Federal Reserve Bank of Cleveland Working Paper Series

Welfare Implications of Asset Pricing Facts: Should Central Banks Fill Gaps or Remove Volatility?

Pierlauro Lopez

Working Paper No. 21-16R

May 2023

Suggested citation: Lopez, Pierlauro. "Welfare Implications of Asset Pricing Facts: Should Central Banks Fill Gaps or Remove Volatility?." Working Paper No. 21-16R. Federal Reserve Bank of Cleveland. https://doi.org/10.26509/frbc-wp-202116r.

Federal Reserve Bank of Cleveland Working Paper Series
ISSN: 2573-7953

Welfare Implications of Asset Pricing Facts: Should Central Banks Fill Gaps or Remove Volatility?

Pierlauro Lopez[a,1]

[a] *Federal Reserve Bank of Cleveland*

Abstract

I find that removing consumption volatility is a priority over filling the gap between consumption and its flexible-price counterpart, or inflation targeting, in a model that matches empirical measures of the welfare costs of consumption fluctuations. Nearly 30 years of financial market data suggest sizable welfare costs of fluctuations that can be decomposed into a term structure that is downward-sloping on average, especially during downturns. This evidence offers guidance in selecting a model to study the benefits of macroeconomic stabilization from a structural perspective. The addition of nonlinear external habit formation to a textbook New Keynesian model can rationalize the evidence, and it offers a framework suitable for studying the desirability of removing fluctuations. The model is nearly observationally equivalent in its quantity implications to a standard New Keynesian model with CRRA utility, but the asset pricing and optimal policy implications are dramatically different. In the model, a central bank that minimizes consumption volatility generates welfare improvements relative to an inflation targeting regime that are equivalent to a 25 percent larger consumption stream.

JEL classification: E32; E44; E61; G12.

Keywords: Welfare cost of business cycles, Macroeconomic priorities, Equity and bond yields, Optimal monetary policy, Financial stability.

Email address: pierlaurolopez@gmail.com (Pierlauro Lopez)

[1] I would like to thank Klaus Adam, David Berger, Jarda Borovicka, John Cochrane, Jordi Galí, Arie Gozluklu, Wouter den Haan, Ayse Imrohoroglu, Urban Jermann, Patrick Kehoe, Ralph Koijen, Ed Knotek, Dirk Krueger, Roberto Marfè, Monika Piazzesi, Ricardo Reis, Juan Rubio-Ramirez, Eric Sims, Fabio Trojani, Harald Uhlig, and Michael Weber for insightful discussions, and seminar participants at several institutions and conferences for very useful comments. The views presented here are those of the author and not necessarily those of the Federal Reserve Bank of Cleveland or the Federal Reserve System.

1. Introduction

There are two distinct perspectives on macroeconomic stability. The first view, originating from work on the welfare cost of fluctuations (Lucas, 1987, 2003), considers reducing consumption fluctuations to alleviate the negative effects of cyclical changes on welfare. Namely, Lucas (1987) defined the welfare cost of fluctuations as the amount of growth people would trade to eliminate the uncertainty around future consumption, and proposed it as an indicator of the priority of stabilization policies. While alternative assumptions on preferences and cash flow processes have widely different implications for the costs of fluctuations,[2] it has been understood at least since Alvarez and Jermann (2004) that a sufficiently complete financial market will reveal the marginal cost of fluctuations directly. Yet, the implications of these measures for the design of stabilization policy have remained elusive. The second view, which is the hallmark of the New Keynesian optimal policy literature (e.g., Clarida et al., 1999; Woodford, 2003), uses general equilibrium foundations to emphasize steering consumption toward an optimal, time-varying level that reflects evolving economic conditions. Notably, a robust result of this literature is that inflation targeting—where policy fills the gap between consumption and its flexible-price counterpart—is optimal, or nearly so.

This paper bridges the two views by arguing that measures of the welfare cost of fluctuations contain powerful information to discriminate among competing general-equilibrium explanations of why fluctuations are costly. In this precise sense, they can have dramatic policy implications. This statement rests on two contributions. First, I use recent asset market data to reveal how short-run fluctuations contribute to the marginal cost of lifetime consumption fluctuations. This term structure information forms a rich set of empirical features to discipline structural models designed to study the benefits of stabilization policy.

Second, I propose a New Keynesian model with realistic discount rates that rationalizes those empirical measures, and carry out a welfare-theoretic analysis of monetary policy by nonlinear solution methods. Besides nominal price rigidities, the model relies on external habit formation and, to match the observed measures of the welfare costs of fluctuations at different horizons, consumers turn out to be sufficiently sensitive to cash flow fluctuations that they value a less volatile consumption path more than less volatile inflation, thereby overturning a pervasive result in the New Keynesian literature. In fact, this model demonstrates that the optimal policy involves a tradeoff between the two views of macroeconomic stability. In practice, the priority lies in reducing consumption fluctuations rather than in filling the consumption gap relative to the flexible-price level.

1.1. Decomposing the cost of fluctuations

The standard definition links the welfare cost of consumption fluctuations to the price of claims to future consumption, which are not observed. While I will rely on a model to estimate such costs, I propose a second definition that can be directly observed. Namely, by focusing on fluctuations in the two main determinants of consumption—wage and equity

[2] See Imrohoroğlu (1989); Atkeson and Phelan (1994); Krusell and Smith (1999); Tallarini (2000); Otrok (2001); De Santis (2007); Krebs (2007); Barro (2009); Ellison and Sargent (2012), among many others.

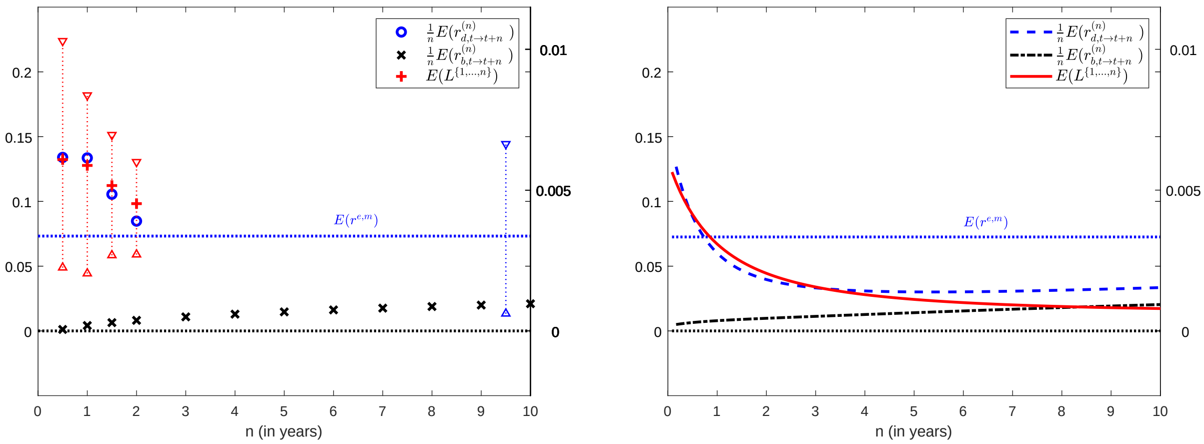

(a) *Point estimates.* (b) *Model-based predictions.*

Figure 1: Average term structures of equity ('∘' and dashed line), interest rates ('×' and dash-dotted line), and welfare costs ('+' and solid line); annualized expected real returns on the left axis for the term structures of equity and interest rates, and on the right axis for welfare costs. (Returns are deflated by the CPI.) Figure (a) plots point estimates during 1994:01-2018:12 (with block-bootstrap 95% confidence interval); $E(r^{e,m})$ is the equity premium (average annualized excess return on a 6-month buy-and-hold strategy). Figure (b) plots model-implied average term structures in the dynamic equilibrium model presented in Section 4. $L^{\{1,\dots,n\}}$ represents the amount of growth in the entire consumption process at an n-period horizon that people would trade against the elimination of consumption risk over the same horizon. The solid curve implies a cost of lifetime uncertainty of 15 hundredths of a percent.

income—and by describing the labor adjustments made by consumers as consumption becomes stabilized, I derive a tight link between the welfare cost of consumption fluctuations at different horizons and the prices of zero-coupon bonds and of claims to single market dividends, so-called dividend strips, which can be observed.[3] The result relies on a standard assumption on the consumption-labor tradeoff, does not require a parametric specification of consumer preferences, and allows for an observable measure of short-run welfare costs since the 1990s.

The empirical analysis finds that such costs are sizable and countercyclical. The point estimates, reported in Figure 1a, suggest a *negatively* sloped term structure of welfare costs, driven *both* by the negative slope of the term structure of equity yields (Binsbergen et al., 2012, 2013) and by the positive slope of the term structure of interest rates. At the margin people would trade an average of 0.5 percentage points of growth in next year's consumption against the elimination of one-year-ahead consumption risk. The volatility of this cost is

[3]A recent literature focuses on risk pricing across maturities (see, for example, Lettau and Wachter, 2007, 2011; Binsbergen et al., 2012, 2013; Borovicka and Hansen, 2014; Belo et al., 2015; Hasler and Marfè, 2016; Binsbergen and Koijen, 2017; Marfè, 2017; Ai et al., 2018; Weber, 2018; Bansal et al., 2021; Gormsen, 2021). The search for a structural explanation for the negative slope of the term structure of equity has only recently received a lot of attention (see Binsbergen and Koijen, 2017, for a survey). While a relatively small sample and potential liquidity issues leave some controversy on the sign of the average slope, as argued by Bansal et al. (2021), the cyclical properties of the term structure of equity are so far less controversial.

similar in size and is countercyclical; the cost increased to 2-3 percentage points during the 2001 and 2007-2009 recessions.

Thus, I complement the analysis of Alvarez and Jermann (2004) through a different take on the two main challenges they face in measuring the marginal cost of fluctuations. First, along with the standard one, I propose a definition of the welfare cost of fluctuations that can be observed in a model-free way. In the thought experiment I propose, people also adjust labor as consumption is stabilized. When the wage equals the marginal rate of substitution between consumption and labor, the positive effect on utility when labor income is marginally stabilized is exactly offset by the effect of the associated adjustment in hours, and hence only claims to equity income are necessary to measure the marginal cost of fluctuations. While this welfare cost measure with *variable labor* is meaningful in its own right, once I build a structural model that matches this welfare cost evidence, I am able to go back to the standard definition of welfare cost with *fixed labor* as in Alvarez and Jermann and estimate it through the model.

Second, the term structure of the cost of fluctuations roughly answers the question: How much compensation do people command to bear n-years-ahead consumption uncertainty? In contrast, the question studied by Alvarez and Jermann is: How much compensation do people command to bear uncertainty at business-cycle frequency in the *entire* consumption process? Their answer depends on the parametric assumptions about the filter that separates the trend and the business-cycle frequencies of the cash flow process. The question I ask is nonparametric and complements their exercise by decomposing the marginal cost of fluctuations in the time domain.

The empirical section uses bond data and extracts evidence about the term structure of equity from index option markets to infer the term structure of the marginal cost of fluctuations. Like Binsbergen et al. (2012) and Golez (2014) this paper relies on option data and no-arbitrage relations to replicate synthetic single market dividend payments by a strategy that avoids the need for an interest-rate proxy and that mitigates measurement error by excluding observations that violate the put-call parity relation.

1.2. Welfare implications of welfare cost measures

Evidence of nontrivial welfare costs is not a sufficient condition to attribute priority to stabilizing fluctuations, which can only be assessed in the context of a structural model. So, what do measures of the welfare costs of fluctuations really say? Since they represent natural evidence about the benefits of stabilization policies, I argue that they represent a set of empirical features that any macroeconomic model that seeks to assess the priority of different stabilization policies should match. In this respect, empirical measures of the welfare costs are highly informative to discriminate across competing models because it is well-known that explaining from a structural perspective the slope and cyclicality of the term structures of equity and interest rates *jointly* is difficult (e.g., Binsbergen and Koijen, 2017; Gormsen, 2021; Lopez et al., 2021).[4]

[4]The online appendix shows the implications of some of the leading consumption-based asset pricing models for the term structure of the welfare cost of fluctuations, and confirms their difficulties in explaining

Accordingly, I develop a structural model that captures the observed term structure of welfare costs and analyze its implications for monetary policy. The model relies on nominal rigidities and on Campbell and Cochrane (1999) *external* habits, and hence implies the desirability of some stabilization policy. I avoid the difficulties of incorporating habit formation in production economies documented by Lettau and Uhlig (2000) by including habit formation in home as well as market consumption. Intuitively, after a bad productivity shock affecting market and home production, both market and home consumption drop close to their habit levels, and hence offset the undesirable effect of habits on the labor choice. Here, note that such an offsetting effect that ensures standard quantity implications can also be achieved if habits were internal rather than external, but only for parameterizations that imply a stochastic discount factor with identical properties as log-utility CRRA preferences. By this argument, I rule out *internal* habit formation, as it is unable to match the asset pricing facts.

In the model, nominal rigidities make equity income highly procyclical but partly mean reverting by its cointegration with consumption; equity income is therefore particularly risky in the short term. Since the slow-moving habit makes people particularly risk averse after a bad productivity shock, it follows that people especially fear instability in fluctuations during downturns. Therefore, as in the data, the model-implied costs of short-run uncertainty are sizable, countercyclical, and substantially larger than the cost of fluctuations over longer horizons. Figure 1b reports the average term structure of welfare costs in the model and in the data, and shows how the model matches the asset pricing properties that are relevant from a welfare perspective. In fact, it turns out that the model captures not only the procyclicality of the slope of the term structures of equity yields documented by Binsbergen et al. (2013), but also the countercyclicality of the slope of the term structure of one-period equity returns documented by Gormsen (2021).[5] This evidence also rules out Epstein-Zin preferences as an alternative ingredient to habits to generate realistic asset prices, as I discuss below.

Conveniently, the model builds on the textbook CRRA New Keynesian model (e.g., Galí, 2008) used to study macroeconomic stabilization. The comparison can be made clear-cut. I can parameterize a version of my model so that quantities and inflation are approximately the same as in the textbook model; a macroeconomist only interested in the dynamics of quantities and inflation would not be able to discriminate between the two models. But the asset pricing and welfare cost implications differ dramatically.

The two externalities—sticky prices and external habits—reconcile the two views of stability in the literature. First, the central bank wants to close the consumption gap (or, equivalently, remove inflation volatility). Second, it wants to remove consumption volatility (or risk premia variation). Achieving both goals is unfeasible in the presence of productivity

the documented facts. I consider the habit formation model of Campbell and Cochrane (1999), the long-run risk model of Bansal and Yaron (2004), the long-run risk model under limited information of Croce et al. (2015), the recursive preferences of Tallarini (2000) and Barillas et al. (2009), and the rare disasters model of Gabaix (2012).

[5] Lopez et al. (2021) study these implications of a similar version of the model in greater detail, and show that it also captures key facts of the term structure of nominal interest rates.

shocks, so the optimal monetary policy trades them off. When comparing the two simple regimes under a parameterization consistent with measured welfare costs, I find that removing uncertainty is a priority over filling the gaps. This analysis exemplifies the potentially dramatic implications of realistic discount-rate variation for macroeconomics (Cochrane, 2011), even in a setup where the dynamics of quantities and inflation are standard.

The model's mechanism emphasizes the role of productivity shocks in matching the asset pricing properties, but abstracts from realistic features of New Keynesian models, including the presence of demand shocks and the zero-lower bound on the interest rate. I address both concerns. First, I solve the model subject to a lower bound on the nominal interest rate. The model's implications for the term structure of welfare costs remain very similar. Second, I extend the model to include demand shocks to match the observed correlation between consumption and inflation.[6] Importantly, even though the presence of demand shocks reduces the procyclicality of the profit share, the properties of the term structure of welfare costs and the priority of removing consumption fluctuations over filling the consumption gap remain similar.

Finally, note that I do not consider additional policy tools to remove the policy tradeoff. For example, Ljungqvist and Uhlig (2000) consider time-varying taxation to address habit externalities. In fact, it is well-known that fiscal policy can be used to mimic the effects of monetary policy or even replace it when monetary policy is constrained (Correia et al., 2013). Rather, the point of the present analysis is twofold. First, it points out how a model consistent with observed welfare cost measures can easily imply that people fear consumption fluctuations so much as to overturn the standard result that stable inflation is the macroeconomic priority. Second, it illustrates the tensions in two classical views of business cycle stabilization.

2. Term Structures of the Welfare Costs of Fluctuations

Identical consumers $j \in [0,1]$ have time-t preferences $E_t U\left(C(j), N(j), X(j)\right)$, where $C \equiv \{C_{t+n}\}_{n=1}^{\infty}$ is consumption, $N \equiv \{N_{t+n}\}_{n=1}^{\infty}$ is labor, and $X \equiv \{X_{t+n}\}_{n=1}^{\infty}$ is any other factor that influences utility.[7] Without loss of generality, let factor $X(j)$ depend on individual consumption and labor only via aggregate consumption and labor, $C = \int_0^1 C(j)dj$ and $N = \int_0^1 N(j)dj$. Since there is a continuum of agents, each of whom has zero mass, this modeling strategy enables me to ask an individual how much consumption growth she would trade against stable consumption and labor streams without having to affect factor X.

I will provide two definitions of the welfare cost of consumption fluctuations. The first definition generalizes the one by Alvarez and Jermann (2004) and has the disadvantage

[6]As shown by Lopez et al. (2021), the inclusion of demand shocks likewise allows for capturing the changing correlation patterns between stock and nominal bond returns emphasized by Campbell et al. (2020), as the extended model easily generates 10-year simulations with negative correlations between stock and bond returns.

[7]The definition of the cost of fluctuations has a meaningful interpretation even if one relaxes the assumption of a representative consumer, by the argument made by Alvarez and Jermann (2004).

of requiring a model to be measured. The second definition changes slightly the thought experiment in that it spells out the labor adjustments made by consumers as consumption becomes stabilized, and it has the important advantage of being directly measurable.

2.1. Welfare cost of fluctuations with fixed labor

Let the conditional expectation of n-period-ahead consumption, $E_t(C_{t+n})$, denote the consumption level that is hypothetically offered to the jth individual at time $t+n$, which I refer to as *stable* consumption. I parameterize stable consumption as

$$\overline{C}_{t+n}(\theta) = \theta E_t(C_{t+n}) + (1-\theta)C_{t+n}$$

where $\theta \in [0,1]$ represents the fraction of consumption volatility that is removed. I assume enough smoothness in preferences to guarantee that $\overline{\mathcal{L}}_t^{\mathcal{N}}$ is a differentiable map in $\theta \in [0,1]$.

