ISSN 0956-8549

The Market Cost of Business Cycle Fluctuations

By

Anisha Ghosh
Christian Julliard
Michael J. Stutzer

DISCUSSION PAPER NO 898

Revised version January 2026

The Market Cost of Business Cycle Fluctuations *

Anisha Ghosh[†] Christian Julliard[‡] Michael J. Stutzer[§]

June 26, 2024

Abstract

We propose a novel approach to measure the cost of aggregate economic fluctuations, that does not require complete specification of investors' risk preferences or their beliefs. With data on consumption and asset prices, an information-theoretic method is used to recover an *information kernel* (I-SDF). The I-SDF accurately prices broad cross-sections of assets, thereby offering a reliable candidate for the measurement of the welfare cost of business cycles. Our method enables the estimation of both the unconditional (or, average) cost of fluctuations as well as the cost conditional on each possible economic state. We find that the cost of fluctuations is strongly time-varying and countercyclical and that the cost of business cycle fluctuations is substantial, accounting for a quarter to a third of the cost of all consumption uncertainty.

Keywords: Aggregate Uncertainty, Business Cycle Risk, Pricing Kernel, Empirical Likelihood, Smoothed Empirical Likelihood.

JEL Classification Codes: JEL: E3, E2, G12, C5.

*We benefited from helpful comments from Jaroslav Borovicka, George Constantinides, Janice Eberly, Deborah Lucas, Sydney Ludvigson, and seminar participants at the NBER's Summer Institute Asset Pricing Workshop and the Capital Markets and the Economy (EFEL) Workshop. All errors and omissions are our own.

[†]Desautels Faculty of Management, McGill University; anisha.ghosh@mcgill.ca

[‡]Department of Finance, London School of Economics; c.julliard@lse.ac.uk

[§]Leeds School of Business, University of Colorado, Boulder; michael.stutzer@colorado.edu.

I Introduction

In his seminal 1987 monograph, Robert E. Lucas Jr. concludes that the welfare benefit of eliminating *all* consumption fluctuations in the U.S. economy is trivially small, hence challenging the desirability of policies aimed at insulating the economy from cyclical fluctuations. As Lucas emphasizes,[1] this result is obtained without taking a stand on the origins of aggregate fluctuations, and it relies solely on the specifications of preferences (a representative agent with time and state separable power utility preferences with a constant coefficient of relative risk aversion) and the data generating process (i.i.d. log-normal aggregate consumption growth rate).

Nevertheless, it is exactly these two assumptions that make Lucas' calculations questionable. This is because evaluating the welfare cost of business cycles is tantamount to *pricing* the risk that households face due to aggregate fluctuations. And, an extensive literature has documented how Lucas' specification grossly underestimates the market price of risk in the U.S. economy: e.g., the average premium on a broad U.S. stock market index over and above short-term Treasury Bills has been about 7% per year over the last century, while Lucas' specification would imply a premium of less than 1%.[2] Lucas' specification also fails to explain the significant cross-sectional differences in average returns between different asset classes (see e.g., Hansen and Singleton (1983), Lettau and Ludvigson (2001), Parker and Julliard (2005), Julliard and Ghosh (2012)).

Indeed, exactly due to the inability of the power utility with log-normal shocks to match households' preferences toward risk revealed by the prices of financial assets, a burgeoning literature, based on modifying the preferences of investors and/or the dynamic structure of the economy, has developed. In these models, the pricing kernel (hereafter referred to as the Stochastic Discount Factor or SDF) can be factored into an observable component consisting of a parametric function of consumption growth, as with power utility, and a model-specific component. That is, the pricing kernel, M, in these models is of the form:

$$M_{t+1} = (C_{t+1}/C_t)^{-\gamma} \, \psi_{t+1}. \tag{1}$$

The Lucas (1987) original setting is nested within this family, corresponding to the case in which ψ_t is a positive constant. Prominent examples of models in this class are: habit formation models (see, e.g., Campbell and Cochrane (1999), Menzly, Santos, and Veronesi (2004)); long run risks models (e.g., Bansal and Yaron (2004)); models with complemen-

[1] *"these calculations rest on assumptions about preferences only, and not about any particular mechanism equilibrium or disequilibrium – assumed to generate business cycles"*, Lucas (1987).

[2] This discrepancy is the so-called Equity Premium Puzzle, first identified by Mehra and Prescott (1985).

tarities in consumption (e.g., Piazzesi, Schneider, and Tuzel (2007), Yogo (2006)); models in which ψ_t captures aggregation over heterogeneous agents who face uninsurable idiosyncratic shocks to their labor income (e.g. Constantinides and Duffie (1996), Constantinides and Ghosh (2017)), as well as solvency constraints (e.g. Lustig and Nieuwerburgh (2005)). While the above models are all based on rational expectations, the multiplicative decomposition of the pricing kernel in Equation (1) also encompasses behavioral models – in such cases, the ψ_t captures deviations from rational expectations (e.g. Basak and Yan (2010), Hansen and Sargent (2010)).

Estimates of the cost of business cycles vary widely across these model specifications (see, e.g., Barlevy (2005) for a survey). Moreover, as with Lucas' original specification, in order for any of the more recent models to constitute a good choice for welfare cost calculations, it should accurately price broad categories of assets and, therefore, reflect agents' attitude towards risk. Ghosh, Julliard, and Taylor (2017) evaluate the empirical pricing performance of several of these models and show that they perform quite poorly, producing large pricing errors and low (and often negative) cross-sectional R^2. Therefore, the shortcomings of using Lucas' specification for welfare cost calculations also apply to the more recent advances.

