DOCUMENT

DE TRAVAIL

N° 521

THE TERM STRUCTURE OF THE WELFARE COST

OF UNCERTAINTY

Pierlauro Lopez

November 2014

BANQUE DE FRANCE

EUROSYSTÈME

DIRECTION GÉNÉRALE DES ÉTUDES ET DES RELATIONS INTERNATIONALES

THE TERM STRUCTURE OF THE WELFARE COST

OF UNCERTAINTY

Pierlauro Lopez

November 2014

The Term Structure of the Welfare Cost of Uncertainty

Pierlauro Lopez[1],[*]

Banque de France

Abstract

The marginal cost of aggregate fluctuations has a term structure that is a simple transformation of the term structures of equity and interest rates. I extract evidence from index option markets to infer a downward-sloping, volatile and procyclical term structure of welfare costs. On average, the gains from greater macroeconomic stability are large, especially in the short run. I estimate that at the margin the elimination of one-year ahead consumption risk is worth around 12 percentage points of additional growth; this number compares to a marginal cost of lifetime uncertainty of 2-3 percentage points. Over time, the term structure of welfare costs varies substantially, predictably and with a volatility that decreases with maturity. These empirical properties of the term structure of welfare costs cannot be easily captured by today's leading dynamic equilibrium models and therefore represent a puzzling piece of evidence with potentially important welfare implications.

Keywords: Welfare cost of business cycles; Macroeconomic priorities; Dividend strips; Return forecastability.

JEL classification: E32; E44; E61; G12.

Résumé

Le coût marginal de fluctuations agrégées a une structure par terme qui est une simple transformation des structures par terme des dividendes et des taux d'intérêt. Cet article utilise l'information implicitement contenue dans les prix d'options pour inférer une structure par terme des coûts en bien-être des fluctuations qui a une pente négative et qui est volatile et procyclique. En moyenne, les bénéfices d'une majeure stabilité macroéconomique sont importants et surtout dans le court terme: l'estimation dans cet article suggère qu'à la marge l'élimination des fluctuations agrégées sur un horizon d'un an a la même valeur en bien-être qu'environ 12 points de pourcentage de croissance additionnelle pendant un an. Ce numéro peut être comparé à un coût marginal de toutes les fluctuations futures de 2-3 points de pourcentage de croissance perpétuelle. La structure par terme des coûts des fluctuations varie substantiellement, prévisiblement et avec une volatilité en relation inverse avec la maturité. Ces propriétés empiriques de la structure par terme des coût des fluctuations ne peuvent pas être capturées facilement par les modéles dynamiques d'équilibre actuels et représentent ainsi un défi pour la modelisation avec des implications en termes de bien-être potentiellement importantes.

Mots clés: coût en bien-être des cycles économiques, priorités macroéconomiques, structure par terme de dividendes, prévisibilité des rendements.

Classification JEL: E32; E44; E61; G12.

Non-technical summary

How much growth are people willing to trade against a stabilization of macroeconomic fluctuations? Lucas (1987) introduced this thought experiment as a measure of the tradeoff between growth and macroeconomic stability, the so-called cost of aggregate uncertainty (also known as the cost of business cycles). Similarly, this paper introduces the term structure of the cost of uncertainty as a measure of the tradeoff between growth and macroeconomic stability at different time horizons.

As long as the appropriate insurance is actually traded the question ceases to be a simple thought experiment detached from the actual tradeoffs that investors face. Recent derivative securities allow for synthesizing the required insurance over some horizons and, therefore, provide new insight into an old question (the tradeoff between growth and stability) as well as evidence to study the new question (the tradeoff between stability at different time horizons). In particular, the term structure of the cost of uncertainty is a simple transformation of the term structures of equity and interest rates. For example, by holding a portfolio short a zero-coupon bond and long a zero-coupon equity security a consumer can buy an increase in consumption growth equal to the ex-ante return on the portfolio for each unit of consumption volatility over the relevant horizon that she accepts to shoulder.

Accordingly, I rely on recent index option market evidence and bond data as well as on current consumption-based asset pricing models to infer a downward-sloping and volatile term structure of welfare costs. The negative average slope can be interpreted as evidence that people value more short-run stability than long-run stability. Conversely, asset markets suggest that marginal increases in aggregate uncertainty are particularly costly, especially in the short run and during downturns such as in the early 2000s and during the recent financial crisis. The large risk premia underlying the term structure suggest that people command a nontrivial amount of growth to shoulder systematic cashflow risk. Over time, the term structure of welfare costs varies substantially and with movements that are revealed by excess return predictors such as equity and bond yields.

This paper is an example of the implications of finance for macroeconomics, both because asset market data reveal people's evaluation of the tradeoff among growth, short-run and long-run stability, and because they suggest that a macroeconomic model that fails to capture the empirical properties of the term structure of welfare costs is likely to miss crucial welfare implications. The negative slope of the estimated term structure of welfare costs (which rests on a downward-sloping term structure of equity and on an upward-sloping term structure of interest rates) cannot be easily captured by today's leading macro-finance models and therefore represents a puzzling piece of evidence. In this sense, 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.

1. Introduction

How much growth are people willing to trade against a marginal stabilization of macroeconomic fluctuations? The marginal cost of aggregate uncertainty (Alvarez and Jermann, 2004) answers this important question in economics, which dates back at least to Lucas (1987). I decompose the cost of uncertainty into a term structure. This decomposition allows for studying how cashflow fluctuations at different horizons contribute to the cost of lifetime uncertainty (proposition 1). The term structure of welfare costs allows for understanding the tradeoff between growth and macroeconomic stability as well as the tradeoff between cashflow stability at different periodicities (for example, between short-run and long-run stability).

Reviving the insight that asset market data can reveal the marginal cost of fluctuations (Alvarez and Jermann, 2004), I then show how the components of the term structure of welfare costs are tightly linked to the risk premia on market dividend strips and on zero-coupon bonds (proposition 2); the term structure of welfare costs is a simple transformation of the term structures of equity and interest rates (e.g., studied in Lettau and Wachter, 2011; Binsbergen, Brandt and Koijen, 2012; Binsbergen, Hueskes, Koijen and Vrugt, 2013).[2] This link allows for a measure of the cost of fluctuations that is directly observable over the last two decades, as asset markets started to trade the relevant insurance for some maturities. Like Alvarez and Jermann (2004), my approach requires only the absence of arbitrage opportunities and does not require a parametric specification of consumer preferences.

Accordingly, in the empirical section I rely on bond data and extract evidence about the term structure of equity from index option markets to infer the cost of uncertainty as a function of its time horizon. Like Binsbergen et al. (2012) and Golez (2014) I rely on option data and no-arbitrage relations to synthetically replicate single market dividend payments, so-called dividend strips. This paper proposes a method to extract dividend claim prices by combining options with different moneyness levels to avoid the need for an interest-rate proxy and to mitigate measurement error by excluding observations that violate the put-call parity relation.

The analysis finds costs that are large, volatile, procyclical and have nontrivial term structure features. The point estimates, reported in figure 1a, suggest a negatively sloped term structure of welfare costs; people command a larger premium to shoulder short-term cashflow uncertainty than to shoulder longer-term uncertainty. The premium at one-year frequency is 12-13 percent on average and its volatility is similar in size. The tradeoff among short-run, long-run stability and

---

[2]A recent and rapidly growing literature is focusing on risk pricing across maturities (see, for example, Lettau and Wachter, 2007; Hansen, Heaton and Li, 2008; Backus, Chernov and Zin, 2014; Borovicka and Hansen, 2014; Dew-Becker and Giglio, 2014).

growth revealed by the term structure of welfare costs is nontrivial and varies substantially over the business cycle.

The evidence of a downward-sloping term structure of welfare costs helps identifying a model to capture and quantify the entire term structure. 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 (2014), the ambiguity averse multiplier preferences of Barillas, Hansen and Sargent (2009), and the rare disasters model of Gabaix (2012). Although these models do not study the term structure of welfare costs directly, they have implications for it and are calibrated to match several other asset pricing facts. Unfortunately, from a structural perspective, replicating a downward-sloping term structure of equity and an upward-sloping term structure of interest rates is problematic (e.g., Binsbergen et al., 2012; Gürkaynak and Wright, 2012).

