

Sticky Expectations and Consumption Dynamics

First version: June 2005; this version: March 6, 2017

Christopher Carroll¹ Jiri Slacalek² Kiichi Tokuoka³ Matthew N. White⁴

Abstract

A model in which consumers have accurate knowledge of their personal circumstances but ‘sticky expectations’ about the macroeconomy can reconcile conflicting micro and macro evidence about the nature of consumption dynamics. Sluggish aggregate spending growth, which has usually been interpreted as reflecting habits, arises here as a consequence of a modest degree of macroeconomic inattention, whose utility cost is calculated to be negligible. The implications of the model closely agree with a simple empirical exercise designed to reproduce the key facts about aggregate and household-level consumption.

Keywords Consumption, Sticky Expectations, Habits, Inattention

JEL codes E21, F41

PDF: <http://econ.jhu.edu/people/ccarroll/papers/cAndCwithStickyE.pdf>

Web: <http://econ.jhu.edu/people/ccarroll/papers/cAndCwithStickyE/>

Archive: <http://econ.jhu.edu/people/ccarroll/papers/cAndCwithStickyE.zip>
(*Contains data and estimation software producing paper's results*)

¹Carroll: Department of Economics, Johns Hopkins University, Baltimore, MD, <http://econ.jhu.edu/people/ccarroll/>, ccarroll@jhu.edu, Phone: (410) 516 7602 ²Slacalek: European Central Bank, Frankfurt am Main, Germany, <http://www.slacalek.com/>, jiri.slacalek@ecb.europa.eu, Phone: +49 69 1344 5047 ³Tokuoka: Ministry of Finance, 3-1-1 Kasumigaseki, Chiyoda-ku, Tokyo 100-8940, Japan, email: kiichi.tokuoka@mof.go.jp, Phone: +81-3-3581-4111 ⁴White: Department of Economics, University of Delaware, Newark, DE 19702, <http://www.lerner.udel.edu/faculty-staff/matthew-n-white/>, email: mwcon@udel.edu, phone: 302-831-1906

Thanks to seminar participants in the NBER Summer Institute, McMaster University, and the University of Michigan for constructive and insightful comments which substantially improved this paper. The views presented in this paper are those of the authors, and should not be attributed to the European Central Bank or to the Japanese Ministry of Finance.

1 Introduction

Starting with Campbell and Deaton (1989), the macroeconomics, finance, and international economics literatures have documented a wide variety of stylized facts suggesting that aggregate consumption is too smooth to be explained using the benchmark Hall (1978) random walk model of consumption. To address this problem, many recent papers have extended the Hall model by incorporating consumption ‘habits’ in the representative agent’s utility function. (See, e.g., results and references in Fuhrer (2000) or Christiano, Eichenbaum, and Evans (2005)).¹

But if habits are a true structural characteristic of utility, their influence should be seen in microeconomic data as well as in macroeconomic data. However, empirical studies using household-level data strongly reject the existence of habits of the magnitude necessary to explain aggregate consumption dynamics.²

At root, the conflict reflects a large difference in the forecastability of spending growth in micro and macro data. Even among micro studies that have found some evidence against the random walk proposition,³ we are not aware of any study that claims to have found more than a few percentage points’ worth of household-level spending growth to be predictable at any horizon. Roughly speaking, the predictability of aggregate spending growth is around an order of magnitude larger than predictability of household-level spending growth (say, 0.4 versus 0.04 in an \bar{R}^2 sense).

This paper proposes a simple solution. In place of habits, we postulate a modest informational friction: Not everybody instantaneously notices all macroeconomic developments. Instead, household macroeconomic expectations are “sticky,” as in Mankiw and Reis (2002) and Carroll (2003, 2006a). Specifically, while each consumer perfectly (‘frictionlessly’) perceives his own personal circumstances (employment status, relative wage rate, etc.), consumers’ information about macroeconomic quantities like productivity is obtained only occasionally (calibrated to match measured stickiness in household-level inflation expectations). Given its trivial utility cost (section 3.4), this small deviation from frictionless updating seems plausible — indeed, personal experience persuades us that even well-informed macroeconomists sometimes do not have perfect knowledge of the most recent macroeconomic statistics.

Our ‘sticky expectations’ approach has connections with a vast body of macroeconomic literature. In addition to the habit formation papers touched upon above,⁴ the model bears some resemblance to the Muth (1960)–Pischke (1995) framework in which consumers face a signal extraction problem in determining whether a shock to income is transitory or permanent, an approach that is in turn closely related to work by Lucas (1973) and Sims (2003); Woodford (2001) and Morris and Shin (2005) pursue a slightly

¹Perhaps the strongest evidence that some modification of the standard framework is necessary is the evident lack of enthusiasm for habit formation in many of these papers, which tend to excuse the assumption as a nose-holding necessity rather than as an intrinsically appealing assumption.

²Dynan (2000) is the best known micro study; others will be cited below. This tension between micro and macro evidence is longstanding; it was an important theme in Deaton (1992)’s canonical book, but little recent literature has focused on it.

³See the detailed discussion in Section 5.1.

⁴See also contributions by Constantinides (1990), Rotemberg and Woodford (1997), Guariglia and Rossi (2002), Carroll, Overland, and Weil (2000), Gabaix and Laibson (2001), Parker (2003), Parker and Preston (2005).

different approach that reaches a similar conclusion. More directly, recent work by Reis (2003) has explored a theoretical model in which consumers optimally choose to be inattentive because of explicit costs of attention; in the present paper, an especially simple specification of inattention permits correspondingly simple and intuitive results (as in Akerlof and Yellen (1985), Cochrane (1991), Calvo (1983) and emphasized recently by Browning and Crossley (2001)).⁵

Although numerical methods are necessary for precise calculation of the cost of stickiness, and for quantitative exactness in simulating consumption dynamics, the principal message for macroeconomists comes through clearly even in our much simpler DSGE model: Sticky aggregate expectations induce sticky aggregate consumption growth, while having little observable effect on household-level consumption dynamics. The model is therefore capable of reconciling the conflict between micro and macro evidence, in a way that closely mirrors recent developments in other branches of macroeconomics.

2 Theory

2.1 Value

Across all our model variants, one element is constant: The value of the consumer's dynamic consumption pattern is the discounted sum of time separable felicity

$$\mathbf{v}_t = \sum_{s=t}^{\infty} \beta^{s-t} \mathbf{u}(\mathbf{c}_s). \quad (1)$$

2.2 Quadratic Utility Benchmark

To fix notation and ideas, we start with the classic Hall (1978) random walk model. Total wealth \mathbf{o} (the sum of human and nonhuman wealth) evolves according to the dynamic budget constraint

$$\mathbf{o}_{t+1} = (\mathbf{o}_t - \mathbf{c}_t)\mathbf{R} + \zeta_{t+1}, \quad (2)$$

where $\mathbf{R} = (1 + r)$ is the interest factor and ζ_{t+1} is a shock to (total) wealth.

2.2.1 Frictionless Expectations

Under frictionless expectations, the usual derivations lead to the usual Euler equation

$$\mathbf{u}'(\mathbf{c}_t) = \mathbf{R}\beta \mathbb{E}_t[\mathbf{u}'(\mathbf{c}_{t+1})], \quad (3)$$

where \mathbb{E}_t denotes an assumption of instantaneous perfect frictionless updating of all information. Quadratic \mathbf{u} and $\mathbf{R}\beta = 1$ lead directly to Hall's random walk proposition:⁶

$$\Delta \mathbf{c}_{t+1} = \epsilon_{t+1}. \quad (4)$$

⁵ Browning and Collado (2001) suggest that the differing results in the micro literature can be resolved if consumers are not perfectly attentive even to all the details of their own personal income processes, an explanation which could easily be interpreted using framework proposed here.

⁶Deaton (1992) provides a lucid exposition.

Consumers spend

$$\mathbf{c}_t = (r/R)\mathbf{o}_t, \quad (5)$$

because this is exactly the amount that maintains expected wealth unchanged:

$$\mathbb{E}_t[\mathbf{o}_{t+1}] = (\mathbf{o}_t - \mathbf{c}_t)R = \mathbf{o}_t. \quad (6)$$

2.2.2 Sticky Expectations

Now suppose consumers update their information, and therefore their behavior, only occasionally. A consumer who updates in period t obtains precisely the same information that the frictionless consumer would have, forms the same expectations, and makes the same choices. Nonupdaters, however, behave as though their expectations had actually come true (since by definition they have obtained no information to the contrary). For example, consider a consumer who updates in periods t and $t + n$ but not between. Designating $\bar{\mathbf{o}}$ as the consumer's *perception* of wealth,

$$\begin{aligned} \bar{\mathbf{o}}_{t+1} &\equiv \mathbb{E}_t[\mathbf{o}_{t+1}] = \mathbf{o}_t \\ &\vdots \\ \bar{\mathbf{o}}_{t+n-1} &\equiv \mathbb{E}_t[\mathbf{o}_{t+n-1}] = \mathbf{o}_t \end{aligned}$$

the consumer spends according to perceived wealth so that

$$\begin{aligned} \mathbf{c}_t &= (r/R)\mathbf{o}_t \\ \mathbf{c}_{t+1} &= (r/R)\bar{\mathbf{o}}_{t+1} = (r/R)\mathbf{o}_t = \mathbf{c}_t \\ &\vdots \\ \mathbf{c}_{t+n-1} &= \mathbf{c}_t. \end{aligned}$$

The dynamics of actual (as distinct from perceived) wealth are given by (2),

$$\begin{aligned} \mathbf{o}_{t+1} &= \underbrace{(\mathbf{o}_t - \mathbf{c}_t)R}_{\mathbf{o}_t} + \zeta_{t+1} \\ \mathbf{o}_{t+2} &= \mathbf{o}_{t+1} + \zeta_{t+2} \\ &= (\mathbf{o}_{t+1} - \mathbf{c}_{t+1})R + \zeta_{t+2} \\ &= (\mathbf{o}_t + \underbrace{\mathbf{o}_{t+1} - \mathbf{o}_t}_{\zeta_{t+1}} - \mathbf{c}_t)R + \zeta_{t+2} \\ &= \mathbf{o}_t + \zeta_{t+1}R + \zeta_{t+2} \\ &\vdots \\ \mathbf{o}_{t+n} &= \mathbf{o}_t + \underbrace{\sum_{s=1}^n R^{n-s} \zeta_{t+s}}_{\equiv \Delta^n \mathbf{o}_{t+n}} \end{aligned}$$

so for a consumer who updates in periods t and $t + n$ but not between, the change in

consumption is

$$\mathbf{c}_{t+n} - \mathbf{c}_t = (r/R)\Delta^n \mathbf{o}_{t+n} \quad (7)$$

where $\Delta^n \mathbf{o}_{t+n}$ is white noise because it is a weighted sum of the white noise errors ζ . Thus, consumption follows a random walk across updating periods; consumers who were only observed during their updating periods would never be seen to deviate from the predictions of Hall (1978).

2.2.3 Aggregation

The economy is populated by consumers indexed by i distributed uniformly along the unit interval; the total population mass is 1, so that aggregate variables are also interpretable as per capita values, e.g. per capita/aggregate consumption is

$$\mathbf{C}_t = \int_0^1 \mathbf{c}_{t,i} di.$$

Whether the consumer at location i updates in period t is determined by the realization of the dichotomous random variable

$$\pi_{t,i} = \begin{cases} 1 & \text{if consumer } i \text{ updates in period } t \\ 0 & \text{if consumer } i \text{ does not update in period } t, \end{cases}$$

and each period's updaters are chosen randomly such that a constant proportion Π update in each period:

$$\begin{aligned} \mathbb{E}_{t,i}[\pi_{t+1,i}] &= \Pi \\ \int_0^1 \pi_{\tau,i} di &= \Pi \quad \forall \tau. \end{aligned}$$

Aggregate consumption is the population-weighted average of per-capita consumption of updaters \mathbf{C}^π and nonupdaters $\mathbf{C}^\not\pi$:

$$\mathbf{C}_{t+1} = \Pi \mathbf{C}_{t+1}^\pi + (1 - \Pi) \underbrace{\mathbf{C}_{t+1}^\not\pi}_{=\mathbf{C}_t} \quad (8)$$

where $\mathbf{C}_{t+1}^\not\pi = \mathbf{C}_t$ because the nonupdaters at time $t + 1$ are a random subset of the population at time t . The first difference of (8) yields

$$\Delta \mathbf{C}_{t+1} = (1 - \Pi)\Delta \mathbf{C}_t + \underbrace{\Pi \Delta \mathbf{C}_{t+1}^\pi}_{\equiv \epsilon_{t+1}} \quad (9)$$

and appendix A shows that ϵ_{t+1} is approximately mean zero.⁷ Thus, in the quadratic utility framework the serial correlation of aggregate per-capita consumption changes is an approximate measure of the proportion of nonupdaters.

⁷Intuitively, this term mainly reflects the behavior of updating consumers, and is therefore unpredictable with respect to sufficiently delayed information.