In this paper, we do not take a stand on either the full specification of investors' preferences, or on the true dynamics of the underlying state variables, or on the latter dynamics as perceived by the investor, i.e. her beliefs. Importantly, as we show formally, our approach does not require investors' beliefs to be rational. The method relies on the insight that asset prices contain information about the stochastic discounting of the different possible future states and, therefore, use observed asset prices to recover the SDF. Specifically, we assume that the underlying SDF has the multiplicative form in Equation (1). We use asset returns and consumption data to extract, non-parametrically, the *minimum relative entropy* estimate of the ψ-component of the pricing kernel M such that the resultant M satisfies the unconditional Euler equations for the assets, i.e. successfully prices the cross-section of assets. This information-theoretic approach, that has its origins in the physical sciences, adds to the standard power utility kernel the *minimum* amount of additional information needed to price assets perfectly, i.e. satisfy the Euler equations. We refer to the estimated M as the *information SDF* (I-SDF) because of the information-theoretic methodology used to recover it. In the absence of knowledge of the true SDF, our framework offers a more robust approach to identifying it, while incorporating the central economic insight that aggregate consumption risk represents an important source of priced risk.

With this I-SDF at hand, we obtain the cost of aggregate consumption fluctuations as the ratio of the (shadow) prices of two hypothetical securities – a claim to a *stabilized* version of the aggregate consumption stream from which certain types of fluctuations (e.g., all fluctuations or fluctuations corresponding to business cycle frequencies only) have been removed, and a claim to the actual aggregate consumption stream. Alvarez and Jermann (2004) show that, in the context of an infinite horizon representative agent economy, the above ratio measures the *marginal cost* of consumption fluctuations, defined as the per unit benefit of a marginal reduction in consumption fluctuations, expressed as a percentage of lifetime consumption. Our approach allows us to estimate the term structure of the cost of fluctuations, i.e. how the cost (or, the welfare benefit of removing fluctuations) rises with the elimination of aggregate fluctuations over each additional future period.

Our information-theoretic approach to the recovery of the SDF corresponds to the Empirical Likelihood (EL) estimator of Owen (2001) and the Exponential Tilting (ET) estimator of Kitamura and Stutzer (1997). Using this methodology to recover the (multiplicative) missing component of the SDF was originally proposed in Ghosh, Julliard, and Taylor (2017). The I-SDF, unlike Lucas' original specification, accurately prices broad cross-sections of assets.[3] It, therefore, offers a more reliable choice for assessing investors' attitude toward risk.

We first apply our methodology to assess the welfare benefits of economic stabilization *on average*, i.e. averaged across all possible states. We find that the cost of business cycle fluctuations in consumption is large and constitutes between a quarter to a third of the cost of *all* consumption fluctuations. For instance, in our baseline 1929–2015 sample, when the I-SDF is extracted using nondurables and services consumption with the excess return on the market portfolio as the sole asset and a utility curvature parameter $\gamma = 10$ in Equation (1), the cost of all fluctuations over a five-year horizon is estimated at 14.4% (11.9%), while the corresponding cost of business cycle fluctuations is 3.6% (3.1%) with the EL (ET) approach. When total (instead of nondurables and services) consumption expenditures is used to recover the I-SDF, the costs of all fluctuations and business cycle fluctuations over a 5-year period are both estimated to be even higher at 19.7% (17.3%) and 5.1% (4.6%), respectively, with the EL (ET) approach. The corresponding costs obtained with Lucas' specification are typically an order of magnitude smaller. These conclusions are robust to the set of test assets used to recover the I-SDF. Our results suggest that economic agents perceive the cost of aggregate fluctuations to be substantial and that business cycles constitute a substantial proportion of this cost.

We next rely on an extension of our methodology – specifically, the Smoothed Empirical Likelihood (SEL) estimator of Kitamura, Tripathi, and Ahn (2004) and the Smoothed Exponential Tilting (SET) estimator (see, e.g., Ghosh and Otsu (2020)) that extend the EL and

[3]See also Ghosh, Julliard, and Taylor (2022) who show that the I-SDF, estimated in a purely out-of-sample fashion, accurately prices the aggregate stock market, broad cross-sections of equity portfolios constructed by sorting stocks on the basis of different observable characteristics (e.g., size, book-to-market-equity, prior returns, industry), as well as currency portfolios and portfolios of commodity futures.

ET estimators, respectively, to a conditional setting – to obtain the cost of all consumption fluctuations in each time period (i.e., in each possible state of the economy). This amounts to calculating the ratio of the time-t prices of the claims to the stabilized consumption stream and the actual risky consumption stream, for each time period t. We find that the cost of consumption fluctuations is strongly time-varying and countercyclical. In our baseline case, the cost of all one-year fluctuations varies from 0.15% to 8.0%. Also, the cost is strongly countercyclical, rising sharply during recessionary episodes. The latter finding also helps explain the high average cost of business cycle fluctuations that we estimate. While the precise magnitudes of the costs are sensitive to the assumed value of the utility curvature parameter (γ in Equation 1), with higher values leading to larger costs, the findings that business cycle costs constitute between a quarter to a third of the costs of all fluctuations, that the cost of fluctuations is strongly countercyclical, and that the costs are substantially higher than those implied by Lucas' specification are robust to perturbations in the value of this parameter.

Comparing our findings to the leading equilibrium models, we find that the model-implied costs of consumption fluctuations are too small at business cycle frequencies in both the habit model of Campbell and Cochrane (1999) and in the long-run risk model of Bansal and Yaron (2004), and the latter implies too large costs in the very long run. Instead, our estimates are more in line with Bryzgalova, Huang, and Julliard (2024a) that documents strong persistence in the consumption process that dies off after the business cycle horizon, and shows that this medium horizon consumption risk is sufficient to rationalize the equity premium puzzle with a low level of relative risk aversion. Furthermore, we derive the implications of our method for the term structure of yields on "consumption strips." Bringing this novel insight to the data we find that the (unconditional) term structure appears to be significantly downward sloping after one quarter, but essentially flat at longer horizons.

