraint associated with the brokerage account.

$$\bar{B} \geq \sum_{t=1}^{\infty} \int \int q(s^t) h(\xi^t) \left(z(s^t, \xi^t) \left[x(s^t, \xi^t) + P(s^t) \xi_t \right] - P(s^{t-1}) (1 - \lambda) y(s^{t-1}) \right) d\xi^t ds^t \quad (33)$$

Equation 33, which is an immediate implication of the availability of a complete set of state-contingent bonds for each household, implies that individual histories are irrelevant beyond current money balances in the bank account. This intuitive property of the model is straightforward, if notationally cumbersome, to prove and is important in our approach to characterizing competitive equilibrium.

Lemma 2 *Given $M(s^{t-1}, \xi^{t-1})$, the decisions $c(s^t, \xi^t)$, $A(s^t, \xi^t)$, $x(s^t, \xi^t)$ and $z(s^t, \xi^t)$ are independent of the history ξ^{t-1} .*

Proof. Equation 33 is equivalent to the following sequence of period by period constraints,

$$B(s^t) = \int_{s_{t+1}} q(s^t, s_{t+1}) B(s^t, s_{t+1}) ds_{t+1} - (1 - \lambda) P(s^{t-1}) y(s^{t-1}) \quad (34) \\ + \int_{\xi^t} h(\xi^t) \left[x(s^t, \xi^t) + P(s^t) \xi_t \right] z(s^t, \xi^t) d\xi^t,$$

and the limit condition

$$\lim_{t \rightarrow \infty} \int_{s^t} q(s^t) B(s^t) ds^t = 0. \quad (35)$$

Let $\xi_t^{t+j} = \{\xi_t, \dots, \xi_{t+j}\}$ be a partial history of individual household shocks, fix (s^t, ξ^t) and define pointwise the continuation value to the household's problem given $(B(s^t), M(s^{t-1}, \xi^{t-1}))$ as

$$V(B_0, M_0; s^t, \xi^t) = \max \sum_{j=0}^{\infty} \beta^j \int_{s_t^{t+j}} \int_{\xi_t^{t+j}} u(c[s^{t+j}, \xi_t^{t+j}]) g(s^{t+j} | s^t) h(\xi_t^{t+j}) ds_t^{t+j} d\xi_t^{t+j} \quad (36)$$

subject to:

$$\begin{aligned} B(s^{t+j}) &= \int_{s_{t+j+1}} q(s^{t+j}, s_{t+j+1}) B(s^{t+j}, s_{t+j+1}) ds_{t+j+1} - (1-\lambda) P(s^{t+j-1}) y(s^{t+j-1}) \\ &+ \int_{\xi_t^{t+j}} h(\xi_t^{t+j}) [x(s^{t+j}, \xi_t^{t+j}) + P(s^{t+j}) \xi_{t+j}] z(s^{t+j}, \xi_t^{t+j}) d\xi_t^{t+j}, \quad \text{for } j = 0, \dots, \end{aligned} \quad (37)$$

$$P(s^{t+j}) c(s^{t+j}, \xi_t^{t+j}) \leq M(s^{t+j-1}, \xi_{t-1}^{t+j-1}) + x(s^{t+j}, \xi_t^{t+j}) z(s^{t+j}, \xi_t^{t+j}) - A(s^{t+j}, \xi_t^{t+j}), \quad \text{for } j = 0, \dots, \quad (38)$$

$$M(s^{t+j}, \xi_t^{t+j}) \leq A(s^{t+j}, \xi_t^{t+j}) + \lambda P(s^{t+j}) y(s^{t+j}) \quad \text{and} \quad A(s^{t+j}, \xi_t^{t+j}) \geq 0, \quad \text{for } j = 0, \dots, \quad (39)$$

$$\text{given } B(s^t) = B_0, \quad M(s^{t-1}, \xi_{t-1}^{t-1}) = M_0 \quad \text{and} \quad \lim_{j \rightarrow \infty} \int_{s_t^{t+j}} q(s^{t+j}) B(s^{t+j}) ds_t^{t+j} = 0.$$