2.3 CRRA Utility

One of the lessons of the consumption literature after Hall (1978) is that his simplifying assumptions (quadratic utility, perfect capital markets, $R\beta = 1$) are far from innocuous; more plausible assumptions can lead to very different conclusions. In particular, a host of persuasive theoretical and empirical considerations has led to the now-standard assumption of constant relative risk aversion utility, $u(\mathbf{c}) = \mathbf{c}^{1-\rho}/(1-\rho)$.

When utility is not quadratic, solution of the model requires specification of the exact stochastic structure of the income and transition processes. We analyze two models that represent the polar extremes of the modern macro literature: A small open economy (or partial equilibrium) model with a rich and empirically realistic calibration of idiosyncratic and aggregate risk but exogenous interest rates and wages; and a dynamic stochastic general equilibrium model closed-economy model that abstracts from idiosyncratic income risk but endogenizes factor returns.

2.4 Partial Equilibrium/Small Open Economy (PE/SOE) Model

2.4.1 Frictionless Expectations

Labor productivity for the individual is determined by the interaction of transitory idiosyncratic (θ), transitory aggregate (Θ), permanent idiosyncratic (p), and permanent aggregate (P) factors:

$$\ell_{t+1} = \overbrace{\theta_{t+1}\Theta_{t+1}}^{\equiv\boldsymbol{\theta}_{t+1}} \underbrace{p_{t+1}P_{t+1}}_{\equiv\mathbf{p}_{t+1}} \quad (10)$$

where θ can be thought of as reflecting, for example, individual unemployment spells; Θ captures transitory aggregate shocks ('hurricanes'); p reflects the individual's permanent labor productivity; and P reflects aggregate permanent labor productivity. (Here and henceforth we drop the i subscripts from the idiosyncratic variables, except when they are useful for clarity; thus, e.g., $\theta_{t+1} \equiv \theta_{t+1,i}$). θ and Θ are iid, and satisfy $\mathbb{E}_t[\theta_{t+n}] = \mathbb{E}_t[\Theta_{t+n}] = 1 \forall n > 0$. The idiosyncratic transitory shock has a minimum possible value of 0 (corresponding to an unemployment spell) which occurs with a small finite probability φ . (This is in essence equivalent to imposing a liquidity constraint, cf. Zeldes (1989b)). For most purposes the aggregate and idiosyncratic transitory shocks can be combined into a single overall transitory shock indicated by the use of boldface, $\boldsymbol{\theta}$, and the aggregate and idiosyncratic levels of permanent income can be combined as \mathbf{p} .

Market resources \mathbf{m} next period reflect the wage rate times the individual's (exogenous) labor supply ℓ plus the capital productivity factor times the individual's 'capital' stock \mathbf{k}_{t+1} ,

$$\mathbf{m}_{t+1} = \underbrace{W_{t+1}\ell_{t+1}}_{\equiv\mathbf{y}_{t+1}} + \underbrace{\mathcal{R}_{t+1}}_{1+r_{t+1}} \mathbf{k}_{t+1} \quad (11)$$

and we will henceforth call \mathbf{y}_{t+1} 'labor income.'

Idiosyncratic and aggregate permanent productivity evolve according to⁸

$$p_{t+1} = p_t \psi_{t+1} \quad (12)$$

$$P_{t+1} = P_t \Psi_{t+1} \quad (13)$$

where both idiosyncratic and aggregate permanent shocks are iid ($\mathbb{E}_t[\psi_{t+n}] = \mathbb{E}_t[\Psi_{t+n}] = 1 \forall n > 0$), so that log idiosyncratic and aggregate permanent productivity follow random walks. The combined permanent shock is boldface $\boldsymbol{\psi}_t \equiv \psi_t \Psi_t$.

The transition process for \mathbf{m} is broken up, for convenience of analysis, into three steps. ‘Assets’ at the end of the period are market resources minus consumption,

$$\mathbf{a}_t = \mathbf{m}_t - \mathbf{c}_t,$$

while next period’s capital is determined from this period’s assets via

$$\mathbf{k}_{t+1} = \mathbf{a}_t \Upsilon / \mathcal{D}, \quad (14)$$

where Υ is the depreciation factor for capital (normally suppressed in micro analysis but necessary here for comparability with the DSGE treatment below); the factor \mathcal{D} , which will be slightly less than one, reflects the individual’s survival probability *a la* Blanchard (1985) as elaborated below.

Defining nonbold variables as the bold version normalized by permanent labor productivity (e.g. $m_t = \mathbf{m}_t / \mathbf{p}_t$), Carroll (2011) shows that, if the aggregate wage rate and capital productivity factors are constant at \mathcal{R} and \mathcal{W} , the solution to this problem can be derived from the solution to the single-state-variable problem

$$\begin{aligned} v(m_t) &= \max_{\{c_t\}} \{u(c_t) + \beta \mathbb{E}_t[\boldsymbol{\psi}_{t+1}^{1-\rho} v(m_{t+1})]\} \\ &\text{s.t.} \\ a_t &= m_t - c_t \\ k_{t+1} &= a_t \Upsilon / \mathcal{D} \boldsymbol{\psi}_{t+1} \\ m_{t+1} &= \mathcal{R} k_{t+1} + \mathcal{W} \boldsymbol{\theta}_{t+1}, \end{aligned} \quad (15)$$

which has a solution if the impatience condition⁹

$$\underbrace{\mathcal{R}}_{\mathcal{R} \Upsilon / \mathcal{D}} \beta \mathbb{E}[\boldsymbol{\psi}^{-\rho}] < 1 \quad (16)$$

holds (where for this equation we drop time subscripts because all shocks are iid).¹⁰

Designating the converged consumption policy function that solves (15) as $c(m)$, the level of consumption for the frictionless consumer can be obtained from

$$\mathbf{c}_t = c(m_t) \mathbf{p}_t.$$

⁸For expositional simplicity, we omit the predictable trend component of permanent productivity growth; it is included below in the model calibration and solution.

⁹Other parametric restrictions are also necessary, but for typical parameterizations are not likely to be binding; see Carroll (2011) for details.

¹⁰Note that the relevant interest factor is the within-period productivity of capital \mathcal{R} adjusted for both of the influences (Υ and \mathcal{D}) that intervene between the amount of assets with which the consumer ends period t and the amount of productive capital owned when capital income is received in period $t + 1$.

2.4.2 Aggregation

Aggregating a productivity process like (12) with infinitely-lived households would generate a nonstationary distribution of idiosyncratic productivity. To avoid this inconvenience, we make the Blanchard (1985) assumption: Each consumer faces a constant probability of mortality in each period. We keep track using another zero-one indicator:

$$\mathfrak{d}_{t+1,i} = \begin{cases} 0 & \text{if consumer at location } i \text{ dies between time } t \text{ and } t+1 \\ 1 & \text{if consumer at location } i \text{ does not die between } t \text{ and } t+1 \end{cases}$$

We refer to this henceforth as a ‘replacement’ event, since the consumer who dies is replaced by an unrelated newborn who happens to inhabit the same location on the numberline. The *ex ante* probability of survival is identical for each consumer, so that the aggregate mass of consumers who survive is time invariant at $\mathfrak{D} = \int_0^1 \mathfrak{d}_{\tau,i} di \forall \tau$.

Under the assumption that newborns start life with permanent income equal to the population mean of one,

$$p_{t+1,i} = \begin{cases} 1 & \text{for newborns} \\ p_{t,i} \psi_{t+1,i} & \text{for survivors,} \end{cases}$$

the population mean of the idiosyncratic component of permanent income is always $\int_0^1 p_{t,i} di = 1$, and appendix D derives, for the idiosyncratic component of permanent income p , the dynamics and steady state of its population variance, which will exist so long as $\mathfrak{D} \mathbb{E}[\psi^2] < 1$ (which imposes a loose restriction on the magnitude of the idiosyncratic permanent shocks).

Households engage in a Blanchardian mutual insurance scheme: Those who survive receive a proportion of the estates of those who die,¹¹

$$\mathbf{k}_{t+1,i} = \begin{cases} 0 & \text{if agent at } i \text{ dies, is replaced by a newborn} \\ \mathbf{a}_{t,i} \mathfrak{T} / \mathfrak{D} & \text{if agent at } i \text{ survives} \end{cases}$$

or more compactly

$$\mathbf{k}_{t+1,i} = (1 - \mathfrak{d}_{t+1,i})0 + \mathfrak{d}_{t+1,i} \mathbf{a}_{t,i} \mathfrak{T} / \mathfrak{D}. \quad (17)$$

Next period’s aggregate capital is the population integral of (17)

$$\begin{aligned} \mathbf{K}_{t+1} &= \int_0^1 (\mathfrak{d}_{t+1,i} \mathfrak{T} \mathbf{a}_{t,i} / \mathfrak{D}) di \\ &= \mathfrak{T} \mathbf{a}_t \\ K_{t+1} &= \mathfrak{T} \mathbf{a}_t / \Psi_{t+1} \end{aligned}$$

where the second equation follows from the first because $\mathfrak{D}^{-1} \int_0^1 \mathfrak{d}_{t+1,i} di = 1$ and $\mathfrak{d}_{t+1,i}$ is independent of $\mathbf{a}_{t,i}$. This explains why the capital transition equation for the individual’s optimization problem, (14), involved \mathfrak{D} : Since the individual receives no utility after

¹¹The constant population permits the analytical convenience of exactly replacing the dying with newborns, but it is important to understand that there is no relationship between successive persons at the same location on the number line; this is not a dynastic model.

death, the optimization problem is implicitly contingent on survival, and survivors' capital is boosted by the proceeds of the insurance scheme.

Since $\int_0^1 \theta_{t,i} = \int_0^1 p_{t,i} = 1$, aggregate labor supply is

$$\mathbf{L}_t = \int_0^1 \boldsymbol{\ell}_{t,i} di = \Theta_t P_t. \quad (18)$$

Aggregate resources can be written as per-capita resources of the survivors times their population mass \mathcal{D} , plus per-capita resources of the newborns times their population mass $1 - \mathcal{D}$:

$$\begin{aligned} \mathbf{M}_{t+1} &= \left(\overbrace{\mathbf{a}_t \mathcal{R} / \mathcal{D} + \Theta_{t+1} P_{t+1} \mathbf{W}}^{\text{per-capita } \mathbf{m} \text{ for survivors}} \right) \mathcal{D} + \overbrace{\Theta_{t+1} P_{t+1} \mathbf{W}}^{\text{per-capita } \mathbf{m} \text{ for newborns}} (1 - \mathcal{D}) \\ M_{t+1} &= K_{t+1} \mathcal{R} + \Theta_{t+1} \mathbf{W}. \end{aligned} \quad (19)$$

2.4.3 Sticky Expectations

Among surviving consumers with sticky expectations, perceptions of aggregate permanent income evolve according to

$$\bar{P}_{t+1,i} = \pi_{t+1,i} P_{t+1} + (1 - \pi_{t+1,i}) \bar{P}_{t,i}. \quad (20)$$

For simplicity, newborns begin life with the mean perceptions prevailing in the existing population at their date of birth, which implies that (20) holds at all i (whether or not a replacement event has occurred at i).¹²

Nonupdaters assume that the aggregate transitory shock takes its mean value of 1,

$$\bar{\Theta}_{t,i} = \begin{cases} \Theta_t & \text{for updaters} \\ 1 & \text{for nonupdaters} \end{cases}$$

so perceived idiosyncratic labor supply is

$$\bar{\boldsymbol{\ell}}_{t,i} = \theta_{t,i} \bar{\Theta}_{t,i} p_{t,i} \bar{P}_{t,i}.$$

Given $\bar{\boldsymbol{\ell}}$, household i 's perceived level of market resources is

$$\bar{\mathbf{m}}_{t,i} = \bar{\mathbf{k}}_{t,i} \mathcal{R} + \bar{\boldsymbol{\ell}}_{t,i} \mathbf{W},$$

which implies that misperceptions of labor productivity translate directly into misperceptions of \mathbf{m} (even for a consumer with a correct perception of beginning-of-period capital, $\bar{\mathbf{k}}_{t,i} = \mathbf{k}_{t,i}$).¹³ Perceived normalized market resources are

$$\bar{m}_{t,i} = \bar{\mathbf{m}}_{t,i} / \bar{P}_{t,i} p_{t,i}.$$

The essence of our approach is to assume that the consumer with sticky expectations behaves according to the decision rule that would be correct for a consumer with

¹²This assumption about newborns' beliefs is numerically inconsequential because the quarterly replacement rate is so low.

¹³A possible interpretation of this misperception (which will turn out to be quantitatively slight) is that it reflects imperfect knowledge of the most recent inflation rate.

frictionless expectations. (Such behavior is suboptimal if costs of information acquisition and processing are zero; below show that it is only a *little* bit suboptimal, so that modest information-gathering costs could justify the assumption.) Since the frictionless consumption function is $c(m)$,

$$\begin{aligned}\bar{c}_{t,i} &= c(\bar{m}_{t,i}) \\ \mathbf{c}_{t,i} &= \bar{c}_{t,i} \bar{P}_{t,i} p_{t,i}\end{aligned}$$

where \mathbf{c} does not need a bar over it because every consumer can always correctly perceive his own level of consumption.

