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. 21-16

August 2021

Suggested citation: Lopez, Pierlauro. 2021. "Welfare Implications of Asset Pricing Facts: Should Central Banks Fill Gaps or Remove Volatility?" Federal Reserve Bank of Cleveland, Working Paper No. 21-16. https://doi.org/10.26509/frbc-wp-202116.

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

More than 20 years of financial market data suggest a term structure of the welfare cost of economic uncertainty 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 monetary model can rationalize the evidence. The model is observationally equivalent in its quantity implications to a standard New Keynesian model with CRRA utility, but the optimal policy prescription is overturned. In the model the central bank should prioritize removing consumption volatility (a targeting of risk premia) over filling the gap between consumption and its flexible-price counterpart (inflation targeting).

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, Urban Jermann, Patrick Kehoe, Ralph Koijen, Ed Knotek, Roberto Marfè, Monika Piazzesi, Eric Sims, Fabio Trojani, and Harald Uhlig for very helpful advice and 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

In seminal contributions, Lucas (1987, 2003) 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, even though the resulting measures have an intuitive relation with welfare, their implications for the design of stabilization policy have remained elusive.

This paper argues that welfare cost measures contain powerful information to discriminate among competing 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 that rationalizes those empirical measures, and carry out a welfare-theoretic analysis of monetary policy. Besides nominal price rigidities, the model relies on external habit formation, and consumers are sufficiently sensitive to cash flow fluctuations to value more a less volatile consumption path than less volatile inflation, thereby overturning a pervasive result in the New Keynesian literature.

1.1. Decomposing the cost of uncertainty

By focusing on fluctuations in the determinants of consumption, in particular wage and equity income, 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 dividend payments, so-called dividend strips.[3] The result relies on no-arbitrage relations and 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 costs that are sizable and countercyclical. The point estimates, reported in Figure 1a, suggest a *negatively* sloped term structure of welfare costs, driven *both*

[2]See Lucas (1987); Atkeson and Phelan (1994); Krusell and Smith (1999); Tallarini (2000); Otrok (2001); De Santis (2007); Galí, Gertler, and López-Salido (2007); Barillas, Hansen, and Sargent (2009); Barro (2009); Croce (2012); Ellison and Sargent (2012); Epstein, Farhi, and Strzalecki (2014), among many others.

[3]A recent literature focuses on risk pricing across maturities (see, for example, Lettau and Wachter 2007, 2011; Hansen et al. 2008; Binsbergen et al. 2012, 2013; Backus et al. 2014; Borovicka and Hansen 2014; Belo et al. 2015; Dew-Becker and Giglio 2016; Binsbergen and Koijen 2017; Marfè 2017; Ai et al. 2018; Weber 2018).

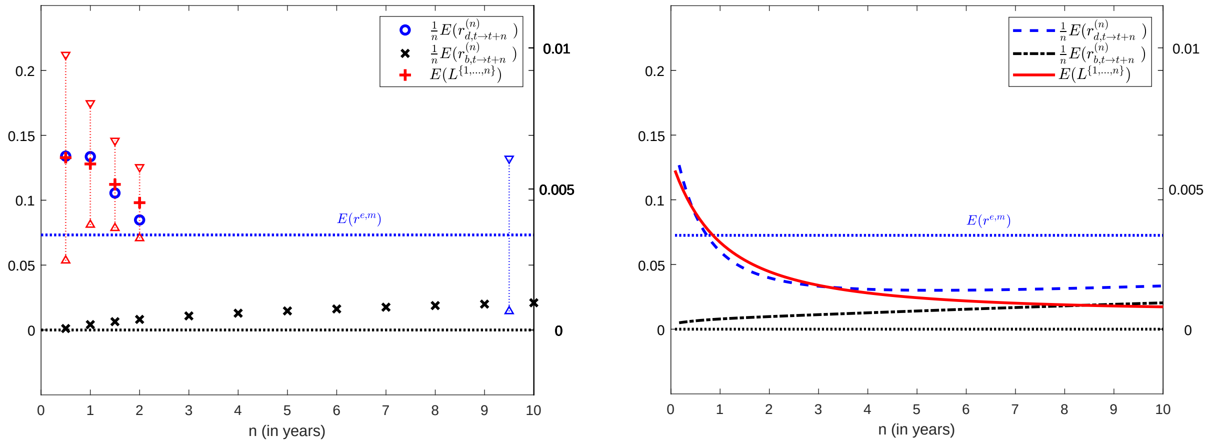

(a) *Point estimates.* (b) *Model-based predictions.*

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 percent 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. $L^{\{1,\ldots,n\}}$ 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.

by the negative slope of the term structure of equity 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 the premium is similar in size; the cost increased to 2-3 percentage points during the 2001 and 2007-2009 recessions while the benefits of long-run stability are smaller.

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 uncertainty. First, I focus on fluctuations in the *determinants* of consumption and consider the endogeneity of the main source of income—labor income. In their calculations Alvarez and Jermann model an exogenous difference between consumption and dividends. In contrast, once I make the difference between consumption and dividends endogenous from a consumer's point of view, the positive effect on utility when labor income is marginally stabilized is offset by the effect of the associated adjustment in hours. In this context, only claims to exogenous determinants of consumption such as equity income are necessary to measure the marginal cost of uncertainty.[4]

[4]The cost of uncertainty can be measured using equity claims as long as claims to the residual determinant

Second, the term structure of the cost of uncertainty 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 uncertainty 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 uncertainty. Like Binsbergen, Brandt, and Koijen (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. The sufficiency of the condition can only be assessed in the context of a structural model. In fact, the literature opened by Lucas (1987) focuses on fluctuations in the level of consumption and is silent as to whether the current fluctuations in the economy are optimal. By contrast, structural macroeconomic models for policy evaluation typically prescribe stabilizing gaps between current and some benchmark consumption levels (see Clarida, Galí, and Gertler 1999, among many others).

What welfare cost measures really say is that their empirical properties are key diagnostics of any macroeconomic model that seeks to assess the priority of different stabilization policies. In this respect, empirical measures of welfare costs are highly informative to discriminate across competing models because it is well-known that explaining *jointly* a downward-sloping term structure of equity and an upward-sloping term structure of interest rates from a structural perspective is hard (e.g., Binsbergen and Koijen 2017).[5]

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

---

of consumption net of equity and wage income command a relatively small risk premium. In any event, the measure can be interpreted as the cost of uncertainty in equity and labor income, which is a meaningful quantity in its own right.

[5]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 uncertainty, and confirms their difficulties in explaining 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, Lettau, and Ludvigson (2015), the recursive preferences of Tallarini (2000) and Barillas et al. (2009), and the rare disasters model of Gabaix (2012).

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 technology shock affecting market and home productions, both market and home consumption drop close to their habit levels, and hence offset the undesirable effect of habits on the labor choice. Figure 1b reports the average term structure of welfare costs in the model. In the model, equity income is highly procyclical but partly mean reverting by its cointegration with consumption; hence, people fear instability especially in the short run. As in the data, the model-implied costs of short-run uncertainty are sizable and substantially larger than the cost of uncertainty over longer horizons, with a cost of lifetime uncertainty of around one-tenth of a percentage point of perpetual growth in consumption.

