

Welfare Implications of the Term Structure of Returns: Should Central Banks Fill Gaps or Remove Volatility?

Pierlauro Lopez^{a,1,*}

^a*Banque de France*

Abstract

The welfare cost of economic uncertainty has a term structure that is a simple transformation of the term structures of the equity premium and interest rates. Twenty years of financial market data suggest a term structure of welfare costs that is downward-sloping on average and during downturns. This evidence offers guidance in selecting a model to study the benefits of greater consumption stability from a structural perspective. A model with nominal rigidities and nonlinear external habits can rationalize the evidence and motivates the competitive level and volatility of consumption as inefficient. The model is observationally equivalent to a standard New Keynesian model with CRRA utility but the optimal policy prescription is overturned; in the model the central bank should focus on removing consumption volatility rather than on filling the gap between consumption and its flexible-price counterpart.

JEL classification: E32; E44; E61; G12.

Keywords: Welfare cost of business cycles, Macroeconomic priorities, Equity and bond yields, Optimal monetary policy, Risk premia targeting.

1. Introduction

How much growth are people willing to trade against the elimination of uncertainty around future consumption? Lucas (1987) introduced this thought experiment as a measure of the welfare cost of lifetime fluctuations and as an indicator of the priority of growth over stabilization policies. At least since Alvarez and Jermann (2004), it has been understood that a sufficiently complete financial market can reveal directly the cost of lifetime fluctuations at the margin. In this context, recent evidence about the term structures of equity and interest rates can reveal how uncertainty at different horizons contribute to the marginal cost of lifetime consumption fluctuations, and forms therefore a rich set of empirical features with implications for welfare analysis.

**Present address:* Macro-Finance Division, Banque de France, 31 rue Croix des Petits Champs, 75001 Paris.

Email address: pierlauro.lopez@banque-france.fr (Pierlauro Lopez)

¹I would like to thank John Cochrane, Martin Eichenbaum, Ralph Koijen, Jordi Galí, Antonio Mele, and Fabio Trojani for very helpful advice and comments. The views presented here are those of the author and do not necessarily represent those of any institution of the Eurosystem.

I show how risk premia on some traded derivative instruments combine to measure the marginal cost of uncertainty around arbitrary sets of consumption coordinates. By focusing on fluctuations in the determinants of consumption, in particular labor and market dividend income, I link directly the cost of uncertainty to the price of claims to market dividends; welfare costs at different horizons are tied by no-arbitrage relations to the term structures of equity and interest rates.² The result relies on a standard assumption on the consumption-labor tradeoff, does not require a parametric specification of consumer preferences, and allows for an observable measure of the cost of uncertainty at several horizons over the last two decades.

The empirical analysis finds costs that are sizable, countercyclical, and have nontrivial term structure features. The point estimates, reported in figure 1a, suggest a negatively sloped term structure of welfare costs, driven both by the negative slope of the term structure of equity and by the positive slope of the term structure of interest rates. At the margin people would trade an average of 0.5-0.75 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 last two recessions and compares to much smaller benefits of long-run stability.

This information improves our understanding of the appropriate model to explain why uncertainty is costly, as calculations of welfare costs can vary by orders of magnitude across alternative assumptions on preferences and cashflow processes.³ I propose a model with nominal rigidities and external habit formation that can rationalize the empirical measures of the welfare cost of uncertainty at different horizons, and carry out an explicit welfare analysis to assess whether interventions that remove risk in the economy are Pareto-improving. In the model, the representative consumer is sufficiently sensitive to cashflow fluctuations to value much more a less volatile consumption path than less volatile inflation and thereby overturn a pervasive result in the New Keynesian literature. This analysis offers an example of the potentially dramatic implications of incorporating realistically large discount-rate variation into macroeconomics (see also Cochrane, 2011b).

1.1. Decomposing the cost of uncertainty

Recent evidence about the term structures of returns allows me to link the marginal cost of uncertainty to a rich set of financial market evidence. 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 et al. (2012) and Golez (2014) this paper relies on option data and no-arbitrage relations to synthetically replicate single market dividend payments, so-called dividend strips. I contribute to this literature by proposing a method to extract dividend claim prices by combining options with different moneyness levels to avoid the

²A recent and rapidly growing literature is focusing on risk pricing across maturities (see, for example, Lettau and Wachter, 2007, 2011; Hansen, Heaton and Li, 2008; Binsbergen, Brandt and Koijen, 2012; Binsbergen, Hueskes, Koijen and Vrugt, 2013; Backus, Chernov and Zin, 2014; Borovicka and Hansen, 2014; Belo, Collin-Dufresne and Goldstein, 2015; Dew-Becker and Giglio, 2013).

³See 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.




Figure 1: Average term structures of equity (‘ $\circ$ ’ and dashed line), interest rates (‘ $\times$ ’ and dash-dotted line) and welfare costs (‘+’ and solid line); annualized premia on the left axis for the term structures of equity and interest rates and on the right axis for welfare costs. The left subfigure plots point estimates during 1994-2013 (with block bootstrapped 95% confidence interval); $E(r^{e,m})$ is the equity premium. The right subfigure plots model-implied average term structures in the DSGE model with Campbell-Cochrane habits presented in section 4. $L^{(1,...,n)}$ represents the amount of growth in the entire consumption process at a n -period horizon that people would trade against the elimination of consumption risk over the same horizon. The solid curve tends to a cost of lifetime uncertainty of 5 hundredths of a percent.

need for an interest-rate proxy and to mitigate measurement error by excluding observations that violate the put-call parity relation.

I complement the analysis by 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 recognize the endogeneity of the main source of income (labor income). In an endowment economy, all consumption determinants are exogenous; for example, they are dividends from a Lucas tree. Consumption and dividends differ in practice, and in their calculations Alvarez and Jermann model implicitly an exogenous difference between them. However, the difference between consumption and dividends is not exogenous from a consumer’s point of view. When an endogenous determinant of consumption such as labor income is marginally stabilized, the positive effect on utility is offset by the effect of the associated adjustment in other decision variables (e.g., labor hours). It follows that only claims to exogenous determinants of consumption such as dividends are necessary to measure the marginal cost of uncertainty.⁴

Second, the term structure of the cost of uncertainty answers roughly the question, ‘how much compensation do people command to bear n -year ahead cashflow uncertainty?’ This question compares to, ‘how much compensation do people command to bear uncertainty at business-cycle frequency in the *entire* cashflow process?’, which is the question they study. Their answer depends a lot on the parametric assumptions about the filter that separates the trend and the business-

⁴The cost of uncertainty can be measured using dividend claims as long as the residual exogenous determinants of consumption net of market and labor income are relatively small or claims to such cashflows command a relatively small risk premium. In any event, the measure can always be interpreted as the cost of uncertainty in market and labor income, which is a meaningful quantity on its own right.

cycle frequencies of the cashflow process. The question I am interested in is nonparametric and complements their exercise by decomposing the marginal cost of uncertainty in the time domain rather than in the frequency domain.

1.2. *Welfare implications of welfare cost measures*

Evidence of substantial welfare costs is only a necessary condition to justify attributing 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 cashflows, as opposed to fluctuations in its distance (gap) from some benchmark level, and is silent as to whether the current level of risk in the economy is Pareto optimal. For example, structural macroeconomic models for policy evaluation typically prescribe stabilizing gaps between current and efficient cashflow levels (see Clarida, Galí and Gertler, 1999; Woodford, 2010, among many others) and not the amount of uncertainty in the economy. What welfare cost measures really say is that the empirical properties of the term structure of welfare costs are not only relevant as a quantitative diagnostic of an asset pricing model that attempts to understand risk pricing across maturities but, also, of any macroeconomic model that seeks to assess the priority of different stabilization policies.

Accordingly, I analyze the welfare implications of a structural model that captures the observed term structure of welfare costs. It is well-known that replicating *jointly* a downward-sloping term structure of equity and an upward-sloping term structure of interest rates from a structural perspective is problematic (e.g., Binsbergen and Koijen, 2015). To the best of my knowledge, the model by Lopez, Lopez-Salido and Vazquez-Grande (2015b) is the only candidate structural model that is able to explain jointly the observed term structure properties. Building on that framework, I rely on Campbell and Cochrane (1999) external habit formation and nominal rigidities and study the policy implications in that setting.

In the model, people fear instability in the short run more than in the long run because market income is highly procyclical but partly mean reverting. Figure 1b reports the average term structure of welfare costs in the model. As in the data, the model-implied costs of short-run uncertainty are sizable and volatile and substantially larger than the cost of uncertainty over long horizons, with a cost of lifetime uncertainty equal to one-half of one-tenth of a percentage point of perpetual growth in consumption, similar to the original calculations by Lucas (1987, 2003).

The model of quantities is observationally equivalent to a textbook New Keynesian model (e.g., Galí, 2008) but the asset pricing implications differ dramatically. Two externalities (nominal rigidities and external habits) drive the model and motivate a role for policy interventions, so I can adopt an explicit welfare-theoretic approach to the design of the optimal monetary policy. The model reconciles the two notions of macroeconomic stability found in the extant literature. First, the central bank wants to stabilize the consumption gap (or, equivalently, inflation); second, it wants to remove consumption uncertainty (or, equivalently, risk premia variation). Since achieving both goals is unfeasible, the optimal monetary policy trades them off to minimize the welfare metric. However, when comparing two simple regimes—one that closes the consumption gap (inflation targeting) and one that removes consumption uncertainty (risk premia targeting)—I find that stabilizing levels is typically a priority over filling the troughs.

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 as to allow people to trade the full set of zero-coupon bonds and equities.

Identical risk-averse consumers indexed by $i \in [0, 1]$ have time- t preferences $U_t = E_t U(C_i, N_i, X_i)$, 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.⁵ Without loss of generality let factor X_i depend on individual consumption and labor only via aggregate consumption and labor, $C = \int_0^1 C_i di$ and $N = \int_0^1 N_i di$. Since there is a continuum of agents each of which 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 thereby having to affect all aggregate quantities, including factor X .

The determinants of consumption, $C_t = D_t + W_t N_t + e_t$, explicitly include market 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 $\bar{C}_{t+n}$ denote the consumption level that is hypothetically offered to the i th individual at time $t + n$, which I refer to as *stable* consumption. Stable consumption consists of a stable stream of dividend income, $E_t(D_{t+n})$, and of labor income, $E_t(W_{t+n}N_{t+n})$, which associates with the labor level $\bar{N}_{t+n} = E_t(W_{t+n}N_{t+n})/W_{t+n}$. I parametrize stable consumption as

$$\bar{C}_{t+n}(\theta) = \theta E_t(D_{t+n} + W_{t+n}N_{t+n}) + (1 - \theta)(D_{t+n} + W_{t+n}N_{t+n}) + e_t$$

where the parameter $\theta \in [0, 1]$ represents the fraction of ex-post uncertainty in the main determinants of consumption that is removed. Accordingly, the associated labor level is

$$\bar{N}_{t+n}(\theta) = \theta \bar{N}_{t+n} + (1 - \theta)N_{t+n}$$

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

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

measures the cost of fluctuations, where the index set $\mathcal{N} \subset \bar{\mathbb{N}} \equiv \{1, \dots, \infty\}$ indicates which coordinates are stabilized and allows for focusing on any window of interest.

Two particularly interesting quantities are the *total cost* $\mathcal{L}_t^{\bar{\mathbb{N}}}(1)$, which measures how much extra growth the elimination of all uncertainty in dividend and labor income is worth, and the *marginal cost* $L_t^{\mathcal{N}} \equiv \frac{\partial}{\partial \theta} \mathcal{L}_t^{\mathcal{N}}(0)$, which represents how much extra growth a marginal stabilization is worth at the current level of uncertainty.⁶

⁵The definition of cost of uncertainty has a meaningful interpretation even if one relaxes the assumption of a representative consumer. The appendix provides the detailed argument.