Definition (Welfare cost of fluctuations, fixed labor)**.** The map $\overline{\mathcal{L}}_t^{\mathcal{N}}(\theta)$ defined by

$$\begin{aligned} &E_t U\Big(\big\{\big(1+\overline{\mathcal{L}}_t^{\mathcal{N}}(\theta)\big)^n C_{t+n}\big\}_{n\in\mathcal{N}}, \{C_{t+n}\}_{n\in\mathbb{N}\setminus\mathcal{N}}, \big\{X_{t+n}\big\}_{n=1}^{\infty}\Big) = \\ &\quad = E_t U\Big(\big\{\theta E_t(C_{t+n}) + (1-\theta)C_{t+n}\big\}_{n\in\mathcal{N}}, \{C_{t+n}\}_{n\in\mathbb{N}\setminus\mathcal{N}}, \{N_{t+n}\}_{n\in\mathbb{N}\setminus\mathcal{N}}, \big\{X_{t+n}\big\}_{n=1}^{\infty}\Big) \end{aligned} \tag{1}$$

measures the cost of consumption fluctuations, where the index set $\mathcal{N} \subset \mathbb{N}$ indicates which coordinates are stabilized.

For example, the *total cost* $\overline{\mathcal{L}}_t^{\mathbb{N}}(1)$ measures how much extra growth the elimination of all uncertainty in consumption is worth, and the *marginal cost* $\overline{L}_t^{\mathcal{N}} \equiv \frac{\partial}{\partial\theta}\overline{\mathcal{L}}_t^{\mathcal{N}}(0)$ represents how much extra growth a marginal stabilization of all coordinates in $\mathcal{N}$ is worth at the current level of uncertainty.[8]

2.2. Welfare cost of fluctuations with variable labor

A second definition of the welfare cost of consumption fluctuations can be defined by stabilizing the two main determinants of consumption. Consumption, $C_t = D_t + W_t N_t + e_t$, includes equity income D_t and labor income $W_t N_t$, with W_t the real wage rate, while e_t denotes any residual income. Stable consumption $\overline{C}_{t+n} = \overline{D}_{t+n} + W_{t+n}\overline{N}_{t+n} + e_{t+n}$ is here defined as a stable stream of dividend income, $\overline{D}_{t+n} = E_t(D_{t+n})$, of the path for labor $\overline{N}_{t+n}$ that ensures a stable labor income $W_{t+n}\overline{N}_{t+n} = E_t(W_{t+n}N_{t+n})$, and of residual income. I parameterize stable consumption by defining

$$\begin{aligned} \overline{C}_{t+n}(\theta) &= \theta\overline{C}_{t+n} + (1-\theta)C_{t+n} \\ &= \theta E_t(D_{t+n} + W_{t+n}N_{t+n}) + (1-\theta)(D_{t+n} + W_{t+n}N_{t+n}) + e_{t+n} \end{aligned}$$

where the parameter $\theta \in [0,1]$ represents the fraction of ex-post uncertainty in equity and labor income that is removed. The associated labor level is $\overline{N}_{t+n}(\theta) = \theta\overline{N}_{t+n} + (1-\theta)N_{t+n}$.

[8]The online appendix discusses the relationship between definition (2) and the definitions by Lucas (1987) and Alvarez and Jermann (2004).

Definition (Welfare cost of fluctuations, variable labor)**.** The map $\mathcal{L}_t^{\mathcal{N}}(\theta)$ defined by

$$E_t U\Big(\big\{\big(1+\mathcal{L}_t^{\mathcal{N}}(\theta)\big)^n C_{t+n}\big\}_{n\in\mathcal{N}}, \{C_{t+n}\}_{n\in\mathbb{N}\setminus\mathcal{N}}, \big\{X_{t+n}\big\}_{n=1}^{\infty}\Big) = $$
$$= E_t U\Big(\big\{\overline{C}_{t+n}(\theta)\big\}_{n\in\mathcal{N}}, \big\{\overline{N}_{t+n}(\theta)\big\}_{n\in\mathcal{N}}, \{C_{t+n}\}_{n\in\mathbb{N}\setminus\mathcal{N}}, \{N_{t+n}\}_{n\in\mathbb{N}\setminus\mathcal{N}}, \big\{X_{t+n}\big\}_{n=1}^{\infty}\Big) \quad (2)$$

measures the cost of fluctuations in equity and labor income, where the index set $\mathcal{N}\subset\mathbb{N}$ indicates which coordinates are stabilized.

The total cost $\mathcal{L}_t^{\mathbb{N}}(1)$ and the marginal cost $L_t^{\mathcal{N}} \equiv \frac{\partial}{\partial\theta}\mathcal{L}_t^{\mathcal{N}}(0)$ can be defined analogously.

Additionally, as consumption and hours are stabilized, I will assume that the wage rate remains consistent with efficiency by the condition[9]

$$W_t = -\frac{\partial U_t/\partial N_t}{\partial U_t/\partial C_t} \quad (3)$$

This assumption will imply that the marginal cost of the adjustment in hours will exactly offset the benefits of a marginal stabilization in labor income.

2.3. Marginal costs of consumption fluctuations

Proposition 1. *The marginal costs of consumption fluctuations within any window of interest $\mathcal{N}$ are*

$$\overline{L}_t^{\mathcal{N}} = \sum_{n\in\mathcal{N}} \frac{E_t(M_{t,t+n})E_t(C_{t+n}) - E_t(M_{t,t+n}C_{t+n})}{\sum_{n\in\mathcal{N}}\ n\ E_t(M_{t,t+n}C_{t+n})} \quad (4)$$

for the definition with fixed labor, and

$$L_t^{\mathcal{N}} = \sum_{n\in\mathcal{N}} \frac{E_t(M_{t,t+n})E_t(D_{t+n}) - E_t(M_{t,t+n}D_{t+n})}{\sum_{n\in\mathcal{N}}\ n\ E_t(M_{t,t+n}C_{t+n})} \quad (5)$$

for the definition with variable labor. Here, $M_{t,t+n} = (\partial U_t/\partial C_{t+n})/(\partial U_t/\partial C_t)$ is the n-period stochastic discount factor, $D_{ct}^{(n)} \equiv E_t(M_{t,t+n}C_{t+n})$ is the price of an n-period consumption strip, $D_{dt}^{(n)} \equiv E_t(M_{t,t+n}D_{t+n})$ is the price of an n-period dividend strip, and $D_{bt}^{(n)} \equiv E_t(M_{t,t+n})$ is the price of an n-period zero-coupon real bond.

For the definition with fixed labor, equation (4) expresses the marginal cost of fluctuations around all coordinates $n\in\mathcal{N}$ as a function of the price of claims to future consumption

[9] There is ample evidence that labor wedges matter (Chari et al., 2007). This observation does not, however, rule out assumption (3), which is consistent with labor wedges generated by distortions between wages and the marginal product of labor on the firm side—including sticky prices—as well as by departures from CRRA utility on the consumption side. Assumption (3) is nevertheless inconsistent with frictions such as sticky wages. Importantly, it is consistent with the model proposed in Section 4.

at those coordinates. These prices are not directly observable and require a model to be estimated. Accordingly, I will construct estimates in Section 4.

In contrast, for the definition with variable labor, equation (5) expresses the numerator of the marginal cost of fluctuations as a function of the price of claims to future dividends at all coordinates $n \in \mathcal{N}$. Importantly, these claims are traded, at least for some maturities, and hence their prices are directly observable. Note also that claims to labor income, which would also be unobserved, do not enter the expression by the assumption that the wage rate equals the marginal rate of substitution between consumption and labor, so the marginal effect of the adjustment in hours offsets exactly the benefits of a marginal stabilization in labor income. Therefore, the definition with variable labor can be much more easily measured in the data.

The expressions in Proposition 1 hold for arbitrary coordinate sets $\mathcal{N}$, including singleton sets of individual coordinates, which give rise to a term structure. The marginal costs $\overline{L}^{\mathcal{N}}$ and $L^{\mathcal{N}}$ can then be computed for any coordinate set $\mathcal{N} \subset \mathbb{N}$ from the term structure components.

Definition (Term structure of the cost of fluctuations)**.** The nth component of the term structure of the welfare cost of fluctuations is the risk premium for holding to maturity a portfolio long in an n-period strip and short in an n-period zero-coupon bond,

$$\bar{l}_t^{(n)} = \frac{1}{n}\left(E_t \overline{R}^{e,(n)}_{t\to t+n} - 1\right), \qquad\qquad l_t^{(n)} = \frac{1}{n}\left(E_t R^{e,(n)}_{t\to t+n} - 1\right) \tag{6}$$

where $\overline{R}^{e,(n)}_{t\to t+n} \equiv C_{t+n} D^{(n)}_{bt} / D^{(n)}_{ct}$ and $R^{e,(n)}_{t\to t+n} \equiv D_{t+n} D^{(n)}_{bt} / D^{(n)}_{dt}$.

As I will show below, $l_t^{(n)}$ can be directly measured, at least for small n.

Proposition 2. *The marginal cost of fluctuations within any window of interest $\mathcal{N}$ is the linear combination of the term structure components,*

$$\overline{L}^{\mathcal{N}}_t = \sum_{n\in\mathcal{N}} \overline{\omega}^{\mathcal{N}}_{nt} \bar{l}_t^{(n)} \tag{7}$$

for the definition with fixed labor, where the weights $\overline{\omega}^{\mathcal{N}}_{nt} \equiv n\ D^{(n)}_{ct} / \sum_{n\in\mathcal{N}} n\ D^{(n)}_{ct}$ are positive and such that $\sum_{n\in\mathcal{N}} \overline{\omega}^{\mathcal{N}}_{nt} = 1$, and

$$L^{\mathcal{N}}_t = \alpha^{\mathcal{N}}_t \sum_{n\in\mathcal{N}} \omega^{\mathcal{N}}_{nt} l_t^{(n)} \approx \alpha \sum_{n\in\mathcal{N}} \omega^{\mathcal{N}}_{n} l_t^{(n)} \tag{8}$$

for the definition with variable labor, with scaling factor $\alpha^{\mathcal{N}}_t \equiv \sum_{n\in\mathcal{N}} n\ D^{(n)}_{dt} / \sum_{n\in\mathcal{N}} n\ D^{(n)}_{ct}$, hence $\alpha^{\mathcal{N}} = \alpha = D/C$ in a nonstochastic stationary state, and where the weights $\omega^{\mathcal{N}}_{nt} \equiv n\ D^{(n)}_{dt} / \sum_{n\in\mathcal{N}} n\ D^{(n)}_{dt}$ are positive and such that $\sum_{n\in\mathcal{N}} \omega^{\mathcal{N}}_{nt} = 1$.

Equation (8) shows how the first-order effect of distinguishing between consumption and dividends is that the cost of fluctuations around any coordinate set $\mathcal{N}$ scales the linear combination of the term structure components by a factor equal to the average dividend-consumption ratio. This factor can be estimated at about 4.6 percent over the 1980-2019 period using data from the US Flow of Funds Accounts on net dividends paid out by nonfarm nonfinancial corporates (Table F.103, line 3) and on personal consumption expenditures in nondurable goods and services (Table F.6, lines 4 and 5). Accordingly, as shown below, $L_t^{\mathcal{N}}$ can be measured directly, at least at short durations.

3. Empirical Measures of the Cost of Fluctuations

The previous section related the welfare cost with variable labor in (2) to zero-coupon bonds and dividend strips. In this section, I use evidence about prices of such claims to measure some of these welfare costs. In contrast, the welfare costs with fixed labor in (1) cannot be directly measured in the absence of consumption strip price data, and will be estimated in Section 4.

Suppose that a full set of zero-coupon real and nominal bonds and a full set of put and call European options whose underlying is an aggregate equity index are traded. In the absence of arbitrage opportunities, put-call parity holds:

$$\mathcal{C}_{t,t+n} - \mathcal{P}_{t,t+n} = \mathscr{P}_t - \sum_{j=1}^{n} D_t^{(j)} - X \; P_{\$bt}^{(n)} \tag{9}$$

where $\mathcal{C}_{t,t+n}$ and $\mathcal{P}_{t,t+n}$ are the nominal prices at time t of call and put European options on the market index with maturity n and nominal strike price X, $P_{\$bt}^{(n)} = E_t M_{t,t+n} P_t / P_{t+n}$ is the nominal price of an n-period nominal zero-coupon bond, $\mathscr{P}_t = P_t E_t \sum_{j=1}^{\infty} M_{t,t+j} D_{t+j}$ is the nominal value of the market portfolio, and $D_t^{(n)} = P_t E_t M_{t,t+n} D_{t+n}$ is the nominal price of the nth dividend strip, where P denotes the price level.

I follow Binsbergen et al. (2012) and Golez (2014) in synthesizing the evidence on dividend claims from put and call European options on the S&P 500 index.[10] Standard index option classes, with 12 monthly maturities of up to one year, and long-term equity anticipation securities (LEAPS), with 10 maturities of up to three years, have been exchange

[10]The literature offers at least three alternatives to extract the term structure of equity. First, index options can be combined with some interest-rate proxy as in the original intra-day approach of Binsbergen et al. (2012). The tick-level approach has the advantage of exploiting information from more data points and avoids asynchronicity issues; however, accuracy can be lost in the choice of an interest-rate proxy. Second, index options can be combined with the interest rate implied in index futures as in Golez (2014). However, CME S&P 500 futures have expiration dates only for eight months in a quarterly cycle over most of the available sample. Finally, the index dividend futures studied by Binsbergen et al. (2013) have the advantage of revealing strip prices without the need for synthetic replication, and they do so for longer maturities. However, for S&P 500 dividend futures exchange trading started only in November 2015, and previously only proprietary data sets starting at the end of 2002 are available covering over-the-counter trades. In this context, the exchange-traded nature of options mitigates concerns that the preferences embedded in their pricing do not reflect those of the average investor, which would complicate the macroeconomic interpretation of the derived costs.

traded on the Chicago Board Options Exchange (CBOE) since 1990. The overall size of the index option market in the US has grown rapidly over the years. During the first year of the sample the average open interest for standard options and LEAPS with maturities of less than six months is around $60 billion and gradually decreases across maturities to less than $400 million for options of two years or more. The corresponding figures in the last year of the sample are an open interest of $2,300 billion for maturities of less than six months and of $50 billion for maturities longer than two years.

Like Golez, but unlike Binsbergen et al. (2012), the main analysis relies on end-of-day option data. I use a data set provided by the CBOE Data Services/Market Data Express containing S&P 500 index option data for CBOE-traded European-style options and running from January 1990 to December 2018. I obtain the daily S&P 500 price and one-day total return indices from Bloomberg and combine them to calculate daily index dividend payouts; I then aggregate the daily payouts to a monthly frequency without reinvestment.

There are three major difficulties when extracting options-implied prices through the put-call parity relation (see also Boguth et al., 2012). First, quotes may violate the law of one price for reasons that include measurement errors such as bid-ask bounce or other microstructural frictions. Second, the synthesized prices are sensitive to the choice of risk-free rate, which multiplies strikes in the put-call parity relation. Since strike prices are large numbers, any error in the interest rate will be magnified in the synthetic prices. Third, end-of-day data quote the closing value of the index, whose components trade on the equity exchange, and the closing prices of derivatives that are exchange-traded on a market that continues to operate for 15 minutes after the equity exchange closes. An asynchronicity of up to 15 minutes may drive a wedge between the reported quotes of the index value and the option prices and bias the synthetic prices.

To address these difficulties, I combine options with different moneyness levels to extract *both* risk-free rates and strip prices in a unique step. Namely, for any date-maturity pair (t, n) with I put-call pairs that differ only in strike price, X_i, $i = 1, ..., I$, define the auxiliary variable $\mathcal{A}_{it}^{(n)} \equiv \mathscr{P}_t - \mathcal{C}_{i,t,t+n} + \mathcal{P}_{i,t,t+n}$. Equation (9) then implies that

$$\mathcal{A}_{it}^{(n)} = \sum_{j=1}^{n} D_t^{(j)} + X_i P_{\$bt}^{(n)} \tag{10}$$

i.e., for each (t, n), a linear cross-sectional regression on a constant and strike prices should fit perfectly and reveal strip prices and the appropriate interest rate for synthetic replication. To the extent that the linear relation fits imperfectly in practice, (10) allows for spotting trade dates that violate the law of one price at some date-maturity pair. When I identify such violations, I drop the associated observations to mitigate the effects of microstructural noise. The appendix details the algorithm for synthetic replication.

Finally, to measure real bond prices I rely on zero-coupon TIPS yields with maturities of up to 10 years from Gürkaynak et al. (2010). However, TIPS yields are either unavailable or unreliable during the 1990s. In fact, there is evidence of sizable liquidity premia in TIPS markets, especially at inception and during the 2007-2009 crisis, that distort measures of real yields extracted from TIPS (D'Amico et al., 2018). In this context, D'Amico et al. find that a TIPS-specific factor not captured by the first three principal components of nominal yields captures the liquidity premium commanded by TIPS. Accordingly, I regress TIPS of up to 10-year maturity on the 1999-2019 period on the first three principal components of nominal yields, and interpret the resulting projection as the real yield net of the TIPS liquidity premium. I then reconstruct real yields on the 1994-1999 period using the same regression coefficients.[11]

3.1. Average costs of fluctuations

The evidence suggests economically and statistically sizable welfare costs of short-term fluctuations. Economically, the evidence points to a downward-sloping term structure of welfare costs, driven *both* by a negatively sloped term structure of equity and by a positively sloped term structure of interest rates.

I follow Binsbergen et al. (2012) and focus on a semestral periodicity; the first strip pays off the next six months of dividends, the second strip the dividends paid out six to 12 months out, and so on. The measure of the hold-to-maturity return on the first semestral strip is the return on a six-month buy-and-hold strategy that pays off the next six months of dividends.[12] Accordingly, I measure the hold-to-maturity return on the n-semester strip as the return for holding for $n-1$ semesters an n-semester strip times the semestral return on the first semestral strip.

To address the concerns raised by Boguth et al. (2012) that microstructural frictions might cause spuriously large arithmetic high-frequency returns on synthetic dividend claims, I report log returns on six-month buy-and-hold strategies and hold-to-maturity returns on strategies with maturities between 0.5 and 2 years, which Boguth et al. advocate as much less biased by microstructural effects related to highly leveraged positions.