In this context, Lettau and Wachter (2011) offer a parsimonious, quasi-structural model designed to capture precisely a downward-sloping term structure of equity, an upward-sloping term structure of interest rates, and time-varying risk premia. The quantitative implications of the model, reported in figure 1b, are a marginal cost of lifetime fluctuations of about 3 percent. Over a horizon of up to ten years, a marginal increase in uncertainty costs more than 10 percentage points of annual growth per unit of uncertainty, as measured by the conditional standard deviation of cashflows over the relevant horizon. These numbers compare with much smaller marginal benefits of long-run stability.
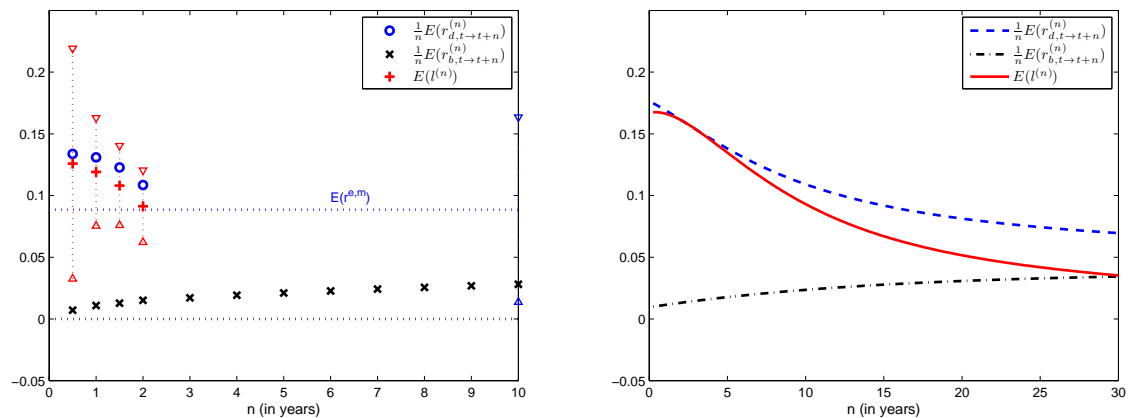
(a) *Point estimates (with block bootstrapped 95% confidence interval). $E(r^{e,m})$ is the equity premium.*

(b) *Model-based estimates in the model of Lettau and Wachter (2011).*

Figure 1: Term structures of equity ('∘' and dashed line), interest rates ('×' and dotted line) and welfare costs ('+' and solid line); annualized premia.

My approach builds on and complements the analysis by Alvarez and Jermann (2004). First, by

focusing on *marginal* welfare costs, I stick to their preference-free setting while linking the cost of fluctuations to a richer set of recent financial market evidence. By not relying on a particular parametrization of preferences the analysis avoids the fragility of model-based estimates of the cost of fluctuations on alternative assumptions on preferences.[3]

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 one studied by Alvarez and Jermann. Their answer depends a lot on the parametric assumptions about the filter that separates the trend and the business-cycle frequencies of the cashflow process. The question I am interested in is nonparametric and complements the exercise of Alvarez and Jermann by decomposing the marginal cost of uncertainty in the time domain rather than in the frequency domain.

This paper also complements the macroeconomic interpretation of equity yields offered by Binsbergen, Hueskes, Koijen and Vrugt (2013) by describing in what sense the evidence about the maturity structure of asset prices maps into evidence about the welfare cost of business cycles. The components of the term structure of uncertainty turn out to be the arithmetic versions of the maturity-specific risk premia in Binsbergen et al.'s decomposition of forward equity yields plus the inflation risk premium over the relevant horizon, and I provide formulas to describe how the term structure components combine to measure the cost of uncertainty around arbitrary sets of cashflow coordinates.[4]

Finally, it is important to remember what welfare costs do and do not say. The literature opened by Lucas (1987) focuses on fluctuations in the level of cashflows, as opposed to fluctuations in its distance (gap) from some efficient level (see Galí, Gertler and López-Salido, 2007), and is silent as to whether the current level of risk in the economy is Pareto optimal.[5] Thus, evidence of large welfare costs is only a necessary condition to justify attributing priority to stabilizing fluctuations;

---

[3]Since the seminal work by Lucas (1987), a large literature has focused on model-based estimates of the costs of lifetime cashflow uncertainty under alternative assumptions on preferences and cashflow processes, with highly dispersed findings that range from virtually zero to more than 20 percent (for example, Atkeson and Phelan, 1994; Krusell and Smith, 1999; Tallarini, 2000; Otrok, 2001; De Santis, 2007; Barillas et al., 2009; Croce, 2012; Ellison and Sargent, 2012, among many others). In this context, Alvarez and Jermann (2004) manage to move the game to a preference-free environment by showing how one can directly use asset market data to measure the cost of fluctuations at the margin.

[4]It follows that time-variation in forward equity yields (corrected by the inflation risk premium) reveals either a time-varying term structure of expected growth (Binsbergen et al., 2013) or a time-varying term structure of the welfare cost of uncertainty, or both.

[5]A recent literature studies the contribution of policy shocks to aggregate uncertainty (e.g., Baker, Bloom and Davis, 2013; Bekaert, Hoerova and Lo Duca, 2013) as well as the potential role of time-varying uncertainty in explaining business cycle fluctuations (e.g., Bloom, 2009; Fernández-Villaverde et al., 2011; Ilut and Schneider, 2014).

the sufficiency of the condition can only be assessed in the context of a structural macroeconomic model. However, such a model should not abstract from the documented properties of welfare costs; the link between the welfare cost of uncertainty and asset prices suggests that the asset pricing implications of a model are likely to have important normative consequences.

In this sense, the empirical properties of the term structure of welfare costs (and hence of the underlying term structures of equity and bond yields) 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. A quantitative assessment of both categories of dynamic equilibrium models based on their implications for the term structure of welfare costs should be part of the calibration process, even though this requirement is particularly demanding (section 4).[6]

2. The term structure of the cost of uncertainty

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 $i \in [0, 1]$ have time-t preferences $U_t = E_t U(C_i, X_i)$, where $C \equiv \{C_{t+n}\}_{n=1}^{\infty}$ is consumption and $X \equiv \{X_{t+n}\}_{n=1}^{\infty}$ is any other factor that influences utility.[7] Without loss of generality let factor X_i depend on individual consumption, C_i, only via aggregate consumption, $C = \int_0^1 C_i di$. Since there is a continuum of agents each of which has zero mass, this modeling strategy allows me to ask an individual how much consumption growth he would trade against a stable consumption stream without thereby having to affect all aggregate quantities, including factor X.

Let $\overline{C}_{t+n}$ denote the consumption level that is hypothetically offered to the ith individual at time $t + n$, which I refer to as *stable* consumption. Then, I parametrize stable consumption as $\overline{C}_{t+n}(\theta) = \theta E_t C_{t+n} + (1 - \theta)C_{t+n}$, where the parameter $\theta \in [0, 1]$ represents the fraction of ex-post uncertainty that is removed.

---

[6]The use of thought experiments of the form, 'how much would you pay for the following gamble?', as a diagnostic of the quantitative predictions of dynamic equilibrium models has been recently emphasized by Swanson (2012) and Epstein, Farhi and Strzalecki (2014).

[7]In section 5 I relax the assumption of a representative consumer.

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

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

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

Two particularly interesting quantities are the *total cost* $\mathcal{L}_t^{\overline{\mathbb{N}}}(1)$, which measures how much extra growth the elimination of all cashflow uncertainty is worth, and the *marginal cost* $L_t^{\mathcal{N}} \equiv \frac{\partial}{\partial \theta} \mathcal{L}_t^{\mathcal{N}}(0)$, which represents the how much extra growth a marginal stabilization is worth at the current level of uncertainty.[8]

I assume enough smoothness in preferences to guarantee that $L_t^{\mathcal{N}}$ is a differentiable map on $\theta \in [0, 1]$. Thus, differentiating (1) with respect to θ,

$$L_t^{\mathcal{N}} = \frac{\sum_{n \in \mathcal{N}} E_t(M_{t,t+n}) E_t(C_{t+n}) - E_t(M_{t,t+n} C_{t+n})}{\sum_{n \in \mathcal{N}} n\, E_t(M_{t,t+n} C_{t+n})} \quad (2)$$

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, $D_t^{(n)} = E_t M_{t,t+n} C_{t+n}$ is the price of a n-period consumption strip and $E_t M_{t,t+n}$ is the price of a n-period zero-coupon bond. It follows that equation (2) expresses the marginal cost of uncertainty around all coordinates $n \in \mathcal{N}$ as a function of the price of a claim to the consumption trend and of the price and duration of a claim to future consumption at all coordinates $n \in \mathcal{N}$.

Definition (Term structure of the cost of uncertainty). Consider the singleton set $\mathcal{N} = \{n\}$, for $n = 1, 2, ...$, and consider the cost $l_t^{(n)} \equiv \frac{\partial}{\partial \theta} \mathcal{L}_t^{\{n\}}(0)$ of a marginal increase in uncertainty in the nth coordinate of consumption. Then, by equation (2), it follows that

$$l_t^{(n)} = \frac{1}{n}\left(\frac{E_t(M_{t,t+n}) E_t(C_{t+n})}{E_t(M_{t,t+n} C_{t+n})} - 1\right)$$

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 1. Given the prices of strips $\{D^{(n)}\}$ and the term structure components $\{l^{(n)}\}$ you can compute the marginal cost $L^{\mathcal{N}}$ for any coordinate set $\mathcal{N} \subset \overline{\mathbb{N}}$.

---

[8]Appendix A discusses the relationship between definition (1) and the definitions by Lucas (1987) and Alvarez and Jermann (2004).