Clearly the optimal choices of $c(s^{t+k}, \xi_t^{t+k})$, $A(s^{t+k}, \xi_t^{t+k})$, $x(s^{t+k}, \xi_t^{t+k})$ and $z(s^{t+k}, \xi_t^{t+k})$, for any $k = 0, 1, \dots$, are independent of ξ^{t-1} given M_0 . Let $c(s^{t+k}, \xi_t^{t+k}; M_0)$, $A(s^{t+k}, \xi_t^{t+k}; M_0)$, $x(s^{t+k}, \xi_t^{t+k}; M_0)$ and $z(s^{t+k}, \xi_t^{t+k}; M_0)$ describe these choices. Next, any solution to the household's problem must also solve the following problem.

$$\begin{aligned} \max \quad & \sum_{j=1}^{t-1} \beta^{j-1} \int_{s^j} \int_{\xi^j} u(c(s^j, \xi^j)) g(s^j) h(\xi^j) ds^j d\xi^j \\ & + \beta^t \int_{s^t} \int_{\xi^t} V(B(s^t), M(s^{t-1}, \xi^{t-1}); s^t, \xi^t) g(s^t) h(\xi^t) ds^t d\xi^t \end{aligned}$$

subject to

$$\begin{aligned} B(s^j) &= \int_{s_{j+1}} q(s^j, s_{j+1}) B(s^j, s_{j+1}) ds_{j+1} - (1-\lambda) P(s^{j-1}) y(s^{j-1}) \\ &+ \int_{\xi^j} h(\xi^j) [x(s^j, \xi^j) + P(s^j) \xi_j] z(s^j, \xi^j) d\xi^j, \end{aligned}$$

$$P(s^j) c(s^j, \xi^j) \leq M(s^{j-1}, \xi^{j-1}) + x(s^j, \xi^j) z(s^j, \xi^j) - A(s^j, \xi^j),$$

$$M(s^j, \xi^j) \leq A(s^j, \xi^j) + \lambda P(s^j) y(s^j) \quad \text{and} \quad A(s^j, \xi^j) \geq 0, \quad \text{for } j = 1, \dots, t-1.$$

In this problem, the optimal choices of $c(s^{t+k}, \xi_t^{t+k})$, $A(s^{t+k}, \xi_t^{t+k})$, $x(s^{t+k}, \xi_t^{t+k})$ and $z(s^{t+k}, \xi_t^{t+k})$, for any $k = 0, 1, \dots$, are given by the functions $c(s^{t+k}, \xi_t^{t+k}; M(s^{t-1}, \xi^{t-1}))$,

$A(s^{t+k}, \xi_t^{t+k}; M(s^{t-1}, \xi^{t-1}))$, $x(s^{t+k}, \xi_t^{t+k}; M(s^{t-1}, \xi^{t-1}))$ and $z(s^{t+k}, \xi_t^{t+k}; M(s^{t-1}, \xi^{t-1}))$, which attained the maximum value $V(B(s^t), M(s^{t-1}, \xi^{t-1}); s^t, \xi^t)$. ■

We proceed to further characterize household behavior by studying the state-contingent plans chosen by households in period 0 when they are all identical. Let Λ denote the multiplier associated with equation 33 and $\nu_0(s^t, \xi^t)$ be the LaGrange multiplier for (4) and $\nu_1(s^t, \xi^t)$ the multiplier for (5). For clarity, the Lagrangian is shown below.