Conveniently, the model builds on the textbook New Keynesian model (e.g., Galí 2008) used to study macroeconomic stabilization. The comparison can be made clear-cut. I can calibrate a version of the model so that quantities and inflation are approximately the same as in the textbook model; a macroeconomist 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 notions of stability in the literature. First, the central bank wants to close the consumption gap (or, equivalently, inflation). Second, it wants to remove consumption uncertainty (or risk premia variation). Achieving both goals is unfeasible in the presence of supply 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).

2. The Term Structure of the Welfare Cost of Fluctuations

People live in a stochastic world, have finite resources, and decide how to allocate them across time. Financial markets are without arbitrage opportunities and sufficiently complete to allow people to trade the full set of zero-coupon real bonds and equities.

Identical risk-averse consumers $j \in [0, 1]$ have time-t preferences $U_t = E_t U\left(C(j), N(j), X(j)\right)$, 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.[6] Without loss of generality, let factor $X(j)$ depend on individual consumption and labor only via aggregate consumption and labor, $C = \int_0^1 C(j)dj$

[6]The definition of cost of uncertainty has a meaningful interpretation even if one relaxes the assumption of a representative consumer, by an argument similar to the one made by Alvarez and Jermann (2004). The online appendix provides the detailed argument.

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 he would trade against stable consumption and labor streams without having to affect factor X.

The determinants of consumption, $C_t = D_t + W_t N_t + e_t$, include 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. Let $\overline{C}_{t+n}$ denote the consumption level that is hypothetically offered to the jth individual at time $t+n$, which I refer to as *stable* consumption. Stable consumption $\overline{C}_{t+n} = \overline{D}_{t+n} + W_{t+n}\overline{N}_{t+n} + e_{t+n}$ is 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 as

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

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

Definition (Welfare cost of uncertainty)**.** The map $\mathcal{L}_t : (\theta, \mathcal{N}) \mapsto \mathcal{L}_t^{\mathcal{N}}(\theta)$ defined by

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

measures the cost of fluctuations, where the index set $\mathcal{N} \subset \mathbb{N}$ indicates which coordinates are stabilized and allows for focusing on any window of interest.

For example, the *total cost* $\mathcal{L}_t^{\mathbb{N}}(1)$ measures how much extra growth the elimination of all uncertainty in equity and wage income is worth, and the *marginal cost* $L_t^{\mathcal{N}} \equiv \frac{\partial}{\partial\theta}\mathcal{L}_t^{\mathcal{N}}(0)$ represents how much extra growth a marginal stabilization is worth at the current level of uncertainty.[7]

I assume enough smoothness in preferences to guarantee that $\mathcal{L}_t^{\mathcal{N}}$ is a differentiable map in $\theta \in [0,1]$. Additionally, suppose that the consumption-labor tradeoff is described by the condition

$$W_t = -\frac{\partial U_t/\partial N_t}{\partial U_t/\partial C_t} \quad (2)$$

[7] The online appendix discusses the relationship between definition (1) and the definitions by Lucas (1987) and Alvarez and Jermann (2004).

i.e., the marginal rate of substitution between consumption and labor equals the relative price. This assumption is a standard optimality condition.[8]

Proposition 1. *The marginal cost of uncertainty within any window of interest* $\mathcal{N}$, $L_t^{\mathcal{N}}$, *is*

$$L_t^{\mathcal{N}} = \sum_{n\in\mathcal{N}} \frac{E_t(M_{t,t+n})E_t(D_{t+n}) - E_t(M_{t,t+n}D_{t+n})}{\sum_{n\in\mathcal{N}} n\ E_t(M_{t,t+n}C_{t+n})} \tag{3}$$

where $M_{t,t+n} = (\partial U_t/\partial C_{t+n})/(\partial U_t/\partial C_t)$ *is the* n*-period stochastic discount factor. Under no-arbitrage,* $D_{ct}^{(n)} = E_t M_{t,t+n} C_{t+n}$ *is the price of an* n*-period consumption strip,* $D_{dt}^{(n)} = E_t M_{t,t+n} D_{t+n}$ *is the price of an* n*-period dividend strip, and* $D_{bt}^{(n)} = E_t M_{t,t+n}$ *is the price of an* n*-period zero-coupon real bond.*

Equation (3) expresses the marginal cost of uncertainty around all coordinates $n \in \mathcal{N}$ as a function of the price of a claim to the trend in dividend income and of the price and duration of claims to future dividends and consumption at all coordinates $n \in \mathcal{N}$. Note that claims to labor income do not enter the expression because people set the marginal rate of substitution between consumption and labor equal to the wage rate, so the marginal effect of the adjustment in hours offsets exactly the benefits of a marginal stabilization in labor income.

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

$$l_t^{(n)} = \frac{1}{n}\left(E_t R_{t\to t+n}^{e,(n)} - 1\right)$$

where $R_{t\to t+n}^{e,(n)} = D_{t+n} D_{bt}^{(n)} / D_{dt}^{(n)}$.

The motivation for calling the map $l_t : n \mapsto l_t^{(n)}$ a term structure of the marginal cost of uncertainty is given by proposition 2. Given the prices of strips $\{D^{(n)}\}$ and the term structure components $\{l^{(n)}\}$ you can compute the marginal cost $L^{\mathcal{N}}$ for any coordinate set $\mathcal{N} \subset \mathbb{N}$.

[8] There is ample evidence that labor wedges matter (Chari, Kehoe, and McGrattan 2007). This observation does not, however, rule out assumption (2), 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 (2) is nevertheless inconsistent with frictions such as sticky wages. Importantly, however, it is consistent with the model proposed in Section 4.

Proposition 2. *The marginal cost of uncertainty within any window of interest $\mathcal{N}$, $L_t^{\mathcal{N}}$, is the linear combination of the term components $\{l_t^{(n)}\}_{n=1}^{\infty}$,*

$$L_t^{\mathcal{N}} = \alpha_t^{\mathcal{N}} \sum_{n\in\mathcal{N}} \omega_{nt}^{\mathcal{N}} l_t^{(n)} \approx \alpha \sum_{n\in\mathcal{N}} \omega_n^{\mathcal{N}} l_t^{(n)} \tag{4}$$

with scaling factor $\alpha_t^{\mathcal{N}} \equiv \sum_{n\in\mathcal{N}} n\ D_{dt}^{(n)} / \sum_{n\in\mathcal{N}} n\ D_{ct}^{(n)}$, hence $\alpha^{\mathcal{N}} = \alpha = D/C$ in a non-stochastic stationary state, and where the weights $\omega_{nt}^{\mathcal{N}} \equiv n\ D_{dt}^{(n)} / \sum_{n\in\mathcal{N}} n\ D_{dt}^{(n)}$ are positive and such that $\sum_{n\in\mathcal{N}} \omega_{nt}^{\mathcal{N}} = 1$.

Equation (4) shows how the first-order effect of distinguishing between consumption and dividends is that the cost of uncertainty 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 1994-2018 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).

There is a powerful intuition behind these formulas. At the margin, people would trade $L_t^{\{n\}}\theta = \alpha_t^{\{n\}} l_t^{(n)}\theta$ points of growth in the nth coordinate of consumption against the elimination of a fraction θ of the aggregate uncertainty around that coordinate, and that coordinate alone. Propositions 1 and 2 show how this tradeoff is precisely the one offered by the financial market. In fact, by holding to maturity a portfolio short in the n-period zero-coupon bond and long an equal amount in the n-period dividend strip, people can experience an average growth rate in n periods of $\frac{1}{n}(E_t R_{t\to t+n}^{e,(n)} - 1)$ by shouldering the uncertainty in n-period ahead dividend income.