Figure 2a illustrates the size of average annualized monthly log returns on six-month strategies over different subsamples by plotting the cumulated real return on an investment strategy that goes long on January 31, 1996, by a dollar in a claim to the next n years of dividends, holds the investment for six months, and then rolls over the position. Monthly average log returns are large and positive for short-duration equities and larger than the real return on the index. Figure 2b plots the analogous cumulated monthly returns on six-month bond strategies long by a dollar on zero-coupon bonds with maturities between six months and 10 years; average returns steadily increase in maturity across bonds, consistent with an upward-sloping average term structure of real interest rates.

Table 1 reports the point estimates for the term structure of welfare costs. The first four term structure components at semestral frequency are 13.3, 12.9, 9.8 and 7.6 percent. These

[11] A simpler proxy for real yields—nominal hold-to-maturity returns on Treasury yields deflated by holding-period inflation—results in similar results.

[12] Golez (2014) raises concerns that equity prices of up to three-month maturity may be biased as a result of firms routinely pre-announcing part of their dividend payouts, which would lower their riskiness. To mitigate such concerns, I roll over three times a two-month buy-and-hold strategy that goes long in the six-month strip rather than hold to maturity a six-month strip.

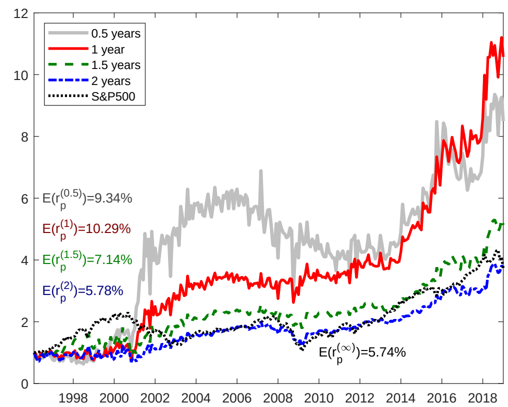

(a) *Monthly cumulated returns on equities.*

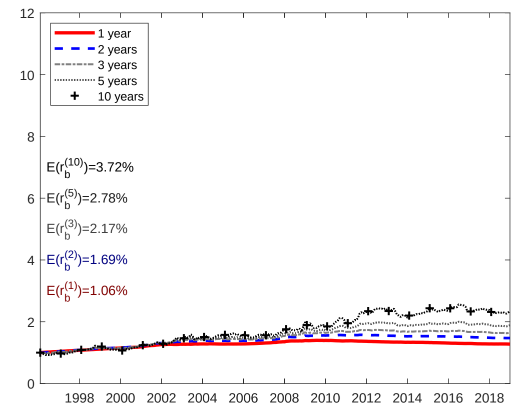

(b) *Monthly cumulated returns on real bonds.*

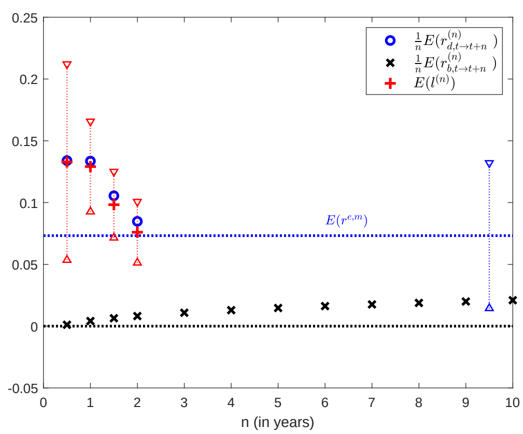

(c) *Average term structures of equity, real rates, and welfare costs (with block-bootstrap 95% intervals).*

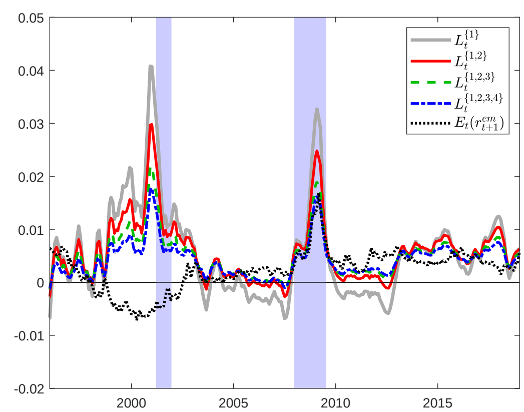

(d) *Semestral term structure of welfare costs and equity premium, conditional on the information set spanned by equity and bond yields.*

Figure 2: Term structures of equity, real interest rates, and welfare costs over the last two decades; annualized real returns. The term structure of equity is synthesized from index options; the term structure of interest rates uses Gürkaynak-Sack-Wright data about Treasuries and TIPS. Real interest rates over 1994-2019 are constructed as the projection of TIPS yields on the first three principal components of the term structure of nominal yields to filter out a liquidity premium and to construct missing data over 1994-1999. Regressors in panel 2d are the first two principal components of semestral equity yields and the first principal component bond yields; the semestral excess return on the index is regressed additionally on the market dividend yield. The shaded areas indicate recessions as declared by the NBER (March-November 2001 and December 2007 to June 2009).

numbers compare to an average equity premium of 7.3 percent. I then rely on proposition 2 to compute the costs of uncertainty around multi-period cash flows. Namely, I estimate average welfare costs of 0.61 percent associated with uncertainty one semester out, of 0.59 percent associated with up to one-year-ahead uncertainty, of 0.52 percent with up to 18-months-ahead uncertainty, and of 0.45 percent with up to two-years-ahead uncertainty.

Additionally, Table 1 shows the welfare cost of one-period-ahead uncertainty for different periodicities, from semestral to biennial. These estimates complement the evidence about the term structure in a way that bypasses the somewhat arbitrary choice of a semestral periodicity of the strips. I find comparable results: the average cost of one-year-ahead uncertainty over the two samples is 0.51 percent, whereas the cost of fluctuations over the next two years is 0.36 percent.

Figures 1a and 2c plot the point estimates for the term structures of equity, interest rates, and welfare costs. The figures also show block-bootstrapped one-sided critical values based on bootstrap-t percentiles corresponding to a 5 percent size for the means of $l_t^{(n)}$ and the six-month market return. (The block size of 10 observations is slightly larger than the number of lags after which the correlogram of the underlying returns becomes negligible.) In both samples I can reject the hypothesis that the term structure components are trivial.

A comparison of the estimated costs of uncertainty with holding-period excess returns on the index reveals the *economic* significance of these estimates. Equation (6) implies that if the cost $l_t^{(1)}$ of one-period-ahead fluctuations is larger than the equity premium over the same period, then strip excess returns must be smaller than the equity premium at some longer duration. It follows that the term structure of welfare costs must slope downward. Furthermore, an upward-sloping real term structure exacerbates this feature. Table 1 confirms this point. Short-term welfare costs are sizable, both economically and statistically.

3.2. Time-variation in the costs of fluctuations

Present-value logic implies that equity yields contain information about expected returns. Since the same states would drive the risk premia that constitute the term structure of welfare costs, it follows that equity yields signal variation in the term structure of welfare costs. Motivated by theoretical models I consider a one- to three-factor specification for these risk premia.[13] To capture the factors I extract the first two principal components of semestral equity yields, which capture 93.8 percent of their volatility, and the first principal component of bond yields, and use them to forecast the hold-to-maturity excess returns whose ex-ante values constitute the welfare cost measures. Table 2 presents the predictive regressions and shows a standard deviation of expected returns about as large as the already sizable level.

The components of the term structure of welfare costs are volatile and countercyclical. Since excess returns are forecastable, the cost of fluctuations varies over time. The cost

[13] For example, the no-arbitrage models of Lettau and Wachter (2011) and Gormsen (2021) as well as the model in Section 4 imply a term structure of welfare costs that can be revealed by the information set spanned by two equity yields and a bond yield.

	0.5 years ($n=1$)	1 year ($n=2$)	1.5 years ($n=3$)	2 years ($n=4$)	0.5 years ($n=1$)	1 year ($n=2$)	1.5 years ($n=3$)	2 years ($n=4$)
	$l_t^{(n)}$				$l_t^{(1)}$, 1 period = n semesters			
Mean	0.1328	0.1291	0.0983	0.0760	0.1328	0.1103	0.0908	0.0777
	(0.0578)	(0.0250)	(0.0168)	(0.0150)	(0.0578)	(0.0358)	(0.0238)	(0.0194)
	$L_t^{\{1,\ldots,n\}}$				$L_t^{\{1\}}$, 1 period = n semesters			
Mean	0.0061	0.0059	0.0052	0.0045	0.0061	0.0051	0.0042	0.0036
	(0.0027)	(0.0015)	(0.0011)	(0.0008)	(0.0027)	(0.0016)	(0.0011)	(0.0009)
$E(r^{e,m})$	0.0732							
	(0.0319)							

Table 1: Options-implied average term structure of the welfare cost of fluctuations, 1994-2019. $l^{(n)}$ is the annualized cost of a marginal increase in uncertainty in n-semesters-ahead cash flows. $L^{\{1,\cdots,n\}}$ is the annualized cost of a marginal increase in uncertainty in 1- to n-semesters-ahead cash flows. The right panel reports the cost of a marginal increase in one-period-ahead uncertainty, $L^{\{1\}}$, for different period lengths. Bootstrap standard errors use block sizes of 10 observations to preserve time-series dependencies. The equity premium $E(r_{em})$ is the average semestral buy-and-hold return on the S&P 500 index in excess of the corresponding risk-free rate (dividends reinvested in real bonds).

	$r^{e,(1)}_{t\to t+1}$		$\frac{1}{2}r^{e,(2)}_{t\to t+2}$		$\frac{1}{3}r^{e,(3)}_{t\to t+3}$		$\frac{1}{4}r^{e,(4)}_{t\to t+4}$		$r^{e,m}_{t+1}$	
const.	0.094	0.094	0.106	0.107	0.078	0.078	0.059	0.060	0.059	0.059
	[0.027]	[0.019]	[0.013]	[0.013]	[0.012]	[0.010]	[0.010]	[0.009]	[0.028]	[0.027]
pc^1_{dt}	0.381	0.367	0.153	0.163	0.083	0.092	0.074	0.082		-0.074
	[0.106]	[0.055]	[0.031]	[0.038]	[0.023]	[0.023]	[0.019]	[0.016]		[0.049]
pc^1_{bt}		-0.015		-0.089		-0.100		-0.151		0.095
		[0.131]		[0.072]		[0.049]		[0.035]		[0.155]
pc^2_{dt}		0.811		-0.050		-0.246		-0.116		
		[0.149]		[0.073]		[0.049]		[0.037]		
dp^m_t									0.399	0.405
									[0.134]	[0.212]
R^2	0.30	0.47	0.20	0.22	0.10	0.23	0.13	0.24	0.13	0.16
$\frac{\sigma(E_t r)}{E(r)}$	2.03	2.54	0.73	0.75	0.55	0.83	0.66	0.90	1.40	1.53

Table 2: Predictive regressions on hold-to-maturity semestral strip returns and on the semestral buy-and-hold market return. Annualized log returns in excess over the riskless return over the holding period. Regressors are the first two principal components of the semestral equity yields, pc^1_d and pc^2_d, the first principal component of up to 10-year bond yields, pc^1_b, and the market dividend yield, dp^m. Monthly data, 1994m1-2018m12. Newey-West standard errors to correct for overlapping.

of short-run cash flow uncertainty is substantial at some junctures of the business cycle. Figure 2d plots the estimated time series of the term structure of the welfare cost of fluctuations over time. The cost of fluctuations rises dramatically to 2-4 percent during the dot-com crash and the period immediately preceding the early 2000s recession as well as during the Great Recession. Moreover, the premium to hedge uncertainty six months out is considerably larger than the premium to hedge longer-run uncertainty. The estimated term structure remains downward-sloping during the downturns, whereas it appears considerably flatter and even upward-sloping in normal times.

4. Term Structure of Welfare Costs in a Simple Model

This section proposes a model that can rationalize the evidence about the term structure of welfare costs described in the previous section. The model can therefore also be used to provide estimates about welfare costs at coordinates that cannot be observed and about the welfare costs of fluctuations with fixed labor. The model combines a textbook New Keynesian model economy (Galí, 2008) with Campbell and Cochrane (1999) habits in consumption to generate realistic asset pricing facts. I include habit formation in home as well as market consumption to offset the undesirable effect of habits on the labor choice documented by Lettau and Uhlig (2000). Importantly, this model rationalizes the term structure evidence of Section 3, including the size and volatility of equity market returns and risk-free rates. In this sense, this model captures key facts about the welfare cost of fluctuations.

4.1. Model

I develop a model that illustrates the point in the simplest possible setup. Namely, I will calibrate the model so that the approximate dynamics of inflation and economic activity under the Campbell-Cochrane pricing kernel coincide with those under a version of the model with CRRA utility, and I will abstract from capital formation. Up to a first-order approximation around the risky steady state, the two models differ only in their asset pricing implications. (Throughout the exposition, lower-case letters denote natural logarithms and hatted variables denote deviations from the mean.)

4.1.1. Firms

Monopolistically competitive firms indexed by $i \in [0, 1]$ maximize intertemporal profits

$$E_0 \sum_{t=0}^{\infty} M_{0,t} \left(\frac{P_t(i)}{P_t} Y_t(i) - (1 - \tau_f) W_t N_t(i) - T_{ft} \right)$$

subject to Calvo-type nominal price stickiness. The t-period stochastic discount factor $M_{0,t}$ is the households'—which own the firms. Firms operate the production technology

$$Y_t(i) = A_t N_t(i)^{1-\alpha}$$

where Y_t is real output; N_t is the labor input, which they acquire at a unit cost equal to the real wage rate W_t; and A_t is aggregate productivity. The government levies lump-sum taxes T_{ft} on each firm to finance an employment subsidy, τ_f, designed to offset any distortions caused by monopolistic competition in the steady state. Corporate profits $D_t(i) = Y_t(i)P_t(i)/P_t - (1-\tau_f)W_tN_t(i) - T_{ft}$ are paid out as dividends on market equity each period to households.

The ith good sells for the nominal price $P_t(i)$. Each firm i can reset prices at any given time only with probability $1-\eta$. If a firm is unable to re-optimize, it adjusts its price from the previous period by the steady-state inflation rate π^*. Individual consumers bundle the continuum of goods via a CES aggregator with elasticity of substitution between goods, ε; their cost-minimizing plan gives rise to the demand curve for the ith good, $Y_t(i) = [P_t(i)/P_t]^{-\varepsilon}Y_t$, which constrains individual firms in their production choices. $P_t \equiv [\int_0^1 P_t(i)^{1-\varepsilon}di]^{1/(1-\varepsilon)}$ is the price index.

In this context, a firm's optimal labor demand schedule implies real marginal costs

$$mc_t(i) = w_t + \ln(1-\tau_f) - \ln(1-\alpha) - y_t(i) + n_t(i) \tag{11}$$

while a nonlinear New Keynesian Phillips curve describes the optimal price-setting behavior of a firm that reset prices at time t as the condition linking inflation and marginal costs,

$$\left(\frac{1-\eta e^{(\varepsilon-1)(\pi_t-\pi^*)}}{1-\eta}\right)^{\frac{1}{1-\varepsilon}} = \frac{E_t\sum_{j=0}^{\infty}(\beta\eta)^j e^{y_{t+j}-\gamma c_{t+j}-\gamma\hat{s}_{t+j}+\varepsilon\sum_{h=1}^{j}(\pi_{t+h}-\pi^*)+mc^*_{t+j}-\ln((\varepsilon-1)/\varepsilon)}}{E_t\sum_{j=0}^{\infty}(\beta\eta)^j e^{y_{t+j}-\gamma c_{t+j}-\gamma\hat{s}_{t+j}+(\varepsilon-1)\sum_{h=1}^{j}(\pi_{t+h}-\pi^*)}} \tag{12}$$

where the equilibrium real marginal cost at date $t+j$ for the firm can be written as

$$mc^*_{t+j} = mc_{t+j} - \frac{\alpha\varepsilon}{(1-\alpha)(1-\varepsilon)}\ln\left(\frac{1-\eta e^{(\varepsilon-1)(\pi_t-\pi^*)}}{1-\eta}\right) + \frac{\alpha\varepsilon}{1-\alpha}\sum_{h=1}^{j}(\pi_{t+h}-\pi^*) - \Delta_{t+j} \tag{13}$$

Note how each resetting firm faces the same problem; so I dropped the i index from equation (12). The aggregate marginal cost in equation (13) integrates equation (11) over i.

Finally, the Calvo structure implies that price dispersion $\Delta_t \equiv \ln\int_0^1[P_t(i)/P_t]^{-\varepsilon/(1-\alpha)}di$ evolves according to the law of motion

$$\Delta_t = \ln\left(\eta\hat{\Pi}_t^{\frac{\varepsilon}{1-\alpha}}e^{\Delta_{t-1}} + (1-\eta)\left[\frac{1-\eta\hat{\Pi}_t^{\varepsilon-1}}{1-\eta}\right]^{\frac{\varepsilon}{(1-\alpha)(\varepsilon-1)}}\right) \tag{14}$$

where $\hat{\Pi}_t \equiv P_t/e^{\pi^*}P_{t-1}$ defines the detrended gross inflation rate.

4.1.2. Households

Identical consumers indexed by $j \in [0,1]$ trade in complete financial markets and choose consumption and labor to maximize intertemporal utility

$$U_0(j) = E_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{[C_t(j) - X_{ct}]^{1-\gamma}}{1-\gamma} + \chi \frac{[H_t(j) - X_{ht}]^{1-\gamma}}{1-\gamma} \right)$$

subject to the present-value budget constraint

$$E_0 \sum_{t=0}^{\infty} M_{0,t} C_t(j) = \frac{B_{-1}(j)}{P_0} + E_0 \sum_{t=0}^{\infty} M_{0,t} \left((1-\tau_h) W_t N_t(j) + D_t + T_{ht} \right)$$

with t-period real contingent claims prices $M_{0,t}$. As in Greenwood and Hercowitz (1991), households derive utility from consumption of two goods: C_t is real consumption purchased in the market, and H_t denotes the consumption produced at home, with production function $H_t = A_t(1 - N_t)$. X_{ct} and X_{ht} represent external habit levels that are slow-moving averages of contemporaneous and past aggregate consumption, described below. Hours worked in market production N_t are remunerated at the real wage rate W_t. B_t denotes holdings of one-period nominal debt issued by a fiscally passive government and with unit price $\exp(-i_t) = E_t M_{t,t+1}/\Pi_{t+1}$. D_t is the equity income households receive from owning the aggregate firm. Parameter β is the subjective discount rate and parameter χ controls steady-state hours.[14] The government levies an income tax $\tau_h = 1 - S_c/S_h$ designed to offset any steady-state distortions caused by the habit externalities, and rebates it in lump-sum fashion, T_{ht}, to households.