Our paper lies at the interface of two, albeit mostly distinct, strands of literature. It contributes to a growing literature that uses an information-theoretic (or, relative-entropy minimizing) alternative to the standard generalized method of moments approach to address a variety of questions in economics and finance. Information-theoretic approaches were first introduced in financial economics by (Stutzer, 1995, 1996) and Kitamura and Stutzer (1997) (see Kitamura (2006) for a survey of these methods). Subsequently, these approaches have been used to assess the empirical plausibility of the rare disasters hypothesis in explaining asset pricing puzzles (see, e.g., Julliard and Ghosh (2012)), construct diagnostics for asset pricing models (see, e.g., Almeida and Garcia (2012), Backus, Chernov, and Zin (2013), Almeida and Garcia (2016)), construct bounds on the SDF and its components and recover the missing component from a candidate SDF (see, e.g., Borovicka, Hansen, and Scheinkman

(2016), Ghosh, Julliard, and Taylor (2017), Sandulescu, Trojani, and Vedolin (2018)), performance evaluation of funds (see, e.g., Almeida, Ardison, and Garcia (2019)), and recover investors' beliefs from observed asset prices (see, e.g., Hansen (2014), Ghosh and Roussellet (2019), Chen, Hansen, and Hansen (2020)).

Our paper also contributes to the literature that tries to assess the welfare costs of aggregate economic fluctuations (see, e.g., Lucas (1987), Imrohoroglu (1989), Atkeson and Phelan (1994), Obstfeld (1994), Pemberton (1996), Dolmas (1998), Tallarini (2000), Beaudry and Pages (2001), Otrok (2001), Storesletten, Telmer, and Yaron (2001), Alvarez and Jermann (2004), Krebs (2007), Martin (2008), Barro (2009), Krusell and Smith (2009), Epstein, Farhi, and Strzalecki (2014), Lustig, Nieuwerburgh, and Verdelhan (2013)). Most of this literature assumes particular parametric forms for preferences as well as the data generating process (DGP). Our paper, on the other hand, is more model-free, not requiring us to fully specify preferences or the DGP.

Our approach is similar in spirit to Alvarez and Jermann (2004) that, to the best of our knowledge, are the first to have used asset prices to infer bounds on the welfare cost of business cycle fluctuations.[4] Our results, however, are in stark contrast to those in Alvarez and Jermann (2004) who argue that, while the cost of all consumption fluctuations is very high (they report a baseline value of 28.6% in an infinite-horizon setting), the cost of business cycle fluctuations in consumption is miniscule, varying from 0.1% to 0.5%. This difference is driven by both our wholly different methodology to the recovery of the SDF that relies on fewer assumptions and approximation results and has well-behaved asymptotics, as well as different approaches to the filtering of the business cycle fluctuations from the historical consumption series. Recent theoretical studies, that aim to explain the behavior of asset prices while simultaneously retaining plausible business cycle dynamics, have argued for very high costs of business cycles – for example, 29% in Bai and Zhang (2020) that develops a general equilibrium model with recursive utility, search frictions, and capital accumulation. Davis and Segal (2020) argue that a small component of the business-cycle can be rationally mistaken to be permanent, thereby understating the importance of business cycle fluctuations. Our estimate of the cost of business cycles is in line with these studies, albeit obtained in a more model-free setting and, therefore, more robust to misspecification.

The reminder of the paper is organized as follows. Section II defines the cost of aggregate consumption fluctuations and describes an information-theoretic methodology to estimate this cost. Section III reports the empirical results. Section IV assesses the sensitivity of our main findings to alternative values of the utility curvature parameter. Finally, Section VI

[4] Hansen, Sargent, and Tallarini (1999) hint at the welfare cost calculation subseqeuntly developed in Alvarez and Jermann (2004).

concludes with suggestions for future research. The Appendix contains simulation evidence on the ability of the methodology to estimate the cost of fluctuations accurately, a data description, and a host of robustness checks.

II Pricing Aggregate Economic Fluctuations

This section defines the welfare costs of fluctuations in aggregate consumption and proposes a novel procedure to measure this cost. Specifically, in Subsection II.1, we follow Alvarez and Jermann (2004) and define the marginal cost of aggregate consumption fluctuations, for two alternative definitions of fluctuations. In Subsection II.2, we propose a novel information-theoretic procedure to measure the costs of these fluctuations. Throughout this section, uppercase letters are used to denote random variables and the corresponding lowercase letters to particular realizations of these variables.

II.1 The Cost of Aggregate Fluctuations

The cost (or, the market price) of consumption fluctuations, ω_0, is defined as the ratio of the prices of two securities: a claim to a *stable* version of the aggregate consumption stream from which certain fluctuations have been removed, and a claim to the actual aggregate consumption stream:

$$\omega_0 = \frac{V_0\left[\left\{C_t^{stab}\right\}_{t\geq 1}\right]}{V_0\left[\left\{C_t\right\}_{t\geq 1}\right]} - 1. \tag{2}$$

In the above equation, $V_0\left[\left\{C_t\right\}_{t\geq 1}\right]$ and $V_0\left[\left\{C_t^{stab}\right\}_{t\geq 1}\right]$ denote the time-0 prices of claims to the future consumption stream and the future stabilized consumption stream, respectively. Therefore, the cost of consumption fluctuations measures how much extra investors would be willing to pay in order to replace the aggregate consumption stream with its stabilized counterpart.

If stabilized consumption, C_t^{stab}, is defined as the expected value of future consumption, i.e. $C_t^{stab} = E_0\left(C_t\right)$, then Equation (2) measures the cost of *all* consumption fluctuations. In other words, it measures the benefit of eliminating all consumption uncertainty. If, on the other hand, stabilized consumption, C_t^{stab}, is defined as the long-term trend consumption, from which fluctuations corresponding to business cycle frequencies (typically defined as lasting for no longer than 8 years) have been removed, then Equation (2) measures the cost of business cycle fluctuations in consumption.