3. Empirical Measures of the Cost of Fluctuations

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:

$$\mathcal{C}_{t,t+n} - \mathcal{P}_{t,t+n} = \mathscr{P}_t - \sum_{j=1}^{n} D_t^{(j)} - X\ P_{\$bt}^{(n)} \tag{5}$$

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.[9] 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 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.

[9] 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.

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 (5) then implies that

$$\mathcal{A}_{it}^{(n)} = \sum_{j=1}^{n} D_t^{(j)} + X_i P_{\$bt}^{(n)} \tag{6}$$

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, (6) 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, Kim, and Wei 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 1999-2018 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.[10]

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

[10] A simpler proxy for real yields—nominal hold-to-maturity returns on Treasury yields deflated by holding-period inflation—results in similar results. In any event, none of the evidence suggests that real yields are significantly larger in absolute magnitude than nominal yields, so the real bond proxy problem is unlikely to affect by much the quantitative estimates of short-term welfare costs because, at the observable end of the curve, the contribution of bond yields is much smaller than the contribution of equity yields.

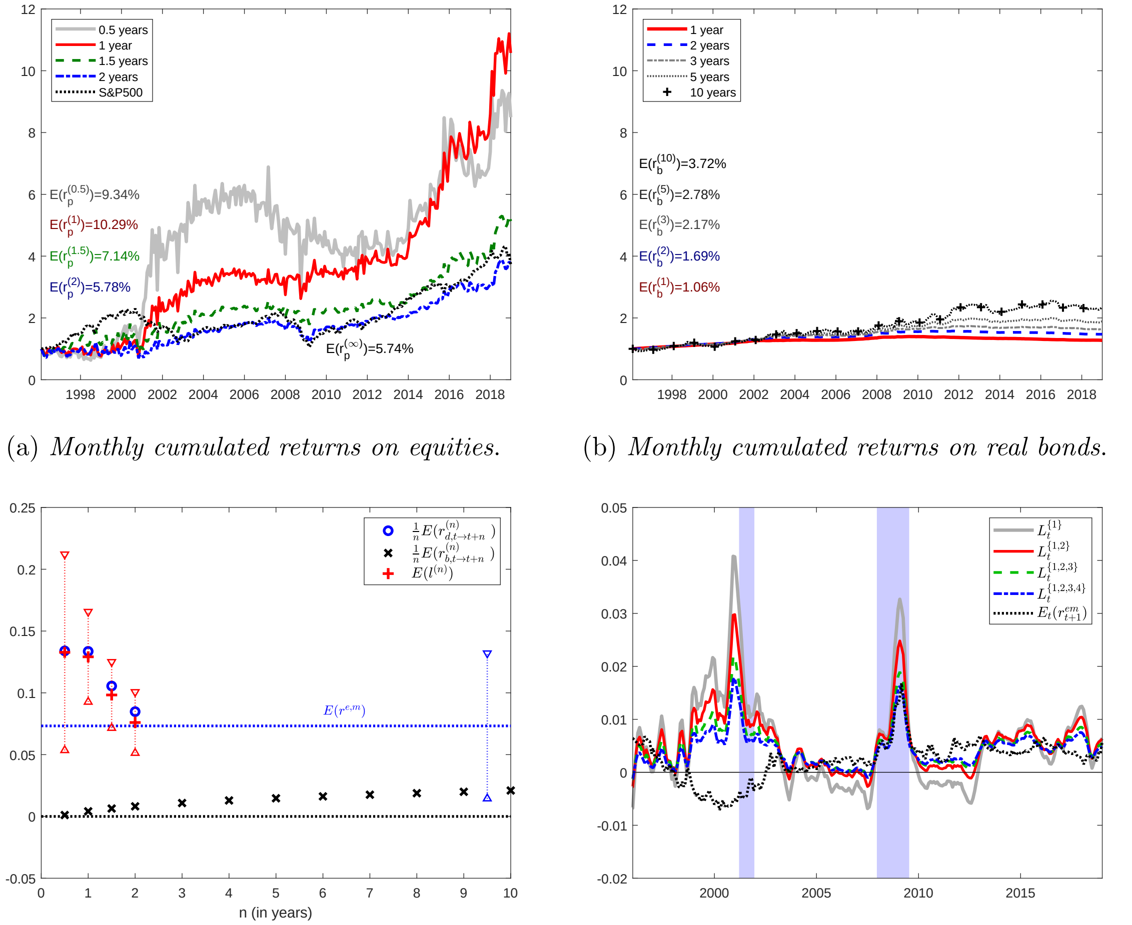

(a) *Monthly cumulated returns on equities.*

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

(c) *Average term structures of equity, real rates, and welfare costs (with block-bootstrap 95 percent 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 1994-2018 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 of nominal 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).

return on a six-month buy-and-hold strategy that pays off the next six months of dividends.[11] 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 of 13.3, 12.9, 9.8 and 7.6 percent. These 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 of 0.51 percent, whereas the cost of uncertainty over the next two years is of 0.36 percent.

A comparison of the estimated costs of uncertainty with holding-period excess returns on the index reveals the *economic* significance of these estimates. To an approximation around

[11] 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.

$\ln R = 0$, the equity premium relates to the term structures of equity and interest rates by

$$\begin{aligned} E_t(\ln R^{e,m}_{t+1}) &\approx \sum_{n=1}^{\infty} \omega_{nt} E_t(\ln R^{e,(n)}_{dt+1}) \\ l^{(n)}_t \approx \frac{1}{n} E_t(\ln R^{(n)}_{d,t\to t+n} - \ln R^{(n)}_{b,t\to t+n}) &= \frac{1}{n}[E_t(\ln R^{(n)}_{dt+1}) + ... + E_t(\ln R^{(1)}_{dt+n})] - y^{(n)}_{bt} \end{aligned} \tag{7}$$

where $y^{(n)}_{bt} = -\ln(E_t M_{t,t+n})/n$ is the nth real yield, and weights $\omega_{nt} \equiv D^{(n)}_{dt}/\mathscr{P}_t$ add up to one. Equation (7) implies that if the cost 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 downwards. 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.

Figure 1a and Figure 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 $l^{(n)}_t$ 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.

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.[12] 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 uncertainty varies over time. The cost 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 uncertainty over time. The cost of uncertainty rises dramatically to 2-4 percent during the dot-com crash

[12]For example, the no-arbitrage model of Lettau and Wachter (2011) 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.

and the period immediately preceding the early 2000s recession as well as during the most recent 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 combines a textbook New Keynesian model economy (Galí 2008) with Campbell and Cochrane (1999) habits in consumption. 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.

This section computes the properties of the model with accurate numerical procedures—the online appendix describes the solution method—but I derive sharper insight from a linearization of the solution of the model around the risky steady state.

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 linearized 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 linearization, the two models differ only in their asset pricing implications. The respective welfare implications turn out to be, however, dramatically different; under CRRA utility the optimal Taylor rule closes the consumption gap, while under habits the priority is to minimize consumption volatility.

(Throughout the exposition, lower-case letters denote natural logarithms and hatted variables deviations from the mean.)

4.1.1. Firms

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

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

	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 uncertainty, 1994-2018. $l^{(n)}$ is the annualized cost of a marginal increase in uncertainty in n-semesters-ahead cash flows. $L^{\{1,\ldots,n\}}$ 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. Boostrap 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, 1996m1-2018m12. Newey-West standard errors to correct for overlapping.