⁶The online appendix discusses the relationship between definition (1) and the definitions by Lucas (1987) and Alvarez and Jermann (2004).

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

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

so the marginal rate of substitution between consumption and labor equals the relative price. This optimality condition is standard and implies that there are no distortions in the consumption side of the economy that generate so-called labor wedges.

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

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

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, $D_{c,t}^{(n)} = E_t M_{t,t+n} C_{t+n}$ is the price of a n -period consumption strip, $D_{d,t}^{(n)} = E_t M_{t,t+n} D_{t+n}$ is the price of a n -period dividend strip, and $D_{b,t}^{(n)} = E_t M_{t,t+n}$ is the price of a n -period zero-coupon 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 exactly offsets the benefits of a marginal stabilization in labor income.

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

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

where $R_{t \rightarrow t+n}^{e,(n)} = D_{t+n} D_{b,t}^{(n)} / D_{d,t}^{(n)}$.

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^N for any coordinate set $\mathcal{N} \subset \overline{\mathbb{N}}$.

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

$$\begin{aligned} L_t^N &= \alpha_t^N \sum_{n \in \mathcal{N}} \omega_{n,t}^N l_t^{(n)} \\ &= \alpha \sum_{n \in \mathcal{N}} \omega_n l_t^{(n)} + O(\|\alpha_t^N \omega_{n,t}^N - \alpha \omega_n^N; l_t^{(n)}\|^2) \end{aligned} \quad (4)$$

with scaling factor $\alpha_t^N \equiv \sum_{n \in N} n D_{d,t}^{(n)} / \sum_{n \in N} n D_{c,t}^{(n)}$, hence $\alpha^N = \alpha = D/C$ in a deterministic steady state, and where the weights $\omega_{n,t}^N \equiv n D_{d,t}^{(n)} / \sum_{n \in N} n D_{d,t}^{(n)}$ are positive and such that $\sum_{n \in N} \omega_{n,t}^N = 1$.

Equation (4) shows how the first-order effect of distinguishing between consumption and dividends is that the cost of uncertainty around any coordinate set N scales the linear combination of the term structure components by a factor equal to the average dividend-consumption ratio. Intuitively, the scaling factor translates points of growth in dividend income into points of growth in consumption. This factor can be estimated at about 4.47% over the 1994-2013 period using data from the U.S. Flow of Funds Accounts on net dividends paid out by nonfarm nonfinancial corporates and on personal consumption expenditures in nondurable goods and services.

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 n th coordinate of consumption against the elimination of a fraction θ of the aggregate uncertainty around it.⁷ 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 the n th coordinate of dividend income of $\frac{1}{n}(E_t R_{t \rightarrow t+n}^{e,(n)} - 1)$ by shouldering a conditional entropy of $\mathcal{V}_t(R_{t \rightarrow t+n}^{e,(n)}) = \mathcal{V}_t(D_{t+n})$, i.e., the uncertainty in n -period ahead dividend income. Since n -period ahead dividend income is expected to finance a fraction $\alpha_t^{(n)}$ of consumption, the experienced growth rate in consumption must be $\frac{1}{n} \alpha_t^{(n)} (E_t R_{t \rightarrow t+n}^{e,(n)} - 1)$.

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 on the market. In absence of arbitrage opportunities, put-call parity holds:

$$C_{t,t+n} - \mathcal{P}_{t,t+n} = \mathcal{P}_t - \sum_{j=1}^n D_t^{(j)} - X P_{b,t}^{(n)} \quad (5)$$

where $C_{t,t+n}$ and $\mathcal{P}_{t,t+n}$ are the nominal prices at time t of a call and a put European options on the market index with maturity n and nominal strike price X , $P_{b,t}^{(n)} = E_t M_{t,t+n} P_t / P_{t+n}$ is the nominal price of a n -period nominal zero-coupon bond, $\mathcal{P}_t = P_t E_t \sum_{j=1}^{\infty} M_{t,t+j} D_{t+j}$ 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 n th dividend strip, where P denotes the price level. Since the only unknowns in equation (5) are the prices of the dividend strips, $\{D_t^{(n)}\}$, one can synthetically replicate them (Binsbergen et al., 2012).

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. Standard index option classes, with twelve monthly maturities of up to one year, and Long-Term Equity Anticipation Securities

⁷Note that the cost L_t^N associated with the singleton set $N = \{n\}$ describes the cost per unit of uncertainty in the n th coordinate of consumption in terms of growth in that coordinate and that coordinate alone. The insurance would therefore change the distribution of the level of the n th coordinate of consumption, keeping the distribution of all other coordinates unchanged.

(LEAPS), with ten maturities of up to three years, are exchange traded on the Chicago Board Options Exchange (CBOE) since 1990. The overall size of the index option market in the U.S. has grown rapidly over the years. During the first year of the sample Options Clearing Corp reports an average open interest of \$60 billion for standard options and LEAPS with maturities of less than six months that gradually decreases across maturities to \$200 million for options of two years or more. The corresponding figures in the last year of the sample are an open interest of \$1,400 billion for maturities of less than six months and of \$40 billion for maturities larger than two years.

Like Golez, but unlike Binsbergen et al. (2012), the main analysis relies on end-of-day option data. I use a dataset provided by Market Data Express containing S&P 500 index option data for CBOE traded European-style options and running from January 1990 to December 2013. 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 the nominal risk-free rate, which multiplies the strikes in the put-call parity relation; since strike prices are large numbers, any error in the interest rate will magnify 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 therefore drive a wedge between the reported quotes of the index value and the option prices and bias the synthetic prices.

To address these difficulties, I combine options with different moneyness levels to extract *both* the risk-free rate and the strip price in a unique step.⁸ This approach produces the appropriate interest rate for synthetic replication as well as it allows for spotting trade dates that violate the law of one price (LOOP) for some maturity and whose associated observations are excluded to mitigate microstructural noise. The appendix describes the algorithm for synthetic replication in detail.

Finally, to measure the real bond prices necessary to compute welfare costs I rely on zero-coupon TIPS yields with maturities of up to ten years from Gürkaynak et al. (2010). Since TIPS yields are either unavailable or unreliable during the 1990s, I also use Treasury yields as a proxy

⁸The literature offers at least three alternative ways to extract the term structure of equity. First, index options can be combined with some interest rate proxy as in the original intra-daily approach of Binsbergen et al. (2012). Their tick-level approach has the clear 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 and thereby maturities of less than two years. Finally, the index dividend futures studied by Binsbergen et al. (2013) have the clear advantage over index options and futures that they directly reveal strip prices without the need for synthetic replication and they do so for longer maturities. However, for S&P 500 dividend futures only proprietary datasets are available covering over-the-counter trades for a relatively short sample (for example, the dataset studied by Binsbergen et al. starts in October 2002). The exchange-traded nature of options mitigates concerns that the preferences embedded in their pricing may not reflect those of the average investor, which would complicate the macroeconomic interpretation of the derived cost of uncertainty.

available over the same sample period as I have dividend strips, in which case hold-to-maturity Treasury returns are corrected by the size of the inflation risk premium to account for the difference between nominal and real yields, as described in the appendix.⁹ Both nominal and real government bonds are computed on the last trading day of the month.

Figure 2a plots the prices of the synthetic dividend claims.

3.1. Average costs of consumption fluctuations

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 twelve months out, and so on. The measure of the hold-to-maturity return on the first semestral strip is the return on a six-month buy-and-hold strategy that pays off the next six months of dividends.¹⁰ Accordingly, I measure the hold-to-maturity return on the n -semester strip as the return for holding for $n - 1$ semesters a n -semester strip times the semestral return on the first semestral strip.

To address the potential concerns raised by Boguth et al. (2012) that microstructural frictions could 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 microstructure effects related to highly levered positions.

The evidence suggests 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.


Figure 2b illustrates the size of average annualized monthly log returns on six-month strategies over different subsamples by plotting the cumulated 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 (close to ten percent for claims to the next semester and year of dividends) and larger than the return on the index.¹¹ Figures 2c and 2d plot the analogous cumulated monthly returns on six-month bond strategies long a dollar on zero-coupon bonds with maturities between six months and ten years; average returns steadily increase in maturity across nominal as well as real government bonds, consistent with an upward-sloping average term structure of interest rates.

The annualized average hold-to-maturity return over the available dataset is of 13.4%, 13.1%, 12.3% and 10.9% for a strategy that goes long in the first to fourth semestral dividend strip,


⁹Low liquidity during the first years after the TIPS market opened in 1997 likely introduced a wedge between TIPS risk premia and real bond premia (see, e.g., Gürkaynak and Wright, 2012). In this regard, Grishchenko and Huang (2013) suggest using TIPS data only after 2000 and taking TIPS data during the recent crisis with caution due to possible mispricing. In any event, none of the evidence suggests that real yields are significantly larger in absolute magnitude than nominal yields (a point stressed by Backus et al., 2014), so the real bond proxy problem is unlikely to affect by much the quantitative estimates of the term structure of welfare costs because at the observable end of the curve the contribution of bond yields is much smaller than the contribution of equity yields.

¹⁰Golez (2014) raises concerns that equity prices of up to three-month maturity may be biased as a result of firms routinely preannouncing 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.


¹¹The economic significance of the large returns on short-term equities becomes even stronger once one recognizes that the initial and final years of the 1994-2013 period are years in which the index performed particularly well.




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




(b) Monthly cumulated returns on equities.




(c) Monthly cumulated returns on nominal bonds.



(d) Monthly cumulated returns on real bonds.



(e) Average term structures of equity, interest rates (mean-adjusted Treasuries) [l.h.s.] and welfare costs [r.h.s.] (with block bootstrapped 95% confidence intervals), 1994-2013.



(f) Average term structures of equity, interest rates (TIPS) [l.h.s.] and welfare costs [r.h.s.] (with block bootstrapped 95% confidence intervals), 2000-2013.

Figure 2: Term structures of equity, interest rates and welfare costs over the last two decades. The term structure of equity is synthesized from index options; the term structure of interest rates uses Gürkaynak-Sack-Wright data.

respectively, and of 2.9%, 3.0%, 3.1% and 3.3% for a strategy that goes long in the first to fourth semestral Treasuries, respectively. Over the 2000-2013 sample period, when TIPS data are available, the hold-to-maturity strategies pay off average returns of 17.0%, 14.0%, 11.5% and 10.5% for long positions in the short-term equities and 0.2%, 0.6%, 0.7% and 0.9% for long positions in the first to fourth semestral TIPS, respectively. These figures compare to an average six-month buy-and-hold equity premium over the two sample periods of 8.8% (1994-2013) and of 5.7% (2000-2013).

Table 1 reports the point estimates for the term structure of welfare costs. The first four term structure components at semestral frequency are of 12.6%, 11.9%, 10.8% and 9.1%, respectively. I then rely on proposition 2 to compute the costs of uncertainty around multi-period cashflows. Namely, I estimate average welfare costs of 0.56% associated with uncertainty one semester out, of 0.53% associated with up to one-year ahead uncertainty, of 0.50% with up to 18-month uncertainty, and of 0.46% with up to two-year ahead uncertainty, respectively. Restricting the attention to the 2000-2013 period, over which TIPS data are available to measure real interest rates, the estimated term structure components and welfare costs are slightly larger.

Additionally, I compute 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 the semestral periodicity of the strips. I find comparable results; the average cost of one-year ahead uncertainty over the two samples is of 0.4% (1994-2013) and 0.5% (2000-2013), whereas the cost of uncertainty over the next two years is of 0.3% and of 0.4%.