In the context of an infinite horizon representative agent economy, Alvarez and Jermann (2004) show that ω_0 in Equation (2) measures the *marginal cost of consumption fluctuations*, defined as the per unit benefit of a marginal reduction in consumption fluctuations, expressed as a percentage of lifetime consumption. Under fairly general conditions, the marginal cost provides an upper bound on the *total cost of consumption fluctuations*, where the latter is defined as the additional lifetime consumption, expressed as a percentage of consumption, that the representative agent would demand in order to be indifferent between the risky consumption stream and the stabilized version of it.

Alvarez and Jermann (2004) (AJ) show that the marginal cost of *all* consumption fluctuations, i.e. the scenario where $C_t^{stab} = E_0\left(C_t\right) = (1+\mu_c)^t C_0$ for $t = 1, 2, ..., \infty$, where μ_c denotes the unconditional mean of consumption growth, is given by:

$$\omega_0 = \frac{r_0 - \mu_c}{y_0 - \mu_c} - 1. \tag{3}$$

In the above equation, y_0 and r_0 denote the yields to maturity on claims to the stabilized sure consumption stream and the risky consumption stream, respectively. AJ calibrate $\mu_c = 2.3\%$, set $y_0 = 3.0\%$ to match the 10-year government bond yield, and estimate the consumption risk premium $r_0 - y_0 \geq 0.2\%$, that implies an estimated cost of at least 28.6%. However, the above equation highlights that the estimate of the cost is very sensitive to the values of y_0, r_0, and μ_c. For instance, reducing the calibrated y_0 by a mere 50bp would yield an estimate of the cost of consumption fluctuations of 350%, and using parameters in the range reported as being reasonable in AJ one can get estimates as high as 1535.7%.

The source of instability in the AJ framework stems from the fact that, as $y_0 \to \mu_c$, we have $\omega_0 \to \infty$, hence the approach breaks down. Blanchard (2019) points out that, in the more recent history, the nominal rate on a 10-year government bond has been close to 2.7%, while the expected nominal growth rate is 4.0%, causing $y_0 - \mu_c$ to be negative, thereby invalidating the use of Equation (3). And this is not just a feature of the US, but also other developed economies such as the UK and the Euro Zone. Moreover, Blanchard (2019) highlights that the current situation is more the norm rather than the exception in the US – the average nominal growth rate and the rate on 1-year government bonds have been 6.3% and 4.7%, respectively, since 1950, and 5.3% and 4.6%, respectively, since 1870, and, in fact, $y_0 - \mu_c$ has been negative in all decades except the 1980s. This outlines the extreme fragility of using Equation (3) to estimate the cost of consumption fluctuations.

Therefore, in this paper, instead of attempting to measure the welfare costs of eliminating consumption fluctuations over an infinite time horizon, we focus on the term structure of finite horizon consumption risk. In other words, we characterize the welfare gains from stabilizing the next $j = 1, ..., J$ periods of consumption uncertainty. This makes our results more robust to the choice of discount rates, since small changes in the latter have large impact at the infinite horizon. In addition, the results are also informative about the persistence of underlying shocks. A further advantage of considering the welfare costs of finite horizon fluctuations is that stabilization of fluctuations can affect the long run mean growth rate in unknown ways (e.g. Barlevy (2004) considers an endogenous growth framework in which shutting down aggregate uncertainty increases annual consumption growth by .35-.40%), so the present endowment-economy exercise that abstracts from this effect may have limited interpretability for long-horizon calculations.

To obtain the term structure, note that the law of one price implies that

$$V_0\left[\{C_t\}_{t=1}^{j}\right] = \sum_{t=1}^{j} V_0\left(C_t\right), \tag{4}$$

for $j \geq 1$, where $V_0\left(C_t\right)$ denotes the time-0 price of a claim to a single payoff equal to the aggregate consumption at time t. Similarly, $V_0\left[\left\{C_t^{stab}\right\}_{t=1}^{j}\right]$ can be written as the sum, over time periods 1, 2, ..., j, of the prices of claims to single payoffs equal to the stabilized consumption in each of these future periods. Therefore, the (cumulative) cost of j-period fluctuations is given by

$$\frac{\sum_{t=1}^{j} V_0\left(C_t^{stab}\right)}{\sum_{t=1}^{j} V_0\left(C_t\right)} - 1. \tag{5}$$

Note that, as $j \to \infty$, the cost of j-period consumption fluctuations in Equation (5) approaches the marginal cost of consumption fluctuations in Equation (2) studied by Alvarez and Jermann (2004).

We provide two types of estimates of the costs of fluctuations. First, we present the *expected* cost of consumption fluctuations, i.e. the average cost over all possible states of the world. This is the ratio of the expected (or, average) prices of claims to a stabilized consumption stream and the actual aggregate consumption stream. For instance, the expected cost of one-period consumption fluctuations is defined as:

$$\frac{\mathbb{E}^{\mathbb{P}}\left[V_t\left(C_{t+1}^{stab}\right)\right]}{\mathbb{E}^{\mathbb{P}}\left[V_t\left(C_{t+1}\right)\right]} - 1, \tag{6}$$

where $\mathbb{E}^{\mathbb{P}}[.]$ refers to the expectation with respect to the (true) physical measure $\mathbb{P}$.

Second, we report how the cost varies over time, i.e. with different possible states of the world. Specifically, the time-t cost of one-period consumption fluctuations is defined as

$$\frac{V_t\left(C_{t+1}^{stab}\right)}{V_t\left(C_{t+1}\right)} - 1, \tag{7}$$

Note that the difference between the average cost in Equation (6) and the time-t cost in Equation (7) is that, while the former involves the evaluation of *unconditional* expectations to obtain the average prices of the consumption claims, the latter requires the computation of the time-t prices of these claims as the *conditional* expectations of their discounted payoffs.