Habit specification. As in Campbell and Cochrane (1999), the law of motion of habits is specified indirectly through the processes for surplus market consumption $S_{ct} \equiv (C_t - X_{ct})/C_t$ and surplus home consumption $S_{ht} \equiv (H_t - X_{ht})/H_t$. The following dynamics for the *logarithms* of aggregate surplus levels ensure nonnegative marginal utilities:

$$\begin{aligned} s_{ct+1} &= (1-\rho_s) s_c + \rho_s s_{ct} + \lambda_{ct} \varepsilon_{ct+1} \\ s_{ht+1} &= (1-\rho_s) s_h + \rho_s s_{ht} + \lambda_{ht} \varepsilon_{ht+1} \end{aligned} \tag{15}$$

where $\varepsilon_{ct} \equiv c_t - E_{t-1} c_t$ and $\varepsilon_{ht} \equiv h_t - E_{t-1} h_t$ denote the innovations in market and home consumption, respectively, with common persistence ρ_s.

It can be readily verified that these preferences imply the intertemporal and static

[14] Let $\chi = \chi_0 (S_c/S_h)^{1-\gamma}$, where χ_0 is the counterpart to χ in the CRRA utility case, i.e., when $X_{ct} = X_{ht} = 0$. I assume a calibration for χ_0 to achieve a level of hours $N = .5$ in the flexible-price steady state.

marginal rates of substitution

$$M_{t+1} = \frac{\partial U_t / \partial C_{t+1}}{\partial U_t / \partial C_t} = \beta \left(\frac{C_{t+1} S_{ct+1}}{C_t S_{ct}} \right)^{-\gamma} \tag{16}$$

$$MRS_t = -\frac{\partial U_t / \partial N_t}{\partial U_t / \partial C_t} = -\frac{\chi C_t^{\gamma}}{H_t^{\gamma}} \frac{\partial H_t}{\partial N_t} \left(\frac{S_{ct}}{S_{ht}} \right)^{\gamma} \tag{17}$$

Habits amplify the consumers' aversion to fluctuations in consumption and hence affect asset prices, including the consumption-saving tradeoff through the Euler equation with stochastic discount factor (16), but they also affect the consumption-labor tradeoff (17). The asset pricing implications and the spillover of habits on the consumption-saving tradeoff are controlled by the process s_{ct}, which affects (16). There are at least two reasons to focus on home consumption rather than standard leisure. First, once it is accepted that people get used to an accustomed market consumption level, it is only natural to assume that people also develop a habit in home consumption. Second, home consumption frees up the elasticity of intertemporal substitution γ by making the economy consistent with balanced growth for any $\gamma > 0$ even though utility is separable in market and home consumption. Consequently, I can accommodate the pricing kernel of Campbell and Cochrane (1999).

The spillover of habits on the consumption-labor tradeoff is controlled instead by the process $s_{ct} - s_{ht}$, which affects (17). In fact, the role of the habit in home consumption is precisely to offset the effect on the labor choice of the habit in market consumption, namely, by generating movements in s_{ht} that offset the movements in s_{ct} on (17). Indeed, the inclusion of only consumption habits in the production economy, as studied by Lettau and Uhlig (2000), would result in a labor process with the wrong cyclicality that is either too volatile or that smooths consumption too much.

I adopt a particular specification for the sensitivity functions λ_{ct} and λ_{ht}. In particular, I let $\sigma_x^2 \equiv var(\varepsilon_{xt})$ and $\sigma_{xy} \equiv cov(\varepsilon_{xt}, \varepsilon_{yt})$ denote the unconditional variance and covariance of innovations to variables x and y, and $\sigma_{xt}^2 \equiv var_t(\varepsilon_{xt+1})$ and $\sigma_{xyt} \equiv cov_t(\varepsilon_{xt+1}, \varepsilon_{yt+1})$ the corresponding conditional variance and covariance. I then consider the sensitivity functions

$$\lambda_{ct} = \frac{\sigma_c}{\sigma_{ct}} \frac{1}{S_c} \sqrt{1 - 2(s_{ct} - s_c)} - 1 \tag{18}$$

if the right-hand side is positive and zero otherwise, and

$$\lambda_{ht} = \frac{S_c}{1 - S_c} \frac{1 - S_h}{S_h} \frac{\sigma_{cht} / \sigma_{ht}^2}{\sigma_{ch} / \sigma_h^2} \lambda_{ct} \tag{19}$$

As shown in the appendix, this choice ensures habits that are predetermined at and around the steady state, reflecting the intuitive notion that households slowly grow used to unanticipated movements in the two types of consumption. The ratios of conditional to unconditional variances and covariances will simplify the resulting expressions for the marginal rates of substitution, as described shortly, and they disappear on average. Those terms likewise disappear if innovations to consumption are homoskedastic, as in the original framework of

Campbell and Cochrane (1999).

Next, I choose the steady-state surplus consumption ratios S_c and S_h to control the spillover of states s_{ct} and s_{ht} on the marginal rates of substitution. Since surplus consumption ratios drive asset prices, it is crucial to gain a handle on such a spillover to avoid the critique by Lettau and Uhlig (2000). In pinning down S_c and S_h, it is convenient to introduce two new free parameters ξ_1 and ξ_2, discussed below, and express the values of parameters S_c and S_h as the following functions of these parameters:

$$S_c = \sqrt{\frac{\gamma\sigma_c^2}{1-\rho_s-\xi_1/\gamma}}, \qquad\qquad S_h = \left(1+\frac{1-S_c}{S_c}(1+\xi_2)\frac{\sigma_{ch}}{\sigma_h^2}\right)^{-1} \tag{20}$$

The motivation for this choice of S_c follows Campbell and Cochrane (1999), and it is made to control the direct spillover of state s_{ct} on the risk-free rate. Namely, the choices in (20) implies that such a spillover is as close as possible to $\xi_1\hat{s}_{ct}$ in a mean-squared sense, and exactly equal to it when consumption innovations are conditionally Gaussian. The choice of S_h is instead motivated to control the direct spillover of surplus consumption on the marginal rate of substitution between consumption and labor. Namely, the value in (20) is such that $S_h = \arg\min_{S_h} var_t[\lambda_{ht}\varepsilon_{ht+1} - (1+\xi_2)\lambda_{ct}\varepsilon_{ct+1}]$; it follows that the term $\hat{s}_{ht} - \hat{s}_{ct}$ in (17) is as close as possible to $\xi_2\hat{s}_{ct}$ in a mean-squared sense. Therefore, through the two free parameters, ξ_1 and ξ_2, the values of S_c and S_h control the spillovers of the habits on the equilibrium allocation. In fact, I will show how a parameterization exists for ξ_1 and ξ_2 that implies an equilibrium allocation that is nearly equivalent to the allocation of a version of the model without habits.

Finally, I assume that the habit levels thus specified are external to any individual consumer. For example, they are driven by *aggregate* market and home consumption, $C_t \equiv \int C_t(j)dj$ and $H_t \equiv \int H_t(j)dj$, and hence each individual consumption level has a trivial impact on these aggregates. Because habits are external, optimizing consumers do not internalize the effect of their individual decisions on habits. In this context, as a final sanity check of the microfoundations of these nonlinear habits, the appendix shows how this specification of the surplus consumption process ensures that welfare increases with consumption for any discrete movement away from the steady state, and therefore avoids Ljungqvist and Uhlig's (2015) critique of Campbell and Cochrane's (1999) internal habit formation.

Households' optimality conditions. Optimality implies the log stochastic real discount factor,

$$m_{0,t} = -t\ln(\beta) - \gamma(c_t - c_0) - \gamma(s_{ct} - s_{c0})$$

the equilibrium consumption-saving equation,

$$i_t = -\ln E_t \beta e^{-\gamma\Delta c_{t+1} - \gamma\Delta s_{ct+1} - \pi_{t+1}} \tag{21}$$

and the labor supply equation,

$$W_t = \frac{\chi}{1-\tau_h} \frac{A_t^{1-\gamma} C_t^{\gamma}}{(1-N_t)^{\gamma}} \left(\frac{S_{ct}}{S_{ht}} \right)^{\gamma} \tag{22}$$

4.1.3. Government

Monetary policy is described by a simple Taylor rule for the nominal interest rate that reacts to inflation and output detrended by the Beveridge-Nelson stochastic trend $c_{Pt} = a_{Pt}$ described below,

$$i_t = i^* + \phi_\pi(\pi_t - \pi^*) + \phi_y[c_t - c_{Pt} - (1-\alpha)(n-\Delta)] \tag{23}$$

for an interest-rate level i^* consistent with a positive steady-state inflation rate π^*.

Fiscal policy runs a balanced budget in that $T_{ht} = \tau_h W_t N_t$ and $T_{ft} = \tau_f W_t N_t$.

4.1.4. Competitive equilibrium

The markets for goods and labor clear, $Y_t = C_t$ and $N_t = \int_0^1 N_t(i)di$, or, combining the individual production functions and demand curves,

$$c_t = a_t + (1-\alpha)(n_t - \Delta_t)$$

As a corollary, aggregate dividends $D_t = \int_0^1 D_t(i)di$ are characterized by the condition

$$D_t = \int_0^1 \left[\frac{P_t(i)}{P_t} Y_t(i) - W_t N_t(i) \right] di = C_t - W_t N_t$$

Log productivity is composed of a permanent component a_{Pt} and a transitory component a_{Tt} such that $a_t = a_{Pt} + a_{Tt}$, with

$$\begin{aligned} a_{Pt+1} &= \mu + a_{Pt} + (1-\theta)e_{t+1} \\ a_{Tt+1} &= \rho_a a_{Tt} + \theta e_{t+1} \end{aligned}$$

with average drift μ and persistence ρ_a, where θ indexes the extent to which a productivity shock $e_t \sim Niid(0, \sigma^2)$ has a permanent effect. For example, $\theta = 0$ is associated with random-walk productivity, while $\theta = 1$ is associated with the typical trend-stationary specification (e.g., Galí, 2008). Consistent with definition (2), θ also indexes the amount of uncertainty in the consumption process that can be removed by monetary policy in that monetary policy in this New Keynesian setup can only have transitory effects.

With this structure in place, I now look for a competitive equilibrium, i.e., a sequence of state-contingent allocations $\{C_t, S_{ct}, H_t, S_{ht}, N_t(i)\}_{t=0}^{\infty}$ and prices $\{W_t, P_t(i)\}$ for each monopolistically competitive firm $i \in (0, 1)$ such that for each date t, *(a)* the choice of prices and labor demand solves the individual firm's problem, *(b)* the choice of market and home consumption solves the individual consumer's problem, *(c)* the goods and labor markets clear, *(d)* market and home consumption habits evolve according to (15), *(e)* price dispersion evolves according to (14), *(f)* the nominal rate follows rule (23), and *(g)* the fiscal authority runs a balanced budget.

4.1.5. Approximate competitive equilibrium

Although I will solve the model with a global projection method, it is useful to gain analytic insight into the model's mechanism by a first-order approximation around the risky steady state as described by Lopez et al. (2020). (Since the asset pricing implications of the model are crucial here, an approximation around the deterministic steady state would be inappropriate.)

It turns out that there are particular values for the free parameters ξ_1 and ξ_2 such that the approximate solutions for consumption and inflation are

$$c_t - a_t = (1-\alpha)n + \psi_c a_{Tt}, \qquad \pi_t = \pi^* + \psi_\pi a_{Tt} \tag{24}$$

to a first-order perturbation around the risky steady state. In words, the equilibrium dependence of consumption and inflation on surplus consumption is zero.[15]

Therefore, by an appropriate choice of free parameters ξ_1 and ξ_2 (or, equivalently, of the levels of surplus market and home consumption S_c and S_h), the model is able to produce a *separation* between risk premia and quantity dynamics that preserves the implications for quantities and inflation of the basic New Keynesian model with CRRA utility. Intertemporally, households balance strong intertemporal substitution and precautionary saving motives, thereby disciplining, via ξ_1, the variation in the real risk-free rate

$$r_t = -\ln E_t m_{t+1} - \frac{1}{2} var_t(m_{t+1}) = -\ln(\beta) + \gamma\mu - \frac{1}{2}\gamma(1-\rho_s - \xi_1/\gamma) + \gamma\psi_c a_{Tt} - \xi_1 \hat{s}_{ct}$$

and hence in households' incentives to save and consume, through (21). For example, a choice of $\xi_1 = 0$ would imply that the risk-free rate is not driven by the surplus consumption ratio, as would be implied by a model without habits.

Statically, households balance a strong aversion to fluctuations in market and home consumption across states, which is key to avoiding excessive variation in the real wage rate (22). As discussed, the choice of the sensitivity function for surplus home consumption minimizes the distance between $\hat{s}_{ht} - \hat{s}_{ct}$ and $\xi_2 \hat{s}_{ct}$, for a free parameter ξ_2. In fact, in equilibrium it implies to a first-order approximation that such a distance is zero, i.e.,

$$\hat{s}_{ht} - \hat{s}_{ct} = \xi_2 \hat{s}_{ct}$$

For example, a choice of $\xi_2 = 0$ would imply a marginal rate of substitution between consumption and labor identical to the one under CRRA utility. This property prevents households from excessively varying labor hours to smooth out consumption, and hence solves the puzzle documented by Lettau and Uhlig (2000) when uniting nonlinear habits and a production economy.

[15] To a first-order approximation price dispersion is a trivial process, $\Delta_t = 0$, so I omit the dependence of the approximation solution on this state.

It follows that the particular parameterization required to achieve an approximate macro-finance separation in my model is very close to the point $\xi_1 = \xi_2 = 0$—though not identical to it because of the nominal rigidities. In fact, a choice of

$$\xi_1 = \frac{\gamma}{S}(1+\psi_c\theta)\psi_\pi\theta\sigma^2, \qquad \xi_2 = \frac{(e^{-L_3}-e^{-L_2})\beta e^{(1-\gamma)\mu}\eta(1-\rho_s)}{(e^{-L_3}-\beta e^{(1-\gamma)\mu}\eta\rho_s)(e^{-L_2}-\beta e^{(1-\gamma)\mu}\eta)}$$

where L_2 and L_3 are small terms proportional to σ^2, implies (24) with the approximate equilibrium coefficients

$$\psi_c = -\frac{[\gamma(1-\rho_a)+\phi_y](1-\beta e^{(1-\gamma)\mu}\rho_a)+\varphi(\phi_\pi-\rho_a)}{[\gamma(1-\rho_a)+\phi_y](1-\beta e^{(1-\gamma)\mu}\rho_a)+\kappa(\phi_\pi-\rho_a)}$$
$$\psi_\pi = -\frac{(\kappa-\varphi)[\gamma(1-\rho_a)+\phi_y]}{[\gamma(1-\rho_a)+\phi_y](1-\beta e^{(1-\gamma)\mu}\rho_a)+\kappa(\phi_\pi-\rho_a)}$$

(The online appendix derives these formulas.) Parameters κ and φ are transformations of other deep parameters and, in the absence of risk ($\sigma = 0$), reduce to $\kappa = (1-\beta e^{(1-\gamma)\mu}\eta)(1-\eta)[\gamma(2-\alpha)+\alpha]/\eta(1-\alpha+\alpha\varepsilon)$ and $\varphi = 0$ as in standard linearizations of the basic New Keynesian model with CRRA utility (e.g., Galí, 2008).

It follows that the equilibrium allocation is approximately observationally equivalent to a model with CRRA households. Figure 3a illustrates this property by comparing the nonlinear solution for consumption, inflation, and dividends in my model and in a version of the model without habits. The solution for these cash flows is virtually identical. An economist interested only in the model's implications for quantities and inflation would be unable to tell apart the model with habits and its CRRA version.

4.2. Quantitative implications

To solve for the competitive equilibrium I use a Smolyak sparse-grid collocation method with an adaptive grid to project the global solution of the model onto the subspace spanned by a basis of Chebyshev polynomials of up to degree sixteen. Following the best practice to solve models with Campbell-Cochrane preferences, I consider a large grid that approaches values close to zero for market surplus consumption, which is the main driver of the stochastic discount factor. Conditional expectations are computed by an accurate Gauss-Hermite quadrature.

Table 3 lists the calibration of the deep parameters of the model. The production side of the economy uses standard values from Galí (2008). Parameters related to the pricing kernel are calibrated as in Campbell and Cochrane (1999), with an EIS of 1/2, a rate of time preference to match the target mean risk-free rate, spillover parameters to produce approximate macro-finance separation, and habit persistence to capture the persistence of observed price-dividend ratios. I estimate the three parameters $[\theta, \rho_a, \sigma]$ that drive the dynamics of productivity by a GMM strategy that minimizes the distance between model-implied moments and observed variances of growth rates over 1- to 5-year horizons for the quantities of interest—consumption growth, dividend growth, and inflation. Figure 3a

	Parameter		Value
Preference block	γ	Inverse EIS	2
	β	Time preference	0.9939
	ρ_s	Habit persistence	0.9917
	ξ_1	Intertemporal spillover	0.0001
	ξ_2	Static spillover	-0.0140
New Keynesian block	$1-\alpha$	Labor share in value added	2/3
	ε	Elasticity of substitution in Dixit-Stiglitz aggregator	6
	$1/(1-\eta)$	Average price duration (in months)	9
	π^*	Average inflation rate (% p.a.)	2
	ϕ_π	Policy response coefficient to inflation movements	1.5
	ϕ_y	Policy response coefficient to output movements	0.5/12
Exogenous block	μ	Mean productivity growth (% p.a.)	1.6
	ρ_a	Persistence of the conditional mean of productivity growth	0.794
	σ	Conditional volatility of productivity	0.0128
	$1-\theta$	Fraction of permanent productivity shocks	0.31

β matches an average real risk-free rate of 0.69% per year.
ρ_s matches a 12-month autocorrelation of price-dividend ratios of .905.
μ matches an average annual per capita real consumption growth of 1.59%.
ρ_a, σ, and θ are estimated by GMM to match the 1- to 5-years-ahead volatility of per capita real consumption growth, real dividend growth, and inflation over 1980-2019.
The calibration for ξ_1 and ξ_2 implies $S_c = 0.078$ and $S_h = 0.295$.