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

$$L_t^{\mathcal{N}} = \sum_{n \in \mathcal{N}} \omega_{n,t} l_t^{(n)}$$

where the weights $\omega_{n,t} \equiv n\, D_t^{(n)} / \sum_{n \in \mathcal{N}} n\, D_t^{(n)}$ are positive and such that $\sum_{n \in \mathcal{N}} \omega_{n,t} = 1$.[9]

Proposition 2 shows how the term structure of welfare costs is a simple transformation of the term structures of equity and interest rates, by which I mean the maps $n \mapsto E_t R_{d,t \to t+n}^{(n)}$ and $n \mapsto E_t R_{b,t \to t+n}^{(n)}$ of ex-ante hold-to-maturity returns on dividend strips and zero-coupon real bonds, respectively, as a function of maturity.[10]

Proposition 2. *The nth component of the term structure of welfare costs 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. The term structure of welfare costs is the transformation of the term structures of equity and interest rates*

$$
\begin{aligned}
l_t^{(n)} &= \frac{1}{n}\left(E_t R_{t \to t+n}^{e,(n)} - 1\right) \\
&= \frac{1}{n}\left(\exp\left\{E_t r_{d,t \to t+n}^{(n)} - E_t r_{b,t \to t+n}^{(n)} + \frac{1}{2}\mathcal{V}_t(R_{d,t \to t+n}^{(n)})\right\} - 1\right)
\end{aligned}
\tag{3}
$$

where $\mathcal{V}_t(X) \equiv 2[\ln E_t X - E_t \ln(X)]$ denotes conditional entropy and $r \equiv \ln(R)$.

Equation (3) provides the link between the term structures of equity and interest rates and the term structure of welfare costs. In particular, since the entropy term is independent of the bond price, it shows clearly how the term structure of welfare costs would unambiguously slope downwards in the case of a downward-sloping term structure of equity and of an upward-sloping term structure of interest rates.

There is a powerful intuition behind these formulas. At the margin, people would trade $l_t^{(n)} \theta$ points of growth against the elimination of a fraction θ of the aggregate cashflow uncertainty around the nth consumption coordinate, C_{t+n}. Proposition 2 shows 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

---

[9]For simplicity, I omit in the notation the dependence of weights $\{\omega_{n,t}\}$ on the coordinate set $\mathcal{N}$.

[10]For now, dividends and consumption are used interchangeably (as in an endowement economy). I discuss in section 5 how dividends rather than consumption can be used to measure the welfare costs of uncertainty in a production economy.

an average growth rate of $\frac{1}{n}(E_t R_{t \to t+n}^{e,(n)} - 1)$ by shouldering a volatility of $\mathcal{V}_t(R_{t \to t+n}^{e,(n)}) = \mathcal{V}_t(C_{t+n})$. Therefore, the cost of n-period ahead uncertainty must be $l_t^{(n)} = \frac{1}{n}(E_t R_{t \to t+n}^{e,(n)} - 1)$.

3. Empirics of the cost of uncertainty

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 as

$$C_{t,t+n} - \mathcal{P}_{t,t+n} = P_t - \sum_{j=1}^{n} D_t^{(j)} - X \, P_{b,t}^{(n)} \tag{4}$$

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, $P_{b,t}^{(n)} = E_t M_{t,t+n} \Pi_t / \Pi_{t+n}$ is the nominal price of a n-period nominal zero-coupon bond, $P_t = \Pi_t E_t \sum_{j=1}^{\infty} M_{t,t+j} D_{t+j}$ is the nominal value of the market portfolio, and $D_t^{(n)} = \Pi_t E_t M_{t,t+n} D_{t+n}$ is the nominal price of the nth dividend strip, where Π denotes the price level. Since the only unknowns in equation (4) 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 (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. grows 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.[11] I use a dataset provided by Market Data Express containing S&P 500 index option data for

---

[11]The 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

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 Binsbergen et al., 2012; Boguth et al., 2012; Golez, 2014). 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 extremely 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 rely on the no-arbitrage relation to extract *both* the risk-free rate and the strip price in a unique step. This approach produces the appropriate interest rate for synthetic replication as well as it allows for spotting violations of the law of one price (LOOP) on any trade date and for any maturity. Moreover, I follow Golez (2014) in using ten days of data at the end of each month to compute options implied prices at monthly frequency. He shows how this strategy reduces substantially the distance between intra-day and end-of-day options implied prices and thereby the potential effect of asynchronicities and other microstructural frictions.[12]

---

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.

[12]I bring further support to his claim by showing 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. (2012); 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).

Finally, to measure the real bond prices necessary to compute welfare costs I rely on zero-coupon TIPS and Treasury yields with maturities of up to ten years from Gürkaynak, Sack and Wright (2007). Since TIPS yields are either unavailable or unreliable during the 1990s, I use Treasury yields as a proxy available over the same sample period as I have dividend strips.[13] I add to hold-to-maturity Treasury returns a constant to match the corresponding TIPS means over the 2000-2013 period; this adjustment corrects for the presence of the inflation risk premium in estimated welfare costs due to the use of nominal rather than real bond yields (see appendix D). Both nominal and real government bonds are computed on the last trading day of the month.

3.1. 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.[14] 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 ith strike price, X_i, $i = 1, ..., I$, define the auxiliary variable

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

where the last equality holds by put-call parity, with $P_{d,t}^{(n)} = \Pi_t \sum_{j=1}^{n} E_t M_{t,t+j} D_{t+j}$ 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 strike prices are available for a given maturity, one can use the no-arbitrage relation to extract $P_{d,t}^{(n)}$ and

---

[13]Low liquidity during the first years after the TIPS market opened in 1997 likely introduced a wedge between TIPS risk premia and real bond premia (D'Amico, Kim and Wei, 2010). 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 (Fleckenstein, Longstaff and Lustig, 2014). 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.

[14]Golez (2014) further refines selection by considering only trades with positive volume or with an open interest of more than 200 contracts. Results are very similar if I add this additional filter.

$P_{b,t}^{(n)}$ as the least absolute deviations (LAD) estimators that minimize expression[15]

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

for a given trade date t and maturity n. The cross-sectional error term accounts for potential measurement error (e.g., because of bid-ask bounce, asynchronicities, or other microstructural 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 (see appendix C for an illustration). 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. In many instances, non-monotonicities in the auxiliary variable are concentrated in deep in- and out-of-the-money options. Whenever I spot non-monotonicities for low and high moneyness levels I restrict the sample to strikes with moneyness levels between 0.7 and 1.1 before running the cross-sectional LAD regression.

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 (2014), I then construct monthly prices using ten days of data at the end of each month. Figure 2a plots the prices of the synthetic dividend claims.
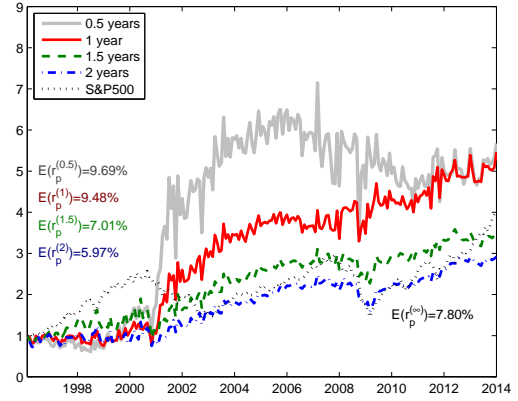
3.2. Mean

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,
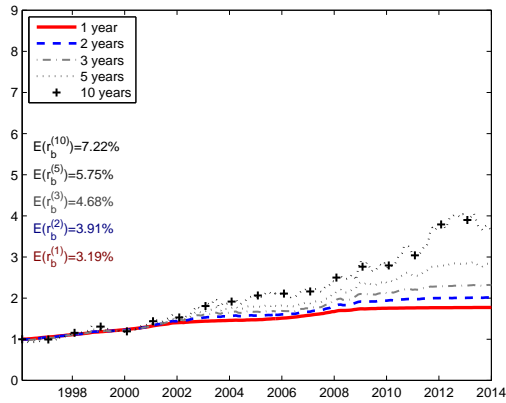
---

[15]A higher-order metric is inappropriate as it attaches much more weight to quotes that violate the LOOP.
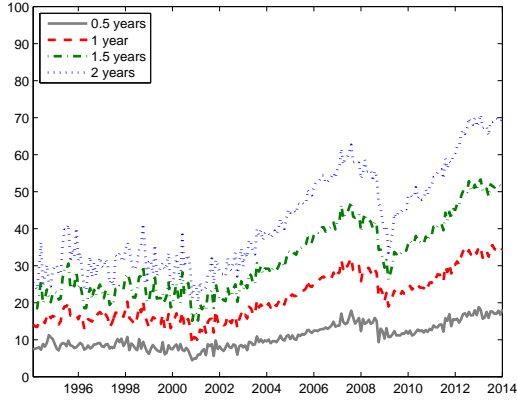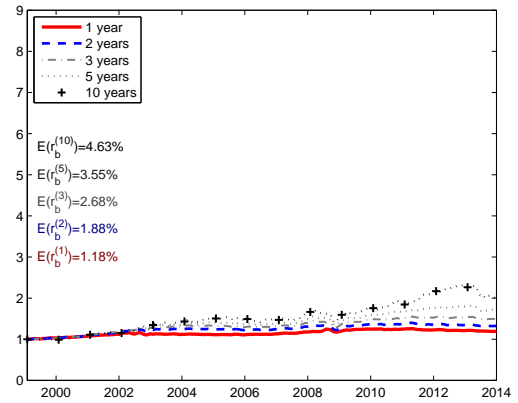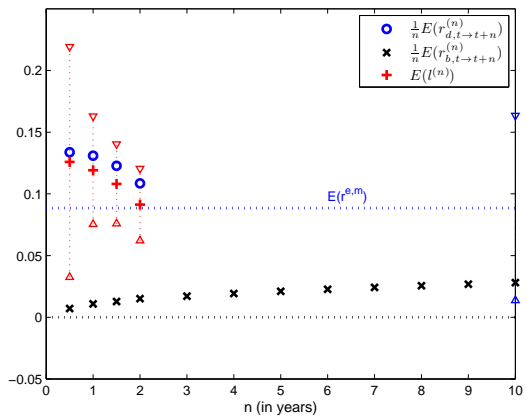
(a) *Price of the next n years of dividends.*