Figures 2e and 2f plot the point estimates for the term structures of equity, interest rates, and welfare costs. The term structures of equity and interest rates report available semestral zero-coupon equities and real bonds and the corresponding per-period hold-to-maturity returns for different maturities. The dotted lines represent the counterfactual case of flat term structures of equity and interest rates, expressed in excess over the first bond yield. The figures also show block-bootstrapped one-sided critical values based on bootstrap- t percentiles corresponding to a five-percent size for the means of $l_t^{(n)}$ and the six-month market return; the block size of ten 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 cost of consumption fluctuations

The components of the term structure of welfare costs are volatile and countercyclical. Present-value logic with time-varying expected returns and expected dividend growth imply that dividend yields contain information about both state variables (Golez, 2014). Since the same states would drive the risk premia that constitute the term structure of welfare costs one could use the semestral equity yields to 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

¹²The descriptive model by Lettau and Wachter (2011) implies a one-factor structure for the term structure of welfare costs that can be revealed by the information set spanned by two equity yields and a bond yield. The external-habit model by Campbell and Cochrane (1999) implies an approximate one-factor structure, while the long-run risk model by Bansal and Yaron (2004) and the external-habit model by Lopez et al. (2015b) imply an approximate two-factor structure.

1994-2013					2000-2013			
	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 ($n = 4$)
$l_t^{(n)}$								
Mean	0.1259 (0.0325) [0.0689]	0.1192 (0.0160) [0.0303]	0.1081 (0.0120) [0.0204]	0.0913 (0.0097) [0.0180]	0.1362 (0.0408) [0.0906]	0.1288 (0.0196) [0.0412]	0.1077 (0.0134) [0.0255]	0.1002 (0.0107) [0.0243]
$L_t^{1,...,n}$								
Mean	0.0056 (0.0015) [0.0031]	0.0053 (0.0008) [0.0018]	0.0050 (0.0006) [0.0012]	0.0046 (0.0005) [0.0009]	0.0061 (0.0018) [0.0040]	0.0058 (0.0011) [0.0024]	0.0053 (0.0008) [0.0017]	0.0050 (0.0006) [0.0013]
$L_t^{(1)}$, 1 period = n semesters								
Mean	0.0056 (0.0015) [0.0031]	0.0040 (0.0008) [0.0020]	0.0035 (0.0006) [0.0014]	0.0032 (0.0004) [0.0011]	0.0061 (0.0018) [0.0040]	0.0050 (0.0010) [0.0026]	0.0046 (0.0007) [0.0017]	0.0043 (0.0005) [0.0014]
Equity premium, semestral								
Mean		0.0885 (0.0169) [0.0398]				0.0575 (0.0207) [0.0471]		

Table 1: Options implied average term structure of the welfare cost of uncertainty. $l^{(n)}$ is the cost of a marginal increase in uncertainty in n -semester ahead cashflows. $L^{1,...,n}$ is the cost of a marginal increase in uncertainty in 1 to n -semester ahead cashflows. The third panel reports the cost of a marginal increase in one-period ahead uncertainty, $L^{(1)}$, for different period lengths. The short sample (2000-2013) uses TIPS yield data to measure the term structure of real interest rates. The full sample (1994-2013) uses Treasury yield data, whose means are adjusted by the average inflation risk premium over the 2000-2013 period, as a proxy for the term structure of real interest rates. The equity premium is the average six-month buy-and-hold return on the S&P 500 index in excess over the six-month risk-free rate. Bootstrapped standard errors use block sizes of one (in parentheses) and ten (in brackets) observations.




Figure 3: Semestral term structure of welfare costs and semestral equity premium, conditional on the information set spanned by equity and bond yields; annualized premia. Regressors are the first two principal components of semestral equity yields and the first principal component of up to ten-year bond yields. The semestral excess return on the index is additionally regressed on the market dividend yield. The shaded areas indicate business-cycle peaks and troughs as declared by the NBER (March-November 2001 and December 2007 to June 2009).

I extract the first two principal components of the semestral equity yields, which capture 95% 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 term structure of welfare costs.

Table 2 presents the predictive regressions and shows a standard deviation of expected returns about as large as the already large level. Note how the forecasting regressions may miss some important return predictors, so the estimates in table 2 represent a lower bound on the actual volatility of the cost of uncertainty. Since excess returns are forecastable, the cost of uncertainty varies over time and considerably so; the cost of short-run cashflow uncertainty is substantial at some junctures of the business cycle.

Figure 3 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% during the dot-com crash and the period immediately preceding the early 2000s recession as well as during the most recent recession as declared by the National Bureau of Economic Research. 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.

The evidence is consistent with people being highly sensitive to cashflow stability, in particular to short-run stability and in bad states of the economy such as during downturns.

	$r_{t \rightarrow t+1}^{e,(1)}$		$\frac{1}{2}r_{t \rightarrow t+2}^{e,(2)}$		$\frac{1}{3}r_{t \rightarrow t+3}^{e,(3)}$		$\frac{1}{4}r_{t \rightarrow t+4}^{e,(4)}$		$r_{t+1}^{e,m}$	
const.	0.083 [0.031]	0.083 [0.022]	0.094 [0.015]	0.094 [0.015]	0.084 [0.014]	0.084 [0.013]	0.069 [0.012]	0.069 [0.012]	0.072 [0.034]	0.072 [0.033]
$pc_{d,t}^1$	0.369 [0.100]	0.366 [0.056]	0.150 [0.028]	0.150 [0.032]	0.077 [0.024]	0.079 [0.023]	0.064 [0.019]	0.066 [0.017]		-0.094 [0.056]
$pc_{b,t}^1$		0.095 [0.145]		-0.017 [0.088]		-0.067 [0.072]		-0.111 [0.057]		0.104 [0.118]
$pc_{d,t}^2$		0.703 [0.144]		-0.063 [0.065]		-0.220 [0.043]		-0.152 [0.037]		
dp_t^m									0.345 [0.152]	0.318 [0.194]
R^2	0.353	0.502	0.254	0.259	0.106	0.229	0.105	0.204	0.098	0.147
$\frac{\sigma(E_t, r)}{E(r)}$	2.68	3.19	1.00	1.01	0.57	0.85	0.57	0.79	1.10	1.35
$E(r)$	0.0832		0.0941		0.0837		0.0694		0.0719	

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_d^1 and pc_d^2 , the first principal component of up to ten-year bond yields, pc_b^1 , and the market dividend yield, dp^m . Monthly data, 1996m1-2013m12. Newey-West standard errors to correct for overlapping.

4. Welfare implications in a simple model

This section investigates the optimal stabilization policy in a model that combines a textbook New Keynesian model economy (Galí, 2008) with Campbell and Cochrane (1999) habits in consumption originally proposed by Lopez, Lopez-Salido and Vazquez-Grande (2015b). This model captures simultaneously a negative slope in the term structure of the equity premium and a positive slope in the term structure of interest rates, and is therefore a natural candidate to explain a downward-sloping term structure of welfare costs.¹³ Coincidentally, the presence of two externalities (nominal rigidities and external habits) breaks the ‘divine coincidence’ of Blanchard and Galí (2007) and motivates nontrivial policy interventions.

4.1. Model

For simplicity I calibrate the model by Lopez et al. (2015b) to full macro-finance separation, so I abstract from capital formation, and focus on the optimal parametrization of a simple Taylor rule. This choice illustrates the point in the simplest possible setup by offering a particularly transparent result. The model of quantities and inflation under the Campbell-Cochrane pricing kernel is the same as under standard CRRA utility, so the two models differ only in their asset pricing implications. The respective welfare implications turn out to be dramatically different, as under CRRA utility the optimal simple Taylor rule closes the consumption gap, while under Campbell-Cochrane habits the optimal rule removes consumption volatility.¹⁴

4.1.1. Firms

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

$$E_0 \sum_{t=0}^{\infty} M_{t,t+h} \left(\frac{P_t(i)}{P_t} Y_t(i) - (1 - \tau) W_t N_t(i) - T_t \right)$$

subject to nominal price stickiness à la Calvo (1983). They 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. Corporate profits are paid out as dividends on market equity each period to households, who own the firms.

¹³The 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 its components 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).

¹⁴The point of this theoretical analysis is to show how the policy prescriptions of the textbook New Keynesian model would change once one incorporates a pricing kernel that produces the observed welfare costs. Like the standard New Keynesian literature, I therefore choose to focus on the unique locally-bounded solution under a Taylor rule despite the recent critiques by Cochrane (2011a) of equilibrium determination with Taylor rules.

The i th good sells for the nominal price $P_t(i)$, with $P_t \equiv [\int_0^1 P_t(i)^{1-\varepsilon} di]^{1/(1-\varepsilon)}$ the price index. 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 i th good, $Y_t(i) = [P_t(i)/P_t]^{-\varepsilon} Y_t$, which constrains individual firms in their production choices. The government levies lump-sum taxes T_t on each firm to finance an employment subsidy, $\tau = 1 - [Z/S](\varepsilon - 1)/\varepsilon \in (0, 1)$, designed to offset any steady-state distortions caused by the monopolistic competition and the habit externality.

The market for goods and labor clear, $Y_t = C_t$ and $N_t = \int_0^1 N_t(i) di$, where real consumption C_t denotes aggregate demand for market produced goods and N_t is aggregate labor effort.

Log technology¹⁵ $a_t \equiv \ln A_t$ is composed of a permanent component a_t^P and a transitory component a_t^T such that $a_t = a_t^P + a_t^T$, with

$$\begin{aligned} a_{t+1}^P &= \mu + a_t^P + (1 - \theta)e_{t+1} \\ a_{t+1}^T &= \rho a_t^T + \theta e_{t+1} \end{aligned}$$

with average drift μ and persistence $|\rho| < 1$, where θ indexes the extent to which a technology shock $e_t \sim Niid(0, \sigma^2)$ has a permanent effect; for example, $\theta = 0$ associates with random-walk technology, while $\theta = 1$ associates with the typical trend-stationary specification in macro (e.g., Galí, 2008). When discussing the optimal Taylor type rule, θ will also index the amount of uncertainty in the consumption process that can be removed via monetary policy, consistent with definition (1).

4.1.2. Households

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

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

subject to the sequence of budget constraints and the terminality condition

$$P_t C_t(j) + \exp(-i_t) B_t(j) \leq W_t N_t(j) + B_{t-1}(j) + P_t D_t, \quad E_t \left(M_{t,t+j} \frac{B_{t+j-1}}{P_{t+j}} \right) \xrightarrow{j \rightarrow \infty} 0$$

with real contingent claims prices $M_{0,t}$.¹⁶ As in Greenwood and Hercowitz (1991), households derive utility from consumption of two goods; C_t is real consumption purchased in the market and H_t denotes the consumption produced at home, with production function $H_t = A_t(1 - N_t)$, with N_t the labor choice. X_t^c and X_t^h represent habit levels that are a nonlinear function of contemporaneous and past consumption. B_t denotes their holdings of one-period nominal debt issued by a fiscally-passive government and with unit price $\exp(-i_t)$, and D_t is the dividend they receive from owning the

¹⁵Lower case letters denote logs and ‘hat’ variables denote deviations from the non-stochastic mean.

¹⁶The budget constraint could be written to include explicitly any asset in zero net supply, which can therefore be priced by no-arbitrage.

aggregate firm. Parameter β is the subjective discount rate and parameter χ controls the steady-state effect of habits. The restriction on the curvature of the utility function in market and home consumption ensures balanced growth (Campbell and Ludvigson, 2001). I assume a calibration for χ to achieve a steady-state level of hours $N = .5$.