Table 3: Deep parameters and their calibration (monthly frequency). Data for real consumption growth use BEA-NIPA data over the period 1980-2019 for per capita personal consumption expenditure in nondurables and services, and are deflated by the CPI. Monthly simulated data are aggregated to an annual frequency and are matched to the corresponding data moments.

reports the associated fit of the model, which captures well the autocorrelation and volatility of cash flows observed over the 1980-2019 period. The figure also reports cash flow volatilities in the model with CRRA utility based on the parameters estimated for the habit model. The implications for dividends of the model with and without habits are nearly identical, while the properties of consumption and inflation are indistinguishable, thereby confirming numerically the equivalence discussed in Section 4.1.5.

The model preserves the successes of Campbell and Cochrane (1999) in solving the equity premium and the equity volatility puzzles. At an annual frequency, the equity premium is 7.25 percent in the model, close to the 7.22 percent in the data reported in Figure 1a. Log price-dividend ratios have a mean and standard deviation of 3.16 and 0.28 in the model vs. 3.75 and 0.38 in the data, and they are strong predictors of future returns. A present-value decomposition of the variance of price-dividend ratios implied by a VAR(1) accounts for 96 percent of the variance to covariance of price-dividend ratios with future returns in the model vs. 92 percent in the data.

Importantly, the model captures the facts documented in Section 3 about the term structure of welfare costs. Figures 1b and 3 report the average term structures of hold-to-maturity equity and real bond returns and of the welfare cost of fluctuations implied by the model, as well as the 5-95 percentile range of the term structure of welfare costs, constructed as in Figure 2d by regressing the corresponding hold-to-maturity excess returns on the first two principal components of equity yields and the first principal component of bond yields. The representative household is particularly sensitive to fluctuations in market income, as the countercyclical labor share exacerbates the procyclicality of corporate profits after a productivity shock, but fears cash flow uncertainty less in the long run because dividends and consumption are cointegrated, and hence have the same riskiness in the long run. Therefore, this simple model is able to capture the main empirical properties of the term structure of welfare costs previously documented, including its countercyclical variation.

Accordingly, Table 4 reports the implications of the model for the marginal costs of fluctuations with fixed labor $\overline{L}_t^{\mathcal{N}}$ and with variable labor $L_t^{\mathcal{N}}$, for uncertainty around all coordinates up to several horizons. Even though the costs of uncertainty of the two definitions are quantitatively similar at very short horizons, with costs of 0.4-0.5 percentage points of growth at a one-year horizon, the costs of lifetime uncertainty in the definition with variable labor are an order of magnitude lower than in the definition with fixed labor, namely, 1.5 percentage points of perpetual growth when labor is fixed and 0.15 percentage points of perpetual growth when labor is variable.

Horizon (years), $\mathcal{N}$ up to	0.5	1	1.5	2	5	10	30	∞
Welfare costs, fixed labor, $\overline{L}^{\mathcal{N}}$ (% p.a.)	0.54	0.52	0.51	0.50	0.50	0.56	1.05	1.53
Welfare costs, variable labor, $L^{\mathcal{N}}$ (% p.a.)	0.51	0.39	0.32	0.27	0.14	0.10	0.11	0.15

Table 4: Model-implied marginal cost of fluctuations at all periodicities $n \in \mathcal{N}$.

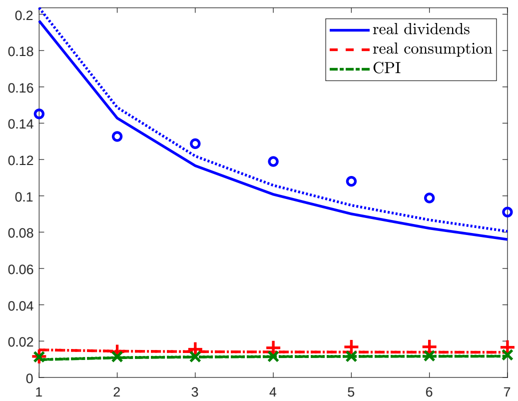

(a) *Average term structures of volatilities in quantities; model (solid and dashed lines), model under CRRA utility (dotted lines), and data (markers).*

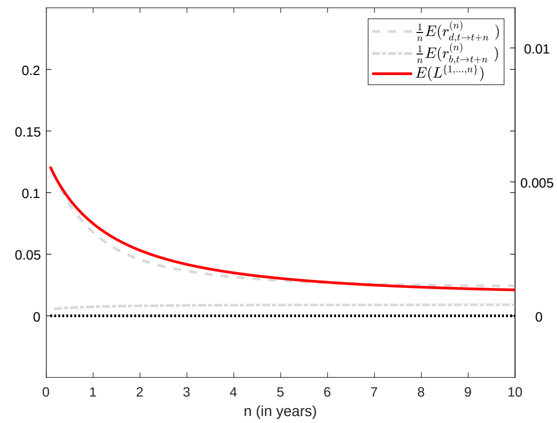

(b) *Average term structures of equity, interest rates, and welfare costs.*

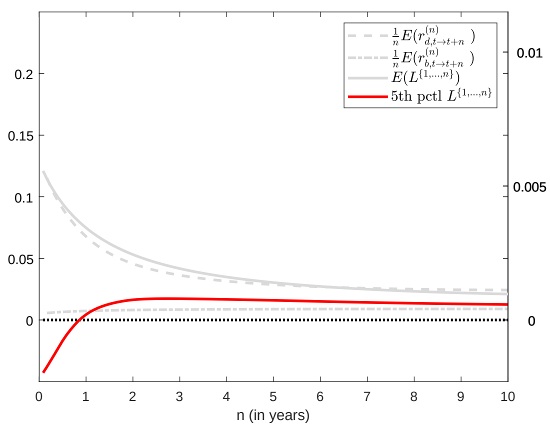

(c) *5th percentile of the term structure of welfare costs.*

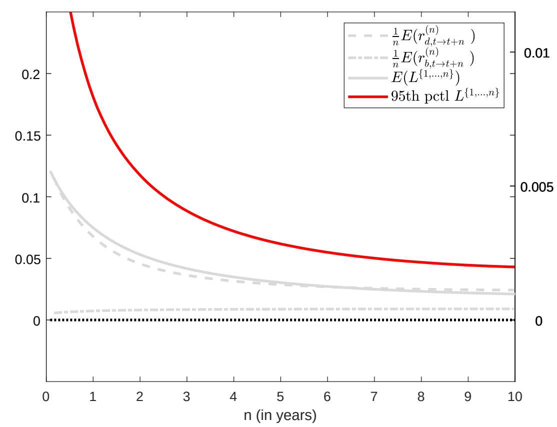

(d) *95th percentile of the term structure of welfare costs.*

Figure 3: Term structures of cash flows and asset prices. The nth cash flow term-structure component is defined as $n^{-1}var[\ln(D_{t+n}/D_t)]$ for a cash flow process D. Model-implied term structures of equity (dashed line), interest rates (dash-dotted line), and welfare costs (solid line); average and 5-95 percentile range. Annualized values on the left axis for the term structures of cash flows, equity, and interest rates and on the right axis for welfare costs.

4.2.1. Inspecting the mechanism

The production economy with sticky prices provides a macroeconomic model that matches the observed volatility and autocorrelation of cash flows (consumption growth, dividend growth, and inflation). Because of the approximate macro-finance separation, the model's quantity and inflation implications and responses to a productivity shock are standard. The central mechanism relies on nominal rigidities that produce a countercyclical labor share after a productivity shock. Therefore, dividend claims pay off badly in a downturn when marginal utility is high, and are therefore risky investments. However, in the periods following the shock, corporate profits increase as more and more firms are able to adjust their prices. The labor share is a stationary process, and hence the payoff of long-duration dividend strips is less procyclical than the payoff of short-duration strips.

The slow-moving external habit formation magnifies the riskiness of dividends into large and countercyclical risk premia. Moreover, for sufficiently long durations, habits imply that both equities and real bonds are expected to pay off a negative return when marginal utility is high; so they are risky in the very long run. In fact, habits make prices drop during a downturn, and the more so the longer the claim's duration, because people will slowly get used to the lower consumption level; therefore, people will want to anticipate consumption and will require compensation for shifting resources in the future. Because of this habit effect, the model produces a term structure of one-period equity returns that is U-shaped, with a negative slope in the short to medium run, driven by the cyclicality of dividends, and a positive slope for longer maturities, driven by this habit effect. The same habit effect generates a positively sloped term structure of real rates.

The nonlinear habits also generate the countercyclicality of costs and the procyclicality of the slope of the term structures we see in the data. In bad times, as consumption falls close to habits and dividends drop, risk premia increase and future dividends are expected to recover, and hence the slope of hold-to-maturity equity returns for the observable maturities becomes more negative, consistent with the evidence in Section 3. Conversely, the slope of the term structure of welfare costs can become positive in normal times. At the same time, the U-shaped term structure of one-period risk premia implies that risk premia also increase for longer horizons and, consequently, it turns out that market equity premia increase by more than short-run equity premia, giving rise to the countercyclicality of the one-period equity term premia documented by Gormsen (2021).[16]

4.3. Robustness: Zero-lower bound on the nominal interest rate

The documented evidence on the term structure of welfare costs covers periods in which the policy rate was constrained by the zero lower bound on the nominal interest rate. Accordingly, I carry out two additional exercises. First, I note that the estimates for the welfare

[16]Since my focus here is on hold-to-maturity returns, consistent with the measures of the welfare costs of fluctuations, I do not explore further these implications for one-period returns beyond highlighting their consistency with the evidence in Gormsen (2021) and their support for a habit specification rather than an Epstein-Zin specification of preferences. The online appendix explores these implications in greater detail; see also Lopez et al. (2021).

costs of short-term fluctuations excluding the 2008m12-2015m11 zero-lower-bound episode have properties similar to those documented for the whole period, as seen in Figure 2d. Namely, in the data, over the unconstrained period the average hold-to-maturity excess equity returns on 6- to 24-month horizons are 15.9, 14.4, 12.3, and 10.7 percent (versus 13.3, 12.8, 11.2, and 9.8 percent over the whole period) and remain well above the equity premium of 5.4 percent (versus 7.4 percent over the whole period). The associated welfare costs of short-term fluctuations are 0.73, 0.66, 0.57, and 0.49 percent; if anything, they are larger than over the full sample (0.61, 0.59, 0.52, and 0.45 percent).

Second, I explore the robustness of the previous results to the additional nonlinearities in the model generated by a zero-lower-bound constraint on the nominal interest rate set by the central bank. Namely, I consider a version of the model in which monetary policy sets the interest rate $i_t = \max(i_t^*, 0)$, where the latent interest rate i_t^* follows rule (23).[17] I maintain the same parameterization and solve the model by the same solution strategy described above. The results remain very similar. In particular, the model spends 13 percent of the time at the zero lower bound; the inflation rate is 3.8 percent lower on average and has 9.8 percent more variance, consistent with lower and more elastic inflation when the interest rate is constrained; and the real risk-free rate is 3.6 percent higher on average and has 14.6 percent less variance, consistent with a central bank that is unable to lower the real risk-free rate sufficiently when constrained. The consumption and dividend growth rates are unchanged on average, and their variances are 4.6 percent lower and 0.4 percent higher, respectively. These changes in cash flows are associated with small changes in the term structures of welfare costs. The term structures of hold-to-maturity equity risk premia shift down between 3 and 20 basis points depending on the maturity, a modest change that affects none of the previous results.

4.4. Robustness: Demand shocks

So far, the model has presented a simple mechanism that focuses on the effects of productivity shocks to capture the evidence about the term structure of welfare costs. Still, a full-fledged model would incorporate more shocks, including demand shocks, to capture the data more comprehensively. For example, the model with only productivity shocks understates the correlation between consumption growth and inflation (-46 percent in the model vs. -38 percent in the 1980-2019 data). Moreover, as argued by Campbell et al. (2020), a mix of demand and supply shocks can capture changing correlation patterns between consumption growth and inflation and between stock and bond returns. Indeed, the implications for the term structure of welfare costs are robust to including demand shocks to capture these correlation patterns. As will be shown in Section 5, the policy implications are likewise similar.

[17] The zero lower bound introduces additional equilibria associated with a deflationary regime, but we focus here on the global solution in the so-called anchored-expectations regime, as defined, for example, in Aruoba et al. (2018), who find limited evidence in favor of a deflationary regime for the US during the 2008m12-2015m11 lower-bound episode.

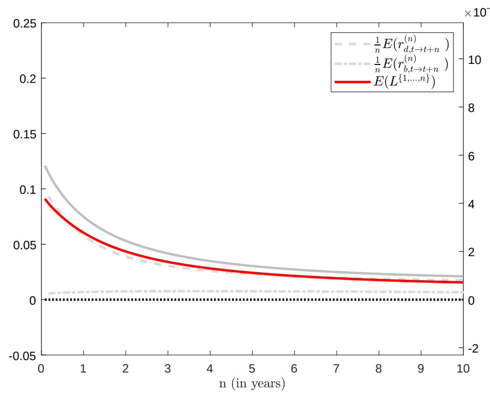

(a) *Average term structures of welfare costs.*

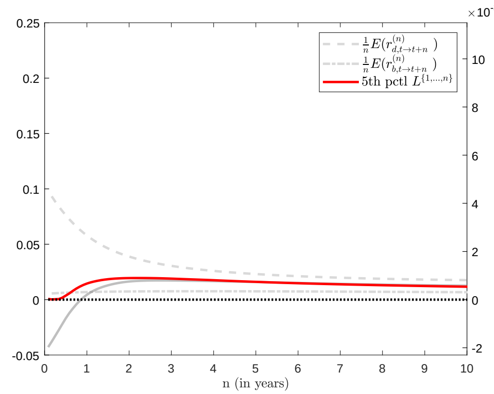

(b) *5th percentile of the term structure of welfare costs.*

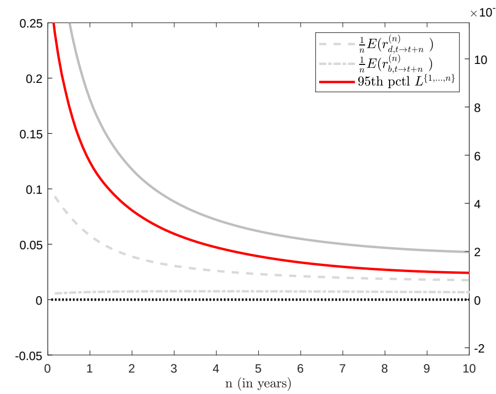

(c) *95th percentile of the term structure of welfare costs.*

Figure 4: Term structures of the welfare costs of fluctuations in the model with demand shocks. Model-implied term structures of equity (dashed line), interest rates (dash-dotted line), and welfare costs (solid line); average and 5-95 percentile range. Annualized values on the left axis for the term structures of cash flows, equity, and interest rates and on the right axis for welfare costs. The gray lines represent the corresponding object in the version with only productivity shocks.

Thus, I extend the model to include demand shocks, Λ_t. Preferences are now described by

$$U_0 = E_0 \sum_{t=0}^{\infty} \beta^t \Lambda_t \left(\frac{[C_t - X_{ct}]^{1-\gamma}}{1-\gamma} + \chi \frac{[H_t - X_{ht}]^{1-\gamma}}{1-\gamma} \right)$$

where Λ_t is an exogenous preference shock, whose log growth rate $\phi_t = \Delta \ln(\Lambda_t)$ follows the autoregressive process

$$\phi_{t+1} = \rho_\phi \phi_t + \gamma(1 + \lambda_{ct})\varepsilon_{\phi t+1} \tag{25}$$

where $\varepsilon_{\phi t}$ is i.i.d. normally distributed with mean 0 and variance σ_ϕ^2 and independent of the productivity shock. Under the specification for the sensitivity function

$$\lambda_{ct} = \max\left(0, \frac{\bar{\sigma}_c^2}{\bar{\sigma}_{ct}^2} \frac{1}{\bar{S}_c} \sqrt{1 - 2\hat{s}_{ct}} - 1\right), \qquad \bar{S}_c = \sqrt{\frac{\gamma \bar{\sigma}_c^2}{1 - \rho_s - \xi_1/\gamma}}$$

where $\bar{\sigma}_{ct}^2 = var_t(\varepsilon_{ct+1} - \varepsilon_{\phi t+1})$, the spillover of s_{ct} on the real rate is once again controlled by the free parameter ξ_1, by analogous logic as in the baseline model.

To pin down the parameters ρ_ϕ and σ_ϕ, I first set the persistence $\rho_\phi = \rho_u$, to preserve the estimated serial correlation of quantities of the baseline model. Furthermore, since the model with only productivity shocks understates the correlation between consumption growth and inflation, I choose the standard deviation σ_ϕ to match the observed correlation between consumption growth and inflation.

In line with the evidence in Campbell et al. (2020), the model can now generate decade-long spells with negative correlations between stock and bond returns. I simulate 10,000 monthly samples with a length of 10 years and compute in each sample the correlation between 5-year nominal bond excess returns and the aggregate stock market excess return. In 18.6 percent of the simulated samples, the model generates a negative correlation. Relatedly, the correlation between the consumption level detrended by its stochastic trend, $c_t - \bar{c}_{Pt}$, and inflation turns positive in 5.7 percent of the samples.

While the inclusion of demand shocks improves the model's ability to fit these empirical regularities, the implications of the model for the term structure of welfare costs remains similar to those in the model with only productivity shocks. Figure 4 shows that average and tail properties of the term structures of welfare costs are preserved, even though the presence of demand shocks slightly reduces the procyclicality of the slope. In this sense, the implications for the asset prices of interest are therefore robust to the inclusion of demand shocks.

4.5. Why not internal habits?

If habits are internal, the flexible-price equilibrium is the Pareto optimum. Contrary to the external-habit specification, however, I am unable to defend the internal-habit specification from the critique by Lettau and Uhlig (2000), and I therefore rule out such a specification. Thus, an internal-habit specification is unable to generate realistic asset pricing implications in the sticky-price production economy without dramatically distorting the positive implications of the model for output and inflation.