(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) and welfare costs (with block bootstrapped 95% confidence intervals).*
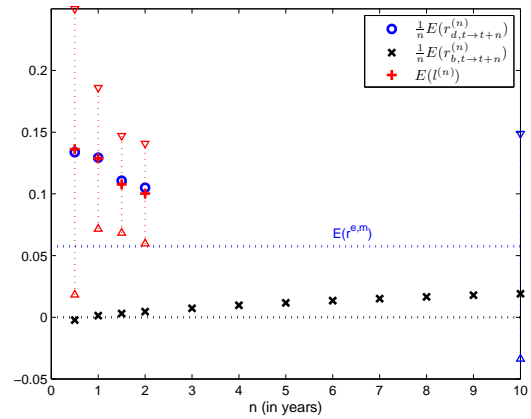
(f) *Average term structures of equity, interest rates (TIPS) and welfare costs (with block bootstrapped 95% confidence intervals), 2000-2013.*

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 use Gürkaynak-Sack-Wright data.

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.[16] 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, Carlson, Fisher and Simutin (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. The 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.

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, 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, in which 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 numbers compare with an average return on an annual buy-and-hold strategy on the market index of 10.2% over the 1994-2013 period and of

---

[16]Golez (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.

6.5% during 2000-2013.

Table 1 reports the point estimates for the term structure of welfare costs. The welfare costs of aggregate uncertainty around semestral cashflows 0.5, 1, 1.5 and 2 years out are of 12.6%, 11.9%, 10.8% and 9.1%, respectively. I then rely on proposition 1 to compute the costs of uncertainty around multi-period cashflows. Namely, I estimate average welfare costs of 11.8% associated with up to one-year ahead uncertainty, of 11.1% with up to 18-month uncertainty, and of 10.4% 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 welfare costs around semestral cashflows 0.5, 1, 1.5 and 2 years out are of 13.6%, 12.9%, 10.8% and 10.0%, respectively; welfare costs associated with up to 1-, 1.5- and 2-year ahead cashflow uncertainty are of 13.0%, 11.8% and 11.1%, 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).

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 9.0% (1994-2013) and 11.2% (2000-2013), whereas the cost of uncertainty over the next two years is of 7.1% and of 9.7%.

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 the 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 zero.

3.3. Volatility

The components of the term structure of welfare costs are volatile and procyclical. 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. In line with Binsbergen et al. (2013) and motivated by theoretical models such as Lettau and Wachter (2007, 2011) I

16

	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	0.1192	0.1081	0.0913	0.1362	0.1288	0.1077	0.1002
	(0.0325)	(0.0160)	(0.0120)	(0.0097)	(0.0408)	(0.0196)	(0.0134)	(0.0107)
	[0.0689]	[0.0303]	[0.0204]	[0.0180]	[0.0906]	[0.0412]	[0.0255]	[0.0243]
$L_t^{\{1,...,n\}}$								
Mean	0.1259	0.1184	0.1113	0.1036	0.1362	0.1299	0.1177	0.1110
	(0.0325)	(0.0190)	(0.0139)	(0.0107)	(0.0408)	(0.0239)	(0.0170)	(0.0129)
	[0.0689]	[0.0398]	[0.0276]	[0.0210]	[0.0906]	[0.0535]	[0.0375]	[0.0295]
$l_t^{(1)}$, 1 period = n semesters								
Mean	0.1259	0.0905	0.0782	0.0711	0.1362	0.1124	0.1032	0.0972
	(0.0325)	(0.0187)	(0.0126)	(0.0097)	(0.0408)	(0.0230)	(0.0154)	(0.0116)
	[0.0689]	[0.0444]	[0.0305]	[0.0243]	[0.0906]	[0.0571]	[0.0384]	[0.0306]
Equity premium, semestral								
Mean		0.0885				0.0575		
		(0.0169)				(0.0207)		
		[0.0398]				[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 to match the corresponding TIPS means 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. Boostrapped standard errors use block sizes of one (in parentheses) and ten (in brackets) observations.

	$r^{e,(1)}_{t\to t+1}$		$\frac{1}{2}r^{e,(2)}_{t\to t+2}$		$\frac{1}{3}r^{e,(3)}_{t\to t+3}$		$\frac{1}{4}r^{e,(4)}_{t\to t+4}$		$r^{e,m}_{t+1}$	
Constant	0.083	0.083	0.094	0.094	0.084	0.084	0.069	0.069	0.072	0.072
	[0.031]	[0.022]	[0.015]	[0.015]	[0.014]	[0.013]	[0.012]	[0.012]	[0.034]	[0.033]
$pc^1_{d,t}$	0.369	0.366	0.150	0.150	0.077	0.079	0.064	0.066		-0.094
	[0.100]	[0.056]	[0.028]	[0.032]	[0.024]	[0.023]	[0.019]	[0.017]		[0.056]
$pc^1_{b,t}$		0.095		-0.017		-0.067		-0.111		0.104
		[0.145]		[0.088]		[0.072]		[0.057]		[0.118]
$pc^2_{d,t}$		0.703		-0.063		-0.220		-0.152		
		[0.144]		[0.065]		[0.043]		[0.037]		
dp^m_t									0.345	0.318
									[0.152]	[0.194]
R^2	0.353	0.502	0.254	0.259	0.106	0.229	0.105	0.204	0.098	0.147
$\frac{std(E_t r)}{E(r)}$	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^1_d and pc^2_d, the first principal component of up to ten-year bond yields, pc^1_b, and the market dividend yield, dp^m. Monthly data, 1996m1-2013m12. Newey-West standard errors to correct for overlapping.

consider a one-factor specification for these risk premia. To capture the factor I extract the first principal component of the semestral equity yields, which captures 82% of their volatility, and use it to forecast the hold-to-maturity excess returns whose ex ante values measure the welfare costs. Motivated by theoretical considerations, I control for the presence of additional factors by including the second principal component of equity yields (these two principal components account for 95% of the variance in equity yields) and the first principal component of bond yields.[17]

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 understate 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 huge 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 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

---

[17]Appendix E shows how the model of Lettau and Wachter (2011) implies that these three variables are sufficient to reveal the factor that drives the term structure of welfare costs.

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.[18]
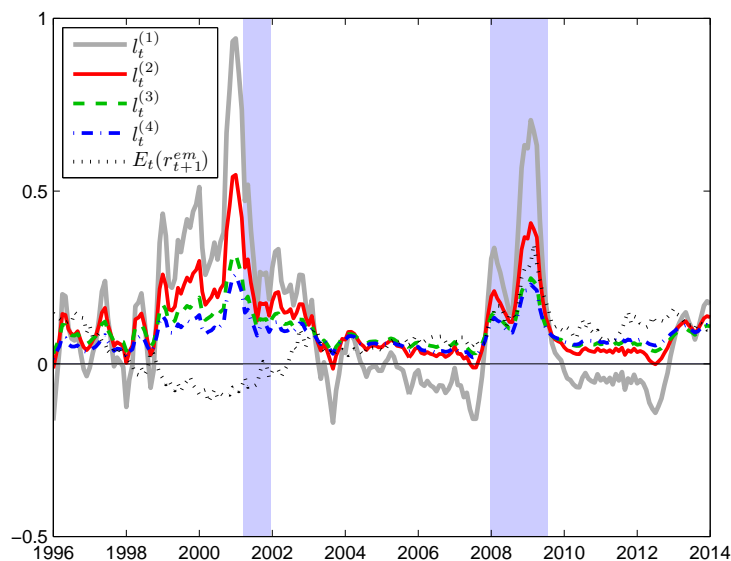

Figure 3: Semestral term structure of welfare costs and semestral equity premium, conditional on the information set spanned by equity and bond yields; annualized premia. Regressors are the first two principal components of semestral equity yields and the first principal component of up to ten-year bond 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).

4. The term structure in some consumption-based asset pricing models

Since the available sample allows for estimating only the first few components of the term structure of welfare costs at semestral frequency, I now turn to a model-based approach to capture and quantify the rest of the term structure.[19] The evidence in section 3 helps to identify a suitable model.

---

[18]When interpreting the huge costs of uncertainty in figure 3, we have to recognize the fact that consumption and market dividends differ. Consistent with section 5, a definition that accounts for the difference between consumption and market dividends scales down the costs of uncertainty in figure 3 by a factor, α.

[19]Such a model-based approach also offers a different way to account for the inflation risk premium in the calculation of welfare costs that mitigates concerns about the reliability of TIPS data.