The law of motion of habits is specified indirectly through the processes for surplus market consumption $S_t \equiv (C_t - X_t^c)/C_t$ and surplus home consumption $Z_t \equiv (H_t - X_t^h)/H_t$. The habit levels thus specified are external to any agent's consumption decisions, as the law of motion of the surplus levels is driven by *aggregate* market and home consumption, $C_t \equiv \int_0^1 C_t(j)dj$ and $H_t \equiv \int_0^1 H_t(j)dj$. To ensure well-behaved marginal utilities, consider the following dynamics for the logarithms of aggregate surplus levels:

$$\begin{aligned}\hat{s}_{t+1} &= \phi \hat{s}_t + \Lambda_c[\hat{s}_t](E_{t+1} - E_t)f_c[C_{t+1}] \\ \hat{z}_{t+1} &= \phi \hat{z}_t + \Lambda_h[\hat{z}_t](E_{t+1} - E_t)f_h[H_{t+1}]\end{aligned}$$

with common persistence $|\phi| < 1$, where $f_c[C_t] = \ln[C_t]$ and $f_h[H_t] = \ln[g_t(H_t)]$, where the market and home production functions and market clearing impose the static restriction $H_t = g_t^{-1}(C_t)$.

The particular calibration for the sensitivity functions (Λ_c and Λ_h) and the steady state levels of the surplus variables

$$\begin{aligned}\Lambda_c[\hat{s}_t] &= \begin{cases} S^{-1} \sqrt{1 - 2\hat{s}_t} - 1, & \hat{s}_t \leq \frac{1}{2}(1 - S^2) \\ 0 & \hat{s}_t > \frac{1}{2}(1 - S^2) \end{cases} & S = \sqrt{\frac{\gamma \text{var}(\varepsilon^c)}{1 - \phi - \xi_1/\gamma}} \\ \Lambda_h[\hat{z}_t] &= (1 + \xi_2)\Lambda_c[\hat{z}_t/(1 + \xi_2)] & Z = \left(1 + \frac{1 - S}{S} \frac{\text{var}(\varepsilon^c)}{\text{cov}(\varepsilon^c, \varepsilon^h)}\right)^{-1}\end{aligned}$$

with $\varepsilon_t^c \equiv (E_t - E_{t-1}) \ln C_t$ and $\varepsilon_t^h \equiv (E_t - E_{t-1}) \ln H_t$, can produce an exact separation between risk premia and quantity dynamics that preserves the ability of the model to fit macroeconomic variables through an appropriate choice of the free parameter $\xi = [\xi_1; \xi_2] \in \mathbb{R}^2$ that controls the spillover of habits dynamics onto the equilibrium quantities. Accordingly, I assume the parametrization, $\xi_2 \propto \xi_1 = \frac{\gamma}{5} \text{cov}(\varepsilon^c, \varepsilon^h)$, described by Lopez et al. (2015b) that ensures such a macro-finance separation in this context. On the one hand, households balance strong intertemporal substitution motives with strong precautionary saving motives, so their marginal utilities vary strongly across states but much less across time, which is key to avoid excessive risk-free rate variation. On the other hand, households balance a strong aversion to fluctuations in market and home consumption across states, which is key to avoid excessive variation in the real wage rate.

Optimality implies the log stochastic real discount factor

$$m_{0,t} = -t \ln(\beta) - \gamma(c_t - c_0) - \gamma(s_t - s_0)$$

the equilibrium consumption-saving equation,

$$i_t = -\ln E_t \beta e^{-\gamma \Delta C_{t+1} - \gamma \Delta S_{t+1} - \pi_{t+1}}$$

and the labor supply equation,

$$W_t = \chi \frac{A_t^{1-\gamma} C_t^\gamma}{(1 - N_t)^\gamma} \left(\frac{S_t}{Z_t} \right)^\gamma$$

Note how the choice of the sensitivity function for surplus home consumption implies

$$\frac{Z_t}{Z} = \left(\frac{S_t}{S} \right)^{1+\xi_2}$$

in equilibrium; this property is crucial to prevent households from varying labor hours to a counterfactually large extent to absorb the shocks (Lettau and Uhlig, 2000; Lopez et al., 2015b).

Finally, the two consumption habits indirectly specified by the surplus processes are a complex nonlinear function of current and past consumption; however, they are approximately two predetermined, slow-moving averages of past consumption,

$$\begin{aligned} x_{t+1}^c &= \ln(1 - S) + c_{t+1} - \sum_{j=0}^{\infty} \phi^j (E_{t-j+1} - E_{t-j}) c_{t-j+1} + O(\|e\|^2) \\ x_{t+1}^h &= \ln(1 - Z) + h_{t+1} - \sum_{j=0}^{\infty} \phi^j (E_{t-j+1} - E_{t-j}) h_{t-j+1} + O(\|e\|^2) \end{aligned}$$

that can be interpreted as a trend in market and home consumption, respectively. Households grow used slowly to unanticipated movements in the two types of consumption. The surplus consumption levels are therefore basically detrended versions of market and home consumption.

4.1.3. Monetary policy

Monetary policy is described by a simple Taylor rule for the nominal interest rate that reacts to inflation $\pi_t \equiv \ln(P_t/P_{t-1})$, and the output gap,

$$i_t = r + \phi_\pi \pi_t + \phi_y (c_t - a_t^P) \quad (6)$$

with $r = -\ln(\beta) + \gamma\mu - .5(1 - \phi - \xi_1/\gamma)$. Monetary policy reacts to economic activity using a realistic description of the output gap as the distance of the level of output from its Beveridge-Nelson trend, $c_t^P = a_t^P$, rather than from the natural (flexible-price) consumption level, $c_t^n = a_t$.

4.1.4. Equilibrium quantities

We solve for the equilibrium allocation of our model using the risk-adjusted loglinearization described by Lopez, Lopez-Salido and Vazquez-Grande (2015a), which correctly captures the first-order component of the equilibrium allocation by accounting for the presence of such components in conditional second moments, which are crucial to correctly characterize the consumption-saving tradeoff in a world with Campbell and Cochrane (1999) habits. Relative to a standard New Keynesian model, the dynamic IS equation and the New Keynesian Phillips curve acquire an additional term that captures time-varying risk.

Accordingly, the risk-adjusted loglinearization of the set of optimality conditions around the zero-inflation deterministic steady state results in the risk-adjusted equation describing the optimal consumption-savings tradeoff,

$$\gamma E_t \Delta(c_{t+1} - a_{t+1}) = i_t - E_t \pi_{t+1} + \xi_1 \hat{s}_t + \frac{1}{2} \text{var}_t(\pi_{t+1}) + x_t \text{cov}_t(c_{t+1}, \pi_{t+1}) \quad (7)$$

where $x_t \equiv \gamma(1 + \Lambda_c[s_t])$ is the price of risk, and in the risk-adjusted New Keynesian Phillips curve

$$\begin{aligned} \pi_t = & \delta E_t \pi_{t+1} + \kappa(c_t - a_t) - \gamma \lambda \xi_2 \hat{s}_t \\ & - \frac{1}{2} \frac{\delta(1 - \alpha + \alpha \varepsilon)}{(1 - \eta)(1 - \alpha)} \text{var}_t(\pi_{t+1}) + \delta \text{cov}_t\left(\pi_{t+1}, \frac{\varepsilon}{1 - \alpha} \pi_{t+1} + (1 - \gamma)e_{t+1} + \ell_{t+1}\right) \end{aligned} \quad (8)$$

up to an irrelevant constant, where $\kappa \equiv \lambda[\gamma(2 - \alpha) + \alpha]/(1 - \alpha)$, with $\lambda \equiv (1 - \eta)(1 - \delta\eta)(1 - \alpha)/\eta(1 - \alpha + \alpha \varepsilon)$, controls the slope of the curve, with $\delta \equiv \beta \exp[(1 - \gamma)\mu]$. Inflation is high when long-run production is expected to be above the flexible-price level, in which case resetting firms adjust prices upwards to realign their markups to the desired level. To this pair of equilibrium conditions adds the equation describing the auxiliary jump variable

$$\begin{aligned} \ell_t = & \delta \eta \left(E_t \ell_{t+1} + \frac{\varepsilon}{1 - \alpha} E_t \pi_{t+1} - (1 - \rho)(1 - \gamma) a_t^T \right) + (1 - \delta \eta) \left(\frac{1 + \gamma}{1 - \alpha} (c_t - a_t) - \gamma(1 + \xi_2) s_t \right) \\ & + \frac{\delta \eta}{2} \text{var}_t \left(\frac{\varepsilon}{1 - \alpha} \pi_{t+1} + (1 - \gamma) e_{t+1} + \ell_{t+1} \right) \end{aligned}$$

Under the determinacy condition $\kappa(\phi_\pi - 1) + (1 - \delta)\phi_y > 0$ with $[\phi_\pi, \phi_y] \in \mathbb{R}_+^2$, it is straightforward to verify how the solution to the set of equations (6) to (8) has the linear form $c_t - a_t = \psi_c a_t^T$, $\pi_t = \psi_\pi a_t^T$, with

$$\begin{aligned} \psi_c = & - \frac{(1 - \delta\rho)[\gamma(1 - \rho) + \phi_y]}{(1 - \delta\rho)[\gamma(1 - \rho) + \phi_y] + \kappa(\phi_\pi - \rho)} \\ \psi_\pi = & - \frac{\kappa[\gamma(1 - \rho) + \phi_y]}{(1 - \delta\rho)[\gamma(1 - \rho) + \phi_y] + \kappa(\phi_\pi - \rho)} \end{aligned}$$

The equilibrium allocation is observationally equivalent to a model with CRRA consumers.

4.1.5. Equilibrium asset prices

The stochastic discount factor $m_{0,t} = \sum_{\tau=0}^{t-1} m_{\tau+1}$, with

$$m_{\tau+1} = -\ln(\beta) - \gamma E_t \Delta c_{\tau+1} + \gamma(1 - \rho_s) \hat{s}_\tau - x_\tau (E_{\tau+1} - E_\tau) c_{\tau+1}$$

is a key ingredient to solve for equilibrium asset prices. To derive the equilibrium term structures I rely on the essentially-affine approximation of Lopez et al. (2015a), as the Campbell-Cochrane price of risk inherits the nonlinear dynamics of surplus consumption and prevents an exact closed-form solution for equilibrium term structures. The online appendix describes the solution method.

Table 3 lists the calibration of the deep parameters of the model. The production side of the

	Parameter		Value
Preference block	β	Subjective discount factor	0.995
	γ	Inverse EIS	2
	ϕ	Habit persistence	0.993
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	0.003
	ρ	Persistence of the conditional mean of technology growth	0.847
	σ	Conditional volatility of technology	0.018
	$1 - \theta$	Fraction of permanent technology shocks	0

β matches an average real interest rate of 0.94% per year.

ϕ matches an average market equity premium of 8.85% per year.

μ matches an average annual consumption growth of 3.60%.

ρ matches an average AR root in an ARMA(1,1) representation of annual consumption growth of 0.136.


σ matches a volatility of annual consumption growth of 2.49%.

Table 3: Deep parameters and their calibration (monthly frequency). Data for real consumption growth use annual BEA-NIPA data over the period 1929-2013 for personal consumption expenditure in nondurables and services, and are deflated by the core PCE price index. Monthly simulated data are aggregated to an annual frequency and are matched to the corresponding data moments.


economy is calibrated using standard values in the New Keynesian literature taken from Galí (2008), while parameters relating with the pricing kernel are calibrated as in Campbell and Cochrane (1999). In particular, the habit persistence coefficient, ϕ , is chosen to match the average equity premium of almost 9% reported in figure 1a, which is nonetheless substantially lower than the premium commanded by short-term equities, which can reach up to 15% on average, consistent with observed strip returns. As in Galí (2008), I choose an $AR(1)$ specification for technology ($\theta = 1$).¹⁷ Finally, to calibrate the parameters that drive technology, I choose a parametrization that matches the persistence and volatility of consumption growth using annual data on personal consumption expenditure in nondurable goods and services over the period 1929-2013 from the Bureau of Economic Analysis, which is expressed in real terms through the core PCE price index.

Figure 4 reports 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. 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


¹⁷This choice, and in particular the one-shock structure, implies a counterfactual perfect correlation between consumption and dividends. Lopez et al. (2015b) relax the one-shock structure and match the empirical correlation between consumption and dividends; however, my present purpose is to describe in the simplest setting possible a model with textbook quantity implications but with empirically appealing welfare implications.