In fact, the only way the internal-habit model can leave quantity and inflation dynamics unchanged relative to a CRRA specification is at $\gamma = 1$. But under a unitary elasticity of intertemporal substitution, one can show that the model reduces to a log utility CRRA specification in all its quantity, asset pricing, and welfare implications. (The appendix derives this result.) The resulting model has trivial asset pricing implications. In this sense, even though in a pure-exchange economy Campbell and Cochrane (1999) had no strong reason to prefer the external- over the internal-habit specification, the Lettau and Uhlig (2000) critique rejects the internal-habit specification in a production economy.

4.6. Why not Epstein-Zin preferences?

The choice of external habits to explain the term structure evidence may seem surprising at first, as some authors document the challenges of the habit framework in producing a downward-sloping term structure of equity (e.g., Binsbergen et al., 2012; Gormsen, 2021). Those results, however, are derived in endowment economies with random walk dividend streams. Once I inject a mean-reverting component in dividends, as endogenously generated by the production economy, I depart from those benchmark models. In particular, by using the model to make the properties of cash flows match their volatility and autocorrelation in the data, I am able to generate the slope and cyclicality of hold-to-maturity equity returns.

Furthermore, the recent evidence by Gormsen (2021) rules out Epstein-Zin preferences as an alternative asset pricing ingredient to habits. In fact, the online appendix and Lopez et al. (2021) explore the implications of a New Keynesian model with habits similar to the one proposed here, and show that it is consistent with the countercyclicality of one-period equity term premia documented in the data by Gormsen (2021). This property is naturally generated by habits, even after the properties of cash flows have been changed by the production economy to generate the right slope and cyclicality of equity yields.

In contrast, in production economies with Epstein-Zin preferences, the change in correlation necessary for flipping the sign of the slope of the equity term structure tends to operate also at long horizons. Indeed, Gormsen shows how recent examples by Hasler and Marfè (2016) and Ai et al. (2018), which are able to generate a downward-sloping equity term structure by changing the dividend process, display as a consequence the wrong cyclicality of the term structure of equity term premia. The extension of those setups to a nominal production economy therefore seems a much more challenging avenue.

5. Is It Desirable to Eliminate Fluctuations?

The welfare cost calculations of Sections 3 and 4.2 provide a measure of the potential benefits of the removal of fluctuations, but they do not address whether such removal is desirable. In fact, I argue that the true relevance of these measures lies in their role as a set of empirical features that any macroeconomic model aiming to evaluate the priority of various stabilization policies should align with. Indeed, the structural model of Section 4, which was able to account for the average welfare costs and their documented cyclicality, as well as to reconstruct welfare costs at all maturities, can be used to take the additional step of asking about the optimal stabilization policy.

As I turn to investigating the optimal stabilization policy in the model laid out in Section 4, note that the presence of two externalities (nominal rigidities and external habits) motivates a nontrivial policy tradeoff.[18] In this context, aggregate (detrended) welfare $V_t \equiv U_t/A_t^{1-\gamma}$ can be written as

$$V_t = \frac{e^{(1-\gamma)\hat{s}_{ct}}}{1-\gamma}\left[e^{(1-\gamma)(c_t-a_t)} + \chi_0 e^{(1-\gamma)(\hat{s}_{ht}-\hat{s}_{ct})}\left(1 - e^{\frac{1}{1-\alpha}(c_t-a_t)+\Delta_t}\right)^{1-\gamma}\right] + \beta E_t e^{(1-\gamma)\Delta a_{t+1}} V_{t+1}$$

In the version of the model with demand shocks, detrended welfare will instead be defined as $V_t = U_t/\Lambda_t A_t^{1-\gamma}$. The constrained-efficient allocation maximizes welfare subject to the optimal price-setting equation (12), to the law of motion of price dispersion (14), and to the definition of habits (15).

In practice, I will study the consequences of two simple monetary policy regimes—one that closes the consumption gap relative to the flexible-price level (or, equivalently, one that minimizes the volatility of inflation) and one that minimizes the volatility of consumption (or, equivalently, one that minimizes risk premia variation). The corresponding solution is constructed by an analogous global projection method as used in Section 4. These two simple regimes can be implemented by an appropriate choice of the reaction parameters in Taylor rule (23). The central bank is able to implement the inflation targeting regime by an extreme anti-inflationary response, $\phi_\pi \to \infty$, while the risk premium targeting regime can be implemented by an extreme response to movements in detrended consumption, $\phi_y \to \infty$. In this sense the two policies are polar opposites.

5.1. Fill the consumption gaps (or: inflation targeting)

Under an inflation targeting regime consumption equals the flexible-price level, and hence inflation is constant ($\pi_t^{it} = \pi^*$). The associated values for market consumption c_t is implicitly defined as the solution of equation

$$0 = \ln\left(\frac{\varepsilon(1-\tau^f)}{\varepsilon-1}\frac{\chi_0}{1-\alpha}\right) + \frac{\gamma(1-\alpha)+\alpha}{1-\alpha}(c_t^{it}-a_t) - \gamma\ln(1-e^{\frac{1}{1-\alpha}(c_t^{it}-a_t)}) + \gamma(\hat{s}_{ct}^{it} - \hat{s}_{ht}^{it})$$

This policy maximizes the level of consumption for a given level of employment in that it minimizes price dispersion, but it also generates a highly volatile surplus consumption ratio, $\hat{s}_{ct+1}^{it} = \rho_s \hat{s}_{ct}^{it} + \lambda_{ct}^{it}\sigma\varepsilon_{t+1}$, and hence volatile risk aversion.

[18] In a different setting Leith et al. (2012) study monetary policy tradeoffs in a New Keynesian model when one-period linear habits in only market consumption are external. Nevertheless, inflation targeting remains very close to the optimal policy in their model. Their result follows from the fast-moving habits they consider, which are unable to generate realistic discount-rate variation and welfare costs of fluctuations.

5.2. Minimize consumption volatility (or: risk premium targeting)

A key property of the New Keynesian model is the stationarity of the labor share. It follows that a regime that removes as much consumption risk as possible has to preserve the property that consumption and productivity are cointegrated. Therefore, consumption has to inherit the permanent component of productivity, which represents the maximum amount of consumption risk that monetary policy can remove.

In the context of the model, removing consumption risk (so $c_t^{rpt} = a_{Pt} + (1-\alpha)(n-\Delta)$ for some steady-state price dispersion $\Delta > 0$) is equivalent to stabilizing surplus consumption. Under such a risk premium targeting regime consumption equals the permanent component of productivity, and hence $\varepsilon_{ct}^{rpt} = (1-\theta)e_t$. Coincidentally, such a policy would likewise stabilize the maximum risk-return tradeoff

$$\sigma_t(m_{t+1}^{rpt}) = \gamma(1+\lambda_{ct}^{rpt})\sigma_t(\varepsilon_{ct+1}^{rpt})$$

that dominates Sharpe ratios $|\ln E_t R_{t+1}^e / \sigma_t(r_{t+1}^e)|$. The policy is therefore equivalent to a regime that targets risk premia by stabilizing the maximum risk-return tradeoff. In this sense, this regime can be viewed as a *risk premium targeting* regime.

This policy minimizes the volatility of surplus consumption $\hat{s}_{ct+1}^{rpt} = \rho_s \hat{s}_{ct}^{rpt} + \lambda_{ct}^{rpt}(1-\theta)\sigma\varepsilon_{t+1}$, and hence of risk aversion, but it generates a strictly positive degree of price dispersion that causes consumption losses for a given level of employment.

5.3. Approximate welfare criterion

Although I will use a global solution to study the effects of the two policies on welfare, an approximation of welfare is useful to gain insight into its determinants. The appendix derives a quadratic approximation around the risky steady state of average (detrended) per-period welfare. Accordingly, I consider average approximate welfare losses $\mathbb{L} \equiv -E(V_t)/(CS_c)^{1-\gamma}$, up to a term independent of policy,

$$\begin{aligned}\mathbb{L} \approx\ & \frac{1}{2}\frac{\eta\varepsilon(1-\alpha+\alpha\varepsilon)}{(1-\eta)^2(1-\alpha)}var(\pi_t) + \frac{1}{2}\frac{\gamma(2-\alpha)+\alpha}{1-\alpha}var\left(c_t - a_t\right) \\ & + \frac{\gamma-1}{2}var(s_{ct}) + \frac{(\gamma-1)(1-\alpha)}{2}var(s_{ht}) - (\gamma-1)cov(c_t - a_t, s_{ht} - s_{ct}) + t.i.p.\end{aligned} \tag{26}$$

The first row of (26) is the standard quadratic expansion of welfare (e.g., Galí, 2008) and reflects the objectives of removing consumption volatility and closing the consumption gap relative to the flexible-price level in the absence of habits, $c_t - a_t$. The second row of expression (26), however, reflects additional objectives. In fact, there are two separate departures from the efficient dynamics of the benchmark RBC model. First, in a sticky-price environment inflation volatility is approximately equivalent to cross-sectional dispersion in prices, which in turn is associated with inefficient employment. Second, when habits are external people fail to internalize the fact that higher consumption *also* has a habit effect that means a higher marginal value of consumption (Ljungqvist and Uhlig, 2000). When facing good news about current and future states, people should not increase their consumption

	Welfare	Cons. Loss, ϱ	Welfare, No habits	Cons. Loss, ϱ
Baseline model				
- Taylor rule	-265.2	12.0%	-222.0	-0.7%
- Fill consumption gap (inflation targeting)	-302.9	25.3%	-221.6	-0.9%
- Minimize consumption volatility (risk premium targ.)	-235.1	0.0%	-223.5	0.0%
Model with demand shocks				
- Taylor rule	-262.5	11.9%	-222.2	-0.7%
- Fill consumption gap (inflation targeting)	-296.3	24.0%	-221.8	-0.8%
- Minimize consumption volatility (risk premium targ.)	-233.0	0.0%	-223.7	0.0%

Table 5: Welfare under different policies and model versions. The consumption loss ϱ is relative to the risk premium targeting regime.

level as much as they want to because they will become addicted to the higher level. It follows that, in an external-habit environment, consumption fluctuates too much.

Therefore, to minimize the welfare loss (26), the central bank wants to stabilize inflation, π_t, and hence the consumption gap relative to the flexible-price level in the absence of habits, $c_t - a_t$, but it also wants to stabilize consumption volatility, $c_t - E_{t-1}c_t$ and $h_t - E_{t-1}h_t$, to stabilize the surplus consumption ratios s_{ct} and s_{ht}. Since achieving both goals is unfeasible, as closing the gap implies that consumption has the same variance as productivity, the optimal monetary policy trades them off to maximize welfare.

5.4. Fill the gaps or remove volatility?

In both the baseline version of the model and in the model with demand and productivity shocks, I find that removing consumption volatility dominates closing the consumption gap. Denoting by V_t^{it} (detrended) welfare under the inflation targeting regime and by V_t^{rpt} welfare under the risk premium targeting regime, I express these welfare differences in an economically interpretable measure by defining the variable ϱ that solves

$$\frac{A_t^{\gamma-1}}{1-\gamma}E\left(\sum_{t=0}^{\infty}\beta^j[(e^{\varrho}C_{t+j}^{it})^{1-\gamma}(S_{ct+j}^{it})^{1-\gamma}+\chi(e^{\varrho}H_{t+j}^{it})^{1-\gamma}(S_{ht+j}^{it})^{1-\gamma}]\right)=E(V_t^{rpt})$$

so that $\varrho = \log[E(V_t^{rpt})/E(V_t^{it})]/(1-\gamma) = \log[E(V_t^{it})/E(V_t^{rpt})]$, where the last equality uses the baseline parameterization $\gamma = 2$. The value ϱ can therefore be interpreted as the percent increase in market and home consumption that would make the household indifferent between the inflation targeting and the risk premium targeting regimes. I define analogously the difference in welfare relative to the equilibrium under the baseline Taylor rule.

Table 5 reports welfare under the baseline Taylor rule parameterization, under the policy regime that fills the consumption gap, and under the policy regime that minimizes consumption volatility. The policy that minimizes consumption volatility improves welfare relative to the baseline equilibrium allocation by an amount equivalent to a 12 percent increase in the lifetime market and home consumption streams. The inflation targeting regime actually reduces welfare relative to the baseline Taylor rule. In fact, removing consumption volatility increases welfare relative to an inflation targeting regime by an amount equivalent to a 25 percent increase in the lifetime market and home consumption streams.

To better understand these results, Table 5 also reports the results in a version of the model without habit formation in either market or home consumption. In such a model, inflation targeting minimizes the output loss due to price dispersion, and hence can achieve a higher consumption level relative to other policies. Indeed, inflation targeting dominates both the baseline and the risk premium targeting regimes. In contrast, while this level effect on consumption also holds in the model with habits, the inefficient volatility in surplus consumption necessary to account for the measured welfare costs of consumption fluctuations outweighs the beneficial effects of minimizing price dispersion. Risk premium targeting ends up dominating inflation targeting as a consequence of the other welfare objectives reflected in equation (26). People hate consumption volatility so much that a policy that achieves the flexible-price equilibrium is suboptimal relative to a policy that removes as much consumption volatility as possible. In fact, a move from the baseline Taylor rule to an inflation targeting regime is welfare decreasing.

The welfare dominance of the risk premium targeting regime over the inflation targeting regime persists in the version of the model with a richer shock structure. Table 5 shows how the model with demand shocks has welfare improvement nearly identical to that of risk premium targeting relative to both the Taylor rule and the inflation targeting regimes.

6. Conclusion

Lucas (1987, 2003) introduced the notion of the cost of aggregate fluctuations as a thought experiment to provide an assessment of the tradeoff between growth and macroeconomic stability. Analogously, the term structure of the cost of fluctuations requires little structure to reveal the tradeoff between growth and macroeconomic stability at different time horizons. Recent derivative securities provide direct information about this term structure and, therefore, new insight into an old question (the tradeoff between growth and stability) and evidence to study a new question (the tradeoff between stability at different horizons).

Asset markets suggest that the potential gains from greater economic stability are not trivial, especially in the short run and during downturns such as in the early 2000s and during the Great Recession.

The result that the marginal cost of fluctuations is a linear combination of risk premia makes one of the main tasks of macroeconomics—the assessment of macroeconomic priorities—inextricably linked to finance. The negative slope and countercyclical variation of the estimated term structure of welfare costs cannot be easily captured by leading asset pricing theories and therefore represents a puzzling piece of evidence with seemingly crucial welfare consequences. For example, I showed how two models that are observationally equivalent in their quantity implications but that differ in their asset pricing implications prescribe fundamentally different policies. The required risk sensitivity to account for measures of the welfare costs of consumption fluctuations can overturn the robust result in New

Keynesian models that inflation targeting is nearly optimal.

Appendix

A. Theoretical results

A.1. Proof of proposition 1

Differentiating (2) with respect to θ and evaluating at $\theta = 0$, it follows that

$$L_t^{\mathcal{N}} = \sum_{n\in\mathcal{N}} \frac{E_t(U_{1t+n})E_t(D_{t+n} + W_{t+n}N_{t+n}) - E_t(U_{1t+n}[D_{t+n} + W_{t+n}N_{t+n}])}{\sum_{n\in\mathcal{N}} n\ E_t(U_{1t+n}C_{t+n})}$$
$$- \sum_{n\in\mathcal{N}} \frac{E_t(-U_{2t+n}/W_{t+n})E_t(W_{t+n}N_{t+n}) - E_t(-U_{2t+n}N_{t+n})}{\sum_{n\in\mathcal{N}} n\ E_t(U_{1t+n}C_{t+n})}$$

where $U_{1t} \equiv \partial U_t/\partial C_t$ and $U_{2t} \equiv \partial U_t/\partial N_t$. Then, under assumption (3), the terms containing labor cancel each other, and hence

$$L_t^{\mathcal{N}} = \alpha_t^{\mathcal{N}} \sum_{n\in\mathcal{N}} \frac{E_t(M_{t,t+n})E_t(D_{t+n}) - E_t(M_{t,t+n}D_{t+n})}{\sum_{n\in\mathcal{N}} n\ E_t(M_{t,t+n}D_{t+n})}$$

with $\alpha_t^{\mathcal{N}} \equiv \sum_{n\in\mathcal{N}} n\ E_t(M_{t,t+n}D_{t+n})/\sum_{n\in\mathcal{N}} n\ E_t(M_{t,t+n}C_{t+n})$.

A.2. Proof of proposition 2

I can rewrite equation (5) as

$$L_t^{\mathcal{N}} = \sum_{n\in\mathcal{N}} \frac{n\ E_t(M_{t,t+n}D_{t+n})}{\sum_{n\in\mathcal{N}} n\ E_t(M_{t,t+n}C_{t+n})} \times \frac{1}{n}\Big(\frac{E_t(M_{t,t+n})E_t(D_{t+n})}{E_t(M_{t,t+n}D_{t+n})} - 1\Big)$$

The absence of arbitrage opportunities ensures positive weights $\{\omega_{nt}^{\mathcal{N}}\}$.

The proposition then follows directly from the expression of the term structure components, $l_t^{(n)}$, and the definition of the hold-to-maturity return on an arbitrary payoff, X, maturing in n periods, $R_{x,t\to t+n}^{(n)} = X_{t+n}/D_{xt}^{(n)}$, with no-arbitrage price $D_{xt}^{(n)} = E_t(M_{t,t+n}X_{t+n})$ at period t. Market equity is characterized by $X = D$ and real bonds by $X = 1$.

A.3. Microeconomic implications

The two habits satisfy the same local condition for a sensible habit as in Campbell and Cochrane (1999). Furthermore, I will argue that these habits are robust to Ljungqvist and Uhlig's (2015) critique of the habit formation mechanism in Campbell and Cochrane (1999).

A.3.1. Extending nonlinear habits to a production economy

The law of motion of surplus consumption assumed by Campbell and Cochrane (1999) in their endowment economy with random-walk consumption with homoskedastic innovations can be cast in two *equivalent* specifications:

$$s_{ct+1} = (1-\rho_s)s_c + \rho_s s_{ct} + \lambda_{ct}(\Delta c_{t+1} - \mu) \tag{A.1}$$

$$= (1-\rho_s)s_c + \rho_s s_{ct} + \lambda_{ct}\varepsilon_{ct+1} \tag{A.2}$$

when $\Delta c_{t+1} \sim Niid(\mu, \sigma_c^2)$. The equivalence of the two specifications, however, breaks down once we allow for a predictable component or heteroskedasticity in consumption growth, consistent with a generic production economy.