We assume, however, that after the period- t consumption decision has been made, the consumer receives a bank statement and other financial reports that reveal the true end-of-period value of $\mathbf{a}_{t,i}$. So any misperception about the level of real resources is not only small (as will emerge in our calibration) but brief (lasting only for the interval between the beginning of the quarter when the consumption decision is made, and the end of the quarter when end-of-period bank balances are revealed).¹⁴

Population-average perceptions follow from the assumptions above. For example, the average perception of the aggregate transitory shock is the population-weighted average of the perceptions of the updaters (Θ_t) and nonupdaters (1):

$$\bar{\Theta}_{t+1} = \Pi \Theta_{t+1} + (1 - \Pi) 1 \quad (21)$$

$$= 1 + (\Theta_{t+1} - 1) \Pi, \quad (22)$$

while average perceptions of aggregate permanent productivity evolve according to

$$\bar{P}_{t+1} = \Pi \overbrace{\bar{P}_{t+1}^\pi}^{P_{t+1}} + (1 - \Pi) \overbrace{\bar{P}_{t+1}^\dagger}^{\bar{P}_t}. \quad (23)$$

Everyone understands that newborns begin with zero capital, so per capita perceived and actual capital among survivors who update in period $t + 1$ will be

$$\begin{aligned}\bar{\mathbf{k}}_{t+1}^\pi &= \Pi^{-1} \int_0^1 (\mathbf{a}_{t,i} \Upsilon / \mathcal{D}) \pi_{t+1,i} di \\ &= \Upsilon \mathbf{a}_t / \mathcal{D}\end{aligned} \quad (24)$$

and, similarly, for $t + 1$ nonupdaters $\bar{\mathbf{k}}_{t+1}^\dagger = \Upsilon \bar{\mathbf{a}}_t / \mathcal{D}$. Average perceived capital is the perceived capital of survivors times their population weight \mathcal{D} plus that of newborns (0) times their weight $1 - \mathcal{D}$:

$$\begin{aligned}\bar{\mathbf{k}}_{t+1} &= \left(\Pi \bar{\mathbf{k}}_{t+1}^\pi + (1 - \Pi) \bar{\mathbf{k}}_{t+1}^\dagger \right) \mathcal{D} + 0 \times (1 - \mathcal{D}) \\ &= (\Pi \mathbf{a}_t + (1 - \Pi) \bar{\mathbf{a}}_t) \Upsilon.\end{aligned} \quad (25)$$

¹⁴Homage is due here to the New Keynesian literature of the 1980s that made the point that suboptimal behavior that has small microeconomic costs (e.g. 'menu costs') can have large macroeconomic consequences.

2.5 Muth (1960)–Lucas (1973)–Pischke (1995) Kalman Filter Model

The principal rival to habit formation as an explanation of sluggishness in aggregate consumption is the Muth–Lucas–Pischke framework in which agents can perceive only the level of their income and not how it is constituted between permanent and transitory components. They must perform an optimal signal extraction problem to decompose income shocks into estimated transitory and permanent components.¹⁵ (This is mathematically equivalent to a simple version of the Kalman filter). In the Muth–Lucas–Pischke framework, each household updates its estimate of permanent income according to an equation of the form

$$\bar{p}_{t+1,i} = \Pi y_{t,i} + (1 - \Pi)\bar{p}_{t,i}. \quad (26)$$

Defining the signal-to-noise ratio $\varphi = \sigma_{\psi}^2/\sigma_{\theta}^2$, if consumers are never allowed to observe the breakdown of shocks between transitory and permanent components, formulae from Muth can be extended to show that the optimal value of Π is

$$\Pi = \left(\frac{1}{1 + 2/(\varphi + \sqrt{\varphi^2 + 4\varphi})} \right), \quad (27)$$

which makes intuitive sense because it says that if shocks are perceived to be completely permanent ($\varphi = \infty$) then perceived permanent income is perceived actual income, while if shocks are perceived to be completely transitory ($\varphi = 0$) then in the limit the current level of income provides no signal about the level of permanent income, so $\Pi = 0$.

The natural way to calibrate this model is to follow Lucas (1973) and assume that individuals cannot distinguish idiosyncratic from aggregate shocks; instead, they simply observe the level of actual income received, and from that extract their estimate of the level of permanent income. For the empirical calibration of the idiosyncratic and aggregate shocks given below, the formula implies that $\Pi \approx 0.025$ so that if consumption is equal to permanent income the Muth–Lucas–Pischke formula implies that the serial correlation in aggregate consumption growth should be approximately $1 - \Pi = 0.975$, in contrast with empirical estimates below that suggest the correct figure is about 0.75. Thus, at least under our reading of the empirical literatures estimating the size of the idiosyncratic and aggregate income processes, the Muth–Lucas–Pischke logic actually goes *too far* in generating aggregate smoothness.¹⁶

One could extend the framework to allow agents to know their idiosyncratic circumstances but perform Kalman filter/Muth–Lucas–Pischke updating with respect to aggregate income data. However, the logic of the model implies that if consumers are updating their estimates of aggregate permanent income ‘correctly’ in the signal extraction sense, aggregate consumption should again follow a random walk. Since the

¹⁵Lucas (1973) is ostensibly about distinguishing idiosyncratic from aggregate shocks, but the mathematical framework is very similar to those of Muth and Pischke, justifying our grouping of these papers together.

¹⁶In order to illustrate the logic of his model, Pischke (1995) calibrates an example which assumes that all aggregate shocks are permanent while all idiosyncratic shocks are transitory. Later, he presents estimates of the magnitude of transitory and permanent shocks obtained from the *Survey of Income and Program Participation*. However, those estimates do not match very well the magnitude of estimated transitory and permanent shocks from the other micro data sources used in the literature cited for calibration purposes here.

whole point of our paper is to explain why aggregate consumption does *not* follow a random walk, this approach is a dead end for our purposes.

Of course, it is likely that there is some calibration of φ which differs from either the value we obtain from our calibration using micro data or the value that characterizes macro data, and which would have implications for the quadratic utility model that might be difficult to distinguish from the implications of the sticky expectations model we advocate. Those who are uncomfortable with assuming sticky expectations directly may take comfort in this alternative possibility.¹⁷

2.6 Dynamic Stochastic General Equilibrium (DSGE) Model

The PE/SOE model takes aggregate wages and interest rates as exogenous. The alternative is to begin with a standard representative agent dynamic stochastic general equilibrium model and modify it to permit sticky expectations. (Idiosyncratic wage and productivity shocks are suppressed for tractability.)¹⁸

With a Cobb-Douglas aggregate production function, aggregate labor supply and output are respectively

$$\begin{aligned}\mathbf{L}_t &= \Theta_t P_t \\ \mathbf{F}(\mathbf{K}_t, \mathbf{L}_t) &= \mathbf{K}_t^\varepsilon \mathbf{L}_t^{1-\varepsilon}\end{aligned}$$

while market resources and the transition are

$$\begin{aligned}\mathbf{M}_t &= \mathbf{K}_t + \mathbf{F}(\mathbf{K}_t, \mathbf{L}_t) \\ \mathbf{A}_t &= \mathbf{M}_t - \mathbf{C}_t \\ \mathbf{K}_{t+1} &= \mathbf{A}_t \bar{\gamma}.\end{aligned}$$

2.6.1 Frictionless Expectations

Normalizing again,

$$\begin{aligned}M_t &= (\mathbf{K}_t + \mathbf{K}_t^\varepsilon (\Theta_t P_t)^{1-\varepsilon}) / P_t \\ &= K_t + K_t^\varepsilon \Theta_t^{1-\varepsilon} \\ &= K_t + \Theta_t^{1-\varepsilon} (\varepsilon K_t^{\varepsilon-1} K_t + (1-\varepsilon) K_t^\varepsilon) \\ &= \underbrace{\mathcal{R}_t}_{\equiv 1 + \varepsilon K_t^{\varepsilon-1} \Theta_t^{1-\varepsilon}} K_t + \underbrace{\mathcal{W}_t}_{\equiv (1-\varepsilon) K_t^\varepsilon \Theta_t^{-\varepsilon}} \Theta_t\end{aligned}\tag{28}$$

where the definition of \mathcal{R} reflects dM/dK and the wage rate is the marginal product of another efficiency unit of labor.

¹⁷A rigorous signal-extraction treatment of the baseline consumption model we are considering here (CRRA utility consumers who see only the level of aggregate income and not how it is decomposed) would in principle require adding at least one extra state variable to the consumer's optimization problem; the consumption problem reduces to the permanent-income-signal-extraction problem only under the full panoply of the original Hall assumptions (quadratic utility, $\mathbf{R}\beta = 1$, etc.) It therefore seems likely that a rigorously correct solution would be computationally very difficult (and at a minimum would be far more complex than our approach).

¹⁸Developing a model with fully realistic idiosyncratic transitory and permanent shocks as well as frictionless expectations general equilibrium dynamics is a formidable problem; see Krusell and Smith (1998) for a version with simple idiosyncratic shocks.

The normalized frictionless representative agent's problem has a structure parallel to that of the frictionless PE/SOE problem (cf. (15))

$$\begin{aligned}
V(M_t) &= \max_{C_t} \{u(C_t) + \beta \mathbb{E}_t[\Psi_{t+1}^{1-\rho} V(M_{t+1})]\} \\
&\text{s.t.} \\
A_t &= M_t - C_t \\
K_{t+1} &= A_t \Upsilon / \Psi_{t+1} \\
M_{t+1} &= \mathcal{R}_{t+1} K_{t+1} + \mathcal{W}_{t+1} \Theta_{t+1}.
\end{aligned} \tag{29}$$

2.6.2 Sticky Expectations

The sticky-expectations representative agent updates perceived transitory and permanent labor productivity according to the same equations (23) and (22) that characterize the evolution of average perceptions in the PE/SOE model. Perceived aggregate capital dynamics are also the same (cf. (25)).

Defining normalized perceived variables as before (e.g. $\bar{M}_{t+1} = \bar{M}_{t+1} / \bar{P}_{t+1}$), we assume that the representative agent's perception of \bar{M} reflects a correct understanding of the production function (28)

$$\bar{M}_t = \bar{K}_t + \bar{K}_t^\varepsilon \bar{\Theta}_t^{1-\varepsilon} \tag{30}$$

and as before the sticky-expectations agent behaves according to the decision rule associated with the frictionless solution,

$$C_t = C(\bar{M}_t) \bar{P}_t \tag{31}$$

and observes the true value of assets after the consumption decision has been made:

$$A_t = P_t (K_t + K_t^\varepsilon \Theta_t^{1-\varepsilon}) - C_t. \tag{32}$$

The sticky-expectations representative agent is assumed to behave according to the consumption rule that would be optimal for a consumer with frictionless expectations, as before.

2.7 The Cost of Stickiness in Theory

Π has so far been taken as exogenous. But the probability of updating presumably depends at least partly on costs and benefits. This section briefly examines the tradeoffs by imagining that newborns make a once-and-for-all choice of their idiosyncratic value of Π ; for a more thorough theoretical examination of the tradeoffs in a related model, see Reis (2003), and for quantitative estimates of the costs see the discussion below.

In the first period of life, we assume that the consumer is employed and experiences no transitory or permanent shocks so that market resources are nonstochastically equal to \mathcal{W} ; value can therefore be written as $v(\mathcal{W})$. There is of course no analytical expression for v ; but, fixing all parameters aside from the variance of the permanent shock, theoretical considerations suggest (and simulations confirm) that the consequences of permanent

uncertainty for value can be well approximated by

$$v(\mathbf{W}) \approx \overleftarrow{v}(\mathbf{W}) - \kappa\sigma_{\Psi}^2 \quad (33)$$

where $\overleftarrow{v}(\mathbf{W})$ is the value that would be generated by a model with no permanent shocks $\sigma_{\Psi}^2 = 0$ and κ is a constant of approximation that captures the cost of aggregate permanent uncertainty.

Suppose now that the effect of sticky expectations is approximately to reduce value by an amount proportional to the inverse of the updating probability:¹⁹

$$\bar{v}(\mathbf{W}) \approx \overleftarrow{v}(\mathbf{W}) - (\kappa/\Pi)\sigma_{\Psi}^2. \quad (34)$$

This assumption has appropriate scaling properties in three senses:

- If $\sigma_{\Psi}^2 = 0$ so that there are no permanent shocks, then the cost of stickiness is zero (given our assumption that initial perceptions are correct)
- If the probability of updating is $\Pi = 1$ so that perceptions are always accurate, value is the same as in the frictionless model
- If expectations never adjust, then $\Pi = 0$ and the utility cost of stickiness is infinite, which is appropriate because consumers would be making choices based on expectations that would eventually be arbitrarily far from the truth

Simulations (available from the corresponding author) confirm that, using the DSGE model as an experimental platform, for variations around the baseline parameter values, (34) is an excellent approximation to the true cost of uncertainty.