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, $\tau_f = 1 - MC$, with MC the real marginal cost, designed to offset any distortions caused by monopolistic competition in the steady state. Corporate profits $D_t(i) = Y_t(i)P_t(i)/P_t - (1-\tau_f)W_t N_t(i) - T_{ft}$ 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 $1-\eta$. 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.

Finally, 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

$$\Delta_t = \ln\left(\eta \Pi_t^{\frac{\varepsilon}{1-\alpha}} e^{\Delta_{t-1}} + (1-\eta)\left[\frac{1-\eta\Pi_t^{\varepsilon-1}}{1-\eta}\right]^{\frac{\varepsilon}{(1-\alpha)(\varepsilon-1)}}\right) \tag{8}$$

where $\Pi_t \equiv P_t/P_{t-1}$ defines the 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 Benhabib et al. (1991) or 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.[13] The government levies an income tax $\tau_h = 1 - S_c/S_h$ designed to offset any steady-state distortions caused by the habit externality, and rebates it in lump-sum fashion, T_{ht}, to households.

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 $S_{ht} \equiv (H_t - X_{ht})/H_t$. To ensure nonnegative marginal utilities, I consider the following dynamics for the *logarithms* of aggregate surplus levels:

$$\begin{aligned} s_{ct+1} &= (1-\rho_s)s_c + \rho_s s_{ct} + \lambda_{ct}\varepsilon_{ct+1} \\ s_{ht+1} &= (1-\rho_s)s_h + \rho_s s_{ht} + \lambda_{ht}\varepsilon_{ht+1} \end{aligned} \tag{9}$$

where $\varepsilon_{ct} \equiv c_t - E_{t-1}c_t$ and $\varepsilon_{ht} \equiv h_t - E_{t-1}h_t$, with common persistence ρ_s.

These preferences imply the intertemporal and static marginal rates of substitution

$$M_{t+1} = \frac{\partial U_t/\partial C_{t+1}}{\partial U_t/\partial C_t} = \beta\left(\frac{C_{t+1}S_{ct+1}}{C_t S_{ct}}\right)^{-\gamma} \tag{10}$$

$$MRS_t = -\frac{\partial U_t/\partial N_t}{\partial U_t/\partial C_t} = -\frac{\chi C_t^\gamma}{H_t^\gamma}\frac{\partial H_t}{\partial N_t}\frac{S_{ct}^\gamma}{S_{ht}^\gamma} \tag{11}$$

Habits amplify the consumers' aversion to fluctuations in consumption and hence affect asset prices, including the consumption-saving tradeoff through the Euler equation, but they also affect the consumption-labor tradeoff. The asset pricing implications and the spillover of habits on the consumption-saving tradeoff are controlled by the process s_{ct}. The spillover of habits on the consumption-labor tradeoff is controlled instead by the process $s_{ct} - s_{ht}$. The role of the habit in home consumption is to offset the effect of the habit in market consumption on the labor choice. In fact, the inclusion of only consumption habits in the production economy, as studied by Lettau and Uhlig (2000), results in a labor process with the wrong cyclicality that is either too volatile or that smooths consumption too much.

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

[13] 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.

Campbell and Cochrane (1999).

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

$$\lambda_{ct} = \frac{\sigma_c}{\sigma_{ct}} \frac{1}{S_c} \sqrt{1 - 2(s_{ct} - s_c)} - 1 \tag{12}$$

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

$$\lambda_{ht} = \frac{S_c}{1 - S_c} \frac{1 - S_h}{S_h} \frac{\sigma_{cht}/\sigma_{ht}^2}{\sigma_{ch}/\sigma_h^2} \lambda_{ct} \tag{13}$$

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 appendix proves these properties.)

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).

Furthermore, 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 Lettau and Uhlig's critique. 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:

$$S_c = \sqrt{\frac{\gamma \sigma_c^2}{1 - \rho_s - \xi_1/\gamma}}, \qquad S_h = \left(1 + \frac{1 - S_c}{S_c}(1 + \xi_2)\frac{\sigma_{ch}}{\sigma_h^2}\right)^{-1} \tag{14}$$

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, which 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; since $S_h = \arg\min_{S_h} var_t[\lambda_{ht}\varepsilon_{ht+1} - (1 + \xi_2)\lambda_{ct}\varepsilon_{ct+1}]$, it follows that the term $\hat{s}_{ht} - \hat{s}_{ct}$ in (11) is as close as possible to $\xi_2 \hat{s}_{ct}$ in a mean-squared error sense.

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, 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.

Thus, 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,

$$i_t = -\ln E_t \beta e^{-\gamma\Delta c_{t+1} - \gamma\Delta s_{ct+1} - \pi_{t+1}} \tag{15}$$

and the labor supply equation,

$$W_t = \frac{\chi}{1-\tau_h}\frac{A_t^{1-\gamma}C_t^{\gamma}}{(1-N_t)^{\gamma}}\left(\frac{S_{ct}}{S_{ht}}\right)^{\gamma} \tag{16}$$

4.1.3. Government

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

$$i_t = i^* + \phi_\pi \pi_t + \phi_y[c_t - c_{Pt} - (1-\alpha)(n-\Delta)] \tag{17}$$

for an interest rate level i^* consistent with zero steady-state inflation, with a real rate $r_t = -\ln E_t M_{t+1}$.

Fiscal policy runs a balanced budget: $T_{ht} = \tau_h W_t N_t$ 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 $N_t = \int_0^1 N_t(i)di$, and hence, 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 technology 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 technology shock $e_t \sim Niid(0, \sigma^2)$ has a permanent effect. For example, $\theta = 0$ is associated with random-walk technology, while $\theta = 1$ is associated with the typical trend-stationary specification (e.g., Galí 2008). Consistent with definition (1), θ also indexes the amount of uncertainty in the consumption process that can be removed by monetary policy, as it cannot have permanent 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 solve the individual firm's problem, *(b)* the choice of market and home consumption solve the individual consumer's problem, *(c)* the goods and labor markets clear, *(d)* market and home consumption habits evolve according to (9), *(e)* price dispersion evolves according to (8), *(f)* the nominal rate follows rule (17), and 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 by a first-order approximation around the risky steady state as described by Lopez, Lopez-Salido, and Vazquez-Grande (2016). (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

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

to a first-order perturbation around the risky steady state. In words, average inflation is zero, and the equilibrium dependence of consumption and inflation on surplus consumption is zero. (To a first-order approximation price dispersion is a trivial process, $\Delta_t = 0$.)

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. The two parameters control the spillover of habit dynamics onto equilibrium quantities and inflation.

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 (15).

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 (16). 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.,

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

and hence ξ_2 disciplines the variation in the marginal rate of substitution between consumption and labor. For example, a choice of $\xi_2 = 0$ would imply a marginal rate of substitution between consumption and labor identical to one under CRRA utility. This property prevents households from varying excessively 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.

Not surprisingly, the particular parameterization required to achieve an approximate macro-finance separation is close to the point $\xi_1 = \xi_2 = 0$:

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

where L_2 and L_3 are second-order moments, and the approximate equilibrium coefficients are

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

(The online appendix derives these formulas.) Parameters κ and φ are transformations of other deep parameters and, in the absence of risk, 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.