To understand what specification we should retain in a production economy, recall the rationale for the specification in Campbell and Cochrane (1999). First, they pick a specification for habits that implies a risk-free rate that adds to the expression under CRRA utility a linear term in surplus consumption, as in

$$r_{ft} = r_f + \gamma E_t(\Delta c_{t+1} - \mu) - \xi_1(s_{ct} - s_c) \tag{A.3}$$

with $r_f = -\ln(\beta) + \gamma\mu - .5\gamma(1-\rho_s - \xi_1/\gamma)$ and ξ_1 a free parameter that controls the volatility of the risk-free rate. Importantly, they pick a specification that breaks the equivalence between the inverse elasticity of intertemporal substitution γ and the risk aversion coefficient that drives the maximum Sharpe ratio, $\gamma(1+\lambda_{ct})$, a required property to generate discount-rate variation without a risk-free rate puzzle.

Second, they require a habit that is a slow-moving average of past consumption that is predetermined, at least at the limit point of the system in which all shocks are zero, and that moves nonnegatively with consumption everywhere with *local* movements in consumption.

Third, I add the requirement that more consumption is socially desirable not just locally around the steady state, but for any discrete movement in consumption, a property that Ljungqvist and Uhlig (2015) pointed out is not necessarily the case in nonlinear habit specifications. In a production economy, these requirements are met by specification (A.2) but not by (A.1).

A.3.2. Implications for the risk-free rate

When consumption innovations ε_{ct} are conditionally normal and heteroskedastic and expected consumption growth $E_t\Delta c_{t+1}$ is time-varying, the implications for the risk-free rate of specification (A.1) is

$$r_{ft} = r_f + \gamma(1+\lambda_{ct})E_t(\Delta c_{t+1} - \mu) - \xi_1(s_{ct} - s_c)$$

while specification (A.2) implies (A.3), as desired. Specification (A.1) fails to separate the inverse elasticity of intertemporal substitution $\partial r_{ft}/\partial E_t\Delta c_{t+1}$ from the price of risk $\gamma(1+\lambda_{ct})$. Based on this result I retain therefore specification (A.2).

A.3.3. Local predeterminedness

The market consumption habit specified indirectly by the surplus consumption process (A.2) is a complex nonlinear function of current and past consumption

$$x_{ct+1} = c_{t+1} + \ln\left(1 - (1 - e^{x_{ct}-c_t})^{\rho_s} e^{(1-\rho_s)s_c + \lambda_{ct}(c_{t+1} - E_t c_{t+1})}\right)$$

However, it is approximately a predetermined, slow-moving average of past consumption, as required for a sensible notion of habit. In fact, a first-order approximation of the expression around $x_{ct} - c_t = \ln(1 - S_c)$ yields

$$\begin{aligned} x_{ct+1} &\approx \ln(1 - S_c) + c_{t+1} + \rho_s(x_{ct} - c_t) - \varepsilon_{ct+1} \\ &= \ln(1 - S_c) + c_{t+1} - \sum_{j=0}^{\infty} \rho_s^j \varepsilon_{ct-j+1} \end{aligned}$$

The habit is predetermined because $x_{ct+1} = E_t x_{ct+1}$, and past consumption shocks receive their full weight only asymptotically. Unanticipated movements in consumption move consumption away from habits; surplus consumption is thus essentially detrended consumption.

Symmetrically, the home consumption habit can be written locally as

$$x_{ht+1} = \ln(1 - S_h) + h_{t+1} - \sum_{j=0}^{\infty} \rho_s^j \varepsilon_{ht-j+1}$$

The home consumption habit is a slow-moving average of past home consumption.

A.3.4. More consumption is welfare increasing

Contrary to specification (A.1), habit specification (A.2) does not produce the welfare gains to discrete-sized endowment destruction documented by Ljungqvist and Uhlig (2015). To prove this statement consider a pure-exchange economy with per-period log endowment y_t with growth rate $\Delta y_{t+1} = \mu + \varepsilon_{yt+1}$, where $\varepsilon_{yt} \sim iid(0, \sigma_y^2)$, and let equilibrium consumption be $c_t = y_t + \psi_t$, with $\psi_t \equiv c_t - y_t \leq 0$ the fraction of the endowment actually consumed; any leftover is destroyed without cost. In this context, there are welfare gains to endowment destruction if welfare

$$U_0(\{\psi_t\}) = E_0 \sum_{t=0}^{\infty} \beta^t \frac{e^{(1-\gamma)(c_t + s_t)}}{1-\gamma} = E_0 \sum_{t=0}^{\infty} \beta^t \frac{e^{(1-\gamma)(y_t + \psi_t + s_t)}}{1-\gamma}$$

is larger when $\psi_t < 0$ for some t than when $\psi_t = 0$ at all dates t. (I disregard here the piece of utility coming from home consumption, as utility is separable. Similarly, I drop the subscript 'c' to denote surplus market consumption.)

Specifically, I will consider a one-time deviation $\psi_0 = h < 0$ of c_0 from y_0 followed by $\psi_t = 0$ for all $t \geq 1$, and whether that deviation increases welfare. The starting point is the steady state $s_t = s$ and I consider a deterministic sequence of endowment innovations $\varepsilon_{yt} = 0, t \geq 0$. Under habit specification (A.1), the case studied by Ljungqvist and Uhlig, we have

$$s_{t+1} = (1-\rho_s)s + \rho_s s_t + \lambda(s_t)(\psi_{t+1} - \psi_t + \varepsilon_{yt+1})$$

and therefore, since $\lambda(s) \geq 0$ and $\lambda'(s) < 0$,

$$s_0(h) = s + \lambda(s)h < s, \qquad s_t(h) = s + \rho_s^{t-1} h[\rho_s \lambda(s) - \lambda(s + \lambda(s)h)] > s, \ t \geq 1$$

Surplus consumption is lower at time 0 when some of the endowment is destroyed but is higher in the subsequent periods. There can therefore be instances in which welfare is actually increased by an endowment destruction. Namely, welfare is

$$U_0(h) = \frac{e^{(1-\gamma)(y_0+h+s_0(h))}}{1-\gamma} + \sum_{t=1}^{\infty} \beta^t \frac{e^{(1-\gamma)(y_0+\mu t+s_t(h))}}{1-\gamma}$$

which, as emphasized by Ljungqvist and Uhlig, can be greater than $U_0(0)$ for some parameterization and choice of h, as the first term is maximized by $h = 0$ but the second term increases as h decreases below 0.

The same is not true under habit specification (A.2), in which case

$$s_{t+1} = (1-\rho_s)s + \rho_s s_t + \lambda(s_t)(\psi_{t+1} - E_t\psi_{t+1} + \varepsilon_{yt+1})$$

and hence an unanticipated endowment destruction occurring at time 0 ($\psi_0 - E_{-1}\psi_0 = h < 0$) implies

$$s_0(h) = s + \lambda(s)h < s, \qquad s_t(h) = s + \rho_s^t h\lambda(s) < s, \ t \geq 1$$

Surplus consumption is always lower after an endowment destruction. It follows that endowment destruction is welfare decreasing because

$$U_0(h) < U_0(0), \quad \text{for all } h < 0$$

The same result is true if the endowment destruction occurring at time 0 is anticipated ($\psi_{t+1} = E_t\psi_{t+1} = h < 0$), in which case $s_t(\psi) = s$ for all $t \geq 0$, and hence

$$U_0(h) = \frac{e^{(1-\gamma)(y_0+h+s)}}{1-\gamma} + \sum_{t=1}^{\infty} \beta^t \frac{e^{(1-\gamma)(y_0+\mu t+s)}}{1-\gamma} < U_0(0), \quad \text{for all } h < 0$$

Thus, specification (A.2) does not give rise to the counterintuitive properties emphasized by Ljungqvist and Uhlig (2015).

A.4. Derivation of the welfare criterion

In the model with external habits and nominal rigidities, aggregate welfare is:

$$U_0 = E_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{\mathscr{C}_t^{1-\gamma}}{1-\gamma} + \chi \frac{\mathscr{H}_t^{1-\gamma}}{1-\gamma} \right)$$

where $\mathscr{C}_t \equiv C_t - X_{ct}$ and $\mathscr{H}_t \equiv H_t - X_{ht}$. Consider the stationary transformation of utility,

$$\frac{U_0}{A_0^{1-\gamma}} = E_0 \sum_{t=0}^{\infty} \beta^t \exp[(1-\gamma)(a_t - a_0)] V_t$$

with per-period detrended utility

$$V_t \equiv \frac{1}{1-\gamma}\left(\frac{\mathscr{C}_t}{A_t}\right)^{1-\gamma} + \frac{\chi}{1-\gamma}\left(\frac{\mathscr{H}_t}{A_t}\right)^{1-\gamma}$$

A second-order expansion around the risky steady state implies:

$$V_t = \widetilde{\mathscr{C}}^{1-\gamma}\left(\ln(\widetilde{\mathscr{C}_t}) + \frac{1-\gamma}{2}\ln(\widetilde{\mathscr{C}_t})^2\right) + \chi\widetilde{\mathscr{H}}^{1-\gamma}\left(\ln(\widetilde{\mathscr{H}_t}) + \frac{1-\gamma}{2}\ln(\widetilde{\mathscr{H}_t})^2\right) + t.i.p.$$

with the detrended surplus consumption levels $\widetilde{\mathscr{C}_t} \equiv \mathscr{C}_t/A_t$ and $\widetilde{\mathscr{H}_t} \equiv \mathscr{H}_t/A_t$, up to a term independent of policy. I can rewrite the ratio of partial derivatives $\chi\widetilde{\mathscr{H}}^{1-\gamma}/\widetilde{\mathscr{C}}^{1-\gamma} = 1-\alpha$, given the efficient employment subsidy and fiscal and monetary policies such that risky and deterministic steady states coincide.

Using the second-order expansion at $N = .5$,

$$\ln(\widetilde{\mathscr{H}_t}) + \frac{1}{2}\ln(\widetilde{\mathscr{H}_t})^2 = \hat{s}_{ht} - \hat{n}_t + \frac{1}{2}\hat{s}_{ht}^2 - \frac{1}{2}\hat{n}_t^2 - \hat{n}_t\hat{s}_{ht}$$

approximate, detrended, per-period average welfare can be rewritten as:

$$\begin{aligned}
E\left(\frac{V_t}{\widetilde{\mathscr{C}}^{1-\gamma}}\right) &= E\left(\ln(\widetilde{\mathscr{C}_t}) + \frac{1-\gamma}{2}\ln(\widetilde{\mathscr{C}_t})^2 + \chi\frac{\widetilde{\mathscr{H}}^{1-\gamma}}{\widetilde{\mathscr{C}}^{1-\gamma}}\left(\ln(\widetilde{\mathscr{H}_t}) + \frac{1-\gamma}{2}\ln(\widetilde{\mathscr{H}_t})^2\right)\right) + t.i.p. \\
&= E\left(\widetilde{c}_t + \hat{s}_{ct} + \frac{1-\gamma}{2}(\widetilde{c}_t^2 + \hat{s}_{ct}^2 + 2\widetilde{c}_t\hat{s}_{ct})\right) \\
&\quad + (1-\alpha)E\left(\hat{s}_{ht} - \hat{n}_t - \frac{1+\gamma}{2}\hat{n}_t^2 + \frac{1-\gamma}{2}\hat{s}_{ht}^2 - (1-\gamma)\hat{n}_t\hat{s}_{ht}\right) + t.i.p. \\
&= -(1-\alpha)E(\Delta_t) - \frac{1}{2}\frac{\gamma(2-\alpha)+\alpha}{1-\alpha}var(\widetilde{c}_t) - \frac{\gamma-1}{2}var(s_{ct}) \\
&\quad + (\gamma-1)cov(\widetilde{c}_t, s_{ht} - s_{ct}) - \frac{(\gamma-1)(1-\alpha)}{2}var(s_{ht}) + t.i.p.
\end{aligned}$$

where $\widetilde{c}_t \equiv c_t - a_t$ and $\Delta_t \equiv \ln\int_0^1 [P_t(i)/P_t]^{-\varepsilon/(1-\alpha)}di$, with the aggregate production relation $(1-\alpha)\hat{n}_t = \widetilde{c}_t + (1-\alpha)\Delta_t$.

By a standard argument, let $S(t) \subset [0,1]$ represent the set of firms not reoptimizing their posted price in period t, recall the definition of the aggregate price level $P_t \equiv (\int_0^1 P_t(i)^{1-\varepsilon}di)^{1/(1-\varepsilon)}$, and recognize that all resetting firms choose an identical price P_t^*. It follows that

$$1 = \int_{S(t)} \left(\frac{P_t(i)}{P_t}\right)^{1-\varepsilon} di + (1-\eta)\left(\frac{P_t^*}{P_t}\right)^{1-\varepsilon} = \eta\hat{\Pi}_t^{\varepsilon-1} + (1-\eta)\left(\frac{P_t^*}{P_t}\right)^{1-\varepsilon}$$
$$e^{\Delta_t} = \int_{S(t)} \left(\frac{P_t(i)}{P_t}\right)^{-\frac{\varepsilon}{1-\alpha}} di + (1-\eta)\left(\frac{P_t^*}{P_t}\right)^{-\frac{\varepsilon}{1-\alpha}} = \eta\hat{\Pi}_t^{\frac{\varepsilon}{1-\alpha}} e^{\Delta_{t-1}} + (1-\eta)\left(\frac{P_t^*}{P_t}\right)^{-\frac{\varepsilon}{1-\alpha}}$$

and hence a second-order expansion around a steady state with $\pi = 0$ implies

$$\Delta_t = \eta\Delta_{t-1} + \frac{1}{2}\frac{\eta\varepsilon(1-\alpha+\alpha\varepsilon)}{(1-\eta)(1-\alpha)^2}\hat{\pi}_t^2$$

Therefore, I can rewrite average, per-period, detrended welfare as:

$$\frac{E(V_t)}{\widetilde{\mathscr{C}}^{1-\gamma}} = -\frac{1}{2}\frac{\eta\varepsilon(1-\alpha+\alpha\varepsilon)}{(1-\eta)^2(1-\alpha)} var(\hat{\pi}_t) - \frac{1}{2}\frac{\gamma(2-\alpha)+\alpha}{1-\alpha} var(\widetilde{c}_t)$$
$$- \frac{\gamma-1}{2} var(s_{ct}) - \frac{(1-\alpha)(\gamma-1)}{2} var(s_{ht}) + (\gamma-1)cov(\widetilde{c}_t, s_{ht} - s_{ct}) + t.i.p.$$

up to a term of at least third order and a term independent of policy.

A.5. Internal habits

The Pareto optimum (flexible prices, internal habits) implies the intertemporal and intratemporal rates of substitution

$$M_{t+1}^{\text{int.}} = \beta\left(\frac{C_{t+1}}{C_t}\right)^{-1} \frac{C_{t+1}^{1-\gamma}S_{ct+1}^{1-\gamma} + \lambda_{ct}(E_{t+1}-E_t)\mathcal{M}_{ct+1}}{C_t^{1-\gamma}S_{ct}^{1-\gamma} + \lambda_{ct-1}(E_t - E_{t-1})\mathcal{M}_t^c}$$
$$-\frac{\partial U_t^{\text{int.}}/\partial N_t}{\partial U_t^{\text{int.}}/\partial C_t} = \frac{C_t}{1-N_t}\frac{\chi A_t^{1-\gamma}(1-N_t)^{1-\gamma}S_{ht}^{1-\gamma} + \lambda_{ht-1}(E_t - E_{t-1})\mathcal{M}_t^h}{C_t^{1-\gamma}S_{ct}^{1-\gamma} + \lambda_{ct-1}(E_t - E_{t-1})\mathcal{M}_t^c}$$

with the shadow values of market and home surplus consumption

$$\mathcal{M}_{ct} = C_t^{1-\gamma}S_{ct}^{1-\gamma} + \beta\rho_s E_t\mathcal{M}_{ct+1}$$
$$\mathcal{M}_{ht} = \chi H_t^{1-\gamma}S_{ht}^{1-\gamma} + \beta\rho_s E_t\mathcal{M}_{ht+1}$$

On the one hand, a positive market (home) consumption shock means a lower marginal value of market (home) consumption; on the other hand, a positive market (home) consumption shock increases the habit level and thereby increases the marginal value of market (home) consumption.

It is straightforward to verify how a unitary elasticity of intertemporal substitution, $\gamma = 1$, produces constant shadow values of surplus market and home consumption. Under this parameterization we have

$$M_{t+1}^{\text{int.}} = \beta \left(\frac{C_{t+1}}{C_t} \right)^{-1}, \qquad -\frac{\partial U_t^{\text{int.}} / \partial N_t}{\partial U_t^{\text{int.}} / \partial C_t} = \frac{\chi C_t}{1 - N_t}$$

so all intertemporal and intratemporal effects of the habit are absent. The condition $\gamma = 1$ is actually necessary to grant a macro-finance separation (even approximately) when habits are internal, for any value of the spillover parameter ξ_2. The macro-financially separate Pareto optimum displays the same low and stable risk premia as a log-utility CRRA specification.

B. Empirical results

B.1. Data selection and synthetic replication

I drop weekly, quarterly, pm-settled, and mini options, whose nonstandard actual expiration dates are not tagged. Index mini options with three-year maturities have been traded since the 1990s but standard classes appear only in the 2000s; for this reason I follow Binsbergen et al. (2012) and focus on options of up to two-year maturity. I eliminate all observations with missing values or zero prices and keep only paired call and put options. I use mid quotes between the bid and the ask prices on the last quote of the day and closing values for the S&P 500 index.

According to equation (10), if there are no arbitrage opportunities and the law of one price holds then the map $X_i \mapsto \mathcal{A}_{it}^{(n)}$ is strictly monotonic and linear. In practice, the relation does not always hold without error across all strike prices available; as long as more than two strikes are available for a given maturity, one can use the no-arbitrage relation to extract $P_{dt}^{(n)} \equiv \sum_{j=1}^{n} D_t^{(j)}$ and $P_{\$bt}^{(n)}$ as the least absolute deviation (LAD) estimators that minimize the expression

$$\sum_{i=1}^{I} \left| \mathcal{A}_{it}^{(n)} - P_{dt}^{(n)} - X_i P_{\$bt}^{(n)} \right|$$

for a given trade date t and maturity n. The cross-sectional error term accounts for potential measurement error (e.g., because of bid-ask bounce, asynchronicities, or other microstructural noise).