$$\begin{aligned} \mathcal{L} = & \sum_{t=1}^{\infty} \beta^{t-1} \int_{s^t} \int_{\xi^t} \left(u(c(s^t, \xi^t)) g(s^t) h(\xi^t) \right. \\ & + \nu_0(s^t, \xi^t) g(s^t) h(\xi^t) \left[M(s^{t-1}, \xi^{t-1}) + x(s^t, \xi^t) z(s^t, \xi^t) - A(s^t, \xi^t) - P(s^t) c(s^t, \xi^t) \right] \\ & + \nu_1(s^t, \xi^t) g(s^t) h(\xi^t) \left[A(s^t, \xi^t) + \lambda P(s^t) y(s^t) - M(s^t, \xi^t) \right] \Big) d\xi^t ds^t \\ & + \Lambda \left(\bar{B} + \sum_{t=1}^{\infty} \int_{s^t} q(s^t) \left[(1-\lambda) P(s^{t-1}) y(s^{t-1}) ds^t \right. \right. \\ & \left. \left. - \sum_{t=1}^{\infty} \int_{s^t} \int_{\xi^t} q(s^t) h(\xi^t) \left(x(s^t, \xi^t) + P(s^t) \xi_t \right) z(s^t, \xi^t) d\xi^t ds^t \right) \right) \end{aligned}$$

Given any choice of $z(s^t, \xi^t)$, the household's choices of $c(s^t, \xi^t)$, $A(s^t, \xi^t)$, $M(s^t, \xi^t)$ and $x(s^t, \xi^t)$, satisfy the following conditions.

$$Du(c(s^t, \xi^t)) - P(s^t) \nu_0(s^t, \xi^t) = 0 \quad (40)$$

$$\begin{aligned} -\nu_0(s^t, \xi^t) + \nu_1(s^t, \xi^t) &\leq 0, \\ &= 0 \text{ if } A(s^t, \xi^t) > 0, \end{aligned} \quad (41)$$

$$\begin{aligned} -\nu_1(s^t, \xi^t) + \beta \int_{s_{t+1}} \int_{\xi_{t+1}} \nu_0(s^t, s_{t+1}, \xi^t, \xi_{t+1}) h(\xi_{t+1}) g(s_{t+1} | s^t) ds_{t+1} d\xi_{t+1} &\leq 0 \\ &= 0 \text{ if } M(s^t, \xi^t) > \lambda P(s^t) y(s^t) \end{aligned} \quad (42)$$

$$\beta^{t-1} \nu_0(s^t, \xi^t) g(s^t) h(\xi^t) - \Lambda q(s^t) h(\xi^t) = 0 \text{ if } z(s^t, \xi^t) = 1. \quad (43)$$

The following lemma shows that households choose the same consumption and subsequent money balances for their bank account whenever they pay transfer costs to access their brokerage accounts. Recall that this result is important for the approach we take to characterizing household behavior in that it implies that heterogeneity across households is limited to periods when they are inactive.

Lemma 3 *For any (s^t, ξ^t) in which $z(s^t, \xi^t) = 1$, $c(s^t, \xi^t)$, $A(s^t, \xi^t)$ and $M(s^t, \xi^t)$ are independent of ξ^t .*

Proof. Given $z(s^t, \xi^t) = 1$, (40) and (43) imply

$$Du(c(s^t, \xi^t)) = \frac{\Lambda q(s^t) P(s^t)}{\beta^{t-1} g(s^t)}. \quad (44)$$

As Λ is the same for all households, this proves that $c(s^t, \xi^t)$ is independent of ξ^t . Next, using lemma 2, notice that strict monotonicity of $u(c(s^t, \xi^t))$, for any (s^t, ξ^t) , implies (38) and (39) will

always bind, hence $V(B(s^{t+1}), M(s^t, \xi^t); s^t, \xi^t)$ is strictly increasing in $M(s^t, \xi^t)$. Without loss of generality, we consider the case where $A(s^t, \xi^t) > 0$ and therefore $M(s^t, \xi^t) > \lambda P(s^t) y(s^t)$. In this case, (41) and (42) together give

$$\nu_0(s^t, \xi^t) = \beta \int_{s_{t+1}} \int_{\xi_{t+1}} \nu_0(s^t, s_{t+1}, \xi^t, \xi_{t+1}) h(\xi_{t+1}) g(s_{t+1} | s^t) ds_{t+1} d\xi_{t+1},$$

which, using (40) and (44) leads to the expression,

$$Du(c(s^t, \xi^t)) = \beta \int_{s_{t+1}} \int_{\xi_{t+1}} Du(c(s^{t+1}, \xi^{t+1})) \left[\frac{P(s^{t+1})}{P(s^t)} \right]^{-1} h(\xi_{t+1}) g(s_{t+1} | s^t) ds_{t+1} d\xi_{t+1}.$$