It follows that the equilibrium allocation is approximately observationally equivalent to a model with CRRA households—an implication confirmed in Figure 3a. 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

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 technology 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 three cash flows of interest—consumption, dividends, and inflation. Figure 3a reports the associated fit of the model, which captures relatively well the autocorrelation and volatility of cash flows observed over the 1994-2018 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. 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.

Figure 1b, Figure 3, and Table 4 report the average term structures of the equity premium, of real interest rates and of the welfare cost of uncertainty implied by the model, as well as the interquartile range of the term structure of welfare costs. The model produces an average equity premium of 7.25 percent, close to the observed value reported in Figure 1a. The market premium is lower than the premium commanded by short-term equities. The representative household is particularly sensitive to fluctuations in market income, as countercyclical marginal costs exacerbate the procyclicality of corporate profits after a technology shock, but fears less uncertainty in the long run because dividends and consumption are cointegrated, and hence have the same riskiness in the long run.

Crucially, this simple model is able to capture the main empirical properties of the term structure of welfare costs documented in Section 3, including its countercyclical variation.

	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
	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 technology growth	0.0013
	ρ_a	Persistence of the conditional mean of technology growth	0.794
	σ	Conditional volatility of technology	0.0128
	$1-\theta$	Fraction of permanent technology 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 1994-2018.
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 1994-2018 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.

horizon (years), $\mathcal{N}$ up to	0.5	1	1.5	2	5	10	30	∞
cost of uncertainty, $L^{\mathcal{N}}$ (% p.a.)	0.41	0.31	0.25	0.20	0.11	0.08	0.10	0.12

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

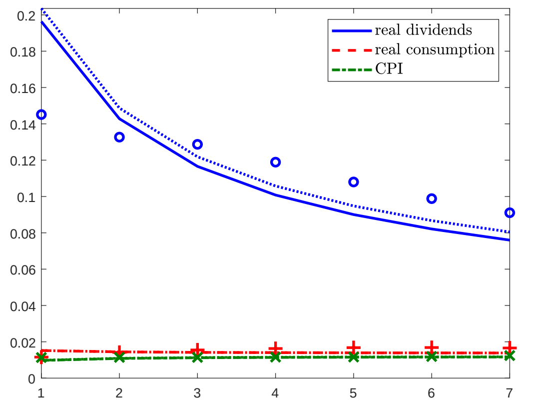

(a) *Average term structures of cash flow volatilities; model (solid and dashed lines), model under CRRA utility (dotted lines), and data (markers).*

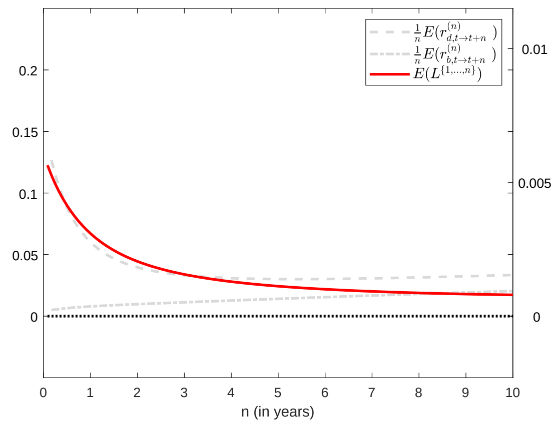

(b) *Average term structures of equity, interest rates, and welfare costs.*

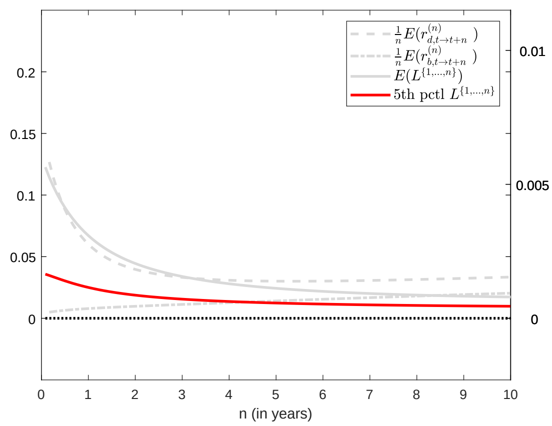

(c) *First quartile of the term structure of welfare costs.*

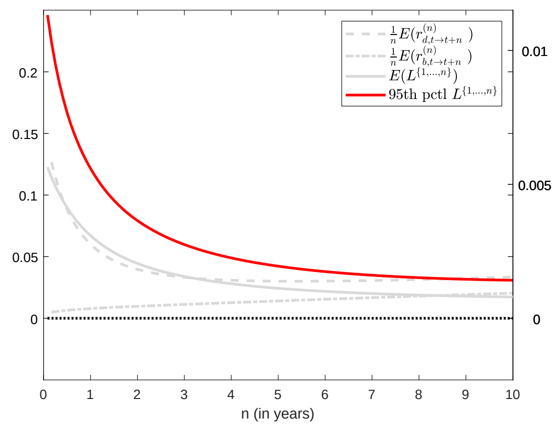

(d) *Third quartile of the term structure of welfare costs.*

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 interquartile range. Annualized values on the left axis for the term structures of casfhlows, equity, and interest rates and on the right axis for welfare costs.

4.2.1. Is the nominal-real covariance plausible?

A recent literature documents that the dynamics of inflation and consumption growth changed precisely in the late 1990s—the period over which we have evidence about the term structure of welfare costs. For example, Song (2017) and Boons et al. (2020) document that predictive regressions of future real consumption growth on lagged inflation are characterized by a negative regression coefficient before the late 1990s and by a positive coefficient thereafter. Furthermore, Campbell et al. (2020) document that the correlation between inflation and the output gap has gone from negative in the 1980s and 1990s to positive in the 2000s.

In this context, the model presented in Section 4 captures correctly the nominal-real covariance that the literature has documented as a distinguishing feature of the last two decades. The model's emphasis on technology shocks as a key driving factor to rationalize the evidence about the term structure of welfare costs implies (i) that high inflation signals good future real growth, and (ii) that high inflation is associated with a large gap between output and potential output. These properties are at odds with the pre-2000s post-war US experience, but are in line with the evidence over the period of interest.

Finally, the model predicts a negative contemporaneous correlation between inflation and real consumption growth. In the data, this correlation coefficient is negative throughout the 1994-2018 period, with the exception of the last quarter of 2008. In fact, the correlation coefficient using annual rates is noisy and even positive when the last quarter of 2008 is included.[14] The point estimate of .04 is particularly imprecise, with a block-bootstrapped p-value of 52 percent for the hypothesis that the correlation coefficient is positive, which I cannot reject. Once I exclude the last quarter of 2008 from the sample, however, the correlation coefficient and the block-bootstrapped p-value drop to $-.36$ and 14 percent.

Therefore, the simple one-shock model presented in Section 4 captures key implications for the nominal-real covariance seen in the data during the period over which evidence about the term structure of welfare costs is available.

5. Stabilization Policy Evaluation

This section investigates the optimal stabilization policy in the model laid out in Section 4. The presence of two externalities (nominal rigidities and external habits) motivates a nontrivial policy tradeoff. I will focus on the optimal parameterization of a simple Taylor rule that nests a policy that fills the consumption gap and a policy that removes consumption uncertainty.[15]

[14] Song (2017) also finds that the contemporaneous correlation between inflation and real consumption growth may have switched sign, turning positive, after 1998 but that the result is not statistically significant.