Now imagine that newborns make a once-and-for-all choice of the value of Π ; a higher Π (faster updating) is assumed to have a linear cost ι in units of normalized value (think of this as utility costs of paying attention to macroeconomic newscasts instead of sportscasts or other more pleasurable news items; since we are examining a model that has been normalized by productivity, this could alternatively be loosely interpreted as a time cost of gathering information). The newborn's objective is therefore to choose the Π that solves

$$\max_{\Pi} \overleftarrow{v}(\mathbf{W}) - (\kappa/\Pi)\sigma_{\Psi}^2 - \iota\Pi. \quad (35)$$

The first order condition is

$$\begin{aligned} 0 &= \Pi^{-2}\kappa\sigma_{\Psi}^2 - \iota \\ \Pi^2 &= (\kappa\sigma_{\Psi}^2)/\iota \end{aligned}$$

which leads to the conclusion that the consumer will pick the Π satisfying

$$\Pi = (\kappa/\iota)^{0.5}\sigma_{\Psi} \quad (36)$$

Thus, the speed of updating should be related directly to the utility cost of permanent uncertainty (κ), inversely to the cost of information (cheaper information = faster updating), and linearly to the standard deviation of permanent shocks.

¹⁹ In principle, \bar{v} will depend on the initial perception error imbibed at birth; think of (33) as the integral over the ergodic distribution of values of the initial perception error.

Of course it is unrealistic to assume that consumers could never reoptimize their choices of Π . It would be straightforward to modify the model to give them occasional opportunities to change Π (the simplest method would be to have an exogenous arrival rate of opportunities to change Π , say on average once every few years). Such an extension is beyond the scope of the present paper, but could be an interesting subject for future work.

3 Calibration and Equilibrium

We begin by calibrating the DSGE model in a traditional way, and next specify the additional parameters necessary for solving the PE/SOE model.

3.1 DSGE Calibration

We assume a coefficient of relative risk aversion of 2, in the middle of the range usually considered plausible. The quarterly depreciation rate is $\Upsilon^4 = 0.94$ at an annual rate. Capital's share in aggregate output takes its usual value of $\varepsilon = 0.36$, and appendix B uses U.S. NIPA data to calibrate the variances of the quarterly transitory and permanent shocks at the approximate values

$$\begin{aligned}\sigma_{\Psi}^2 &= 0.00004 \\ \sigma_{\Theta}^2 &= 0.000004,\end{aligned}$$

which are consistent with standard estimates in the literature (see the references in the appendix for details). Finally, the perfect foresight steady state aggregate capital/output ratio is 12 (where income is measured quarterly; this corresponds to the usual ratio of 3 for capital divided by annual income). These assumptions imply values for the other steady-state characteristics of the model

$$\begin{aligned}K &= 12^{1/(1-\varepsilon)} \\ \mathcal{W} &= (1 - \varepsilon)K^\varepsilon.\end{aligned}$$

The model implies a net-of-depreciation between-period interest factor of $(1 + \varepsilon K_t^{\varepsilon-1})\Upsilon$; if there were no aggregate shocks, the DSGE model would have a deterministic steady state satisfying the Euler equation with $C_{t+1} = C_t$:

$$\overbrace{(1 + \varepsilon K^{\varepsilon-1})\Upsilon}^{\mathcal{R}} \beta = 1 \tag{37}$$

$$\beta = \mathcal{R}^{-1}. \tag{38}$$

We use the endogenous gridpoints method described in Carroll (2006b) to solve for the policy function $C(M)$. The target M_t ratio, \check{M} , is calculated as the level satisfying

$$M_t = \mathbb{E}_t[M_{t+1}] \tag{39}$$

$$= \mathbb{E}_t[K_{t+1} + K_{t+1}^\varepsilon \Theta_{t+1}^{1-\varepsilon}] \tag{40}$$

Table 1 Calibration

DSGE Model		
Calibrated Parameters		
ρ	2.	Coefficient of Relative Risk Aversion
$\bar{\gamma}$	0.94 ^{1/4}	Quarterly Depreciation Factor
K/K^ε	12	Perf Foresight SS Capital/Output Ratio
σ_Θ^2	0.00001	Variance Qtrly Tran Agg Pty Shocks
σ_Ψ^2	0.00004	Variance Qtrly Perm Agg Pty Shocks
Steady State Solution of Model With $\sigma_\Psi = \sigma_\Theta = 0$		
$K = 12^{1/(1-\varepsilon)}$	≈ 48.55	Steady State Quarterly K/P Ratio
$M = K + K^\varepsilon$	≈ 52.6	Steady State Quarterly M/P Ratio
$\mathcal{W} = (1 - \varepsilon)K^\varepsilon$	≈ 2.59	Quarterly Wage Rate
$\mathcal{R} = 1 + \varepsilon K^{\varepsilon-1}$	$= 1.03$	Quarterly Gross Capital Income Factor
$\mathbf{R} = \mathcal{R}\bar{\gamma}$	≈ 1.014	Quarterly Between-Period Interest Factor
$\beta = \mathbf{R}^{-1}$	≈ 0.986	Quarterly Time Preference Factor
Partial Equilibrium/Small Open Economy (PE/SOE) Model Parameters		
Calibrated Parameters		
σ_ψ^2	0.012	Variance Annual Perm Idiosyncratic Shocks (PSID)
σ_θ^2	0.03	Variance Annual Tran Idiosyncratic Shocks (PSID)
\wp	0.05	Quarterly Probability of Unemployment Spell
Π	0.25	Quarterly Probability of Updating Expectations
$(1 - \Omega)$	0.005	Quarterly Probability of Mortality
Calculated Parameters		
$\beta = 0.99\Omega/E[(\psi)^{-\rho}]\mathbf{R}$	0.969	Satisfies Impatience Condition: $\beta < \Omega/E[(\Psi\psi)^{-\rho}]\mathbf{R}$
σ_ψ^2	0.004	Variance Qtrly Perm Idiosyncratic Shocks ($=\sigma_\psi^2/4$)
σ_θ^2	0.12	Variance Qtrly Tran Idiosyncratic Shocks ($=4\sigma_\theta^2$)

where

$$K_{t+1} = (M_t - C(M_t)) \Upsilon / \Psi_{t+1} \quad (41)$$

and target K is the value of K_t which, if $\Theta_t = 1$, produces $M_t = \check{M}$.

3.2 PE/SOE Calibration

We fix the aggregate interest factor \mathcal{R} and wage rate W to the values obtained at the perfect foresight steady state of the DSGE model. The annual-rate idiosyncratic transitory and permanent shocks are assumed to be

$$\begin{aligned} \sigma_\psi^2 &= 0.012 \\ \sigma_\theta^2 &= 0.03. \end{aligned}$$

These figures are conservative in comparison with standard raw estimates from the micro data; using data from the *Panel Study of Income Dynamics*, for example, Carroll and Samwick (1997) estimate $\sigma_\psi^2 = 0.0217$ and $\sigma_\theta^2 = 0.0440$; Storesletten, Telmer, and Yaron (2004) estimate $\sigma_\psi^2 \approx 0.017$, with varying estimates of the transitory component. But recent work by Low, Meghir, and Pistaferri (2005) suggests that controlling for job mobility and participation decisions reduces estimates of the permanent variance somewhat; and using very well-measured Danish administrative data, Nielsen and Vissing-Jorgensen (2006) estimate $\sigma_\psi^2 \approx 0.005$ and $\sigma_\theta^2 \approx 0.015$, which presumably constitute lower bounds for plausible values for the truth in the U.S. (given the comparative generosity of the Danish welfare state).

Since the variance of the annual permanent innovation is 4 times the variance of the quarterly innovation, this calibration implies that the variance of the idiosyncratic permanent innovations is about 100 times the variance of the aggregate permanent innovations (4×0.00004 divided by 0.012). This is a point worth emphasizing: Idiosyncratic uncertainty is *approximately two orders of magnitude larger* than aggregate uncertainty. While reasonable people could differ a bit from our calibration of either the aggregate or the idiosyncratic risk, no plausible calibration of either magnitude will change the fundamental point that the aggregate component of risk is tiny compared to the idiosyncratic component. This is why assuming that people do not pay close attention to the macroeconomic environment is plausible.

We assume the probability of unemployment is 5 percent per quarter. This approximates the historical mean unemployment rate in the U.S., but model unemployment differs from real unemployment in (at least) two important ways. First, the model does not incorporate unemployment insurance, so labor income of the unemployed is zero. Second, model unemployment shocks last only one quarter, so their duration is shorter than the typical U.S. unemployment spell (about 6 months). The idea of the calibration is that a single quarter of unemployment with zero benefits is roughly as bad as two quarters of unemployment with an unemployment insurance payment of half of permanent labor income (a reasonable approximation to the typical situation facing unemployed workers). The model could be modified to permit a more realistic treatment of unemployment spells; this is a promising topic for future research, but would

involve a major increase in model complexity because realism would require adding the individual's employment situation as a state variable.

The probability of mortality is set at $1 - \delta = 0.005$ which implies an expected working life of 50 years; results are not sensitive to plausible alternative values of this parameter, so long as the life length is short enough to permit a stationary distribution of idiosyncratic permanent income.

We calibrate the probability of updating at $\Pi = 0.25$ per quarter, for several reasons. First, this is the parameter value assumed for the speed of expectations updating by Mankiw and Reis (2002) in their analysis of the consequences of sticky expectations for inflation. They argue that an average frequency of updating of once a year is intuitively plausible. Second, Carroll (2003) estimates an empirical process for the adjustment process for household inflation expectations in which the point estimate of the corresponding parameter is 0.27 for inflation expectations and 0.32 for unemployment expectations; the similarity of these figures suggests 0.25 is a reasonable benchmark, and provides some insulation against the charge that the model is *ad hoc*: It is calibrated in a way that corresponds to estimates of the stickiness of expectations in fundamentally different context. Finally, empirical results presented below will also suggest a speed of updating for US consumption dynamics of about 0.25 per quarter.

β remains to be tied down. The β obtained from the perfect foresight DSGE calibration cannot be used, because the extra uncertainty found in idiosyncratic data would imply that such a calibration would not satisfy the impatience condition (16) and the precautionary saving motive would cause wealth to diverge to infinity.

We could seek a value of β which causes the PE/SOE model to match some empirical measure of observed wealth holdings in micro data. However, there is considerable ambiguity about which measure of wealth the model should match; the answer depends in part on whether housing equity should be viewed as part of the precautionary buffer stock, the age range of the households being matched, the measure of permanent income, and many other extraneous issues. Given this ambiguity, rather than seeking a β which reproduces a specific target level of wealth, we choose a simple calibration in which the quarterly value of β is 0.99 times the impatience criterion (16); if there were no uncertainty and no growth, this would translate into about a 4 percent annual discount rate.

As discussed below, these two parameterizations are chosen so that they will span the full range of calibrations found in the micro and macro literatures. Further alternative calibrations are possible, but experimentation has indicated that results are not very sensitive to such choices.

3.3 Equilibrium Characteristics

This section briefly characterizes some of the equilibrium characteristics of the solutions to the models under the parameters specified above. Results are reported in Table 2.

Note first the considerable difference between the mean level of assets in the DSGE and PE/SOE models (first row of the table). As indicated above, this reflects our goal of presenting results that span the full range of the micro and macro literatures; the micro

Table 2 Equilibrium Statistics

	PE/SOE Economy		DSGE Economy	
	Frictionless	Sticky	Frictionless	Sticky
Means				
A	6.650	6.648	49.382	49.371
C	2.684	2.684	3.290	3.289
Standard Deviations				
Aggregate Time Series ('Macro')				
$\log \mathbf{A}$	0.089	0.091	0.085	0.085
$\Delta \log \mathbf{C}$	0.005	0.002	0.003	0.001
$\Delta \log \mathbf{Y}$	0.008	0.003	0.005	0.002
Individual Cross Sectional ('Micro')				
$\log \mathbf{a}$	1.273	1.273		
$\log \mathbf{c}$	1.207	1.207		
$\log \mathbf{p}$	1.221	1.221		
$\log \mathbf{y} \mathbf{y} > 0$	0.846	0.846		
$\Delta \log \mathbf{c}$	0.151	0.149		
Cost Of Stickiness	0.31×10^{-4}		0.53×10^{-5}	

Notes: The cost of stickiness is calculated as the proportion by which the permanent income of a frictionless consumer would need to be reduced in order to achieve the same reduction of expected value associated with forcing them to become a sticky expectations consumer.

literature has often focused on trying to explain the wealth holdings of the median household, which are much smaller than wealth holdings of the representative agent. In the spirit of reconciliation between the two camps, we urge macroeconomists to be open minded about the micro approach of matching the median household; reciprocally, microeconomists are urged to tolerate the reporting of DSGE model results.