4.1. Structural approach

Table 3 and figure 4 show the implications of some of today's 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. (2014), 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. Note that alternative calibrations of the models may capture the term structure evidence of section 3; however, I am not aware of such calibrations. The online 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.

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 4.4% 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 7.7% and an equity premium of 4.9%.

20

Croce et al. (2014) 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 therefore 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 8.8%, 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 fluctuations. 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 2.0% and an equity premium of about the same size.

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 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 than than short-duration equities, which produces a negative slope in the term structure of hold-to-maturity returns. The model implies a marginal cost of all fluctuations of 6.9% and a slightly larger equity premium.[20]

4.2. Quasi-structural approach

I turn to the quasi-structural exponential-Gaussian model of Lettau and Wachter (2011), which is designed to capture a downward-sloping term structure of equity and an upward-sloping term

---

[20]Recently, Wachter (2013) and Nakamura, Steinsson, Barro and Ursúa (2013) have made progress in modifying the rare-disaster framework to also account for the evidence by Binsbergen et al. (2012, 2013); however, their framework based on recursive utility has recently been challenged by Epstein et al. (2014), who show how it implies a cost of late resolution that seems implausible based on introspection.
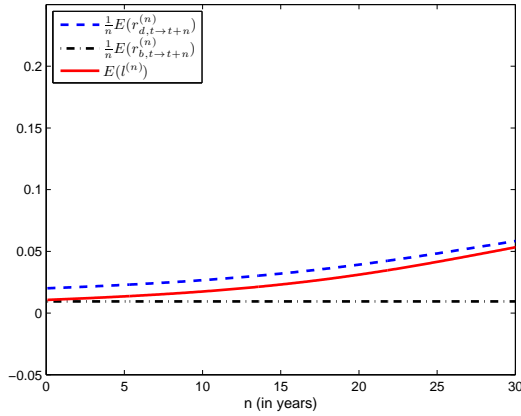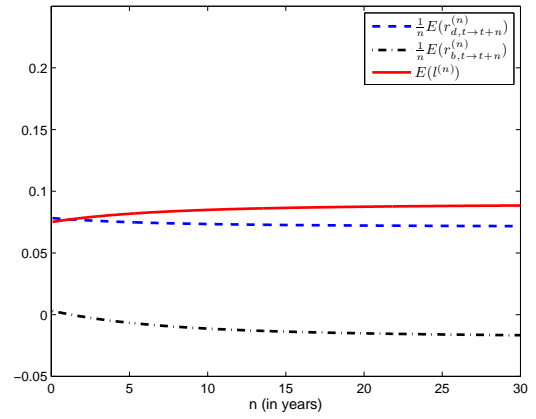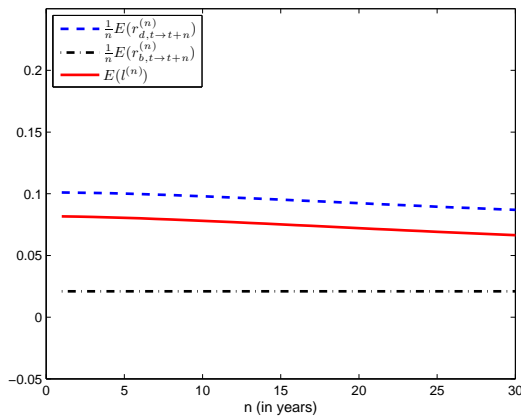
(a) *Habit formation (Campbell and Cochrane, 1999).*
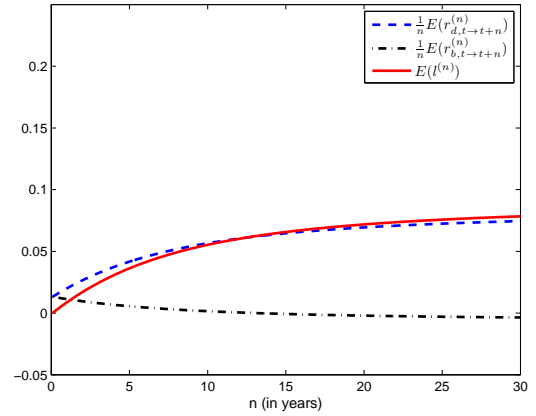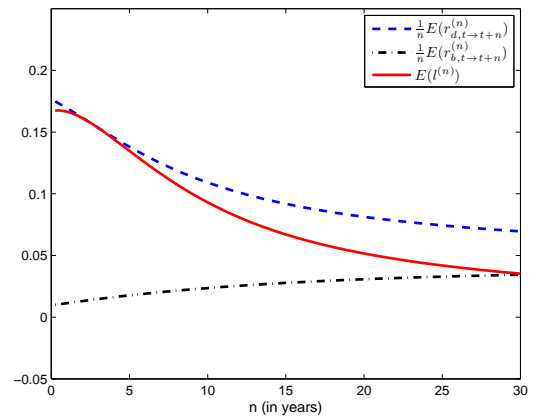
(b) *Long-run risk (Bansal and Yaron, 2004).*
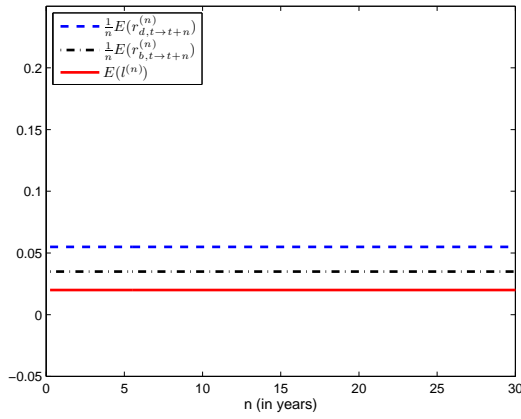
(c) *Ambiguity aversion (Barillas et al., 2009), Epstein-Zin-Weil log utility (Tallarini, 2000).*

(d) *Long-run risk (limited information) (Croce et al., 2014).*

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

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

Figure 4: The term structures of equity, interest rates and welfare costs of uncertainty in some consumption-based asset pricing models.

	$L^{\overline{\mathbb{N}}}$	$\ln E(R^{e,m})$
Campbell and Cochrane (1999)	4.42	6.81
Bansal and Yaron (2004)	7.68	4.90
Croce et al. (2014)	8.82	6.56
Barillas et al. (2009)	2.00	1.92
Gabaix (2012)	6.87	7.89
Lettau and Wachter (2011)	2.79	7.18

Table 3: Mean marginal cost of lifetime uncertainty and equity premium (in percent per year) for different consumption-based asset pricing models.

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 load of long-run dividend strips on the state that drives the price of risk offsets the long-run insurance effect (see also Lettau and Wachter, 2007).

Finally, since only short-run cashflow shocks are priced, Lettau and Wachter (2011) 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 2.8% and an equity premium of 7.2%. Table 4 reports the cost of short- and long-run fluctuations over different coordinate sets as captured by the quasi-structural model of Lettau and Wachter (2011). An increase in consumption uncertainty by a fraction θ over a ten-year period has a marginal cost of more than 12θ 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, reported in figure 5a, 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

$\mathcal{N}$	$L^{\mathcal{N}}$	$L^{\overline{\mathbb{N}}\backslash\mathcal{N}}$
up to 1 year	16.70	2.77
up to 2 years	16.44	2.75
up to 3 years	16.04	2.71
up to 5 years	15.00	2.61
up to 10 years	12.31	2.35
up to 20 years	8.67	1.93
$\overline{\mathbb{N}}$	2.79	0

Table 4: Marginal cost of fluctuations at all periodicities $n \in \mathcal{N}$. Lettau and Wachter (2011) model-based estimates

time-varying market price of risk, as shown in appendix E. Figure 5b shows the impulse-response of the term structure of the cost of uncertainty after a one standard deviation discount-rate shock. News to the price of risk forecast a transitory change in the state of the economy that makes people temporarily increase the price of risk, and hence the cost of uncertainty. The positive initial effect then decays across maturities and over time.
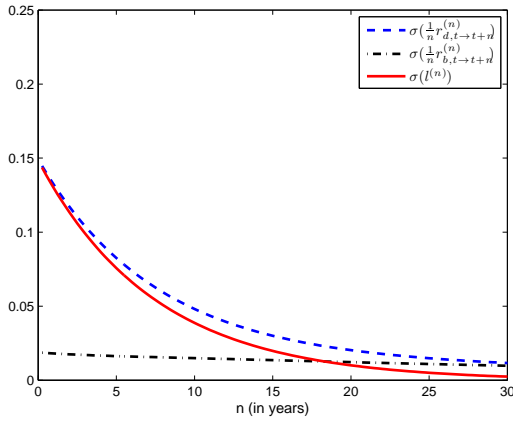
(a) *Model-based estimates of volatilities in the model of Lettau and Wachter (2011).*
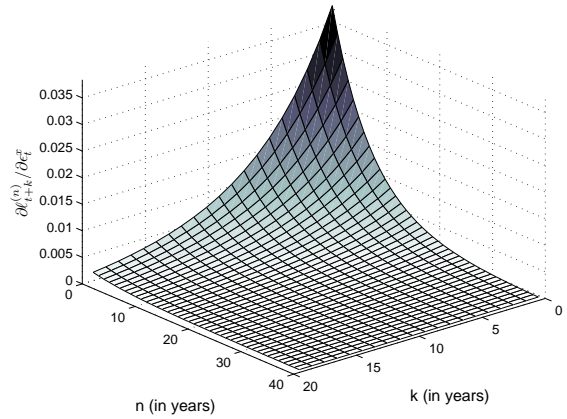
(b) *Impulse response $\partial \ell_{t+k}^{(n)}/\partial \varepsilon_t^x$ to a one-standard deviation discount-rate shock.*

Figure 5: Time-variation in the term structure of welfare costs in the model of Lettau and Wachter (2011).