[15] A characterization of the optimal policy regime under discretion or commitment is beyond the scope of this paper. Similarly, I do not consider additional policy tools to remove the policy tradeoff—for example, Ljungqvist and Uhlig (2000) consider time-varying taxation to address the habit externality. The purpose of the present analysis is to point out how a model consistent with observed welfare cost measures can easily

Aggregate (scaled) welfare $V_t \equiv U_t e^{(\gamma-1)[(1-\alpha)n+s_c]}$ can be written as:

$$V_t = (N^{1-\alpha}S_c)^{\gamma-1}E_t\sum_{j=0}^{\infty}\beta^j\left(\frac{(C_{t+j}S_{ct+j})^{1-\gamma}}{1-\gamma}+\chi\frac{(H_{t+j}S_{ht+j})^{1-\gamma}}{1-\gamma}\right)$$
$$= \frac{e^{(1-\gamma)(a_t+\hat{s}_{ct})}}{1-\gamma}\left[e^{(1-\gamma)(c_t-a_t-(1-\alpha)n)}+(1-\alpha)\left(e^{\hat{s}_{ht}-\hat{s}_{ct}}\frac{1-Ne^{\hat{\Delta}_t+\frac{c_t-a_t-(1-\alpha)(n-\Delta)}{1-\alpha}}}{1-N}\right)^{1-\gamma}\right]+\beta E_t V_{t+1} \tag{18}$$

where the second equality used the property of the efficient steady state, $(1-\alpha)(CS_c)^{1-\gamma} = \chi(HS_h)^{1-\gamma}$. (See the appendix for a derivation.)

The constrained-efficient allocation maximizes welfare subject to the definition of habits (9), to the law of motion of price dispersion (8), and to the price setting equations,

$$L_t = E_t\beta\eta e^{(1-\gamma)[\mu+(1-\theta)\sigma\varepsilon_{t+1}]}\Pi_{t+1}^{\frac{\varepsilon}{1-\alpha}}L_{t+1}+\frac{\varepsilon(1-\tau_f)\chi_0}{(\varepsilon-1)(1-\alpha)}e^{\frac{\widetilde{c}_t}{1-\alpha}+(1-\gamma)a_{Tt}-\gamma\hat{s}_{ct}}\left(1-e^{\frac{\widetilde{c}_t}{1-\alpha}+\Delta_t}\right)^{-\gamma}$$
$$L_t = E_t\beta\eta e^{(1-\gamma)[\mu+(1-\theta)\sigma\varepsilon_{t+1}]}\Pi_{t+1}^{\varepsilon-1}L_{t+1}\left(\frac{1-\eta\Pi_{t+1}^{\varepsilon-1}}{1-\eta\Pi_t}\right)^{\Theta}+e^{(1-\gamma)(\widetilde{c}_t+a_{Tt})-\gamma\hat{s}_{ct}}\left(\frac{1-\eta}{1-\eta\Pi_t^{\varepsilon-1}}\right)^{\Theta} \tag{19}$$

where $\Theta \equiv (1-\alpha+\alpha\varepsilon)/(1-\alpha)(\varepsilon-1)$ and $\widetilde{c}_t = c_t - a_t$ is the market consumption gap, and L_t is an auxiliary variable.

By examining the structural equations we can study the consequences of two simple regimes—inflation targeting and risk premium targeting. Implementation will be done by Taylor rule (17).

5.1. Inflation targeting (or: fill the consumption gap)

Under an inflation targeting regime consumption equals the flexible-price level, and hence inflation is zero ($\pi_t = 0$). We have:

$$c_t = a_t + (1-\alpha)n, \qquad n_t = \bar{n}, \qquad \varepsilon_{ct+1} = e_{t+1} \sim Niid(0,\sigma^2)$$

imply that people fear consumption fluctuations so much as to overturn the standard result that stable inflation is the macroeconomic priority. Similarly, in a different setting Leith, Moldovan, and Rossi (2012) study monetary policy trade offs 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.

It follows that welfare (18) under this regime is:

$$V_t^{it} = \frac{e^{(1-\gamma)(a_{Pt}+\hat{s}_{ct})}}{1-\gamma}\left[1+(1-\alpha)e^{(1-\gamma)(\hat{s}_{ht}-\hat{s}_{ct})}\right]e^{(1-\gamma)a_{Tt}} + \beta E_t V_{t+1}^{it} \tag{20}$$

with highly volatile surplus consumption $\hat{s}_{ct+1}^{it} = \rho_s \hat{s}_{ct}^{it} + \lambda_{ct}^{it}\sigma\varepsilon_{t+1}$.

5.2. Risk premium targeting (or: remove consumption volatility)

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

In the context of the model, removing consumption risk—i.e., $c_t = a_{Pt} + (1-\alpha)(n-\Delta)$—is equivalent to stabilizing surplus consumption, and hence the maximum risk-return tradeoff

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

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.

Under such a risk premium targeting regime consumption equals the permanent component of technology, and hence $\varepsilon_{ct} = (1-\theta)e_t$. If follows that welfare (18) under this regime is:

$$V_t^{rpt} = \frac{e^{(1-\gamma)(a_{Pt}+\hat{s}_{ct})}}{1-\gamma}\left[e^{(\gamma-1)(1-\alpha)\Delta} + (1-\alpha)\left(e^{a_{Tt}+\hat{s}_{ht}-\hat{s}_{ct}}\frac{1-Ne^{-\frac{a_{Tt}}{1-\alpha}+\hat{\Delta}_t}}{1-N}\right)^{1-\gamma}\right] + \beta E_t V_{t+1}^{rpt} \tag{21}$$

with minimally volatile surplus consumption $\hat{s}_{ct+1}^{rpt} = \rho_s \hat{s}_{ct}^{rpt} + \lambda_{ct}^{rpt}(1-\theta)\sigma\varepsilon_{t+1}$. The only remaining unknown is price dispersion Δ_t, which depends on the dynamics of equilibrium inflation driven by the price setting equations (19).

5.3. Approximate welfare criterion

Although the results use a global solution, 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 welfare losses $\mathbb{L} \equiv -E(V_t)/(CS_c)^{1-\gamma}$,

$$\begin{aligned}\mathbb{L} = {} & \frac{1}{2}\frac{\eta\varepsilon(1-\alpha+\alpha\varepsilon)}{(1-\eta)^2(1-\alpha)}var(\pi_t) + \frac{1}{2}\frac{\gamma(2-\alpha)+\alpha}{1-\alpha}var(\widetilde{c}_t) \\ & + \frac{\gamma-1}{2}var(s_{ct}) + \frac{(\gamma-1)(1-\alpha)}{2}var(s_{ht}) - (\gamma-1)cov(\widetilde{c}_t, s_{ht}-s_{ct}) + t.i.p.\end{aligned} \tag{22}$$

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

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 level as much as they would 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 (22), the central bank wants to stabilize the consumption gap, $c_t - a_t$—and hence to stabilize inflation in the process—as well as consumption uncertainty, $c_t - E_{t-1}c_t$—to stabilize the price of risk, and hence risk premia. Since achieving both goals is unfeasible, the optimal monetary policy trades them off to maximize welfare.