The table hints at a broad generalization: With respect to either cross section statistics, mean outcomes, or idiosyncratic consumption dynamics, the frictionless expectations and sticky expectations models are virtually indistinguishable. The principal difference between the models is in the dynamic behavior of aggregate income and consumption, which are substantially less volatile in the sticky than in the frictionless economy.

3.4 The Utility Costs of Sticky Expectations

Our measure of the cost of stickiness is the answer to the following question: How much would the typical perfectly informed frictionless consumer be willing to pay to avoid being forced to switch to the behavior of the sticky expectations agent?

The answer will depend, in principle, on the initial level of resources. Formally, the normalized end-of-period value functions are

$$\begin{aligned} \mathbf{v}(a_t) &= \beta \mathbb{E}_t[\boldsymbol{\psi}_{t+1}^{1-\rho} \mathbf{v}_{t+1}(m_{t+1})] \\ \bar{\mathbf{v}}(a_t, 1) &= \beta \mathbb{E}_t[\boldsymbol{\psi}_{t+1}^{1-\rho} \bar{\mathbf{v}}_{t+1}(m_{t+1}, Q_{t+1})] \end{aligned}$$

where the second argument in the sticky-expectations value function is the ratio of perceived to actual aggregate permanent income (which is a state variable for the purposes of computing the sticky expectations consumer's value). If the initial level of permanent productivity for the frictionless consumer is normalized to 1, we seek the value of actual permanent income $\hat{\mathbf{p}}$ such that

$$\hat{\mathbf{p}}^{1-\rho} \mathbf{v}(a_t/\hat{\mathbf{p}}, \hat{\mathbf{p}}) = \mathbf{v}(a_t),$$

and the cost of stickiness can be measured by $\hat{\mathbf{p}} - 1$, because this is the increment to permanent income (relative to a benchmark of $\mathbf{p} = 1$) such that, given initial a_t , expected value from behaving as a sticky expectations consumer exactly matches expected value from behaving as a frictionless expectations consumer.

As the definition makes clear, the cost of stickiness will depend on the frictionless consumer's end-of-period assets. The bottom row of Table 2 reports the cost of stickiness for a frictionless consumer whose a_t is equal to the target value of a implied by the model. As the table shows, the compensation that a frictionless consumer would demand for being forced to behave in the manner of a sticky expectations consumer is very small for either version of the model, but smaller for the DSGE model than for the PE/SOE model.²⁰

4 Empirical Benchmark

This section presents an empirical benchmark that will guide our investigation of the implications of the model.

The random walk model provides the organizing framework around which both micro and macro literatures have been organized. Reinterpreted to incorporate CRRA utility and permit time-varying interest rates, the random walk proposition has frequently been formulated (in both micro and macro literatures) as a claim that $\mu = 0$ in regressions of the form

$$\Delta \log \mathbf{c}_{t+1} = \varsigma + \vartheta \mathbb{E}_t[\mathbf{r}_{t+1}] + \mu X_t + \epsilon_{t+1} \quad (42)$$

where X_t is any variable whose value was known to consumers at the time the period- t consumption decision was made, and ϵ_{t+1} is white noise.

For macroeconomic models, simulation analysis shows that the relationship between the normalized asset stock A_t and the expected interest rate $\mathbb{E}_t[\mathbf{r}_{t+1}]$ is very nearly linear, so (42) can be reformulated with no loss of statistical power as

$$\Delta \log \mathbf{c}_{t+1} = \varsigma + \alpha A_t + \mu X_t + \epsilon_{t+1}.$$

²⁰ Arguably it would be more appropriate to report the mean value of $\mathbf{p} - 1$ for the population of consumers distributed according to the ergodic distribution for a . When this is done, the conclusion is qualitatively unchanged.

This reformulation is convenient because the literatures on precautionary saving and liquidity constraints have argued that the effects of capital market imperfections can be captured by incorporating a lagged measure of resources like A_t in consumption growth regressions. (See below for further discussion).

Campbell and Mankiw (1989) famously proposed a modification of this model in which a proportion η of income goes to rule-of-thumb consumers who spend $C = Y$ in every period. They argued that η can be estimated by incorporating the predictable component of income growth as an additional regressor. Finally, Dynan (2000) and Sommer (2001) show that in some habit formation models, the size of the habit formation parameter can be captured by including lagged consumption growth as a regressor. These considerations lead to a benchmark specification of the form

$$\Delta \log \mathbf{c}_{t+1} = \varsigma + \chi \Delta \log \mathbf{c}_t + \eta \mathbb{E}_t[\Delta \log \mathbf{Y}_{t+1}] + \alpha A_t + \epsilon_{t+1} \quad (43)$$

There is of course an extensive existing literature on aggregate consumption dynamics, but Sommer (2001) is the only paper we are aware of that estimates an equation of precisely this form in aggregate data. Sommer (2001) interprets the serial correlation of consumption growth as reflecting habit formation.²¹ However, Sommer's choice of instruments, estimation methodology, and tests do not correspond precisely to our purposes for this paper. We therefore conducted a simple empirical exercise along the lines of Sommer's work and modified to correspond to the testable implications of our model for aggregate data. The results are reported in table 3.²²

Rather than attempting here a full recapitulation of the existing empirical literature, and describing the many choices that underlie the regressions reported in the table, we refer the reader to the careful and detailed discussion in Sommer (2001). Three points are worth emphasizing here.

First, while the existing empirical literature has tended to focus on spending on nondurables and services, there are reasons to be skeptical about the interpretation of quarterly dynamics (or lack of such dynamics) in large portions of the services component of measured spending.²³ Hence, we report results both for the traditional measure of nondurables and services spending, and for the more restricted category of nondurables spending alone. Fortunately, as the table shows, our results are robust to the measure of spending (indeed, similar results hold even when the measure of spending is total personal consumption expenditures, or an even stricter version of nondurables spending).²⁴

Second, Sommer (2001) emphasizes the importance of taking account of the effects of measurement error and transitory shocks (e.g., hurricanes) on high frequency consumption data. In principle, transitory shocks to the level of consumption could lead

²¹Weber (2002) makes a similar point using a different methodology.

²²An archive that produces this table, along with a variety of alternative estimations demonstrating the robustness of the main results, is available as part of the archive associated with this paper on the first author's website.

²³In particular, imputed rent on housing is the largest component of services spending, but even a careful examination of the published documentation on the construction of this data series leaves it unclear what causes quarterly variation in the imputations, since some of the main data sources seem to be annual; the same is true of many other types of services spending.

²⁴See Wilcox (1992) for a detailed discussion of the data measurement issues and their connection to theoretical questions.

Table 3 Aggregate Consumption Dynamics in US Data

$$\Delta \log \mathbf{C}_{t+1} = \varsigma + \chi \Delta \log \mathbf{C}_t + \eta \mathbf{E}[\Delta \log \mathbf{Y}_{t+1}] + \alpha A_t$$

Consumption Series	χ	η	α	Method		IV F p -val	
				OLS/IV	\bar{R}_2^2	IV	OID
Nondurables and Services							
	0.358*** (0.066)			OLS	0.123		
		0.577*** (0.118)		IV	0.172	0.000	0.702
			0.0006 (0.0006)	OLS	0.002		
	0.826*** (0.147)			IV	0.143	0.000	0.714
	0.731*** (0.230)	0.071 (0.118)	0.0000 (0.0003)	IV	0.135		0.482
Memo:	For instruments \mathbf{Z} , $\Delta \log \mathbf{C}_{t+1} = \mathbf{Z}\zeta$, $\bar{R}^2 = 0.168$						
Nondurables							
	0.247*** (0.072)			OLS	0.056		
		0.789*** (0.215)		IV	0.113	0.000	0.338
			0.0015* (0.0009)	OLS	0.016		
	0.736*** (0.162)			IV	0.087	0.000	0.685
	0.498** (0.226)	0.286* (0.153)	0.0005 (0.0006)	IV	0.093		0.868
Memo:	For instruments \mathbf{Z} , $\Delta \log \mathbf{C}_{t+1} = \mathbf{Z}\zeta$, $\bar{R}^2 = 0.122$						

Instruments: L(2/3).diffcons L(2/3).wyRatio L(2/3).bigTheta L(2/3).dfedfunds L(2/3).ics
Time frame: 1960Q1-2004Q3, $\sigma_{\Psi}^2 = .0000429$, $\sigma_{\Theta}^2 = .0000107$

Notes: Robust standard errors are in parentheses. The penultimate column reports the \bar{R}^2 from a regression of the dependent variable on the RHS variables (instrumented, when indicated); the final column reports two tests of instrument validity: The p -value from the Shea (1997) test of first-stage instrument validity (top), and the p -value from the Sargan overidentification test (bottom). {*, **, ***} = Statistical significance at {10, 5, 1} percent.

to a severe downward bias in the estimated serial correlation of measured consumption growth as distinct from ‘true’ consumption growth. The simplest solution to this problem is the classic response to measurement error in any explanatory variable: Instrumental variables estimation. This point is illustrated in the fact that instrumenting drastically increases the estimated serial correlation of consumption growth.

Finally, we needed to balance the desire for the empirical exercise to match the theory with the need for sufficiently powerful instruments. This would not be a problem if, in empirical work, we could use once-lagged instruments as is possible for the theoretical model. However, empirical consumption data are likely to be subject to time aggregation bias (Working (1960); Campbell and Mankiw (1989)), which can be remedied by lagging the time-aggregated instruments an extra period. But over the sample period available, two lags of the instruments generated by the theoretical model have only borderline statistical significance; hence the instrument set was augmented by two variables that have been shown to have strong predictive power for consumption growth: The Federal Funds rate and the University of Michigan’s Index of Consumer Sentiment (cf. Carroll, Fuhrer, and Wilcox (1994)). (An extensive literature has found a broad range of other variables with predictive power for spending growth; our experience is that results similar to those in the table can be obtained with any collection of instruments with a statistically robust predictive capacity for consumption growth).

The table demonstrates three principal points. First, when lagged consumption growth is excluded from the regression equation, the classic Campbell and Mankiw (1989) result holds: Consumption growth is strongly related to predictable income growth.

Second, when predictable income growth is excluded but lagged consumption growth is included, the serial correlation of consumption growth is estimated to be in the range of $0.7 \sim 0.8$, very far indeed from the benchmark random walk coefficient of zero.

Finally, in the ‘horserace’ regression in which predictable income growth and lagged consumption growth are both included, lagged consumption growth retains its statistical significance and large point estimate, while the predictable income growth term becomes statistically insignificant and economically small.

5 Simulated Empirical Tests

This section of the paper considers the results that an econometrician could expect to obtain from estimating an equation like (43) using data generated by the models described above.

5.1 Simulated Micro Empirical Results

Zeldes (1989a) pointed out long ago that the Euler equation on which the random walk proposition is based fails to hold for consumers who are liquidity constrained; if consumers with low levels of wealth relative to permanent income are more likely to be constrained, then low wealth consumers will experience systematically faster consumption growth than otherwise-similar high-wealth consumers. Precautionary saving

motives have a similar impact on consumption growth; households who have low levels of wealth will have felt the need to cut back sharply on their spending, and as wealth recovers, such households will experience a predictable recovery in their consumption.

These points are captured here by following Zeldes’s example and incorporating a dummy variable to measure low wealth status: $\underline{a}_{t,i}$ is equal to 0 if household i ’s level of a in period t is in the bottom 1 percent of the distribution, and $\underline{a}_{t,i} = 1$ for all other households.

Details of the simulation procedure are recounted in appendix C. The result of the simulations is a dataset analogous to the empirical datasets on which empirical microeconomic research has been conducted (the most directly comparable data source would be a consumer expenditure survey). On these simulated data, we estimate regressions of the form

$$\Delta \log \mathbf{c}_{t+1,i} = \varsigma + \chi \Delta \log \mathbf{c}_{t,i} + \eta \mathbb{E}_{t,i}[\Delta \log \mathbf{y}_{t+1,i}] + \alpha \underline{a}_{t,i} \quad (44)$$

Results are presented in Table 4.

For our purposes, the key conclusion from the table is that the predictable component of idiosyncratic consumption growth is very modest, even for the ‘sticky expectations’ version of the model. \bar{R}^2 ’s are below 1 percent for all specifications; and the predictive power of lagged consumption growth is always negligible.

In more detail, consider first the frictionless expectations version of the model in the top panel. Note first that the coefficient on expected income growth is zero. This contrasts with Carroll (2001)’s finding that when the instruments capture permanent differences in income growth across consumers, theory predicts $\eta \approx 1$. But the contrast is not a contradiction; instead, it illustrates the point that η is not a structural parameter: Its value will depend on the extent to which the instruments used are capturing transitory or permanent variations in predictable income growth. Here, they are capturing purely transitory variation; in Carroll (2001) they are capturing purely permanent differences.