Since $c(s^t, \xi^t)$ is independent of ξ^t for all households with $z(s^t, \xi^t) = 1$, it follows that $c(s^{t+1}, \xi^{t+1})$ is not a function of ξ^t . Strict monotonicity of the value function, alongside strict concavity of the utility function, requires that $M(s^t, \xi^t)$ be independent of ξ^t , which in turn implies the same for $A(s^t, \xi^t)$. ■

Our final result shows that households follow threshold cost policies with respect to the costs of accessing their brokerage accounts

Lemma 4 *For any (s^t, ξ^{t-1}) , $Z = \{\xi_t | z(s^t, \xi^t) = 1\}$ is a convex set bounded below by 0.*

Proof. Assume not, let $\tilde{Z} = \{\xi_t | z(s^t, \xi^t) = 1\}$ and $m_{\tilde{Z}} = \int_{\tilde{Z}} h(\xi) d\xi$. Define $\bar{\xi}$ implicitly using the equation $\int_0^{\bar{\xi}} h(\xi) d\xi = m_{\tilde{Z}}$. By the Axiom of Choice there exists a 1-1 function $S: \tilde{Z} \rightarrow [0, \bar{\xi}]$. Given $B(s^t)$ and $M(s^{t-1}, \xi^{t-1})$, for each $\xi_0 \in \tilde{Z}$ and $\xi = S(\xi_0)$, construct an alternate continuation plan that assigns the original continuation plan associated with ξ_0 to ξ . In other words, $c^a(s^t, \xi^{t-1}, \xi) = c(s^t, \xi^{t-1}, \xi_0)$, $A^a(s^t, \xi^{t-1}, \xi) = A(s^t, \xi^{t-1}, \xi_0)$ and so on. This alternative plan satisfies (3) - (6) and is thus feasible. Moreover, by construction, it offers the same expected lifetime utility as the original plan. Any solution to the household's problem, solving (2) subject to (3) - (6), must satisfy (33) with equality. Since the original plan solved the household's problem by assumption, the alternative plan satisfies (33) as a strict inequality since

$$\begin{aligned} & \int_{\xi^{t-1}} q(s^t) h(\xi^{t-1}) \int_0^{\bar{\xi}} (x^a(s^t, \xi^t) + P(s^t) \xi_t) d\xi d\xi^{t-1} \\ & < \int_{\xi^t} q(s^t) h(\xi^t) (x(s^t, \xi^t) + P(s^t) \xi_t) z(s^t, \xi^t) d\xi^t ds^t. \end{aligned}$$

This contradicts optimality of the original plan. ■

FIGURE 1A: Endogenous Segmentation Steady-State (mean duration: 4.818)




FIGURE 1B: Fixed Duration Steady-State (mean duration: 5)