(a) Average term structures of equity and interest rates.



(b) Average term structure of welfare costs.



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



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

Figure 4: Model-implied term structures of equity (dashed line), interest rates (dash-dotted line) and welfare costs (solid line); average and interquartile range. Annualized premia on the left axis for the term structures of equity and interest rates and on the right axis for welfare costs. $E(r^{e,m})$ is the equity premium.

countercyclical variation.

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. In fact, the dividend share gradually increases after a technology shock as the price level increases with a lag, so the payoffs of short-duration dividend strips are more procyclical than the payoff of long-duration claims.

4.2. Welfare objective

The appendix derives a quadratic approximation of welfare that is appropriate to study a linear-quadratic approximate optimal policy problem (Benigno and Woodford, 2012). Accordingly, I

consider the average welfare loss criterion, $\mathbb{L} \equiv -(1 - \delta)E(U_t)/(CS)^{1-\gamma}$,

$$\mathbb{L} = \frac{1}{2} \frac{\varepsilon}{\kappa} \text{var}(\pi_t) + \frac{1}{2} \text{var}(c_t - c_t^n) + \frac{1}{2} \frac{\nu \Lambda^2}{1 - \phi^2} \text{var}(\varepsilon_t^c) + \frac{\nu \Lambda}{1 - \delta \phi} \left(1 - \frac{\theta \phi (1 - \rho)}{1 - \phi \rho} \right) \text{cov}(\varepsilon_t^c, e_t) \quad (9)$$

with $\nu \equiv (\gamma - 1)(2 - \alpha)(1 - \alpha)/[\gamma(2 - \alpha) + \alpha]$, where $\varepsilon_t^c \equiv (E_t - E_{t-1})c_t$ denotes consumption innovations.

Welfare criterion (9) compares with the corresponding criterion in the standard CRRA utility case; the volatility of consumption shocks and their positive covariance with technology represent additional sources of welfare losses from a macroeconomic perspective.¹⁸ There are two separate sources of strategic complementarities that generate departures from the efficient dynamics of the benchmark RBC model. On the one hand, in a sticky-price environment inflation volatility is approximately equivalent to cross-sectional dispersion in prices, which in turn associates with inefficient employment.

On the other hand, 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 bad news about current and future states, people should not cut their consumption level as much as they would want to because they will grow used to the lower consumption level, so their marginal utilities will be relatively lower tomorrow.

4.3. Optimal stabilization policy

To maximize the intertemporal welfare criterion (9), a central bank would want to stabilize the consumption gap, $c_t - a_t$, and hence stabilize inflation in the process, as well as consumption uncertainty, $(E_t - E_{t-1})c_t$, to stabilize the price of risk, and hence risk premia. Since achieving both goals is unfeasible, the optimal monetary policy trades them off to minimize the welfare metric.

A characterization of the optimal policy regime under discretion or commitment is beyond the scope of this paper. The purpose of the present analysis is to point out how a model with a pricing kernel sufficiently volatile to be consistent with the observed term structure of the welfare cost of uncertainty can easily imply that people hate consumption fluctuations so much as to overturn the standard result that removing inflation volatility is the macroeconomic priority. This result holds despite the fact that the model of quantities under the two different pricing kernels is observationally equivalent. I therefore focus on a well-understood interest-rule, a Taylor rule in inflation and in the business cycle component of output, and show how the optimal policy under the two pricing kernels are the polar opposite of each other—one fills the troughs, the other removes volatility.

4.3.1. Inflation targeting (close the consumption gap)

Under an inflation targeting regime consumption equals the flexible-price level, and hence inflation is zero. It follows that the welfare loss function that associates with this monetary policy

¹⁸Note that when people have log utility in consumption they do not care about fluctuations in log consumption and the flexible price equilibrium coincides with the Pareto optimum to a first-order approximation. However, an elasticity of intertemporal substitution (EIS) smaller than unity has more empirical support, and a small deviation from a unit EIS can have dramatic welfare implications by the large volatility of surplus consumption in the external-habit framework.

regime can be written as:

$$\mathbb{L}^\pi = \frac{\nu}{2} \left[\frac{\Lambda}{1 - \phi^2} + \frac{2}{1 - \delta\phi} \left(1 - \frac{\theta\phi(1 - \rho)}{1 - \phi\rho} \right) \right] \Lambda\sigma^2$$

4.3.2. Risk premia targeting (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 be cointegrated. In any event, therefore, consumption has to inherit the permanent component of technology, which represents the maximum amount of consumption risk that a simple monetary policy regime can remove.

In the context of the model, removing consumption risk is equivalent to stabilizing surplus consumption, and hence the price of risk by the relation $\hat{s}_t = -\frac{1}{2}(x_t^2 - x^2)/x^2$. Empirical measures of the Hansen and Jagannathan (1991) bound, which in the model takes the simple expression

$$\left| \frac{\ln E_t R_{t+1}^e}{\sqrt{\text{var}_t(r_{t+1}^e)}} \right| \leq x_t \sqrt{\text{var}(\varepsilon_t^c)}$$

for all available excess returns, would in turn reveal the price of risk. This policy is therefore equivalent to a regime that targets risk premia by stabilizing the maximal risk-return tradeoff.

Under such a risk premia targeting regime consumption equals the permanent component of technology, and hence $\pi_t = -[\kappa/(1 - \delta\rho)]a_t^T$ and $\varepsilon_t^c = (1 - \theta)e_t$. The associate welfare loss function can then be written as:

$$\mathbb{L}^x = \frac{1}{2} \left(1 + \frac{\kappa\varepsilon}{(1 - \delta\rho)^2} \right) \frac{\theta^2\sigma^2}{1 - \rho^2} + \frac{\nu}{2} \left[\frac{\Lambda(1 - \theta)}{1 - \phi^2} + \frac{2}{1 - \delta\phi} \left(1 - \frac{\theta\phi(1 - \rho)}{1 - \phi\rho} \right) \right] \Lambda(1 - \theta)\sigma^2$$

4.3.3. Stabilizing gaps vs. levels

Under the baseline calibration risk premia targeting dominates inflation targeting, as $\mathbb{L}^\pi \geq \mathbb{L}^x$ for any value of $\theta \in (0, 1)$. People hate so much consumption volatility that a policy that achieves the flexible-price equilibrium is suboptimal relative to a policy that removes as much consumption volatility as possible. Moreover, the inequality is strict for any $\theta \in (0, 1]$, with equality if and only if $\theta = 0$, as random-walk technology negates any role to nominal rigidities; 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 premia targeting over inflation targeting is deeply linked to the model's ability to explain the empirical term structure properties.

Table 4 exemplifies this result numerically under the baseline calibration. The welfare dominance of the risk premia targeting regime is overwhelming in the trend-stationary case. Average welfare losses relative to the Pareto optimum are up to 25 times larger under inflation targeting than under risk premia targeting. The welfare weight attached to the inflation objective dominates the welfare weights attached to consumption and risk aversion stabilization, however, the volatility of

θ	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
$\mathbb{L}^x/\mathbb{L}^\pi$	1	0.84	0.69	0.55	0.42	0.31	0.21	0.14	0.08	0.05	0.04

Table 4: Per period welfare losses under a simple policy regime that targets inflation (or the consumption gap), $\mathbb{L}^\pi$, and a simple policy regime that targets risk premia (or consumption uncertainty), $\mathbb{L}^x$. Only ratios are reported as the absolute values of per-period losses are meaningless. θ denotes the fraction of uncertainty in the consumption process that can be removed by monetary policy.

surplus consumption is so large that removing its fluctuations becomes the priority.

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 the simple Taylor rule (6). The central bank is able to implement the inflation targeting regime by choosing an extreme anti-inflationary response,

$$\phi_\pi \rightarrow \infty$$

while the risk premia targeting regime can be implemented by choosing an extreme response to movements in detrended consumption,

$$\phi_y \rightarrow \infty$$

In this sense, therefore, the two policies are polar opposite.

5. Conclusion

Lucas (1987) introduced the notion of 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. A main result that justifies looking at this theoretical object is that recent derivative securities provide direct information about the term structure and, therefore, new insight into an old question (the tradeoff between growth and stability) and evidence to study the new question (the tradeoff between stability at different time horizons).

Asset markets suggest that 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 inscribes into a burgeoning literature that finds high and time-varying short-maturity risk premia as a pervasive phenomenon across different asset classes (Binsbergen and Koijen, 2015).

The result that the marginal cost of uncertainty is a linear combination of risk premia makes one of the main tasks of macroeconomics—that of assessing the 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 consumption-based asset pricing theories and therefore represents a puzzling piece of evidence with seemingly crucial welfare consequences. For example, I showed how the optimal monetary policy intervention in two models that are observationally equivalent in their quantity implications but differ in their asset pricing implications prescribe polar opposite policies. Accordingly, dynamic equilibrium models

able to rank different stabilization policies while rationalizing empirical measures of the welfare cost of uncertainty should be high on the macro-finance research agenda.

Appendix

A. Theoretical results

A.1. Proof of proposition 1

Differentiating (1) with respect to θ and evaluating at $\theta = 0$, it follows that

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

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

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

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

A.2. Proof of proposition 2

I can rewrite equation (3) as

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

The absence of arbitrage opportunities ensures positive weights $\{\omega_{n,t}^N\}$.

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, $R_{x,t \rightarrow t+n}^{(n)} = X_{t+n}/D_{x,t}^{(n)}$, with no-arbitrage price $D_{x,t}^{(n)} = E_t(M_{t,t+n}X_{t+n})$ at period t . Market equity is characterized by $X = D$ and real bonds by $X = 1$.

A.3. Extension to heterogeneous consumers

The term structure of marginal costs of uncertainty remains well-defined even in a heterogeneous-agent, incomplete-market setting. Consider agents with heterogeneous preferences, U_i , and an idiosyncratic consumption component, $\varepsilon_{i,t}$, driving their consumption stream, $C_{i,t} = D_t + W_t N_t + e_t + \varepsilon_{i,t}$. Agents can only trade the entire term structures of dividend strips and zero-coupon bonds and may therefore face uninsurable idiosyncratic risk.

In a heterogeneous-agent, incomplete-market context the cost of fluctuations can be defined as

$$E_t U_i \left(\left\{ (1 + \mathcal{L}_{i,t}^N(\theta))^n C_{t+n} + \varepsilon_{i,t+n} \right\} \right) = E_t U_i \left(\left\{ \bar{C}_{t+n}(\theta) + \varepsilon_{i,t+n} \right\}, \left\{ \bar{N}_{t+n}(\theta) \right\} \right)$$

so $\mathcal{L}_i$ measures the cost of uncertainty around the systematic part of the i th agent's consumption (Alvarez and Jermann, 2004). By the absence of arbitrage opportunities, the projection of the marginal rates of substitution on the payoff space is the same across people, so their valuations of available assets are equal. Since all agents have access to the entire term structures of strips and zero-coupon bonds, they end up equalizing their valuation of welfare costs. Therefore, the marginal cost of uncertainty around all coordinates $n \in \mathcal{N}$ is constant across agents $i \in (0, 1)$,

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

A.4. Derivation of the welfare criterion

In the macro-financially separate model with external habits and nominal rigidities, welfare of the representative household is

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

where $\mathcal{C}_t \equiv C_t - X_t^c$ and $\mathcal{H}_t \equiv H_t - X_t^h$. Consider the stationary transformation of utility,

$$\begin{aligned} \frac{U_0}{A_0^{1-\gamma}} &= E_0 \sum_{t=0}^{\infty} \delta^t \exp[(1-\gamma)(a_t - a_0 - t\mu)] V(\widetilde{\mathcal{C}}_t, \widetilde{\mathcal{H}}_t) \\ \text{with } V_t &\equiv \frac{1}{1-\gamma} \left(\frac{\mathcal{C}_t}{A_t} \right)^{1-\gamma} + \frac{\chi}{1-\gamma} \left(\frac{\mathcal{H}_t}{A_t} \right)^{1-\gamma} - \frac{1+\chi}{1-\gamma} A_t^{\gamma-1} \end{aligned}$$

with the detrended surplus consumption levels $\widetilde{\mathcal{C}}_t \equiv \mathcal{C}_t/A_t$ and $\widetilde{\mathcal{H}}_t \equiv \mathcal{H}_t/A_t$.