Since neither of the two assets that characterize the marginal cost of consumption fluctuations – namely, the claims to aggregate consumption or its stabilized counterpart – is directly traded in financial markets, their prices are not directly observed. Therefore, the values of these claims need to be estimated in order to obtain the cost of consumption fluctuations. Historically, this has involved taking a stance on investors' preferences, i.e. their stochastic discounting of the various possible future states of the world, and the dynamics of the data generating process (DGP), i.e. the likelihood of the states being realized. The resultant estimates of the cost of economic fluctuations have proven to be quite sensitive to these two assumptions (see, e.g., Barlevy (2005)). The following subsection outlines a novel econometric methodology for estimating the cost of consumption fluctuations, that does not require any specific functional-form assumptions either about investors' preferences or the dynamics of the DGP.

II.2 Measuring the Cost of Aggregate Fluctuations

The (shadow) value of a claim to the aggregate consumption next period can be generally expressed as

$$V_t\left(C_{t+1}\right) = \mathbb{E}^{\mathbb{P}}\left[M_{t+1}C_{t+1}|\mathcal{F}_t\right], \tag{8}$$

where M_t is the SDF, $\mathcal{F}_t$ denotes the investors' information set at time t, and $\mathbb{E}^{\mathbb{P}}\left[.|\mathcal{F}_t\right]$ refers to the expectation with respect to the physical measure $\mathbb{P}$ conditional on the investors' time-t information set. The existence of a (strictly positive) SDF is guaranteed by the assumption of the absence of arbitrage opportunities.

Dividing Equation (8) by C_t to make both sides stationary, we have

$$\tilde{pc}_{1,t} := \frac{V_t\left(C_{t+1}\right)}{C_t} := \mathbb{E}^{\mathbb{P}}\left[M_{t+1}\frac{C_{t+1}}{C_t}|\mathcal{F}_t\right]. \tag{9}$$

$\tilde{pc}_{1,t}$ can be interpreted as the time-t price (expressed as a fraction of current consumption) of an asset with a single payoff equal to the aggregate consumption next period.

Similarly, the (shadow) value of a claim to a *stabilized* version of the aggregate consumption next period can be expressed as

$$V_t\left(C^{stab}_{t+1}\right) = \mathbb{E}^{\mathbb{P}}\left[M_{t+1}C^{stab}_{t+1}|\mathcal{F}_t\right]. \tag{10}$$

implying that

$$\tilde{pc}_{1,t}^{stab} := \frac{V_t\left(C_{t+1}^{stab}\right)}{C_t} := \mathbb{E}^{\mathbb{P}}\left[M_{t+1}\frac{C_{t+1}^{stab}}{C_t}|\mathcal{F}_t\right]. \tag{11}$$

If the true underlying model were known, i.e. the SDF M and the physical measure $\mathbb{P}$ were known, then the prices of the claims to the aggregate consumption and the stabilized aggregate consumption next period could be determined using Equations (9) and (11), respectively. Therefore, the time-t cost of one-period consumption fluctuations, defined in Equation (7) could be obtained as

$$\frac{V_t\left(C_{t+1}^{stab}\right)}{V_t\left(C_{t+1}\right)} - 1 = \frac{\tilde{pc}_{1,t}^{stab}}{\tilde{pc}_{1,t}} - 1. \tag{12}$$

And, the average (over all possible states of the world) cost of one-period fluctuations, defined in Equation (6), would then obtain as

$$\frac{\mathbb{E}^{\mathbb{P}}\left[V_t\left(C_{t+1}^{stab}\right)\right]}{\mathbb{E}^{\mathbb{P}}\left[V_t\left(C_{t+1}\right)\right]} - 1 \approx \frac{\mathbb{E}^{\mathbb{P}}\left(\tilde{pc}_{1,t}^{stab}\right)}{\mathbb{E}^{\mathbb{P}}\left(\tilde{pc}_{1,t}\right)} - 1 \equiv \frac{\tilde{pc}_1^{stab}}{\tilde{pc}_1} - 1. \tag{13}$$

The costs of multi-period fluctuations can be similarly obtained.

For instance, assuming a representative agent endowed with power utility preferences with a constant CRRA, $\tilde{pc}_1$ can be estimated as $\frac{1}{T}\sum_{t=1}^{T}\delta\left(\Delta C_t\right)^{1-\gamma}$, where γ denotes the CRRA and δ the subjective discount factor. Moreover, assuming log-normality of the aggregate consumption growth as in Lucas (1987):

$$\tilde{pc}_1 = \mathbb{E}^{\mathbb{P}}\left[\delta(\Delta C_t)^{1-\gamma}\right] = e^{\ln(\delta)+(1-\gamma)\mathbb{E}^{\mathbb{P}}[\ln(\Delta C_t)]+.5(1-\gamma)^2 Var^{\mathbb{P}}[\ln(\Delta C_t)]}.$$

The price of a claim to sure consumption next period, $C_{t+1}^{stab} = (1+\mu_c)\,C_t$, is, similarly, given by

$$\tilde{pc}_1^{stab} = \mathbb{E}^{\mathbb{P}}\left[\delta(\Delta C_t)^{-\gamma}(1+\mu_c)\right] = (1+\mu_c)\,e^{\ln(\delta)-\gamma\mathbb{E}^{\mathbb{P}}[\ln(\Delta C_t)]+.5\gamma^2 Var^{\mathbb{P}}[\ln(\Delta C_t)]}.$$

Using calibrated (or estimated) values of the first two moments of log consumption growth and the preference parameters, we can obtain $\tilde{pc}_1$ and $\tilde{pc}_1^{stab}$ and, therefore, the price of one-period consumption fluctuations.

However, in practice, neither the pricing kernel M nor the physical measure $\mathbb{P}$ is directly observable and, therefore, need to be estimated. In this paper, we do not make any strong assumptions either about the functional-form of preferences, or the dynamics of the DGP. Instead, our methodology is based on the observation that, albeit not directly observable, information about M is available in financial markets. Specifically, we assume that the pricing kernel, M, has the form in Equation (1). As discussed in the introduction, this multiplicative decomposition of the SDF encompasses virtually all representative agent consumption-based asset pricing models proposed in the literature, including Lucas' original specification, and even certain heterogeneous agents incomplete markets models. Different models offer different economic interpretations of the ψ-component.