5. Robustness

This section examines what features of the cost of uncertainty are robust across models as well as it relaxes two important modeling assumptions—the theoretical equality between consumption and market dividends and the existence of a representative agent.

5.1. Robustness across models

Even though I can only turn to the quasi-structural model of Lettau and Wachter (2011) to capture the term structure features I am interested in, I can draw some lessons that are robust across all asset pricing models considered.

Table 3 shows how the marginal cost of uncertainty in the entire consumption process is above two percent in each of today's leading asset pricing models. This robustness across models is in line with the result by Alvarez and Jermann (2004) that a model that is consistent with the main stylized asset pricing facts must increase the original estimates by Lucas (1987) by two orders of magnitude.

5.2. Market equity or consumption equity?

The definition of cost of fluctuations is in terms of consumption equity. The evidence I consider is in terms of market equity. These two notions of equity may not perfectly substitute. In fact, along with the choice of the moving-average filter, the other main empirical challenge of Alvarez and Jermann (2004) is to find a proxy for the price of a claim to the entire consumption stream.

The critique does not bite in an endowement economy, in which $C_t = D_t$, so the theoretical definition of cost of uncertainty can be stated in terms of dividends in the first place. However, in a production economy the equality breaks down. Fortunately, under mild general equilibrium assumptions one can save the link between the cost of fluctuations and market equity.

A natural approach in a production economy is to consider preferences that depend also on labor effort, N_t. Moreover, to model theoretically a difference between consumption and dividends I rely on the following assumption.

Assumption 1. The condition
$$W_t = -\frac{\partial U_t/\partial N_t}{\partial U_t/\partial C_t}$$
describes the consumption-labor tradeoff, i.e., the marginal rate of substitution between consumption and labor equals the relative price—the real wage rate, W_t. Market dividends correspond to the period profits of the aggregate firm, $D_t = Y_t - W_t N_t - I_t$, where Y_t is total output and I_t are gross investments made by the firm. The market-clearing condition $Y_t = C_t + I_t$ is satisfied.

The definition of market dividends as current aggregate profits is standard in the Q literature (e.g., Kogan and Papanikolaou, 2012). The optimality condition is standard (e.g., Swanson, 2012) and implies that there are no distortions in the consumption-side of the economy that generate so-called labor wedges (such as in Galí et al., 2007).

Definition (Marginal cost of uncertainty, production economy). Define the cost of fluctuations as

$$E_t U\Big(\{(1 + \mathcal{L}_t^N(\theta))^n C_{t+n}\}, \{N_{t+n}\}, \{X_{t+n}\}\Big) =$$
$$= E_t U\Big(\{\theta E_t C_{t+n} + (1 - \theta)C_{t+n}\}, \{\theta\overline{N}_{t+n} + (1 - \theta)N_{t+n}\}, \{X_{t+n}\}\Big) \quad (5)$$

where stable hours worked $\overline{N}_t$ are defined as $W_{t+n}\overline{N}_{t+n} \equiv E_t(W_{t+n}N_{t+n})$, so stable hours provide a stable labor income. In a production economy, $\mathcal{L}_t$ measures the cost of aggregate uncertainty around consumption and labor income, i.e., how much people would pay to hedge against uncertainty in consumption *and* in labor income.

Proposition 1′. *Under assumption 1, the marginal cost of uncertainty around all coordinates $n \in \mathcal{N}$ is*

$$L_t^N = \alpha_t \sum_{n\in\mathcal{N}} \omega_{n,t} l_t^{(n)}$$

with weights $\{\omega_{n,t}\}$ and term structure components $\{l_t^{(n)}\}$ defined by

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

with a scaling factor $\alpha_t = [\sum_{n\in\mathcal{N}} n\, E_t(M_{t,t+n}D_{t+n})]/[\sum_{n\in\mathcal{N}} n\, E_t(M_{t,t+n}C_{t+n})]$.

In this production economy, as I distinguish between market equity and consumption equity, only the maturity-independent scaling factor α_t depends on consumption equity.[21] The weights and the term structure components remain functions of market equities, which therefore exclusively control the slope of the term structure of welfare costs. Via the scaling factor, an unobservable consumption-equity component appears in the level of the cost of uncertainty, L_t^N. However, the downward slope in the term structure is robust to the departure from the equality between market and consumption equity.

---

[21] For example, under a frictionless production side that operates with a Cobb-Douglas production function $Y_t = K_t^\alpha(A_t N_t)^{1-\alpha}$, $\alpha \in (0, 1)$, with technology A and capital K, equilibrium dividends as defined in section 3 relate with consumption as $D_t = \alpha C_t$, so the scaling factor is $\alpha_t = \alpha < 1$. In this case, the term structure of welfare costs estimated from market equity data would overstate the level of welfare costs by a factor $1/\alpha$.

The term structure of marginal costs of uncertainty remains well-defined even in a heterogeneous-agent incomplete-market setting, by an argument analogous to one already made by Alvarez and Jermann (2004). Consider agents with heterogeneous preferences, U_i, and an idiosyncratic consumption component, $\varepsilon_{i,t}$, driving their consumption stream, $C_{i,t} = C_t + \varepsilon_{i,t}$. Agents can only trade the entire term structures of dividend strips and zero-coupon bonds and therefore face possibly uninsurable idiosyncratic risk.

Definition (Marginal cost of uncertainty, heterogenenous-agent incomplete-market economy)**.** Define the cost of fluctuations as

$$E_t U_i\Big(\{(1 + \mathcal{L}_{i,t}^N(\theta))^n C_{t+n} + \varepsilon_{i,t+n}\}\Big) = E_t U_i\Big(\{\theta(E_t C_{t+n} + \varepsilon_{i,t+n}) + (1 - \theta)C_{i,t+n}\}\Big)$$

so $\mathcal{L}_t$ measures the cost of uncertainty around the systematic part of the ith agent's consumption, i.e., how much people would pay to hedge against uncertainty in their systematic consumption component.

By the absence of arbitrage opportunities, the projection of their marginal rates of substitution on the payoff space is the same across people, so their valuations of available asset prices are equal. Since all agents have access to the entire term structures of dividend strips and zero-coupon bonds, they end up equalizing their valuation of welfare costs. Proposition 1″ formalizes this result.

Proposition 1″. *The marginal cost of uncertainty around all coordinates $n \in \mathcal{N}$ for agent i is*

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

and is constant across agents.

6. Conclusion

Lucas (1987, 2003) introduced the notion of cost of aggregate uncertainty as a thought experiment to provide an assessment of the tradeoff between growth and macroeconomic stability. Similarly, the term structure of the marginal cost of uncertainty requires little structure to reveal the tradeoff between growth and macroeconomic stability at different time horizons; as long as such an insurance is actually traded the question ceases to be a simple thought experiment detached from the actual tradeoffs investors face. Recent derivative securities allow for synthesizing the required insurance and, therefore, provide 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). The finding of large 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 et al., 2012, 2013; Berg, 2010; Duffee, 2010; Aït-Sahalia et al., 2012; Palhares, 2013).

Asset markets suggest that marginal increases in aggregate uncertainty are particularly costly, especially in the short run and during downturns such as in the early 2000s and during the recent financial crisis. However, the term structure of welfare costs can vary substantially across the business cycle with movements that are revealed by innovations in the information set made by excess return predictors.

The result that the welfare 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 of the estimated term structure of welfare costs (which rests on a downward-sloping term structure of equity *and* on an upward-sloping term structure of interest rates) cannot be easily captured by today's leading consumption-based asset pricing models and therefore represents a puzzling piece of evidence with seemingly crucial welfare consequences. This evidence is diagnostic more than conclusive as the question of whether the volatility level is Pareto optimal remains open. Accordingly, a structural model able to rank different stabilization policies while explaining the stylized evidence about the term structures of welfare costs should be high on the research agenda.

Appendix

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

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

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

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 allows for a direct focus on the relevant periodicity of economic fluctuations.

B. Proofs

Proof of proposition 1. I can rewrite equation (2) as

$$
\begin{aligned}
L_t^{\mathcal{N}} &= \frac{\sum_{n \in \mathcal{N}} E_t(M_{t,t+n}) E_t(C_{t+n}) - E_t(M_{t,t+n} C_{t+n})}{\sum_{n \in \mathcal{N}} n \, E_t(M_{t,t+n} C_{t+n})} \\
&= \sum_{n \in \mathcal{N}} \frac{n \, E_t(M_{t,t+n} C_{t+n})}{\sum_{n \in \mathcal{N}} n \, E_t(M_{t,t+n} C_{t+n})} \times \frac{1}{n}\Big(\frac{E_t(M_{t,t+n}) E_t(C_{t+n})}{E_t(M_{t,t+n} C_{t+n})} - 1\Big) \\
&= \sum_{n \in \mathcal{N}} \omega_{n,t} l_t^{(n)}
\end{aligned}
$$

The absence of arbitrage opportunities guarantees that the weights are positive.