Expanding around the zero-inflation steady state, the detrended per-period utility function can be written as

$$V_t = \widetilde{\mathcal{C}}_t^{1-\gamma} \left(\ln(\widetilde{\mathcal{C}}_t) + \frac{1-\gamma}{2} \ln(\widetilde{\mathcal{C}}_t)^2 \right) + \chi \widetilde{\mathcal{H}}_t^{1-\gamma} \left(\ln(\widetilde{\mathcal{H}}_t) + \frac{1-\gamma}{2} \ln(\widetilde{\mathcal{H}}_t)^2 \right) + t.i.p. + O(\|\widetilde{\mathcal{C}}_t, \widetilde{\mathcal{H}}_t\|^3)$$

up to a term independent of policy. I can rewrite the ratio of partial derivatives $\chi \widetilde{\mathcal{H}}_t^{1-\gamma} / \widetilde{\mathcal{C}}_t^{1-\gamma} = 1 - \alpha$, given the efficient employment subsidy.

Using the second-order expansion,

$$\ln(\widetilde{\mathcal{H}}_t) + \frac{1}{2} \ln(\widetilde{\mathcal{H}}_t)^2 = \hat{z}_t - \hat{n}_t + \frac{1}{2} \hat{z}_t^2 - \frac{1}{2} \hat{n}_t^2 - \hat{n}_t \hat{z}_t + O(\|\widetilde{\mathcal{H}}_t\|^3)$$

it follows that

$$\begin{aligned}\frac{V_t}{\bar{\mathcal{C}}^{1-\gamma}} &= \ln(\bar{\mathcal{C}}_t) + \frac{1-\gamma}{2} \ln(\bar{\mathcal{C}}_t)^2 + \chi \frac{\bar{\mathcal{H}}^{1-\gamma}}{\bar{\mathcal{C}}^{1-\gamma}} \left(\ln(\bar{\mathcal{H}}_t) + \frac{1-\gamma}{2} \ln(\bar{\mathcal{H}}_t)^2 \right) + t.i.p. + O(\|\bar{\mathcal{C}}_t, \bar{\mathcal{H}}_t\|^3) \\ &= (2-\alpha)\hat{s}_t - (1-\alpha)\Delta_t - \frac{1}{2} \frac{\gamma(2-\alpha) + \alpha\bar{c}_t^2}{1-\alpha} - \frac{(\gamma-1)(2-\alpha)}{2} \hat{s}_t^2 + t.i.p. + O(\|\bar{\mathcal{C}}_t, \bar{\mathcal{H}}_t\|^3)\end{aligned}$$

where $\bar{c}_t \equiv c_t - a_t$ and $\Delta_t \equiv \ln \int_0^1 [P_t(i)/P_t]^{-\varepsilon/(1-\alpha)} di$, with the aggregate production relation $(1-\alpha)\hat{n}_t = \bar{c}_t + (1-\alpha)\Delta_t$.¹⁹

Therefore, I can rewrite intertemporal welfare as

$$\begin{aligned}\frac{U_0}{\bar{\mathcal{C}}^{1-\gamma}} &= \frac{V_0}{\bar{\mathcal{C}}^{1-\gamma}} + \beta E_0 \left(\exp[(1-\gamma)\Delta a_1] \frac{U_1}{\bar{\mathcal{C}}^{1-\gamma}} \right) + t.i.p. + O(\|\bar{\mathcal{C}}_0, \bar{\mathcal{H}}_0\|^3) \\ &= -\frac{1}{2} E_0 \sum_{t=0}^{\infty} \delta^t \left(2(1-\alpha)\Delta_t + \frac{\gamma(2-\alpha) + \alpha\bar{c}_t^2}{1-\alpha} + (\gamma-1)(2-\alpha)\hat{s}_t^2 + 2(\gamma-1)(2-\alpha)(\widehat{a_t - a_0})\hat{s}_t \right) \\ &\quad + t.i.p. + O(\|e\|^3)\end{aligned}$$

where I used the exact property $E_0 s_t = \phi^t s_0$ to eliminate the purely linear term (see Benigno and Woodford, 2012), and where $e_t \sim Niid(0, \sigma^2)$ denotes the underlying shock that drives the economy.

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 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

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

and hence a second-order expansion around the zero-inflation steady state implies

$$\Delta_t = \eta \Delta_{t-1} + \frac{1}{2} \frac{\eta \varepsilon (1-\alpha + \alpha \varepsilon)}{(1-\eta)(1-\alpha)^2} \pi_t^2 + O(\|\pi_t\|^3)$$

Thus, I rely on the properties

$$\begin{aligned}\sum_{t=0}^{\infty} \delta^t (1-\alpha) \Delta_t &= \frac{\varepsilon}{2\lambda} \sum_{t=0}^{\infty} \delta^t \pi_t^2 + O(\|e\|^3) \\ E_0 \sum_{t=0}^{\infty} \delta^t (a_t - a_0 - t\mu) \hat{s}_t &= \hat{s}_0 + \frac{1}{1-\delta\phi} E_0 \sum_{t=1}^{\infty} \delta^t \hat{s}_t \widehat{\Delta a_t}\end{aligned}$$

¹⁹ $N_t = \int_0^1 N_t(i) di = \int_0^1 [Y_t(i)/A_t]^{1/(1-\alpha)} di = (Y_t/A_t)^{1/(1-\alpha)} \int_0^1 [P_t(i)/P_t]^{-\varepsilon/(1-\alpha)} di$, with the clearing condition $Y_t = C_t$.

to rewrite intertemporal welfare as

$$\frac{U_0}{\mathcal{C}^{1-\gamma}} = -\frac{1}{2}E_0 \sum_{t=0}^{\infty} \delta^t \left(\frac{\varepsilon}{\lambda} \pi_t^2 + \frac{\gamma(2-\alpha) + \alpha}{1-\alpha} \tilde{c}_t^2 + (\gamma-1)(2-\alpha) \hat{s}_t^2 + 2 \frac{(\gamma-1)(2-\alpha)}{1-\delta\phi} \hat{s}_t \widehat{\Delta a}_t \right)$$

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

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 are 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.

On any date t , consider all available put-call pairs that differ only in strike price. For the i th strike price, X_i , $i = 1, \dots, I$, define the auxiliary variable

$$\begin{aligned} \mathcal{A}_{it}^{(n)} &\equiv \mathcal{P}_t - C_{t,t+n} + P_{t,t+n} \\ &= P_{d,t}^{(n)} + X_i P_{b,t}^{(n)} \end{aligned}$$

where the last equality holds by put-call parity, with $P_{d,t}^{(n)} = \sum_{j=1}^n D_t^{(j)}$ the no-arbitrage price of the next n periods of dividends. Therefore, if there are no arbitrage opportunities and the LOOP holds then the map $\mathcal{A}_t^{(n)} : X_i \mapsto \mathcal{A}_{it}^{(n)}$ 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 $P_{d,t}^{(n)}$ and $P_{b,t}^{(n)}$ as the least absolute deviations (LAD) estimators that minimize expression

$$\sum_{i=1}^I |\mathcal{A}_{it}^{(n)} - P_{d,t}^{(n)} - X_i P_{b,t}^{(n)}|$$

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 frictions).

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 LOOP. 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 that associate 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% of the extracted dividend claim

price.²⁰

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. 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 that associate with prices $P_{d,t}^{(n)}$ that are nonincreasing in maturity, as they would represent arbitrage opportunities.

Overall, my selection method based on LOOP 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 thereby the potential effect of asynchronicities and other microstructural frictions, I then construct monthly prices using ten days of data at the end of each month.²¹

B.2. Errors in synthetic replication

Figure B.5 plots the auxiliary map $\mathcal{A}_t^{(n)} : X \mapsto P_{d,t}^{(n)} + X P_{b,t}^{(n)}$ for selected trading days and maturities. As shown by the lower part of figure B.5, the typical map towards 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 that associate with each cross-sectional data point.

Figure B.6 box-plots the size of the LOOP violations present in the sample which, for the most part, concentrate around errors of less than 1%; larger violations associate with the first years of the sample—probably because of a relatively low liquidity—and to years of greater volatility such as 2001 or the last years of the sample. Data previous to 1994 are more problematic by this metric (see also Golez, 2014) and I therefore exclude them altogether from the sample.

²⁰In 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.

²¹I 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 of .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 ten days is nearly imperceptible. Since my approach extracts very similar prices, I bring additional robustness to the synthetic prices extracted by Binsbergen et al.; the nearly white-noise deviation between their estimate and mine over the comparable sample are likely a mixture of asynchronicities and different proxies for the interest rate (I find options implied interest rates with nearly perfect correlations with the corresponding LIBOR and Treasury rates but with different levels that lie about halfway between the two proxies).




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




Figure B.6: Violations are in percent of the associated options implied dividend claim price; only observations not filtered out by my data selection approach are included.

B.3. Accounting for the inflation risk premium

In the empirical section, when I consider the long sample (1994-2013), the excess hold-to-maturity returns whose expectations equal the cost of uncertainty are computed relative to the term structure of interest rates extracted from nominal government bonds as a proxy for the term structure of real interest rates. The associated empirical hold-to-maturity excess returns are

$$\begin{aligned}
 \psi_{t,t+n} &\equiv \ln \frac{P_{t+n} D_{t+n} E_t M_{t,t+n} / \Pi_{t,t+n}}{P_t E_t M_{t,t+n} D_{t+n}} \\
 &= \ln \frac{D_{t+n} E_t M_{t,t+n}}{E_t M_{t,t+n} D_{t+n}} + \ln \frac{\Pi_{t,t+n} E_t M_{t,t+n} / \Pi_{t,t+n}}{E_t M_{t,t+n}} \\
 &= r_{t \rightarrow t+n}^{e,(n)} + \pi_{t,t+n} - \ln E_t \Pi_{t,t+n} - irp_t^{(n)}
 \end{aligned}$$

where $\Pi_{t,t+n} \equiv P_{t+n}/P_t$ is the inflation rate and where $irp_t^{(n)} \equiv -\frac{1}{2}\mathcal{V}_t(M_{t,t+n}/\Pi_{t,t+n}) + \frac{1}{2}\mathcal{V}_t(M_{t,t+n}) - \frac{1}{2}\mathcal{V}_t(\Pi_{t,t+n})$ defines the inflation risk premium over a n -period horizon. It follows that

$$l_t^{(n)} = \frac{1}{n} \left(\exp \left[E_t \psi_{t,t+n} + irp_t^{(n)} + \frac{1}{2} \mathcal{V}_t(\Pi_{t,t+n}) \right] - 1 \right)$$

so the welfare cost measure is biased by the presence of the inflation risk premium.²² The Jensen's term $\frac{1}{2}\mathcal{V}_t(\Pi_{t,t+n})$ can be safely disregarded (see also Grishchenko and Huang, 2013). To account for the bias, I compute average inflation risk premia during the 2000-2013 period, over which TIPS data are available, as

$$E(irp_t^{(n)}) = E(\ln P_{0,t}^{(n)}) - E(\ln P_{b,t}^{(n)}) - E(\pi_{t,t+n})$$

²²Note that the maturity-specific risk premium studied by Binsbergen et al. (2013) is $E_t \psi_{t,t+n}$.