II.2.1 Average Cost of Fluctuations

Given the assumed form of the SDF, for any vector of excess returns $\mathbf{R}_t^e \in \mathbb{R}^N$ on N traded assets, the following set of *unconditional* Euler equations must hold in the absence of arbitrage opportunities:

$$\mathbf{0} = \mathbb{E}^{\mathbb{P}}\left[M_t \mathbf{R}_t^e\right] = \int_{\mathbf{z}} (\Delta C(\mathbf{z}))^{-\gamma} \psi(\mathbf{z}) \mathbf{R}^e(\mathbf{z}) d\mathbb{P}(\mathbf{z}) = \int_{\mathbf{z}} (\Delta C(\mathbf{z}))^{-\gamma} \mathbf{R}^e(\mathbf{z}) d\mathbb{F}(\mathbf{z}), \tag{14}$$

where $\mathbf{0}$ is an N-dimensional vector of zeros and $\mathbf{z}$ denotes the (latent) state vector. The third equality follows from a change of measure from $\mathbb{P}$ to $\mathbb{F}$, with an associated Radon-Nikodym derivative of $\frac{d\mathbb{F}(\mathbf{z})}{d\mathbb{P}(\mathbf{z})} = \frac{\psi(\mathbf{z})}{\mathbb{E}^{\mathbb{P}}(\psi(\mathbf{z}))}$.

Note that the representation in Equation (14) holds even if investors have biased beliefs as long as there are no arbitrage opportunities under the subjective expectation measure $\tilde{\mathbb{P}}$, since in this case we would have

$$\mathbf{0} = \mathbb{E}^{\tilde{\mathbb{P}}}\left[M_t \mathbf{R}_t^e\right] = \int_{\mathbf{z}} (\Delta C(\mathbf{z}))^{-\gamma} \tilde{\psi}(\mathbf{z}) \mathbf{R}^e(\mathbf{z}) d\mathbb{P}(\mathbf{z}) = \int_{\mathbf{z}} (\Delta C(\mathbf{z}))^{-\gamma} \mathbf{R}^e(\mathbf{z}) d\mathbb{F}(\mathbf{z}) \tag{15}$$

where $\tilde{\psi}(\mathbf{z}) \propto \psi(\mathbf{z}) \times \frac{d\tilde{\mathbb{P}}(\mathbf{z})}{d\mathbb{P}(\mathbf{z})}$. Hence, in this case, the unobservable component that our method recovers would capture both the missing component of the SDF and the belief distortions $d\tilde{\mathbb{P}}(\mathbf{z})/d\mathbb{P}(\mathbf{z})$. All that is required for the above change of measure to be valid is that the subjective beliefs should be absolutely continuous with respect to the physical probability measure, i.e., the subjective and objective measures should have the same zero probability set.[5]

Using asset returns and consumption data, we can estimate the $\mathbb{F}$ distribution. Suppose that $p(\mathbf{z})$ and $f(\mathbf{z})$ denotes the pdfs associated with the measures $\mathbb{P}$ and $\mathbb{F}$, respectively. The $\mathbb{F}$ distribution can be estimated to minimize the Kullback-Leibler Information Criterion (KLIC) divergence (or the relative entropy) between the $\mathbb{P}$ and $\mathbb{F}$ measures:

$$\min_{\mathbb{F}} \int \log\left(\frac{d\mathbb{P}}{d\mathbb{F}}\right) d\mathbb{P} \equiv \int_{\mathbf{z}} \log\left(\frac{p(\mathbf{z})}{f(\mathbf{z})}\right) p(\mathbf{z}) d\mathbf{z}, \text{ s.t. } \mathbf{0} = \int_{\mathbf{z}} \mathbf{R}^e(\mathbf{z}) (\Delta C(\mathbf{z}))^{-\gamma} f(\mathbf{z}) d\mathbf{z}. \tag{16}$$

Since relative entropy is not symmetric, we can reverse the roles of $\mathbb{P}$ and $\mathbb{F}$ in Equation

[5] Modelling subjective beliefs via a change of measure is, e.g., used in Basak and Yan (2010).

(16) to obtain an alternative divergence criterion between these two measures, which can be minimized to recover an alternative estimate of the measure, $\mathbb{F}$:

$$\min_{\mathbb{F}} \int \mathbf{log}\left(\frac{d\mathbb{F}}{d\mathbb{P}}\right) d\mathbb{F} = \int \mathbf{log}\left(\frac{f(\mathbf{z})}{p(\mathbf{z})}\right) f(\mathbf{z}) d\mathbf{z} \text{ s.t. } \mathbf{0} = \int \mathbf{R}^e(\mathbf{z}) \left(\Delta C(\mathbf{z})\right)^{-\gamma} f(\mathbf{z}) d\mathbf{z}, \quad (17)$$

Equations (16) and (17) are, respectively, the Empirical Likelihood (EL) estimator of Owen (2001) and the Exponentially-Tilted (ET) estimator of Kitamura and Stutzer (1997) (see also Schennach (2005)), originally proposed in Ghosh, Julliard, and Taylor (2017) to recover the multiplicative missing component of the pricing kernel. Once the $\mathbb{F}$-measure, or, from the expression for the Radon-Nikodym derivative, the missing component, ψ, of the pricing kernel, is estimated, the pricing kernel, M, can be obtained using Equation (1). We refer to this kernel as the *Information*-SDF, or I-SDF, because of the information-theoretic approach used to recover it. In the main text, we focus on estimates based on the EL method, and show in the Appendix that almost identical results are obtained with the ET approach.

Ghosh, Julliard, and Taylor (2017) point out several reasons why relative entropy minimization is an attractive criterion for recovering the pricing kernel. Some of these are restated here for convenience.

First, the use of relative entropy, due to the presence of the logarithm in the objective function in Equations (16) and (17), naturally imposes the non-negativity of the SDF.