Proof of proposition 2. The proposition follows directly from the expression of the term structure components, $l_t^{(n)}$, and the definition of the hold-to-maturity return on an arbitrary payoff, X, maturing in n periods,

$$
R_{x,t \to t+n}^{(n)} = \frac{X_{t+n}}{D_{x,t}^{(n)}}
$$

with no-arbitrage price $D_{x,t}^{(n)} = E_t(M_{t,t+n} X_{t+n})$ at period t, where consumption equity is characterized by $X = C$ and real bonds by $X = 1$.

Proof of proposition 1'. Differentiating definition (5) with respect to θ, it follows that

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

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 1, equation (B) becomes

$$
\begin{aligned}
L_t^{\mathcal{N}} &= \frac{\sum_{n \in \mathcal{N}} E_t(M_{t,t+n}) E_t(D_{t+n}) - E_t(M_{t,t+n} D_{t+n})}{\sum_{n \in \mathcal{N}} n \, E_t(M_{t,t+n} C_{t+n})} \\
&= \alpha_t \sum_{n \in \mathcal{N}} \omega_{n,t} l_t^{(n)}
\end{aligned}
$$

C. Errors in synthetic replication

Figure A 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 A, the typical map towards the end of the sample is virtually perfectly linear and monotonic as we move 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 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.

D. Accounting for the inflation risk premium

In the empirical section, I compute the excess hold-to-maturity returns whose expectations equal the cost of uncertainty relative to the term structure of interest rates extracted from nominal government bonds as a proxy for the term structure of real interest rates. Accordingly, consider the empirical hold-to-maturity excess returns under no-arbitrage pricing

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

with the inflation rate $\Pi_{t,t+n} \equiv \Pi_{t+n} / \Pi_t$ 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.[22] It follows that

$$
\ell_t^{(n)} = \frac{1}{n} E_t \psi_{t,t+n} + \frac{1}{n} irp_t^{(n)} + \frac{1}{2n} \mathcal{V}_t(\Pi_{t,t+n})
$$

so the welfare cost measure is biased by the presence of the inflation risk premium.[23] As shown by Ang, Bekaert and Wei (2008) the Jensen's term $\frac{1}{2}\mathcal{V}_t(\Pi_{t,t+n})$ can be safely disregarded (see also Grishchenko and Huang, 2013). During the 2000-2013 period, over which TIPS data are available,

---

[22]In a Gaussian setting the inflation risk premium reduces to the more familiar definition as the covariance between inflation and the nominal discount factor, $irp_t^{(n)} = cov_t(m_{t,t+n} - \pi_{t,t+n}, \pi_{t,t+n})$ (e.g., Ang et al., 2008).

[23]Note that the maturity-specific risk premium studied by Binsbergen et al. (2013) is $\frac{1}{n} E_t \psi_{t,t+n}$.
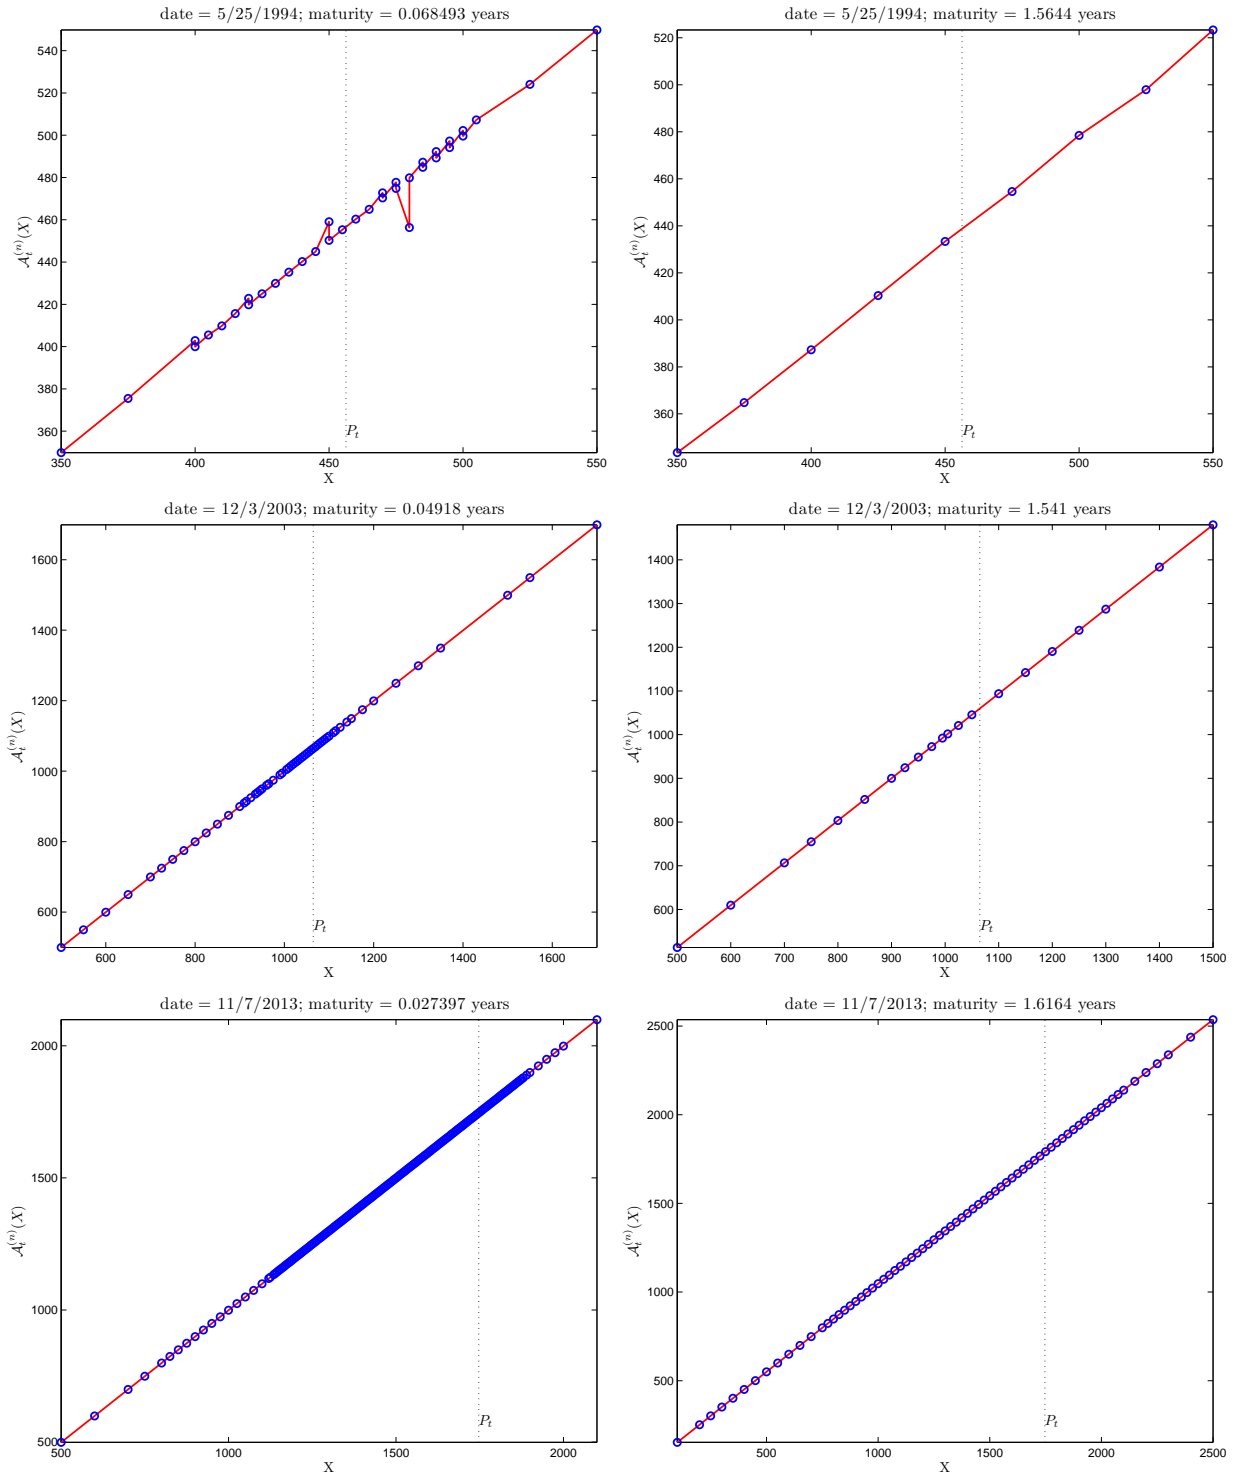
Figure A: Auxiliary map $\mathcal{A}_t^{(n)} : X \mapsto P_{d,t}^{(n)} + X P_{b,t}^{(n)}$ on given trading days and for given maturities. P_t is the value of the index at each respective trading day.
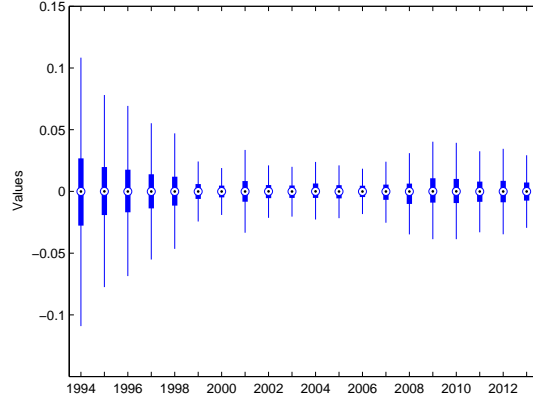
Figure B: Law-of-one-price violations, in percent of the associated options implied dividend claim price; only observations not filtered out by my data selection approach are included.