where $P_{0,t}^{(n)}$ denotes the price of the n th real zero-coupon bond, and correct welfare costs accordingly. The maintained assumption is that the mean inflation risk premium over the 2000-2013 sample is the same as during the 1994-2013 period. To correct for the inflation term I use CPI inflation (BEA-NIPA database).²³

References

- Alvarez, F., Jermann, U.J., 2004. Using asset prices to measure the cost of business cycles. *Journal of Political Economy* 112, 1223–56.
- Atkeson, A., Phelan, C., 1994. Reconsidering the costs of business cycles with incomplete markets. *NBER Macroeconomics Annual* 9, 187–218.
- Backus, D., Chernov, M., Zin, S.E., 2014. Sources of entropy in representative agent models. *Journal of Finance* 69, 51–99.
- Bansal, R., Yaron, A., 2004. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59, 1481–1509.
- Barillas, F., Hansen, L.P., Sargent, T.J., 2009. Doubts or variability? *Journal of Economic Theory* 144, 2388–418.
- Barro, R., 2009. Rare disasters, asset prices, and welfare costs. *American Economic Review* 99, 243–64.
- Belo, F., Collin-Dufresne, P., Goldstein, R.S., 2015. Dividend dynamics and the term structure of dividend strips. *Journal of Finance* forthcoming.
- Benigno, P., Woodford, M., 2012. Linear-quadratic approximation of optimal policy problems. *Journal of Economic Theory* 147, 1–42.
- Binsbergen, J.H.v., Brandt, M., Koijen, R.S.J., 2012. On the timing and pricing of dividends. *American Economic Review* 102, 1596–1618.
- Binsbergen, J.H.v., Hueskes, W.H., Koijen, R.S.J., Vrugt, E.B., 2013. Equity yields. *Journal of Financial Economics* 110, 503–19.
- Binsbergen, J.H.v., Koijen, R.S.J., 2015. The term structure of returns: Facts and theory. Manuscript.
- Blanchard, O.J., Galí, J., 2007. Real wage rigidities and the New Keynesian model. *Journal of Money, Credit and Banking* 39, 35–65.
- Boguth, O., Carlson, M., Fisher, A., Simutin, M., 2012. Leverage and the limits of arbitrage pricing: Implications for dividend strips and the term structure of equity risk premia. Manuscript.
- Borovicka, J., Hansen, L.P., 2014. Examining macroeconomic models through the lens of asset pricing. *Journal of Econometrics* 183, 67–90.
- Calvo, G.A., 1983. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12, 383–98.
- Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107, 205–51.
- Campbell, J.Y., Ludvigson, S., 2001. Elasticities of substitution in Real Business Cycle models with home production. *Journal of Money, Credit, and Banking* 33, 847–75.
- Clarida, R., Galí, J., Gertler, M., 1999. The science of monetary policy: A New Keynesian perspective. *Journal of Economic Literature* 37, 1661–1707.
- Cochrane, J.H., 2011a. Determinacy and identification with Taylor rules. *Journal of Political Economy* 119, 565–615.
- Cochrane, J.H., 2011b. Discount rates. *Journal of Finance* 66, 1047–1108.
- Croce, M.M., 2012. Welfare costs in the long run. Manuscript.
- Croce, M.M., Lettau, M., Ludvigson, S.C., 2015. Investor information, long-run risk, and the term structure of equity. *Review of Financial Studies* forthcoming.
- De Santis, M., 2007. Individual consumption risk and the welfare cost of business cycles. *American Economic Review* 97, 1488–1506.
- Dew-Becker, I., Giglio, S., 2013. Asset pricing in the frequency domain: Theory and empirics. Manuscript.

²³Using the core PCE price index produces virtually identical results.

- Ellison, M., Sargent, T.J., 2012. Welfare cost of business cycles in economies with individual consumption risk. Manuscript.
- Epstein, L.G., Farhi, E., Strzalecki, T., 2014. How much would you pay to resolve long-run risk? *American Economic Review* 104, 2680–97.
- Gabaix, X., 2012. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics* 127, 645–700.
- Galí, J., 2008. *Monetary Policy, Inflation, and the Business Cycle. An Introduction to the New Keynesian Framework*. Princeton University Press, Princeton.
- Galí, J., Gertler, M., López-Salido, J.D., 2007. Markups, gaps, and the welfare cost of business fluctuations. *Review of Economics and Statistics* 89, 44–59.
- Golez, B., 2014. Expected returns and dividend growth rates implied in derivative markets. *Review of Financial Studies* 27, 790–822.
- Greenwood, J., Hercowitz, Z., 1991. The allocation of capital and time over the business cycle. *Journal of Political Economy* 99, 1188–1214.
- Grishchenko, O.V., Huang, J.Z., 2013. The inflation risk premium: Evidence from the TIPS market. *Journal of Fixed Income* 22, 5–30.
- Gürkaynak, R.S., Sack, B., Wright, J.H., 2010. The TIPS yield curve and inflation compensation. *American Economic Journal: Macroeconomics* 2, 70–92.
- Gürkaynak, R.S., Wright, J.H., 2012. Macroeconomics and the term structure. *Journal of Economic Literature* 50, 331–67.
- Hansen, L.P., Heaton, J.C., Li, N., 2008. Consumption strikes back? Measuring long-run risk. *Journal of Political Economy* 116, 260–302.
- Hansen, L.P., Jagannathan, R., 1991. Implications of security market data for models of dynamic economies. *Journal of Political Economy* 99, 225–62.
- Krusell, P., Smith, A.A., 1999. On the welfare effects of eliminating business cycles. *Review of Economic Dynamics* 2, 245–72.
- Lettau, M., Uhlig, H., 2000. Can habit formation be reconciled with business cycle facts? *Review of Economic Dynamics* 3, 79–99.
- Lettau, M., Wachter, J.A., 2007. Why is long-horizon equity less risky? A duration-based explanation of the value premium. *Journal of Finance* 62, 55–92.
- Lettau, M., Wachter, J.A., 2011. The term structures of equity and interest rates. *Journal of Financial Economics* 101, 90–113.
- Ljungqvist, L., Uhlig, H., 2000. Tax policy and aggregate demand management under catching up with the Joneses. *American Economic Review* 90, 356–66.
- Lopez, P., Lopez-Salido, D., Vazquez-Grande, F., 2015a. Exponential-affine approximations of macro-finance models. Manuscript.
- Lopez, P., Lopez-Salido, D., Vazquez-Grande, F., 2015b. Nominal rigidities and the term structures of equity and bond returns. FEDS Board of Governors of the Federal Reserve System, 2015-064.
- Lucas, R.E.J., 1987. *Models of Business Cycles*. Oxford University Press, Oxford.
- Lucas, R.E.J., 2003. Macroeconomic priorities. *American Economic Review* 93, 1–14.
- Otrok, C., 2001. On measuring the welfare cost of business cycles. *Journal of Monetary Economics* 47, 61–92.
- Tallarini, T.D.J., 2000. Risk-sensitive real business cycles. *Journal of Monetary Economics* 45, 507–32.
- Woodford, M., 2010. Optimal monetary stabilization policy, in: Friedman, B., Woodford, M. (Eds.), *Handbook of Monetary Economics*. Elsevier, New York. volume 3B, pp. 723–828.

ONLINE APPENDIX

I. Relationship with the definitions of Lucas (1987) and Alvarez and Jermann (2004)

Definition (1) is slightly different from the one studied by Alvarez and Jermann (2004). First, they measure the cost of fluctuations by the uniform compensation Ω_t in

$$\begin{aligned} E_t U(\{(1 + \Omega_t(\theta))C_{t+n}\}_{n=1}^{\infty}, \{X_{t+n}\}_{n=1}^{\infty}) &= \\ &= E_t U(\{\theta E_t C_{t+n} + (1 - \theta)C_{t+n}\}_{n=1}^{\infty}, \{X_{t+n}\}_{n=1}^{\infty}) \end{aligned}$$

and focus on Lucas's (1987) total cost, $\Omega_t(1)$, and on the marginal cost, $\frac{\partial}{\partial \theta} \Omega_t(0)$. The definition I consider measures instead the cost of fluctuations by a compounded compensation, so it can be interpreted as the tradeoff between *growth* and macroeconomic stability (the working paper version of Alvarez and Jermann, 2004, makes this point).

Second, I allow for considering the stabilization of only some coordinates of consumption—the set $\mathcal{N}$ in definition (1)—rather than of the whole stochastic process. This flexibility grants a direct focus on the relevant periodicity of economic uncertainty.

Finally, I address the proxy problem by stabilizing the determinants of consumption rather than consumption itself.

II. Essentially-affine approximation of equilibrium asset prices

The practical approximation proceeds in three steps. See Lopez, Lopez-Salido and Vazquez-Grande (2015a) for a treatment in greater generality and detail, and for a comparison of its quality with alternative solution methods.

II.1. First step

Cashflows. After a risk-adjusted loglinearization, the approximate equilibrium quantity dynamics of the model of section 4 can be written as:

$$\begin{aligned} \Delta c_{t+1} &= \mu_c + C_c \zeta_t + D_c \varepsilon_{t+1} + O(\|\zeta_t, \varepsilon_{t+1}\|^2) \\ \Delta d_{t+1} &= \mu_d + C_d \zeta_t + D_d \varepsilon_{t+1} + O(\|\zeta_t, \varepsilon_{t+1}\|^2) \end{aligned}$$

where c is log consumption, d is the log of an arbitrary cashflow process, and where the state ζ_t that drives quantities follows the autoregressive process

$$\zeta_{t+1} = A\zeta_t + B\varepsilon_{t+1}$$

with $\varepsilon_t \sim \text{Niid}(0, I)$ a vector of shocks.

Discount rates. The stochastic discount factor is

$$\begin{aligned} m_{t+1} &= \ln(\beta) - \gamma\mu_c - \gamma C_c \zeta_t + \gamma(1 - \phi)\hat{s}_t - x_t D_c \varepsilon_{t+1} + O(\|\zeta_t, \varepsilon_{t+1}\|^2) \\ &= -r_t - \frac{1}{2} x_t^2 \|D_c\|^2 - x_t D_c \varepsilon_{t+1} + O(\|\zeta_t, \varepsilon_{t+1}\|^2) \end{aligned}$$

where the residual term comes from the approximate equation for consumption growth and the last equality is by the no-arbitrage relation $r_t = -\ln E_t M_{t+1}$. The time-varying price of risk follows a nonlinear process $x : \hat{s}_t \mapsto x(\hat{s}_t) = (\gamma/S)(1 - 2\hat{s}_t)^{1/2}$ that is responsible for the absence of a closed-form solution to the problem, which would otherwise take an exponential-affine form. Thus, approximate the endogenous and nonlinear dynamics of the price of risk as

$$x_t = x(0) + x'(0)\hat{s}_t + O(\|\hat{s}_t\|^2) \tag{II.1}$$

$$x_t^2 = x(0)^2 + 2x(0)x'(0)\hat{s}_t + O(\|\hat{s}_t\|^2) \tag{II.2}$$

where, since in the Campbell-Cochrane specification $x''(0)x(0) + x'(0)^2 = 0$, the residual in equation (II.2) is exactly zero.

Thus approximated, the price of risk has an essentially-affine form and thereby allows for an exponential-affine solution for equilibrium yields, since all sources of stochastic volatility owe to the time-varying price of risk and since the risk-free rate is exactly affine in the state vector.

II.2. Second step

Guess the exponential-affine solution for yields $y_{d,t}^{(n)} \equiv -\frac{1}{n} \ln(P_t^{(n)}/D_t)$,

$$y_{d,t}^{(n)} = -\frac{1}{n}A^{(n)} - \frac{1}{n}B_\zeta^{(n)}\zeta_t - \frac{1}{n}B_s^{(n)}\hat{s}_t + O(\|\zeta_t, \hat{s}_t, \varepsilon_{t+1}\|^2)$$

and verify it by the fundamental no-arbitrage pricing formula $0 = \ln E_t(M_{t+1}R_{t+1}^i)$.