Precautionary saving motives are captured by the robust statistical significance of α , which indicates that consumers with low levels of wealth in period t exhibit faster-than-average consumption growth in the next period. This reflects the fact that these consumers have had to depress their consumption for precautionary reasons in period t , but on average their income and wealth circumstances will improve sharply between t and $t+1$, relaxing the precautionary motive. However, despite the statistical significance of this effect, the \bar{R}^2 indicates that the magnitude of consumption predictability remains very modest.

When all three terms are combined, the results are what would be expected from the one-by-one regressions: Only the low-wealth indicator variable a is statistically significant, and even this variable has only very modest forecasting ability.

The lower panel contains results from estimating the same regressions on the sticky expectations version of the model. These results are virtually indistinguishable from those obtained for the frictionless expectations model. As before, aside from the precautionary component captured by a , idiosyncratic consumption growth is essentially unpredictable.

These results are broadly consistent with the substantial literature that has examined empirical micro data for habit formation effects. Most papers have found no evidence of habits (Naik and Moore (1996), Meghir and Weber (1996), Dynan (2000); Flavin and Nakagawa (2005));²⁵ using similar methods but arguably better data (from a multi-year expenditure survey in Spain), Carrasco, Labeaga, and López-Salido (2005) have recently found some evidence for habits in some commodities. But, taken as a whole, even this last paper (which seems to be the strongest evidence for habits in micro data) finds habit formation effects that are much smaller than those commonly used in the macroeconomic literature.

It should be admitted that a related literature, with prominent contributions by Souleles (1999), Parker (1999), and Johnson, Parker, and Souleles (2006), finds that household spending growth is related to predictable components of income growth (though see Hsieh (2003) or Coulibaly and Li (2006) for counterexamples). A possible route to reconciliation would be to admit some degree of predictability to spending growth, but to argue that measurement error or other noise biases the estimate of χ downward; alternatively, Browning and Collado (2001), who summarize the literature nicely, argue that the best way to reconcile the varying microeconomic findings is to suppose that consumers are not always fully aware of the predictable components of their incomes, an explanation that is consistent with the theme (if not the details) of this paper. Our broader point, however, is that only a very small proportion of the changes in micro consumption expenditures seems to be predictable, in contrast with the robust predictability of macroeconomic consumption growth. This point would survive even if the most generous estimates of predictability in the micro literature were adopted.

5.2 Simulated Small Open Economy Empirical Estimation

Our small open economy model tracks the aggregate dynamics of an economy filled with consumers behaving according to exactly the same model examined at the micro level in the previous section.

Now, however, we simulate the model over a longer time frame, 20 quarters. In principle, an appropriate procedure would be to simulate over many subsamples corresponding to, say, 25 years, then to report the variation in estimates across the subsamples; in practice, the amount of variation in model estimation results across such subsamples is small, and so the gains from such an extended Monte Carlo exercise are not worth the expositional cost.

The second modification is a consequence of the first: Given the longer time frame, and given that the idiosyncratic shocks to income are washed away by the law of large numbers, it is now feasible to use instrumental variables techniques to obtain the coefficient on the expected growth term. This is the appropriate procedure for comparison with empirical results in any case, since instrumental variables estimation is the standard way of estimating the benchmark Campbell–Mankiw model. As instruments, we use the appropriate lags of $\Delta \log \mathbf{C}_t$, $\Delta \log \mathbf{Y}_t$, A_t , and, following standard practice since

²⁵The maximum point in Dynan's 95 percent confidence interval for the habits parameter is 0.15.

Campbell and Deaton (1989) and Campbell (1987), the aggregate personal saving rate $s_t = (\mathbf{Y}_t + (r/R)\mathbf{B}_t - \mathbf{C}_t)/\mathbf{C}_t$ where we normalize by consumption rather than income because consumption is smoother than income.²⁶

Finally, for comparison to empirical results we take into account Sommer (2001)'s argument (based on Wilcox (1992)) that both transitory components of aggregate spending (e.g. hurricanes) and measurement problems introduce transitory components in measured NIPA spending compared to the level of spending that would prevail if the model as presented were literally true.²⁷ We adopt Sommer (2001)'s specification in which the level of consumption is disturbed by a white noise error with a variance about 1/2 as large as the estimated variance of 'true' consumption innovations.²⁸ That is, we suppose that measured consumption

$$\tilde{\mathbf{c}}_{t+1} = \mathbf{c}_{t+1}\epsilon_{t+1} \quad (45)$$

where $\mathbb{E}_t[\epsilon_{t+n}] = 0 \forall n$ and $\text{var}(\log \epsilon)$ is calibrated from Sommer (2001)'s empirical estimates.

We present results for the frictionless model assuming that consumption is measured without error, and for the sticky expectations model assuming that it is measured with error. (The frictionless model with error has the counterfactual prediction that consumption growth has strongly negative serial correlation; the sticky model without error implies an excessively large serial correlation coefficient of 0.75 in *measured* consumption growth). Thus, for the two models we estimate respectively the two equations

$$\Delta \log \mathbf{c}_{t+1} = \varsigma + \chi \Delta \mathbb{E}[\log \mathbf{c}_t] + \eta \mathbb{E}[\Delta \log \mathbf{Y}_{t+1}] + \alpha \mathbb{E}[A_t] \quad (46)$$

$$\Delta \log \underbrace{\tilde{\mathbf{c}}_{t+1}}_{=\epsilon_{t+1}\mathbf{c}_{t+1}} = \varsigma + \tilde{\chi} \Delta \mathbb{E}[\log \tilde{\mathbf{c}}_t] + \eta \mathbb{E}[\Delta \log \bar{\mathbf{Y}}_{t+1}] + \alpha \mathbb{E}[\bar{A}_t] \quad (47)$$

Table 5 reports the results. We simulate the model 20 periods (after a "presample" period with no aggregate shocks to let the model reach the ergodic state).

Note that the (tiny) standard errors reflect the long sample simulation period; in the limit if we were to simulate an arbitrarily large number of periods the standard errors would go to zero. Thus, these errors are not appropriate for comparison with the standard errors in an empirical table like 3, where the standard errors reflect a much smaller sample size. An appendix table reports the results of a Monte Carlo exercise where we divide our full sample into nonoverlapping subsamples of size approximately equal to the size of the empirical dataset, and compute regression statistics for each sample in the same manner in which they are computed in empirical data. The conclusion from that analysis is that the empirical estimates are in line with the predictions of the model.

For the frictionless version of the model, estimation on aggregate data largely con-

²⁶Estimation of the sticky expectations model always uses the corresponding appropriate variables, e.g. \bar{s}_t rather than s_t .

²⁷It is worth pointing out here that in Friedman (1957)'s original statement of the Permanent Income Hypothesis, transitory shocks to expenditures were given equal billing with transitory shocks to income. The subsequent literature deemphasized expenditure shocks, perhaps inappropriately.

²⁸Sommer (2001), Table 3.

firms the results from idiosyncratic data: Consumption growth remains almost entirely unpredictable, as per the random walk framework.

We examine the sticky expectations model in two ways.

First, we simply compute the serial correlation coefficient on lagged consumption growth (first row, second panel of the table). We find that this coefficient is very close to the $(1 - \Pi) = 0.75$ figure that reflects the proportion of consumers who do not adjust their expectations in any period. The bottom line is that the intuition derived from the ‘toy’ quadratic utility version of the model survives all the subsequent complications and elaborations.

Next, we introduce measurement error in consumption *a la* Sommer (2001). This is necessary because Sommer finds that the measurement error produces a severe downward bias in the estimate of the serial correlation in consumption growth, relative to the ‘true’ serial correlation. This point is confirmed in our context by a comparison of the second row of the panel, which presents an OLS estimate of the serial correlation coefficient, and the third row, which presents an IV estimate designed to correct for the measurement error problem. In fact, the regression coefficient after instrumenting is actually slightly larger than the ‘true’ serial correlation coefficient for consumption growth reported in the first row of the panel.

The next row reflects what would have been found by Campbell and Mankiw had they estimated their model on data produced by the simulated economy. The coefficient on predictable component of perceived income growth term is large and highly statistically significant (though note that its predictive power for the second stage regression is far lower than the \bar{R}^2 for the first model).²⁹ The last row of the table presents the ‘horserace’ between the CM model and the sticky expectations model; the results indicate that the dynamics of consumption are dominated by the serial correlation in the predictable component of consumption growth stemming from the stickiness of expectations. This can be seen not only from the magnitude of the coefficients, but also by comparison of the second-stage \bar{R}_2^2 ’s, which indicate that the contribution of predictable income growth to the predictability of consumption growth is negligible. Note, however, that when the predictable income growth term is included in the regression, the coefficient on the lagged consumption growth term is no longer indistinguishable from $(1 - \Pi)$. This reflects a modest predictable component of aggregate consumption growth associated with the consumption out of misperceived transitory income by the nonupdating consumers; details are uninteresting, but the lesson for empirical work is that it may be better to omit the predictable component of income growth altogether than to include it in estimation efforts of this kind.

5.3 Simulated DSGE Model Estimation

To generate simulated data, the DSGE model was initialized in period 0 with the representative agent assumed to own a capital stock equal to the target capital stock

²⁹The predictability in perceived income growth reflects both the ‘catching up’ of perceived income to permanent innovations, and the disappearance of the transitory shock to the level of income. Most of the predictability in consumption growth stems from its relationship to the first of these terms.

implied by the converged consumption rule. A sequence of 100 transitory and permanent aggregate shocks was drawn according to the distributional assumptions specified in table 1, and the representative agent was assumed to behave according to the converged consumption rule, thereby generating model-consistent sequences of $\{c, Y, A\}$ and other model variables.

The results of estimating (46) and (47) on these data are reported in table 6.

Results are substantially the same as for the PE/SOE version of the model: The model with frictionless expectations implies aggregate consumption growth that is essentially a random walk, while the model with sticky expectations implies a serial correlation coefficient of consumption growth of about 0.75. The horserace regression again indicates that the success of the Campbell-Mankiw specification reflects the correlation of predicted current income growth with instrumented lagged consumption growth.

6 Conclusions

A model in which consumers perceive their idiosyncratic circumstances frictionlessly but update their expectations about macroeconomic events only occasionally has several advantages over standard models of aggregate consumption dynamics.

Perhaps the most appealing feature of the model is its connection with many other macroeconomic literatures that have found that for many variables, macroeconomic dynamics are more sluggish than implied by benchmark rational expectations representative agent macroeconomic models. This insight has been at the center of many papers outside of the consumption literature over the past few years (Sims (2003); Woodford (2001); Mankiw and Reis (2002)), while the macro consumption literature has tended to attribute consumption sluggishness to habits. Sticky expectations present a unified, tractable, simple methodology for generating sluggishness that appears to work well both in the consumption context and elsewhere.

Another, related, appeal of the model is its ability to reconcile the micro and macro data on consumption dynamics. It does so by emphasizing a point that will be familiar to anyone who has worked with both micro and macro data: Idiosyncratic variation is vastly greater than aggregate variation. This means that small imperfections like the one proposed here in macroeconomic perceptions have very modest utility consequences. Indeed, the sticky expectations assumption is arguably more attractive in the consumption context than in other areas where it has been proposed (and increasingly used) precisely because in the consumption context there is a well-defined utility-based metric for calculating how costly the stickiness is (in contrast, say, with models in which households' inflation expectations are sticky).

A final attraction of the model is tractability. In comparison with habit formation models or frameworks in which agents perform an optimal signal extraction problem, the basic building block of this model is the solution to the frictionless rational expectations model, for which a vast literature exists and time-tested solution methods are available.

Table 4 Typical Micro Consumption Estimation on Simulated Data

$$\Delta \log \mathbf{c}_{t+1,i} = \varsigma + \chi \Delta \log \mathbf{c}_{t,i} + \eta \mathbf{E}_{t,i}[\Delta \log \mathbf{y}_{t+1,i}] + \alpha \underline{a}_{t,i}$$

Model of Expectations	χ	η	α	\bar{R}^2	nobs
Frictionless					
	0.083 (0.077)			0.007	76020
		0.003 (0.004)		-0.000	76020
			-0.111 (0.052)	0.000	76020
	0.083 (0.004)	0.009 (0.004)	-0.059 (0.024)	0.007	76020
Sticky					
	0.084 (0.077)			0.007	76020
		0.003 (0.004)		-0.000	76020
			-0.111 (0.051)	0.000	76020
	0.083 (0.004)	0.009 (0.004)	-0.059 (0.024)	0.007	76020

Notes: $\mathbf{E}_{t,i}$ is the expectation from the perspective of person i in period t ; \underline{a} is a dummy variable indicating that agent i is in the top 99 percent of the a distribution. Heteroskedasticity-robust standard errors are in parentheses. Standard tests detect no serial correlation in the residuals. Sample is restricted to households with positive income in period t .