I compute average inflation risk premia as

$$E(irp_t^{(n)}) = E(\ln P_{0,t}^{(n)}) - E(\ln P_{b,t}^{(n)}) - E(\pi_{t,t+n}) - \frac{1}{2}E[\mathcal{V}_t(\Pi_{t,t+n})]$$

where $P_{0,t}^{(n)}$ denotes the price of the nth real zero-coupon bond, and correct welfare costs accordingly. To correct for the inflation term I use CPI inflation (BEA-NIPA database).

When TIPS are available, the empirical hold-to-maturity excess returns are

$$\frac{\Pi_{t+n}D_{t+n}E_t M_{t,t+n}}{\Pi_t E_t M_{t,t+n}D_{t+n}} = \frac{\Pi_{t+n}}{\Pi_t}R_{t\to t+n}^{e,(n)}$$

by the nominality of observed dividends. Accordingly, I rely on CPI inflation to correct the computed hold-to-maturity returns.

E. The term structure of welfare costs in the model of Lettau and Wachter

The variation of the term structure of welfare costs over the business cycle is driven by the time-varying components of $\{\ell_t^{(n)}\}$, where $\ell_t^{(n)} \equiv \ln(1 + nl_t^{(n)})^{1/n}$ is the continuously compounded marginal cost of uncertainty associated with $l_t^{(n)}$. In the model of Lettau and Wachter (2011), the online appendix shows how this component is driven entirely by movements in the market price of risk x_t as, up to a constant term,

$$\ell_t^{(n)} = \frac{1}{n}\left[\frac{1-\phi_x^n}{1-\phi_x}\sigma_d + \left(\frac{1-\phi_x^n}{1-\phi_x} - \frac{\phi_x^n - \phi_z^n}{\phi_x - \phi_z}\right)\frac{\sigma_z}{1-\phi_z}\right]\frac{\sigma_d'}{\|\sigma_d\|}x_t$$

32

where the price of risk follows the autoregressive process $x_{t+1} = (1 - \phi_x)x + \phi_x x_t + \sigma_x \varepsilon_{t+1}$, with x the average discount rate. Vectors σ'_d, σ'_z and σ'_x are respectively the loadings of short-run and long-run cashflows and of the price of risk on the reduced-form shocks that drive the system, and ϕ_z is the persistence of the predictable component of cashflows, z_t. The term structure of the costs of uncertainty starts from a level of $\|\sigma_d\|x$, which is about 17 percent per year in the calibration of Lettau and Wachter, to then decay with a slope determined by the persistence coefficients ϕ_z and ϕ_x.

The model suggests that a sufficient information set to reveal news to the market price of risk is made by the first two dividend yields, $e_t^{(1)}$ and $e_t^{(2)}$, and the first bond yield, $y_t^{(1)}$. In fact,

$$e_t^{(1)} = -z_t + \|\sigma_d\|x_t + r_t$$

$$e_t^{(2)} = -\frac{1 + \phi_z}{2}z_t + \frac{1 + \phi_r}{2}r_t + \frac{1}{2}\left(1 + \phi_x + \frac{(\sigma_z - \sigma_r)\sigma'_d}{\|\sigma_d\|^2}\right)\|\sigma_d\|x_t$$

$$y_t^{(1)} = r_t$$

where the variables are in log-deviations from the mean and r_t is the exogenous state that drives bond yields with autocorrelation coefficient ϕ_r. Therefore, the projection of hold-to-maturity excess strip returns on the information set induced by the history of the observable vector $[e_t^{(1)}; e_t^{(2)}; y_t^{(1)}]$ is the true cost of uncertainty $\ell_t^{(1)} = \|\sigma_d\|x_t$. The residuals of the Wold representation of the information set then reveal discount-rate news.

References

Alvarez, F., Jermann, U.J., 2004. Using asset prices to measure the cost of business cycles. Journal of Political Economy 112, 1223–56.

Ang, A., Bekaert, G., Wei, M., 2008. The term structure of real rates and expected inflation. Journal of Finance 63, 797–849.

Aït-Sahalia, Y., Karaman, M., Mancini, L., 2012. The term structure of variance swaps, risk premia and the expectations hypothesis. Manuscript.

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 forthcoming.

Baker, S.R., Bloom, N., Davis, S., 2013. Measuring economic policy uncertainty. Manuscript.

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.

Bekaert, G., Hoerova, M., Lo Duca, M., 2013. Risk, uncertainty and monetary policy. Journal of Monetary Economics 60, 771–88.

Berg, T., 2010. The term structure of risk premia: New evidence from the financial crisis. ECB working paper series 1165.

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 forthcoming.

Bloom, N., 2009. The impact of uncertainty shocks. Econometrica 77, 623–85.

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 forthcoming.

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.

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

Croce, M.M., Lettau, M., Ludvigson, S.C., 2014. Investor information, long-run risk, and the term structure of equity. Review of Financial Studies forthcoming.

D'Amico, S., Kim, D.H., Wei, M., 2010. Tips from TIPS: The informational content of Treasury Inflation-Protected Security prices. Federal Reserve Board Finance and Economics Discussion Series.

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., 2014. Asset pricing in the frequency domain: Theory and empirics. Manuscript.

Duffee, G.R., 2010. Sharpe ratios in term structure models. Manuscript.

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 forthcoming.

Fernández-Villaverde, J., Guerrón-Quintana, P., Rubio-Ramírez, J.F., Uribe, M., 2011. Risk matters: The real effects of volatility shocks. American Economic Review 101, 2530–61.

Fleckenstein, M., Longstaff, F.A., Lustig, H., 2014. The TIPS-Treasury bond puzzle. Journal of Finance forthcoming.

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., 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 forthcoming.

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., 2007. The U.S. treasury yield curve: 1961 to the present. Journal of Monetary Economics 54, 2291–2304.

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.

Ilut, C., Schneider, M., 2014. Ambiguous business cycles. American Economic Review forthcoming.

Kogan, L., Papanikolaou, D., 2012. Economic activity of firms and asset prices. Annual Review of Financial Economics 4, 361–84.

Krusell, P., Smith, A.A., 1999. On the welfare effects of eliminating business cycles. Review of Economic Dynamics 2, 245–72.

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.

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.

Nakamura, E., Steinsson, J., Barro, R., Ursúa, J., 2013. Crises and recoveries in an empirical model of consumption disasters. American Economic Journal: Macroeconomics 5, 35–74.

Otrok, C., 2001. On measuring the welfare cost of business cycles. Journal of Monetary Economics 47, 61–92.

Palhares, D., 2013. Cash-flow maturity and risk premia in CDS markets. Manuscript.

Swanson, E.T., 2012. Risk aversion and the labor margin in dynamic equilibrium models. American Economic Review 102, 1663–91.

Tallarini, T.D.J., 2000. Risk-sensitive real business cycles. Journal of Monetary Economics 45, 507–32.

Wachter, J.A., 2013. Can time-varying risk of rare disasters explain aggregate stock market volatility? Journal of Finance 68, 987–1035.

Documents de Travail

510. G. Bazot, M. D. Bordo and E. Monnet, "The Price of Stability. The balance sheet policy of the Banque de France and the Gold Standard (1880-1914) ," October 2014

511. K. Istrefi and A. Piloiu, "Economic Policy Uncertainty and Inflation Expectations," October 2014

512. C. Jardet and A. Monks, "Euro Area monetary policy shocks: impact on financial asset prices during the crisis?," October 2014

513. J. Dugast and T. Foucault, "False News, Informational Efficiency, and Price Reversals," October 2014

514. G. Cette, J. Lopez and J. Mairesse, "Product and Labor Market Regulations, Production Prices, Wages and Productivity," October 2014

515. L. Ferrara and C. Marsilli, "Nowcasting global economic growth: A factor-augmented mixed-frequency approach," October 2014

516. F. Langot and M. Lemoine, "Strategic fiscal revaluation or devaluation: why does the labor wedge matter? ," October 2014

517. J.-P. Renne, "Fixed-Income Pricing in a Non-Linear Interest-Rate Model," October 2014

518. J.-P. Renne, "Options Embedded in ECB Targeted Refinancing Operations," October 2014

519. A. Penalver, "Pre-crisis credit standards: monetary policy or the savings glut?," November 2014

520. C. Marsilli, "Variable Selection in Predictive MIDAS Models," November 2014

521. P. Lopez, "The Term Structure of the Welfare Cost of Uncertainty," November 2014