5.4. Fill the gaps or remove volatility?

I denote by $\mathbb{L}^{\pi}$ the average per-period welfare loss under the inflation targeting regime and by $\mathbb{L}^{x}$ the loss under the risk premium targeting regime. Under the baseline calibration $\mathbb{L}^{\pi} \geq \mathbb{L}^{x}$, i.e., risk premium targeting dominates inflation targeting. As shown in Table 5, this result is robust to varying the value of $\theta \in (0, 1)$, and hence the size of the transitory component of consumption that can be removed by monetary policy.

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, Table 5 shows that a move from the baseline Taylor rule to an inflation targeting regime is welfare decreasing. Although the welfare weight attached to the inflation objective dominates the welfare weights attached to consumption and risk aversion stabilization, the volatility of surplus consumption necessary to explain asset prices is sufficiently large that removing its fluctuations becomes the priority.

The welfare dominance of the risk premium targeting regime over the inflation targeting regime is strict for any $\theta \in (0, 1)$, with equality only when $\theta = 0$, that is, when technology is a pure random walk that negates any role for monetary policy and nominal rigidities. When all technology shocks are permanent, consumers cannot forecast future movements in the real rate, so they choose a random-walk consumption path, while firms cannot forecast future movements in marginal costs, so inflation is zero. Since some mean reversion is necessary to let nominal rigidities model a component in dividends that captures the downward-sloping term structure of welfare costs, it follows that the dominance of risk premium targeting over inflation targeting is deeply linked to the model's ability to explain the empirical term structure properties.

Fraction of removable uncertainty, θ	0	1/3	1/2	2/3	Baseline	3/4	1
Ratio of welfare losses under risk premia vs. inflation targeting, $\mathbb{L}^x/\mathbb{L}^\pi$	1	0.85	0.72	0.61	0.60	0.59	0.58
Welfare improvement relative to baseline parameterization and Taylor rule under:							
- risk premium targeting					+8%		
- inflation targeting					−53%		

Table 5: Ratio of welfare losses $\mathbb{L}^x/\mathbb{L}^\pi$. $\mathbb{L}^\pi$: Welfare losses under a simple policy regime that targets inflation (or the consumption gap). $\mathbb{L}^x$: Welfare losses under a simple policy regime that targets risk premia (or consumption uncertainty). θ denotes the fraction of uncertainty in the consumption process that can be removed by monetary policy.

The extent of the difference between the two policies can be made most evident by noticing how the simple regimes can be implemented by an appropriate choice of the reaction parameters in Taylor rule (17). 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.

5.5. What if habits were internal?

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). It follows that 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.

6. Conclusion

Lucas (1987, 2003) introduced the notion of the cost of aggregate uncertainty as a thought experiment to provide an assessment of the tradeoff between growth and macroeconomic stability. Analogously, the term structure of the cost of uncertainty 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 recent financial crisis. The finding of sizable and volatile costs imposed by an increase in short-run uncertainty fits in with a recent literature that finds high and time-varying short-maturity risk premia as a pervasive phenomenon across different asset classes (Binsbergen and Koijen 2017).

The result that the marginal cost of uncertainty is a linear combination of risk premia makes one of the main tasks of macroeconomics—the assessment of macroeconomic priorities (Lucas 2003)—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 differ in their asset pricing implications prescribe fundamentally different policies.

Appendix

A. Theoretical results

A.1. Proof of proposition 1

Differentiating (1) 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 (2), the terms cointaining 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 (3) 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 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

$$
\begin{aligned}
1 &= \int_{S(t)} \left(\frac{P_t(i)}{P_t}\right)^{1-\varepsilon} di + (1-\eta)\left(\frac{P_t^*}{P_t}\right)^{1-\varepsilon} = \eta\Pi_t^{\varepsilon-1} + (1-\eta)\left(\frac{P_t^*}{P_t}\right)^{1-\varepsilon} \\
e^{\Delta_t} &= \int_{S(t)} \left(\frac{P_t(i)}{P_t}\right)^{-\frac{\varepsilon}{1-\alpha}} di + (1-\eta)\left(\frac{P_t^*}{P_t}\right)^{-\frac{\varepsilon}{1-\alpha}} = \eta\Pi_t^{\frac{\varepsilon}{1-\alpha}} e^{\Delta_{t-1}} + (1-\eta)\left(\frac{P_t^*}{P_t}\right)^{-\frac{\varepsilon}{1-\alpha}}
\end{aligned}
$$

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

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

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

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

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 (6), 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 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.[16]

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.4d, 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.[17]

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

[17] I find options-implied interest rates with nearly perfect correlations with the corresponding LIBOR and

B.2. Errors in synthetic replication

Figure B.5 plots the auxiliary map $X \mapsto P^{(n)}_{dt} + X\ P^{(n)}_{\$bt}$ for selected trading days and maturities. As shown by the lower part of Figure B.5, 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.4c 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 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.4a plots the prices of the synthetic dividend claims. Figure B.4b shows how synthesized dividend strip prices are leading indicators of subsequent dividends, in line with their interpretation as risk-neutral expectations of dividends.

Treasury rates but with different levels that lie about halfway between the two proxies.

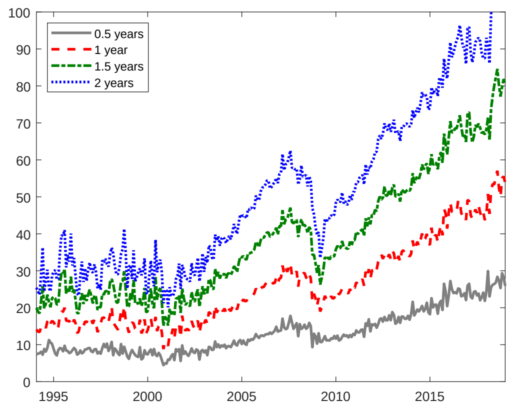

(a) *Price of the next n years of dividends.*

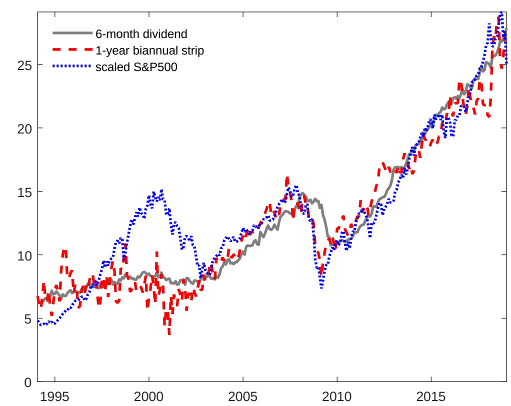

(b) *Dividend strips and realized dividends.*

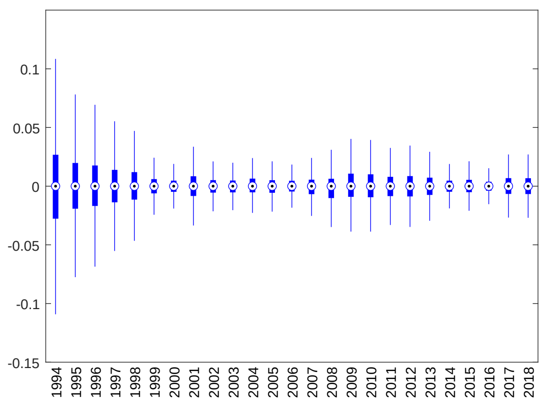

(c) *Law-of-one-price violations.*

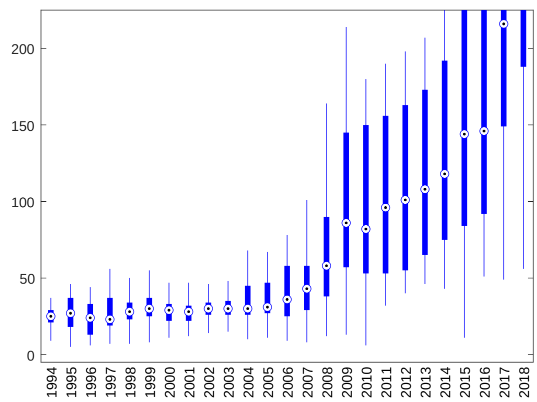

(d) *Number of strikes per trade date.*

Figure B.4: 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.