Table 5 Small Open Economy Aggregate Consumption Dynamics

Expectations:Dep Var			OLS or IV	2nd Stage \bar{R}^2	IV F p -val IV OID
Independent Variables					
Frictionless: $\Delta \log \mathbf{C}_{t+1}$					
$\Delta \log \mathbf{C}_t$	$\Delta \log \mathbf{Y}_{t+1}$	A_t			
0.040 (0.082)			OLS	-0.005	
	0.123 (0.155)		IV	-0.004	0.001 0.484
		0.0035 (0.0042)	OLS	-0.002	
0.064 (0.078)	0.077 (0.171)	0.0041 (0.0041)	IV	-0.004	
Sticky					
$\Delta \log \bar{\mathbf{C}}_t$	$\Delta \log \bar{\mathbf{Y}}_{t+1}$	\bar{A}_t			
0.741 (0.055)			OLS	0.543	
$\Delta \log \tilde{\mathbf{C}}_t$					
0.328 (0.082)			OLS	0.089	
0.795 (0.118)			IV	0.325	0.000 0.510
	1.222 (0.126)		IV	0.332	0.000 0.391
		0.0153 (0.0055)	OLS	0.042	
0.360 (0.282)	0.739 (0.425)	-0.0029 (0.0048)	IV	0.330	0.334
Memo: For instruments \mathbf{Z}_t , $\Delta \log \mathbf{C}_{t+1} = \mathbf{Z}_t \zeta$, $R^2 =$				0.329	

Notes: Model was simulated for 20 periods (quarters); to generate results comparable to the roughly 40 year span of U.S. empirical data, the table reports mean outcomes across nonoverlapping 16 period subsamples. Bars indicate the sticky expectations model data, and \sim indicates the presence of introduced measurement error as discussed in the text. ‘IV%(all)’ indicates instruments that include lags of $\Delta \log \mathbf{C}_t$, $\Delta \log \mathbf{Y}_t$, A_t and Θ_t (resp. $\Delta \log \bar{\mathbf{C}}_t$, $\Delta \log \bar{\mathbf{Y}}_t$, \bar{A}_t and $\bar{\Theta}_t$). The average robust standard across the simulations is presented in parentheses. The penultimate column reports the \bar{R}^2 from a regression of the dependent variable on the RHS variables (instrumented, when indicated); the final column reports two tests of instrument validity: The p -value from the Shea (1997) test of first-stage instrument validity (top), and the p -value from the Sargan overidentification test (bottom).

Table 6 Aggregate Consumption Dynamics in DSGE Model

Expectations:Dep Var Independent Variables			OLS or IV	2nd Stage \bar{R}^2	IV F p -val IV OID
Frictionless: $\Delta \log \mathbf{C}_{t+1}$					
$\Delta \log \mathbf{C}_t$	$\Delta \log \bar{\mathbf{Y}}_{t+1}$	A_t			
0.010 (0.032)			OLS	-0.001	
	0.184 (0.050)		IV	0.007	0.000 0.001
		-0.0002 (0.0001)	OLS	0.010	
-0.019 (0.027)	0.152 (0.052)	-0.0002 (0.0001)	IV	0.007	
Sticky					
$\Delta \log \bar{\mathbf{C}}_t$	$\Delta \log \bar{\mathbf{Y}}_{t+1}$	\bar{A}_t			
0.823 (0.018)			OLS	0.677	
$\Delta \log \tilde{\mathbf{C}}_t$					
0.387 (0.030)			OLS	0.141	
0.845 (0.042)			IV	0.422	0.000 0.210
	0.815 (0.025)		IV	0.395	0.000 0.000
		-0.0004 (0.0000)	OLS	0.115	
0.750 (0.148)	0.065 (0.146)	-0.0001 (0.0000)	IV	0.423	0.126
Memo: For instruments \mathbf{Z}_t , $\Delta \log \mathbf{C}_{t+1} = \mathbf{Z}_t \zeta$, $R^2 =$				0.425	

See notes for previous table.

References

- AKERLOF, GEORGE A., AND JANET L. YELLEN (1985): “A Near-Rational Model of the Business Cycle, with Wage and Price Intertia,” *The Quarterly Journal of Economics*, 100(5), 823–38.
- BLANCHARD, OLIVIER J. (1985): “Debt, Deficits, and Finite Horizons,” *Journal of Political Economy*, 93(2), 223–247.
- BOLDRIN, MICHELE, LAWRENCE J. CHRISTIANO, AND JONAS D. FISHER (2001): “Habit Persistence, Asset Returns and the Business Cycle,” *American Economic Review*, 91(1), 149–66.
- BROWNING, MARTIN, AND M. DOLORES COLLADO (2001): “The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates,” *American Economic Review*, 91(3), 681–692.
- BROWNING, MARTIN, AND THOMAS F. CROSSLEY (2001): “The Life-Cycle Model of Consumption and Savings,” *Journal of Economic Perspectives*, 15(3), 3–22.
- CALVO, GUILLERMO A. (1983): “Staggered Contracts in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12, 383–98.
- CAMPBELL, JOHN Y. (1987): “Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis,” *Econometrica*, 55, 1249–73.
- CAMPBELL, JOHN Y., AND ANGUS S. DEATON (1989): “Why Is Consumption So Smooth?,” *Review of Economic Studies*, 56, 357–74, <http://ideas.repec.org/a/bla/restud/v56y1989i3p357-73.html>.
- CAMPBELL, JOHN Y., AND N. GREGORY MANKIW (1989): “Consumption, Income, and Interest Rates: Reinterpreting the Time-Series Evidence,” in *NBER Macroeconomics Annual, 1989*, ed. by Olivier J. Blanchard, and Stanley Fischer, pp. 185–216. MIT Press, Cambridge, MA, <http://ideas.repec.org/p/nbr/nberwo/2924.html>.
- CARRASCO, RAQUEL, JOSÈ M. LABEAGA, AND J. DAVID LÒPEZ-SALIDO (2005): “Consumption and Habits: Evidence from Panel Data,” *Economic Journal*, 115(500), 144–165, available at <http://ideas.repec.org/a/ecj/econjl/v115y2005i500p144-165.html>.
- CARROLL, CHRISTOPHER D. (2001): “Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation),” *Advances in Macroeconomics*, 1(1), Article 6, <http://econ.jhu.edu/people/ccarroll/death.pdf>.
- (2003): “Macroeconomic Expectations of Households and Professional Forecasters,” *Quarterly Journal of Economics*, 118(1), 269–298, <http://econ.jhu.edu/people/ccarroll/epidemiologyQJE.pdf>.

- (2006a): “The Epidemiology of Macroeconomic Expectations,” in *The Economy as an Evolving Complex System, III*, ed. by Larry Blume, and Steven Durlauf. Oxford University Press, <http://econ.jhu.edu/people/ccarroll/epidemiologySFI.pdf>.
- (2006b): “The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems,” *Economics Letters*, pp. 312–320, <http://econ.jhu.edu/people/ccarroll/EndogenousArchive.zip>.
- (2011): “Theoretical Foundations of Buffer Stock Saving,” *Manuscript, Department of Economics, Johns Hopkins University*, <http://econ.jhu.edu/people/ccarroll/papers/BufferStockTheory.pdf>.
- CARROLL, CHRISTOPHER D., JEFFREY C. FUHRER, AND DAVID W. WILCOX (1994): “Does Consumer Sentiment Forecast Household Spending? If So, Why?,” *American Economic Review*, 84(5), 1397–1408, <http://econ.jhu.edu/people/ccarroll/SentAERC CarrollFuhrerWilcox.pdf>.
- CARROLL, CHRISTOPHER D., JODY R. OVERLAND, AND DAVID N. WEIL (2000): “Saving and Growth with Habit Formation,” *American Economic Review*, 90(3), 341–355, <http://econ.jhu.edu/people/ccarroll/AERHabits.pdf>.
- CARROLL, CHRISTOPHER D., AND ANDREW A. SAMWICK (1997): “The Nature of Precautionary Wealth,” *Journal of Monetary Economics*, 40(1), 41–71, <http://econ.jhu.edu/people/ccarroll/nature.pdf>.
- CHARI, V. V., PATRICK J. KEHOE, AND ELLEN R. MCGRATTAN (2005): “A Critique of Structural VARs Using Real Business Cycle Theory,” working paper 631, Federal Reserve Bank of Minneapolis.
- CHRISTIANO, LAURENCE J., MARTIN EICHENBAUM, AND CHARLES L. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113(1), 1–45.
- COCHRANE, JOHN H. (1991): “The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives,” *American Economic Review*, 79, 319–37.
- CONSTANTINIDES, GEORGE M. (1990): “Habit Formation: A Resolution of the Equity Premium Puzzle,” *Journal of Political Economy*, 98, 519–43.
- COULIBALY, BRAHIMA, AND GENG LI (2006): “Do Homeowners Increase Consumption after the Last Mortgage Payment? An Alternative Test of the Permanent Income Hypothesis,” *The Review of Economics and Statistics*, 88(1), 10–19.
- DEATON, ANGUS S. (1992): *Understanding Consumption*. Oxford University Press, New York.

- DYNAN, KAREN E. (2000): “Habit Formation in Consumer Preferences: Evidence from Panel Data,” *American Economic Review*, 90(3).
- FLAVIN, MARJORIE J., AND SHINOBU NAKAGAWA (2005): “A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence,” *Manuscript UCSD*.
- FRIEDMAN, MILTON A. (1957): *A Theory of the Consumption Function*. Princeton University Press.
- FUHRER, JEFFREY C. (2000): “An Optimizing Model for Monetary Policy: Can Habit Formation Help?,” *American Economic Review*, 90(3).
- GABAIX, XAVIER, AND DAVID LAIBSON (2001): “The 6D Bias and the Equity Premium Puzzle,” *NBER Macroeconomics Annual*, pp. 257–312.
- GUARIGLIA, ALESSANDRA, AND MARIACRISTINA ROSSI (2002): “Consumption, Habit Formation and Precautionary Saving: Evidence from the British Household Panel Survey,” *Oxford Economic Papers*, 54, 1–19.
- HALL, ROBERT E. (1978): “Stochastic Implications of the Life-Cycle/Permanent Income Hypothesis: Theory and Evidence,” *Journal of Political Economy*, 96, 971–87, Available at <http://www.stanford.edu/~rehall/Stochastic-JPE-Dec-1978.pdf>.
- HSIEH, CHANG-TAI (2003): “Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund,” *American Economic Review*, 99, 397–405.
- JERMANN, URBAN J. (1998): “Asset Pricing in Production Economies,” *Journal of Monetary Economics*, 41(2), 257–75.
- JOHNSON, DAVID S., JONATHAN A. PARKER, AND NICHOLAS S. SOULELES (2006): “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 96(5), 1589–1610.
- KRUSELL, PER, AND ANTHONY A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), 867–896.
- LOW, HAMISH, COSTAS MEGHIR, AND LUIGI PISTAFERRI (2005): “Wage Risk and Employment Over the Life Cycle,” *Manuscript, Stanford University*.
- LUCAS, ROBERT E. (1973): “Some International Evidence on Output-Inflation Tradeoffs,” *American Economic Review*, 63, 326–334.
- MANKIW, N. GREGORY, AND RICARDO REIS (2002): “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 117(4), 1295–1328.

- MEGHIR, COSTAS, AND GUGLIELMO WEBER (1996): “Intertemporal Non-Separability or Borrowing Restrictions? A Disaggregate Analysis Using a U.S. Consumption Panel,” *Econometrica*, 64(5), 1151–82.
- MORRIS, STEPHEN, AND HYUNG SOHN SHIN (2005): “The Inertia of Forward Looking Expectations,” *Manuscript, London School of Economics*.
- MUTH, JOHN F. (1960): “Optimal Properties of Exponentially Weighted Forecasts,” *Journal of the American Statistical Association*, 55(290), 299–306.
- NAIK, NARAYAN, AND MICHAEL MOORE (1996): “Habit Formation and Intertemporal Substitution in Individual Food Consumption,” *Review of Economics and Statistics*, 78(2), 321–328.
- NIELSEN, HELENA SKYT, AND ANNETTE VISSING-JORGENSEN (2006): “The Impact of Labor Income Risk on Educational Choices: Estimates and Implied Risk Aversion,” *Manuscript*.
- PARKER, JONATHAN (1999): “The Reaction of Household Consumption to Predictable Changes in Social Security Taxes,” *American Economic Review*, Vol 89(4), 959–73.
- PARKER, JONATHAN A. (2003): “Consumption Risk and Expected Stock Returns,” *American Economic Review Papers and Proceedings*, 93(2), Available at <http://ideas.repec.org/p/nbr/nberwo/9548.html>.
- PARKER, JONATHAN A., AND BRUCE PRESTON (2005): “Precautionary Saving and Consumption Fluctuations,” *American Economic Review*, 95(4), 1119–1143.
- PISCHKE, JÖRN-STEFFEN (1995): “Individual Income, Incomplete Information, and Aggregate Consumption,” *Econometrica*, 63(4), 805–40.
- REIS, RICARDO (2003): “Inattentive Consumers,” *Manuscript, Harvard University*.
- ROTEMBERG, JULIO J., AND MICHAEL WOODFORD (1997): “An Optimization-Based Econometric Model for the Evaluation of Monetary Policy,” in *NBER Macroeconomics Annual, 1997*, ed. by Benjamin S. Bernanke, and Julio J. Rotemberg, vol. 12, pp. 297–346. MIT Press, Cambridge, MA.
- SHEA, JOHN (1997): “Instrument Relevance in Multivariate Linear Models: A Simple Measure,” *Review of Economics and Statistics*, 49(2), 348–352.
- SIMS, CHRISTOPHER A. (2003): “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50(3), 665–690, available at <http://ideas.repec.org/a/eee/moneco/v50y2003i3p665-690.html>.
- SOMMER, MARTIN (2001): “Habits, Sentiment and Predictable Income in the Dynamics of Aggregate Consumption,” *Manuscript, The Johns Hopkins University*, <http://econ.jhu.edu/pdf/papers/wp458.pdf>.