FIGURE 2: Temporary Money Growth Shock in 5-Period Fixed Duration Model






FIGURE 3A: Inflation Rate Responses to Money Growth Shock




FIGURE 3B: Interest Rate Responses in the Endogenous Segmentation Model




FIGURE 3C: Aggregate Velocity Responses to Money Growth Shock




FIGURE 4A: Endogenous Segmentation Response in Total Active Households




FIGURE 4B: Endogenous Segmentation Response in Consumption of Active Households

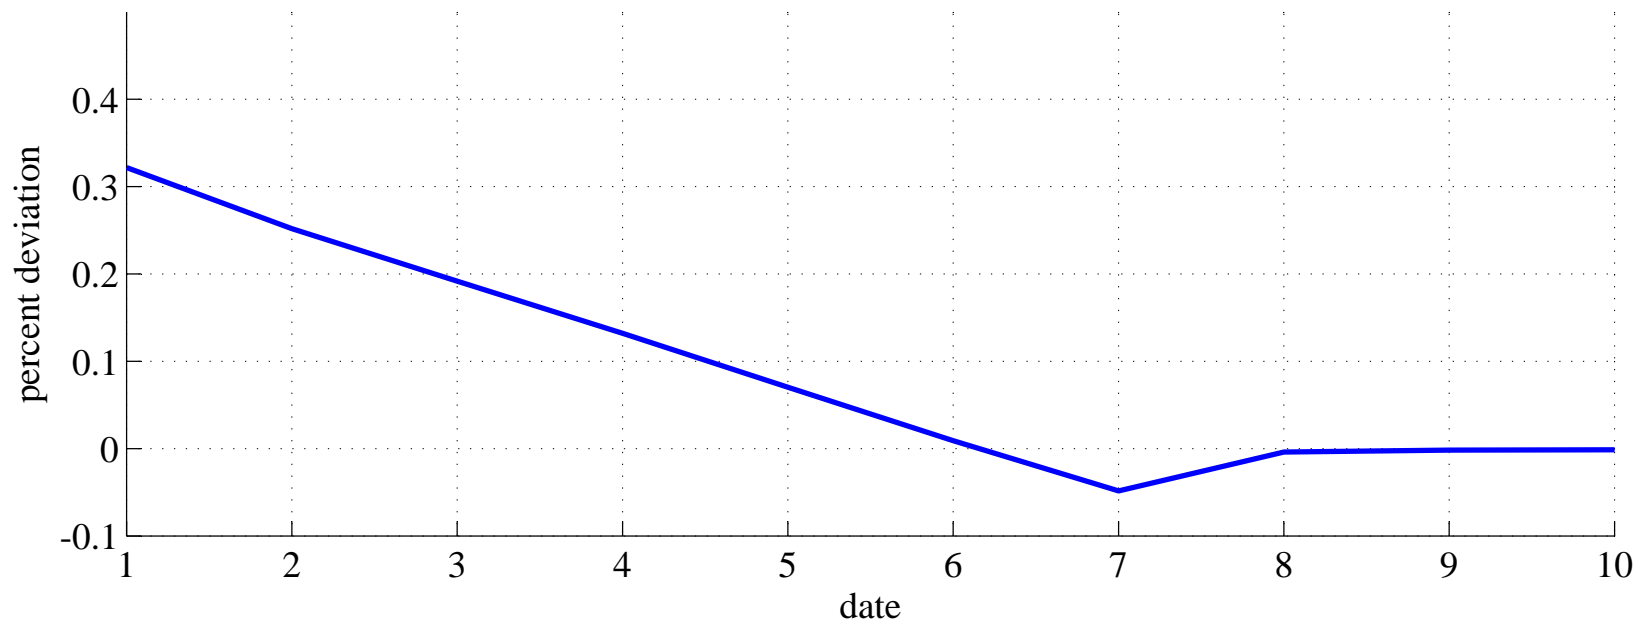


FIGURE 5A: Role of State-Dependent Activity Rates in Inflation Rate Response




FIGURE 5B: Interest Rate Responses with Changes in Activity Rates Suppressed




FIGURE 5C: Role of State-Dependent Activity Rates in Velocity Response




FIGURE 6: Temporary Money Growth Shock in High Transfer Costs Model




FIGURE 7: Temporary Money Growth Shock in High Mean Duration Model




FIGURE 8: Money and Prices with a Persistent Money Growth Shock




FIGURE 9A: Inflation and Interest Rates with a Persistent Endowment Shock and a Taylor Rule




FIGURE 9B: Interest Rates in Endogenous Segmentation Model




FIGURE 9C: Interest Rates in Time-Dependent Model




FIGURE 10: Employment, Production and Inflation under a Taylor Rule




FIGURE 11: Production Economy Interest Rates under a Taylor Rule




FIGURE 12: Market versus Euler-Implied Real Interest Rates in the Production Economy