In fact, given the Gaussianity of log returns $r_{d,t+1}^{(n)} \equiv \Delta d_{t+1} - (n-1)y_{d,t+1}^{(n-1)} + ny_{d,t}^{(n)}$, it follows that

$$\begin{aligned} 0 &= E_t m_{t+1} + dp_t^{(n)} - E_t dp_{t+1}^{(n-1)} + E_t \Delta d_{t+1} + \frac{1}{2} \text{var}_t(m_{t+1} - dp_{t+1}^{(n-1)} + \Delta d_{t+1}) \\ &= \ln(\beta) + \mu_d - \gamma\mu_c - A^{(n)} + A^{(n-1)} + [C_d - \gamma C_c - B_\zeta^{(n)} + B_\zeta^{(n-1)}A]\zeta_t + [\gamma(1-\phi) - B_s^{(n)} + B_s^{(n-1)}\phi]\hat{s}_t \\ &\quad + \frac{1}{2}\|D_c\|^2 x_t^2 + \frac{1}{2}V_{n-1,t}V'_{n-1,t} - x_t D_c V'_{n-1,t} + O(\|\zeta_t, \hat{s}_t, \varepsilon_{t+1}\|^2) \end{aligned}$$

where $V_{n-1,t} = D_d + B_\zeta^{(n-1)}B - B_s^{(n-1)}D_c + x_t B_s^{(n-1)}D_c/\gamma$. Therefore, using equations (II.1) and (II.2),

$$\begin{aligned} 0 &= \ln(\beta) + \mu_d - \gamma\mu_c - A^{(n)} + A^{(n-1)} + \frac{1}{2}\|V_{0,n-1} - x(0)(D_c - V_{1,n-1})\|^2 \\ &\quad + [\gamma(1-\phi) - B_s^{(n)} + B_s^{(n-1)}\phi + x(0)x'(0)\|D_c - V_{1,n-1}\|^2 - x'(0)V_{0,n-1}(D_c - V_{1,n-1})']\hat{s}_t \\ &\quad + [C_d - \gamma C_c - B_\zeta^{(n)} + B_\zeta^{(n-1)}A]\zeta_t + O(\|\zeta_t, \hat{s}_t, \varepsilon_{t+1}\|^2) \end{aligned}$$

which identifies the exponential-affine solution as the solution to the Riccati equations

$$\begin{aligned} A^{(n)} &= A^{(n-1)} + \ln(\beta) + \mu_d - \gamma\mu_c + \frac{1}{2}\|V_{0,n-1} - x(0)(D_c - V_{1,n-1})\|^2 \\ B_\zeta^{(n)} &= B_\zeta^{(n-1)}A + C_d - \gamma C_c \\ B_s^{(n)} &= B_s^{(n-1)}\phi + \gamma(1-\phi) + x(0)x'(0)\|D_c - V_{1,n-1}\|^2 - x'(0)V_{0,n-1}(D_c - V_{1,n-1})' \end{aligned}$$

with

$$\begin{aligned} V_{0,n-1} &= D_d + B_\zeta^{(n-1)}B - B_s^{(n-1)}D_c \\ V_{1,n-1} &= \frac{1}{\gamma}B_s^{(n-1)}D_c \end{aligned}$$

These closed-form expressions allow for computing the entire term structure of yields, $y_{d,t}^{(n)}$, from a simulated path of the state vector $[\zeta_t; \hat{s}_t]$ up to a remainder of order at least $O(\|\zeta_t, \hat{s}_t, \varepsilon_{t+1}\|^2)$.

II.3. Third step

Finally, use the lognormal no-arbitrage pricing formula to compute

$$\begin{aligned} \ln E_t R_{t+1}^{e,(n)} &= x_t D_c V'_{n-1,t} \\ E_t r_{t+1}^{e,(n)} &= x_t D_c V'_{n-1,t} - \frac{1}{2}V_{n-1,t}V'_{n-1,t} \end{aligned}$$

and simulate a sample path for risk premia and return volatilities by using the exact dynamics $x(\hat{s}_t)$.

III. Term structures in some consumption-based asset pricing models

Table III.5 and figure III.7 show the implications of some of the leading consumption-based asset pricing models for the three term structures and for the welfare cost of uncertainty and the equity premium. I consider the habit formation model of Campbell and Cochrane (1999), the long-run risk model of Bansal and Yaron (2004), the long-run risk model under limited information of Croce et al. (2015), the recursive preferences of Tallarini (2000) and Barillas et al. (2009), the rare disasters model of Gabaix (2012), and the quasi-structural model of Lettau and Wachter (2011). In studying the term structures in the different asset pricing models, I consider the original calibrations, which the authors choose to match some asset pricing facts. The appendix works out the details of each model. I refer to the original writings for a list of the stylized asset pricing facts each model replicates.

III.1. Structural approach

The habit formation model of Campbell and Cochrane (1999) predicts a flat term structure of interest rates and an upward-sloping term structure of equity. The term structure of interest rates is driven by a particular calibration of the time-varying risk aversion that produces a constant risk-free rate. The term structure of equity is instead driven by the positive correlation between the pricing factor—shocks to consumption growth—and dividend growth, and by the perfectly negative correlation between the pricing factor and the shocks to the price of risk, which decreases as consumption grows away from the external habit. Since dividend strips load negatively on shocks to the price of risk, and the more so the longer the maturity, people command a greater risk premium to bear long-run dividend strip risk. Under the baseline calibration, the model of Campbell and Cochrane predicts a marginal cost of all fluctuations of 0.20% and an equity premium of 6.8%.

The long-run risk model of Bansal and Yaron (2004) generates an upward-sloping term structure of equity and a downward-sloping term structure of interest rates. Bansal and Yaron introduce rich dynamics in consumption growth, which is driven both by shocks to expected consumption growth and to consumption volatility. Epstein-Zin-Weil utility then makes all shocks to the consumption opportunity set show up as pricing factors. In the calibration of Bansal and Yaron, long-run dividend strips load more heavily on the shocks to the consumption opportunity set and therefore are more risky, as long as the elasticity of intertemporal substitution is larger than one. In the model the risk-free rate is driven by shocks to the predictable component of consumption, which is positively priced; since long-run zero-coupon bonds load less on this state than the risk-free rate, they provide long-run insurance. This property explains the downward-sloping term structure of interest rates. The quantitative implications of the long-run risk model is a marginal cost of all fluctuations of 0.34% and an equity premium of 4.9%.

Croce et al. (2015) consider the long-run risk model of Bansal and Yaron (2004) and change the information structure. Under limited information, not all shocks to the cashflow opportunity set are observable; the shocks that are priced are a linear combination of both short-run and long-run cashflow shocks. Then, since long-run shocks have a relatively small volatility, long-run dividend strips load less on the shocks that are priced under limited information than short-run dividend strips. This strategy allows for generating a downward-sloping term structure of equity; however, the curvature is not enough quantitatively, at least under the baseline calibration, it still predicts a downward-sloping term structure of interest rates, and it works in a world in which risk premia are not time-varying. The model predicts a marginal cost of all fluctuations of 0.39%, against a predicted equity premium of 6.6%.

The ambiguity averse multiplier preferences in Barillas et al. (2009) and the recursive preferences of Tallarini (2000) yield two flat term structures which imply the equality between the equity premium and the cost of uncertainty around any coordinate set $\mathcal{N} \in \bar{\mathbb{N}}$, up to a scale factor. The unitary elasticity of intertemporal substitution that characterizes the recursive preferences of Tallarini (2000) and the robust control literature implies constant dividend yields, as discount-rate effects exactly offset cashflow effects in pricing equity claims; the random walk in consumption in turn implies constant interest rates and thereby a flat bond term structure. The multiplier preferences of Barillas et al. (2009) and the observationally equivalent model of Tallarini (2000) predict a marginal cost of all fluctuations of 0.09% and an equity premium of about 2.0%.

Finally, the rare disasters model of Gabaix (2012) produces two flat term structures of holding-period returns but a slightly downward-sloping term structure of hold-to-maturity equity returns. The intuition behind the flat term structure

	$L^{\bar{N}}$	$\ln E(R^{e,m})$
Campbell and Cochrane (1999)	0.20	6.81
Bansal and Yaron (2004)	0.34	4.90
Croce et al. (2015)	0.39	6.56
Barillas et al. (2009)	0.09	1.92
Gabaix (2012)	0.31	7.89
Lettau and Wachter (2011)	0.12	7.18

Table III.5: Mean marginal cost of lifetime uncertainty and equity premium (percent per year).

$\mathcal{N}$	$L^{\mathcal{N}}$	$L^{\bar{N} \mathcal{N}}$
up to 1 year	0.746	0.124
up to 5 years	0.671	0.117
up to 10 years	0.550	0.105
up to 20 years	0.388	0.086
$\bar{N}$	0.125	0

Table III.6: Marginal cost of fluctuations at all periodicities $n \in \mathcal{N}$. Lettau and Wachter (2011) model-based estimates

of holding period returns in the rare disasters model is that different dividend strips have the same exposure to the disaster event, whose probability is independent of the cashflow shocks that are priced. However, the mean-reversion in the state that drives equity prices makes long-duration equities load slightly less on it if held to maturity than short-duration equities. The model implies a marginal cost of all fluctuations of 0.31% and an equity premium of 7.0%.


III.2. Descriptive approach

I turn to the exponential-Gaussian no-arbitrage model of Lettau and Wachter (2011), which is designed to capture a downward-sloping term structure of equity and an upward-sloping term structure of interest rates. Without micro-founding it, Lettau and Wachter directly specify a stochastic discount factor, whose existence is guaranteed by the no-arbitrage theorems. They assume a single conditional pricing factor perfectly related to short-run cashflow shocks and a single state driving the price of risk. To match the downward-sloping term structure of equity, they assume that the predictable component of cashflows is negatively related to the priced shocks. Long-run dividend strips thus contain a component that provides long-run insurance. They then assume a zero correlation between cashflow and discount-rate shocks to avoid that the negative loading of long-run dividend strips on the state that drives the price of risk offsets the long-run insurance effect.


Finally, since only short-run cashflow shocks are priced, Lettau and Wachter manage to capture an upward-sloping term structure of interest rates by assuming that shocks to the state driving the risk-free rate are negatively correlated with the priced shocks. Since long-run zero-coupon bonds are less exposed to this state than short-run bonds are, the assumption generates a positive bond risk premium as the maturity increases.

The model predicts a marginal cost of total uncertainty of 0.12% and an equity premium of 7.2%. Table III.6 reports the cost of short- and long-run fluctuations over different coordinate sets. An increase in consumption uncertainty by a fraction θ over a ten-year period has a marginal cost of more than 0.5θ percentage points of growth per year during the decade. These numbers are in line with the options implied estimates in table 1 and compare to smaller yet nontrivial marginal benefits of long-run stability, which tend to zero as the stabilization becomes asymptotic.


The volatility of the term structure as captured by the model of Lettau and Wachter supports the evidence in table 2 and figure 3. The standard deviation of the cost of fluctuations at short periodicities is large and decays over long horizons. In their model, the term structure of welfare costs follows a one-factor structure driven entirely by movements in the time-varying market price of risk, as shown in the online appendix.




(a) *Habit formation (Campbell and Cochrane, 1999).*




(b) *Long-run risk (Bansal and Yaron, 2004).*




(c) *Ambiguity aversion (Barillas, Hansen and Sargent, 2009), Epstein-Zin-Weil log utility (Tallarini, 2000).*



(d) *Long-run risk (limited information) (Croce, Lettau and Ludvigson, 2015).*



(e) *Rare disasters (Gabaix, 2012).*



(f) *Quasi-structural (Lettau and Wachter, 2011).*

Figure III.7: The term structures of equity, interest rates and welfare costs of uncertainty in some consumption-based asset pricing models.