Second, by construction, our approach to recover the ψ_t component adds the *minimum amount of information* – in the Shannon sense – needed for the pricing kernel to price assets. That is, from an information-theoretic standpoint, it satisfies the Occam's razor, or law of parsimony. To provide some intuition, suppose that the consumption growth component of the pricing kernel, $(\Delta C_t)^{-\gamma}$, were sufficient to price assets perfectly. Then $\psi_t \equiv 1, \forall t$, and we have that $\mathbb{F} \equiv \mathbb{P}$, delivering a KLIC divergence $\int \log\left(\frac{d\mathbb{P}}{d\mathbb{F}}\right) d\mathbb{P} = 0$ in Equation (16) (the same holds for Equation (17)). However, if the consumption growth component is not sufficient to price assets (as is the case in reality), then the estimated measure $\mathbb{F}$ is distorted relative to the physical measure $\mathbb{P}$, i.e. the KLIC divergence is positive. And, the estimator searches for a measure $\mathbb{F}$ that is as close as possible, in an information-theoretic sense, to the physical measure $\mathbb{P}$. In other words, the approach distorts the physical probabilities as little as possible in order to satisfy the Euler equation restrictions. And the estimator is non-parametric in the sense that it does not require any parametric functional-form assumptions about the ψ-component of the kernel or the physical distribution $\mathbb{P}$.

Third, as implied by the work of Brown and Smith (1990), the use of entropy is desirable if we think that tail events are an important component of the risk measure.[6]

[6]Brown and Smith (1990) develop what they call "a Weak Law of Large Numbers for rare events;" that

Fourth, this approach is numerically simple to implement. Given a history of excess returns and consumption growth $\{\mathbf{r}_t^e, \Delta c_t\}_{t=1}^T$, Equation (16) can be made operational by replacing the expectation with a sample analogue, as is customary for moment based estimators, and using the Radon-Nikodym derivative to rewrite the criterion function in terms of the ψ-component of the SDF:[7]

$$\arg\max_{\{\psi_t\}_{t=1}^T} \frac{1}{T}\sum_{t=1}^{T}\log\psi_t \quad \text{s.t.} \quad \frac{1}{T}\sum_{t=1}^{T}(\Delta c_t)^{-\gamma}\,\psi_t \mathbf{r}_t^e = \mathbf{0}. \tag{18}$$

An application of Fenchel's duality theorem to the above problem (see, e.g., Csiszár (1975), Owen (2001)), delivers the estimates (up to a positive constant scale factor):

$$\hat{\psi}_t = \frac{1}{T(1+\hat{\theta}(\gamma)'\mathbf{r}_t^e\,(\Delta c_t)^{-\gamma})} \quad \forall t, \tag{19}$$

where $\hat{\theta} \in \mathbb{R}^N$ is the vector of Lagrange multipliers that solves the unconstrained dual problem:

$$\hat{\theta}(\gamma) = \arg\min_{\theta} -\sum_{t=1}^{T}\log(1+\theta'\mathbf{r}_t^e\,(\Delta c_t)^{-\gamma}). \tag{20}$$

The solution to Equation (17) is similarly simple to implement and is presented in Appendix A.

Fifth, and perhaps most importantly, the I-SDF successfully prices assets. Note that this result is not surprising *in sample*, because the I-SDF is constructed to price the test assets in-sample (see Equation (16)). However, Ghosh, Julliard, and Taylor (2022) show that the good pricing performance of the I-SDF also obtains out-of-sample for broad cross-sections of assets, including domestic and international equities, currencies, and commodities. The *out-of-sample* performance of the I-SDF is superior to not only the single factor CAPM and the Consumption-CAPM, but also the Fama-French 3 and 5 factor models as well as the 4 factor model of Hou, Xue, and Zhang (2014). We show this formally in Table A.I of Appendix B.1. Furthermore, as shown in Appendix C, the I-SDF also successfully prices multi-period returns. These features suggest that the I-SDF is a robust approach for capturing the relevant sources of priced risk and, therefore, offers a reliable candidate kernel with which to measure the cost of aggregate economic fluctuations.

Finally, we show via simulations in Appendix D that the methodology is quite successful in estimating the cost of fluctuations in hypothetical economies for empirically realistic

is, they show that the empirical distribution observed in a very large sample converges to the distribution that minimizes the relative entropy.

[7]This amounts to assuming ergodicity for both the pricing kernel and asset returns.

sample sizes.

Nevertheless, two important limitations of the method are worth mentioning. First, since our method approximates the distribution of the possible states that the economy might encounter with a multinomial distribution over the historical observations, excluding extreme, yet salient, events from the sample (as, e.g., the Great Depression) can potentially have a large impact on the estimates. Consequently, as the baseline sample, we use annual data going as far back as the Great Depression (instead of quarterly data that is available only in the postwar sample).

Second, since our minimum entropy method is moment based, its finite sample performance quickly deteriorates when the number of moment restrictions becomes large relative to the number of time series data points. This constraints the size of the cross-section of base assets that can be used in the recovery of the I-SDF – especially in the annual data setting that we use as the baseline. Hence, one should choose a small cross-section, that is nevertheless informative about the main priced risk factors in the market. Consequently, in our empirical analysis, we include as base asset a small number of portfolios motivated by the previous literature. In the smallest set, we consider only the market index, since a) it predicts future consumption growth (Parker (2001)), and b) it is very close to the first principal component of asset returns and the latter, as shown in Bryzgalova, Huang, and Julliard (2024a), captures the overwhelming majority of priced shocks in the time series of consumption growth. As an extended set of base assets, we consider the six size and book-to-market sorted Fama-French portfolios, since these capture the well established value and size anomalies, and have predictive power for both future economic conditions and consumption growth (see, e.g., Liew and Vassalou (2000) and Parker and Julliard (2003))

With the recovered ψ-component, the I-SDF is obtained (up to a positive scale factor, κ) as

$$\left\{\widehat{M}_t\right\}_{t=1}^{T} = \left\{\kappa \left(\Delta c_t\right)^{-\gamma} \widehat{\psi}_t\right\}_{t=1}^{T}. \tag{21}$$

The proportionality constant, κ, can be recovered from the Euler equation for the risk free rate. Specifically, $\kappa = \frac{\frac{1}{T}\sum_{t=1}^{T} r_{f,t}}{\frac{1}{T}\sum_{t=1}^{T}(\Delta c_t)^{-\gamma}\widehat{\psi}_t}$, where $\{r_{f,t}\}_{t=1}^{T}$ are the realized returns on the risk free asset in the historical sample. This ensures that, although the I-SDF is recovered using excess returns, it also satisfies the Euler equation for the risk free rate.