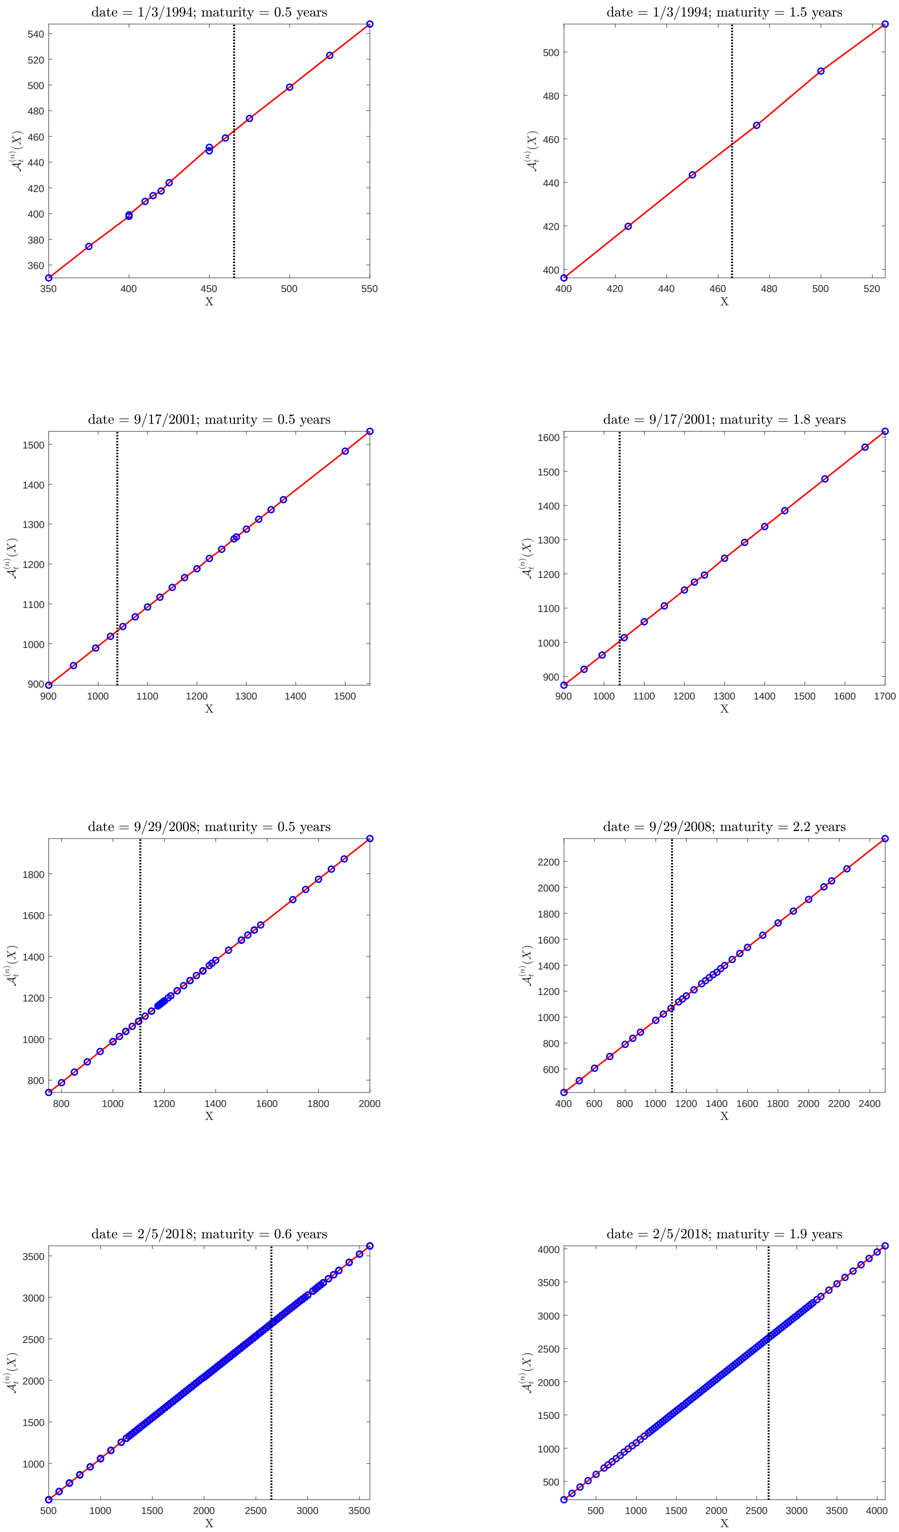

Figure B.5: Auxiliary map $X \mapsto P_{dt}^{(n)} + X\ P_{\$bt}^{(n)}$ on given trading days and for given maturities. The dotted lines indicate the value of the index at each respective trading day.

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.

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.

D. Backus, M. Chernov, and S. E. Zin. Sources of entropy in representative agent models. *Journal of Finance*, 69(1):51–99, 2014. doi: 10.1111/jofi.12090.

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.

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. Benhabib, R. Rogerson, and R. Wright. Homework in macroeconomics: Household production and aggregate fluctuations. *Journal of Political Economy*, 99(6):1166–1187, 1991. doi: 10.1086/261796.

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.

M. Boons, F. Duarte, F. De Roon, and M. Szymanowska. Time-varying inflation risk and stock returns. *Journal of Financial Economics*, 136(2):444–470, 2020. doi: 10.1016/j.jfineco.2019.09.012.

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.

M. M. Croce. Welfare costs in the long run. Manuscript, 2012.

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.

I. Dew-Becker and S. Giglio. Asset pricing in the frequency domain: Theory and empirics. *Review of Financial Studies*, 29(8):2029–68, 2016. doi: 10.1093/rfs/hhw027.

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.

L. G. Epstein, E. Farhi, and T. Strzalecki. How much would you pay to resolve long-run risk? *American Economic Review*, 104(9):2680–97, 2014. doi: 10.1257/aer.104.9.2680.

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.

J. Galí, M. Gertler, and J. D. López-Salido. Markups, gaps, and the welfare cost of business fluctuations. *Review of Economics and Statistics*, 89(1):44–59, 2007. doi: 10.1162/rest.89.1.44.

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.

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.

L. P. Hansen, J. C. Heaton, and N. Li. Consumption strikes back? Measuring long-run risk. *Journal of Political Economy*, 116(2):260–302, 2008. doi: 10.1086/588200.

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. Risk-adjusted linearizations of dynamic equilibrium models. Manuscript, 2016.

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.

D. Song. Bond market exposures to macroeconomic and monetary policy risks. *Review of Financial Studies*, 30(8):2761–2817, 2017. doi: 10.1093/rfs/hhx039.

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.