Over most of the sample the strikes and the auxiliary variables are in a nearly perfect linear relation except for a few points that violate the law of one price. The LAD estimator is particularly appropriate to attach little weight to those observations as long as their number is small relative to the sample size of the cross-sectional regression. Accordingly, I drop all trade dates and maturities associated with a linear relation between X_i and $\mathcal{A}_{it}^{(n)}$ that fails to fit at least a tenth of the cross-sectional size (with a minimum of five points) up to an error that is less than 1 percent of the extracted dividend claim price.[19]

[19] In many instances, non-monotonicities in the auxiliary variable are concentrated in deep in- and out-

The procedure results in a finite number of matches, which I combine to calculate the prices of options-implied dividend claims and nominal bonds by using the put-call parity relation. As shown in Figure B.5b, the number of cross-sectional observations available to extract the options-implied prices of bonds and dividend claims increases over time (from medians of around 25 observations per trading day up to more than 100) as the market grows in size and declines with the options' maturity. Of the resulting extracted prices I finally discard all trading days associated with prices $P_{dt}^{(n)}$ that are nonincreasing in maturity, as they would represent arbitrage opportunities.

Overall, my selection method based on law-of-one-price violations excludes almost a fifth of the available put-call pairs. Finally, to obtain monthly implied dividend yields with constant maturities, I follow Binsbergen et al. (2012) and Golez (2014) and interpolate between the available maturities. As advocated by Golez to reduce the distance between intra-day and end-of-day options-implied prices and hence the potential effect of asynchronicities and other microstructural frictions, I then construct monthly prices using 10 days of data at the end of each month.

I find large correlations with the intra-day options-implied prices extracted by Binsbergen et al. (2012) over the 1996-2009 period; the correlation of the 6-, 12-, 18- and 24-month equity prices with Binsbergen-Brandt-Koijen data are .91, .95, .95 and .94, respectively, with a mean-zero difference in levels. End-of-month data using a one-day window have slightly higher volatilities and correlations between .80 and .95. The median or the mean over a three-day window centered on the end-of-month trading day increases correlations to .87-.95; the marginal increase in correlations for window widths of more than 10 days is nearly imperceptible. The nearly white-noise deviation between the estimates by Binsbergen et al. and mine over the comparable sample is likely a mixture of asynchronicities and different proxies for the interest rate.[20]

B.2. Errors in synthetic replication

Figure B.6 plots the auxiliary map $X \mapsto P_{dt}^{(n)} + X\ P_{\$bt}^{(n)}$ for selected trading days and maturities. As shown by the lower part of Figure B.6, the typical map toward the end of the sample is virtually perfectly linear and monotonic as one moves along the strike prices, so cross-sectional errors are immaterial. In the middle of the sample, the relationship still holds with almost no error despite a lower number of strike prices available relative to the last years of the sample. However, note how the cross-sectional errors are clearly visible during the first years of the sample, in which the strikes available are relatively few. (The figure also reports the index price to better gauge the moneyness of the put and call options associated with each cross-sectional data point.)

Figure B.5a box-plots the size of the law-of-one-price violations present in the sample, which, for the most part, are concentrated around errors of less than 1 percent; larger

of-the-money options. Whenever I spot non-monotonicities for low and high moneyness levels, I restrict the sample to strikes with moneyness levels between 0.7 and 1.1 before running the cross-sectional LAD regression.

[20] I find options-implied interest rates with nearly perfect correlations with the corresponding LIBOR and Treasury rates but with different levels that lie about halfway between the two proxies.

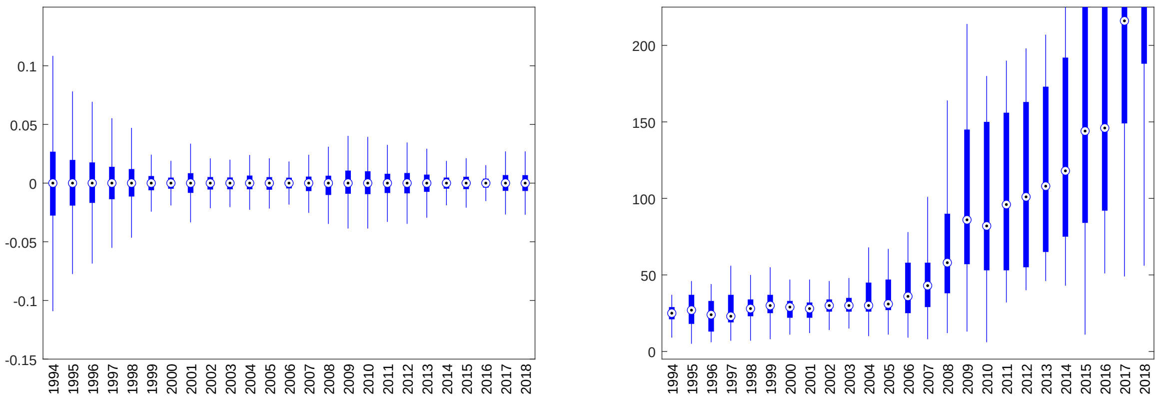

(a) *Law-of-one-price violations.* (b) *Number of strikes per trade date.*

Figure B.5: Implied dividend claim prices, index, dividends, and sample statistics. Law-of-one-price violations are expressed in percent of the associated options-implied dividend claim price. Only observations not filtered out by the data selection criteria are included.

violations are associated with the first years of the sample—probably because of a relatively low liquidity—and with years of greater volatility such as 2001 or 2009-10. Data previous to 1994 are more problematic by this metric (see also Golez, 2014) and I therefore exclude them altogether from the sample.

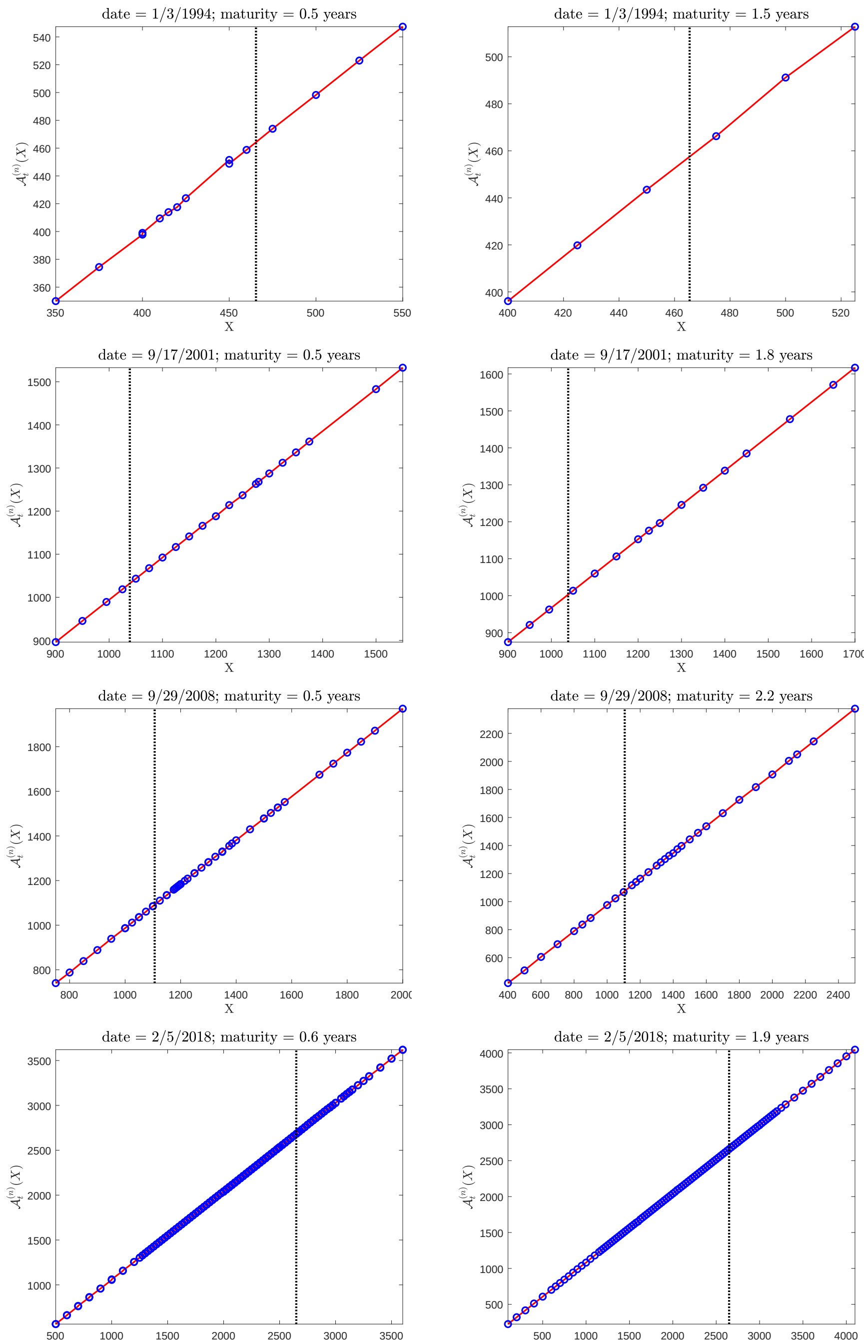

Figure B.6: Auxiliary map $X \mapsto P_{dt}^{(n)} + X\ P_{\$bt}^{(n)}$ on given trading days and for given maturities. The dotted lines indicate the value of the index at each respective trading day.

References

H. Ai, M. M. Croce, A. M. Diercks, and K. Li. News shocks and the production-based term structure of equity returns. *Review of Financial Studies*, 31(7):2423–67, 2018. doi: 10.1093/rfs/hhy015.

F. Alvarez and U. J. Jermann. Using asset prices to measure the cost of business cycles. *Journal of Political Economy*, 112(6):1223–56, 2004. doi: 10.1086/424738.

B. S. Aruoba, P. Cuba-Borda, and F. Schorfheide. Macroeconomic dynamics near the ZLB: A tale of two countries. *The Review of Economic Studies*, 85(1):87–118, 2018. doi: 10.1093/restud/rdx027.

A. Atkeson and C. Phelan. Reconsidering the costs of business cycles with incomplete markets. *NBER Macroeconomics Annual*, 9:187–218, 1994. doi: 10.1086/654247.

R. Bansal and A. Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509, 2004. doi: 10.1111/j.1540-6261.2004.00670.x.

R. Bansal, S. Miller, D. Song, and A. Yaron. The term structure of equity risk premia. *Journal of Financial Economics*, 2021.

F. Barillas, L. P. Hansen, and T. J. Sargent. Doubts or variability? *Journal of Economic Theory*, 144: 2388–418, 2009. doi: 10.1016/j.jet.2008.11.014.

R. Barro. Rare disasters, asset prices, and welfare costs. *American Economic Review*, 99(1):243–64, 2009. doi: 10.1257/aer.99.1.243.

F. Belo, P. Collin-Dufresne, and R. S. Goldstein. Dividend dynamics and the term structure of dividend strips. *Journal of Finance*, 70(3):1115–60, 2015. doi: 10.1111/jofi.12242.

J. H. v. Binsbergen and R. S. J. Koijen. The term structure of returns: Facts and theory. *Journal of Financial Economics*, 124:1–21, 2017. doi: 10.1016/j.jfineco.2017.01.009.

J. H. v. Binsbergen, M. Brandt, and R. S. J. Koijen. On the timing and pricing of dividends. *American Economic Review*, 102(4):1596–1618, 2012. doi: 10.1257/aer.102.4.1596.

J. H. v. Binsbergen, W. H. Hueskes, R. S. J. Koijen, and E. B. Vrugt. Equity yields. *Journal of Financial Economics*, 110:503–19, 2013. doi: 10.1016/j.jfineco.2013.08.017.

O. Boguth, M. Carlson, A. Fisher, and M. Simutin. Leverage and the limits of arbitrage pricing: Implications for dividend strips and the term structure of equity risk premia. Manuscript, 2012.

J. Borovicka and L. P. Hansen. Examining macroeconomic models through the lens of asset pricing. *Journal of Econometrics*, 183:67–90, 2014. doi: 10.1016/j.jeconom.2014.06.010.

J. Y. Campbell and J. H. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–51, 1999. doi: 10.1086/250059.

J. Y. Campbell, C. Pflueger, and L. M. Viceira. Macroeconomic drivers of bond and equity risks. *Journal of Political Economy*, 128(8):3148–3185, 2020. doi: 10.1086/707766.

V. V. Chari, P. J. Kehoe, and E. R. McGrattan. Business cycle accounting. *Econometrica*, 75(3):781–836, 2007. doi: 10.1111/j.1468-0262.2007.00768.x.

R. Clarida, J. Galí, and M. Gertler. The science of monetary policy: A New Keynesian perspective. *Journal of Economic Literature*, 37(4):1661–1707, 1999. doi: 10.1257/jel.37.4.1661.

J. H. Cochrane. Discount rates. *Journal of Finance*, 66(4):1047–1108, 2011. doi: 10.1111/j.1540-6261.2011.01671.x.

I. Correia, E. Farhi, J. P. Nicolini, and P. Teles. Unconventional fiscal policy at the zero bound. *American Economic Review*, 103(4):1172–1121, 2013. doi: 10.1257/aer.103.4.1172.

M. M. Croce, M. Lettau, and S. C. Ludvigson. Investor information, long-run risk, and the term structure of equity. *Review of Financial Studies*, 28(3):706–42, 2015. doi: 10.1093/rfs/hhu084.

S. D'Amico, D. H. Kim, and M. Wei. Tips from TIPS: The informational content of Treasury Inflation-Protected Security prices. *Journal of Financial and Quantitative Analysis*, 53(1):395–436, 2018. doi: 10.1017/S0022109017000916.

M. De Santis. Individual consumption risk and the welfare cost of business cycles. *American Economic Review*, 97(4):1488–1506, 2007. doi: 10.1257/aer.97.4.1488.

M. Ellison and T. J. Sargent. Welfare cost of business cycles with idiosyncratic consumption risk and a preference for robustness. *American Economic Journal: Macroeconomics*, 7(2):40–57, 2012. doi: 10.1257/mac.20130098.

X. Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics*, 127(2):645–700, 2012. doi: 10.1093/qje/qjs001.

J. Galí. *Monetary Policy, Inflation, and the Business Cycle. An Introduction to the New Keynesian Framework.* Princeton University Press, Princeton, 2008.

B. Golez. Expected returns and dividend growth rates implied in derivative markets. *Review of Financial Studies*, 27(3):790–822, 2014. doi: 10.1093/rfs/hht131.

N. J. Gormsen. Time variation of the equity term structure. *The Journal of Finance*, 76(4):1959–1999, 2021. doi: 10.1111/jofi.13020.

J. Greenwood and Z. Hercowitz. The allocation of capital and time over the business cycle. *Journal of Political Economy*, 99(6):1188–1214, 1991. doi: 10.1086/261797.

R. S. Gürkaynak, B. Sack, and J. H. Wright. The TIPS yield curve and inflation compensation. *American Economic Journal: Macroeconomics*, 2(1):70–92, 2010. doi: 10.1257/mac.2.1.70.

M. Hasler and R. Marfè. Disaster recovery and the term structure of dividend strips. *Journal of Financial Economics*, 122:116–134, 2016. doi: 10.1016/j.jfineco.2015.11.002.

A. Imrohoroğlu. Cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political Economy*, 97(6):1364–1383, 1989. doi: 10.1086/261658.

T. Krebs. Job displacement risk and the cost of business cycles. *American Economic Review*, 97(3):664–686, 2007. doi: 10.1257/aer.97.3.664.

P. Krusell and A. A. Smith. On the welfare effects of eliminating business cycles. *Review of Economic Dynamics*, 2:245–72, 1999. doi: 10.1006/redy.1998.0043.

C. Leith, I. Moldovan, and R. Rossi. Optimal monetary policy in a New Keynesian model with habits in consumption. *Review of Economic Dynamics*, 15(3):416–35, 2012. doi: 10.1016/j.red.2012.03.001.

M. Lettau and H. Uhlig. Can habit formation be reconciled with business cycle facts? *Review of Economic Dynamics*, 3:79–99, 2000. doi: 10.1006/redy.1998.0035.

M. Lettau and J. A. Wachter. Why is long-horizon equity less risky? A duration-based explanation of the value premium. *Journal of Finance*, 62(1):55–92, 2007. doi: 10.1111/j.1540-6261.2007.01201.x.

M. Lettau and J. A. Wachter. The term structures of equity and interest rates. *Journal of Financial Economics*, 101:90–113, 2011. doi: 10.1016/j.jfineco.2011.02.014.

L. Ljungqvist and H. Uhlig. Tax policy and aggregate demand management under catching up with the Joneses. *American Economic Review*, 90(3):356–66, 2000. doi: 10.1257/aer.90.3.356.

L. Ljungqvist and H. Uhlig. Comment on the Campbell-Cochrane habit model. *Journal of Political Economy*, 123(5):1201–13, 2015. doi: 10.1086/676119.

P. Lopez, D. Lopez-Salido, and F. Vazquez-Grande. Accounting for risk in a linearized solution: How to approximate the risky steady state and perturb around it. Federal Reserve Bank of Cleveland Working Paper 22-14, 2020. doi: 10.26509/frbc-wp-202214.

P. Lopez, D. Lopez-Salido, and F. Vazquez-Grande. Nominal rigidities and the term structures of equity and bond returns. Federal Reserve Bank of Cleveland Working Paper 23-11, 2021. doi: 10.26509/frbc-wp-202311.

R. E. Jr. Lucas. *Models of Business Cycles.* Oxford University Press, Oxford, 1987.

R. E. Jr. Lucas. Macroeconomic priorities. *American Economic Review*, 93(1):1–14, 2003. doi: 10.1257/000282803321455133.

R. Marfè. Income insurance and the equilibrium term structure of equity. *Journal of Finance*, 72(5): 2073–2130, 2017. doi: 10.1111/jofi.12508.

C. Otrok. On measuring the welfare cost of business cycles. *Journal of Monetary Economics*, 47:61–92, 2001. doi: 10.1016/S0304-3932(00)00052-0.

T. D. Jr. Tallarini. Risk-sensitive real business cycles. *Journal of Monetary Economics*, 45:507–32, 2000. doi: 10.1016/S0304-3932(00)00012-X.

M. Weber. Cash flow duration and the term structure of equity returns. *Journal of Financial Economics*, 128:486–503, 2018. doi: 10.1016/j.jfineco.2018.03.003.

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