- SOULELES, NICHOLAS S. (1999): "The Response of Household Consumption to Income Tax Refunds," *American Economic Review*, 89(4), 947–958.
- STORESLETTEN, KJETIL, CHRIS I. TELMER, AND AMIR YARON (2004): "Consumption and Risk Sharing Over the Life Cycle," *Journal of Monetary Economics*, 51(3), 609–633.
- WEBER, CHRISTIAN E. (2002): "Intertemporal Non-Separability and 'Rule-Of-Thumb' Consumption," *Journal of Monetary Economics*, 49, 293–308.
- WILCOX, DAVID W. (1992): "The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work," *American Economic Review*, 82(4), 922–941.
- WOODFORD, MICHAEL (2001): "Imperfect Common Knowledge and the Effects of Monetary Policy," *NBER Working Paper Number 8673*.
- WORKING, HOLBROOK (1960): "Note on the Correlation of First Differences of Averages in a Random Chain," *Econometrica*, 28(4), 916–18.
- ZELDES, STEPHEN P. (1989a): "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 97, 305–46.
- (1989b): "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, 104(2), 275–298.

A Quadratic Utility Consumption Dynamics

This appendix derives the equation (9) asserted in the main text. Start with the definition of consumption for the updaters,

$$\mathbf{C}_t^\pi \equiv \Pi^{-1} \int_0^1 \pi_{t,i} \mathbf{c}_{t,i} di \quad (48)$$

$$= \Pi^{-1} \int_0^1 \pi_{t,i} (r/R) \mathbf{o}_{t,i} di \quad (49)$$

$$= \Pi^{-1} (r/R) \int_0^1 \pi_{t,i} \mathbf{o}_{t,i} di \quad (50)$$

$$= \Pi^{-1} (r/R) \Pi \mathbf{O}_t \quad (51)$$

$$= (r/R) \mathbf{O}_t \quad (52)$$

where the penultimate line follows from the fact that the updaters are chosen randomly among members of the population so that the average per capita value of \mathbf{o} among updaters is equal to the average per capita value of \mathbf{o} for the population as a whole.

The text asserts (cf. (9)) that

$$\mathbf{C}_{t+1} = \Pi \Delta \mathbf{C}_{t+1}^\pi + (1 - \Pi) \Delta \mathbf{C}_t \quad (53)$$

$$\approx (1 - \Pi) \Delta \mathbf{C}_t + \xi_{t+1} \quad (54)$$

To see this, write

$$\mathbf{C}_{t+1}^\pi = (r/R) [M_{t+1} + H_{t+1}] \quad (55)$$

$$\mathbf{C}_t^\pi = (r/R) [M_t + H_t] \quad (56)$$

$$\mathbf{C}_{t+1}^\pi - \mathbf{C}_t^\pi = (r/R) [M_{t+1} - M_t + H_{t+1} - H_t] \quad (57)$$

$$\mathbf{C}_{t+1}^\pi - \mathbf{C}_t^\pi = (r/R) [R(Y_t + M_t - \mathbf{C}_t) - M_t + H_{t+1} - H_t]. \quad (58)$$

What theory tells us is that *if aggregate consumption were chosen frictionlessly in period t* then this expression would be white noise; that is, we know that

$$(r/R) [R(Y_t + M_t - \mathbf{C}_t) - M_t + H_{t+1} - H_t] = \xi_{t+1} \quad (59)$$

for some white noise ξ_{t+1} . The only difference between this expression and the RHS of (58) is the Π superscript on the \mathbf{C}_t . Thus, substituting, we get

$$\mathbf{C}_{t+1}^\pi - \mathbf{C}_t^\pi = (r/R) [R(Y_t + M_t - (\mathbf{C}_t + \mathbf{C}_t^\pi - \mathbf{C}_t^\pi)) - M_t + H_{t+1} - H_t] \quad (60)$$

$$\begin{aligned} \mathbf{C}_{t+1}^\pi - \mathbf{C}_t^\pi &= (r/R) [R(Y_t + M_t - \mathbf{C}_t) - M_t + H_{t+1} - H_t] + (r/R) (\mathbf{C}_t^\pi - \mathbf{C}_t) \\ &= \xi_{t+1} + (r/R) (\mathbf{C}_t^\pi - \mathbf{C}_t). \end{aligned} \quad (61)$$

So equation (9) can be rewritten as

$$\Delta \mathbf{C}_{t+1} = (1 - \Pi) \Delta \mathbf{C}_t + \Pi ((r/R) (\mathbf{C}_t^\pi - \mathbf{C}_t) + \xi_{t+1}) \quad (62)$$

where ξ_{t+1} is a white noise variable. Thus,

$$\Delta \mathbf{C}_{t+1} = (1 - \Pi) \left(1 + \underbrace{(r/R)}_{\approx 0} \right) \Delta \mathbf{C}_t + \underbrace{\Pi \xi_{t+1}}_{\equiv \epsilon_{t+1}} \quad (63)$$

for a white noise variable ϵ_{t+1} , and $(r/R) \approx 0$ for plausible quarterly interest rates. (63) leads directly to (9).

B Calibration of Aggregate Income Process

We calibrate our aggregate income process using U.S. NIPA labor income, constructed as wages and salaries plus transfers minus personal contributions for social insurance.

We calibrate our aggregate income process using U.S. NIPA labor income, constructed as wages and salaries plus transfers minus personal contributions for social insurance.

This appendix describes how we construct the aggregate transitory shock Θ_t , used as one of the instruments. Following the theoretical structure of the model, we assume that the log of labor income follows the process

$$\log \mathbf{Y}_{t+1} = \log P_{t+1} + \log \Theta_{t+1}, \quad (64)$$

$$\log P_{t+1} = \log \Lambda + \log P_t + \log \Psi_{t+1}. \quad (65)$$

To extract Θ_t we apply the Kalman filter algorithm. Due to difficulties in estimating the parameters Λ , σ_Θ^2 and σ_Ψ^2 directly by maximum likelihood, we use the following procedure. We first calibrate the signal-to-noise ratio $\tau \equiv \sigma_\Psi^2 / \sigma_\Theta^2$ so that the first autocorrelation of the process, generated using (64)–(65), is 0.96.³⁰ Differencing equation (64) and expressing the second moments yields

$$\text{var}(\Delta \log \mathbf{Y}_{t+1}) = \sigma_\Psi^2 + 2\sigma_\Theta^2 \quad (66)$$

$$= (\tau + 2)\sigma_\Theta^2. \quad (67)$$

Given $\text{var}(\Delta \log \mathbf{Y}_{t+1})$ and τ we identify $\sigma_\Theta^2 = \text{var}(\Delta \log \mathbf{Y}_{t+1}) / (\tau + 2)$ and $\sigma_\Psi^2 = \tau \sigma_\Theta^2$. The strategy yields the following estimates: $\tau = 4$, $\sigma_\Psi^2 = 4.29 \times 10^{-5}$ and $\sigma_\Theta^2 = 1.07 \times 10^{-5}$. Finally, $\log \Lambda$ is estimated as 0.515%, the average quarterly growth of income. For these values we extract an estimate of transitory component of income by the Kalman filter, $\Theta_t = \Theta_{t|t}$.

This parametrization of the aggregate income process has two attractive features. First, our model of income is consistent with assumptions in papers by Jermann (1998), Boldrin, Christiano, and Fisher (2001), and Chari, Kehoe, and McGrattan (2005), considered standard exercises in the RBC literature. These authors model the state of technology as persistent AR(1) process (or random walk), whose autocorrelation properties closely correspond to our specification of income. Second, the transitory components of income Θ_t are strongly correlated for different choices of τ . This in turn

³⁰ We generate 10,000 replications of a process with 180 observations, which corresponds to 45 years of quarterly observations used in Table 6 in the main text. The mean and median first autocorrelations (across replications) of such process with $\tau = 4$ are 0.956 and 0.965, respectively. In comparison, the mean and median of sample first autocorrelations of pure random walk are 0.970 and 0.977 (with 180 observations), respectively.

implies that our IV estimates are robust to the choice of τ (as long as the variance of the permanent component σ_{Θ}^2 is large enough).

C Simulation Procedures

This appendix describes details of the simulation procedure used to create the dataset on which the results in Table 2 and Table 4 are based.

A population with members indexed by i is constructed in which households are randomly assigned initial (period 0) levels of idiosyncratic permanent income $p_{0,i}$ from a lognormal distribution whose variance matches the steady-state population variance of p . Initial $k_{0,i}$ for every household is equal to the target k implied by the converged consumption rule for the model as described above. Given the levels of permanent income and initial k , idiosyncratic and aggregate shocks are drawn as per the model's assumptions above, and each household obeys the converged consumption rule to determine consumption and saving behavior.

Results are kept track of in the variables $c_{0,i}, c_{1,i}$, and so on for a total simulation period of 40 quarters (ten years, a long time compared to most micro panel data on consumption). One detail is worth emphasizing: When making their forecasts of expected income growth, households are assumed to forecast that the transitory component of income will grow by the factor $1/\theta_{t,i}$, which is the forecast implied by their beliefs about the levels of the idiosyncratic and aggregate transitory components of income. This assumption is made in order to avoid the small sample problems that would arise in attempting to estimate a forecast of income growth using instrumental variables over such a short panel (only 40 periods of time series variation). Substantively, this point reflects the real-world fact that essentially all of the predictable variation in income growth at the household level comes from idiosyncratic components of income.

D Population Variance of Idiosyncratic Permanent Income

This appendix computes dynamics and steady state of the square of the idiosyncratic component of permanent income (from which the variance can be derived).

$$p_{t+1,i} = p_{t,i}\psi_{t+1,i}\omega_{t+1,i} + (1 - \omega_{t+1,i}) \quad (68)$$

$$p_{t+1,i}^2 = (p_{t,i}\psi_{t+1,i}\omega_{t+1,i})^2 + 2p_{t,i}\psi_{t+1,i}\underbrace{(1 - \omega_{t+1,i})\omega_{t+1,i}}_{=0} + (1 - \omega_{t+1,i})^2 \quad (69)$$

and since $\mathbb{E}_t[(1 - \omega_{t+1,i})^2] = \mathbb{E}_t[(1 - 2\omega_{t+1,i} + \omega_{t+1,i}^2)] = (1 - \Omega)$ we have

$$\mathbb{E}_t[p_{t+1,i}^2] = \mathbb{E}_t[(p_{t,i}\psi_{t+1,i}\omega_{t+1,i})^2] + (1 - \Omega) \quad (70)$$

$$= p_{t,i}^2\Omega\mathbb{E}[\psi^2] + (1 - \Omega) \quad (71)$$

so defining the mean operator $\mathbb{M}[\bullet_t] = \int_0^1 \bullet_{t,t} d\iota$, we have

$$\mathbb{M}[p_{t+1}^2] = \mathbb{M}[p_t^2] \Omega \mathbb{E}[\psi^2] + (1 - \Omega) \quad (72)$$

so that the steady state expected level of $\mathbb{M}[p^2] \equiv \lim_{t \rightarrow \infty} \mathbb{M}[p_t^2]$ can be found from

$$\mathbb{M}[p^2] = (1 - \Omega) + \Omega \mathbb{E}[\psi^2] \mathbb{M}[p^2] \quad (73)$$

$$\mathbb{M}[p^2] = \left(\frac{1 - \Omega}{1 - \Omega \mathbb{E}[\psi^2]} \right) \quad (74)$$

Finally, note the relation between p^2 and the variance of p :

$$\sigma_p^2 = \mathbb{M}[(p - \mathbb{M}[p])^2] \quad (75)$$

$$= \mathbb{M}[p^2 - 2p\mathbb{M}[p] + (\mathbb{M}[p])^2] \quad (76)$$

$$= \mathbb{M}[p^2] - 1 \quad (77)$$

where the last line follows because under the other assumptions we have made, $\mathbb{M}[p] = 1$.

Of course for the preceding derivations to be valid, it is necessary to impose the parameter restriction $\Omega \mathbb{E}[\psi^2] < 1$. This requires that income does not spread out so quickly among consumers who survive as to overcome the compression of the distribution that arises because of death.