Armed with the I-SDF, we can now estimate the welfare benefits of eliminating consumption fluctuations. Specifically, the value of eliminating *all* consumption fluctuations in the next period alone is obtained as:

$$\widehat{\tilde{pc}_1^{stab}/\tilde{pc}_1} - 1 = \frac{\sum_{t=1}^{T}\widehat{M}_t\left(1+\mu_c\right)}{\sum_{t=1}^{T}\widehat{M}_t \Delta c_t} - 1. \tag{22}$$

Note that under the subjective beliefs formulation in Equation (15), the estimated ψ component would contain both the missing part of the SDF as well as the belief distortions. Consequently, in this case, the above would yield the subjective value of eliminating all consumption fluctuations and replacing them with the constant growth rate μ_c.

The value of eliminating business cycle fluctuations in the next period is:

$$\widehat{p\tilde{c}_1{}^{stab}}/p\tilde{c}_1 - 1 = \frac{\sum_{t=1}^{T} \widehat{M}_t \Delta c_t^{stab}}{\sum_{t=1}^{T} \widehat{M}_t \Delta c_t} - 1, \tag{23}$$

where Δc_t^{stab} denotes a time-varying stabilized consumption growth from which the business cycle variations have been removed. This stabilized version of consumption can be obtained by an application of a smoothing filter to the original consumption series.

We will soon see that estimates obtained by using the I-SDF differ markedly from estimates obtained by Lucas' method (summarized on p.10 herein). To help explain this, recall that Lucas' method presumes a complete markets exchange economy in which the (unique) SDF M_t is the marginal rate of substitution (MRS) from the discounted power utility functional. The MRS depends only on consumption growth and model parameters. Under the complete markets assumption, all assets must satisfy the pricing condition $E[M_t \mathbf{R_t^e}] = \mathbf{0}$, including the risk free asset. Yet the Equity Premium Puzzle and the variance and entropy bounds literatures cited herein all establish that the excess returns of popular equity indices will *not* satisfy these constraints when the Lucas SDF is specified with economically plausible parameters. In light of this, subsequent work has proposed other consumption-based asset pricing models, but Ghosh, Julliard, and Taylor (2017) show that these are similarly problematic when the returns of broad cross-sections of equity factor portfolios are included in $\mathbf{R^e}$.

In contrast, the I-SDF satisfies these pricing constraints *by construction* while still including consumption growth in its makeup. This provides a method of pricing consumption fluctuations in a way that is consistent with the pricing of equity portfolios, albeit without the theoretical desideratum of first specifying an exchange or other economic model from which it was derived. Theorists who maintain the complete markets assumption can view our approach as a data-driven procedure to estimate the unknown unique SDF, with the aforementioned desirable properties.

Finally, note that Equations (22) and (23) represent the costs of all consumption fluctuations and business cycle fluctuations, respectively, for one period alone. It is straightforward to extend the analysis to obtain the cost of fluctuations for multiple periods. For instance, the (shadow) value of a claim to the aggregate consumption j periods into the future can be expressed as

$$V_t\left(C_{t+j}\right) = \mathbb{E}_t^{\mathbb{P}}\left[M_{t:t+j}C_{t+j}\right],$$

where $M_{t:t+j}$ denotes the j-period SDF. Thus, the expected price-consumption ratio of a security that delivers a single payoff equal to the aggregate consumption j periods into the future is given by

$$\tilde{pc}_j := \mathbb{E}^{\mathbb{P}}\left[\frac{V_t\left(C_{t+j}\right)}{C_t}\right] = \mathbb{E}^{\mathbb{P}}\left[M_{t:t+j}\frac{C_{t+j}}{C_t}\right].$$

The one-period I-SDF, recovered in Equation (21), can be compounded to recover the j-period discount factor:[8]

$$M_{t:t+j} = \prod_{i=1}^{j} M_{t+i}.$$

Using $M_{t:t+j}$, we can estimate the price-consumption ratio $\tilde{pc}_j$ for a single consumption claim j periods in the future. And this can be done for any $j = 2, 3, 4, ...$ Using the estimated price-consumption ratios of the claims to single future payoffs, we can estimate the price-consumption ratio of an asset that delivers the stochastic consumption in each of the next J periods i.e. $\tilde{pc}_{1:J} := \sum_{j=1}^{J}\tilde{pc}_j$. Hence, it is straightforward to compute the value of removing all, or only business cycle, fluctuations in consumption over J periods with expressions analogous to the ones in Equations (22)-(23).

II.2.2 Time-Varying Cost of Fluctuations

Here we describe an extension of the EL approach, namely the smoothed empirical likelihood (SEL) estimator of Kitamura, Tripathi, and Ahn (2004), that we use to recover the time-varying cost of fluctuations. The corresponding extension of the ET approach – the smoothed exponential tilting (SET) estimator – can be similarly used (see, e.g., Ghosh and Otsu (2020)). To our knowledge, this is the first attempt to provide quantitative estimates of the time-variation in the welfare costs of aggregate fluctuations, without fully specified preferences and DGP.

Recall that the EL and ET approaches recover a pricing kernel (the I-SDF) that prices assets unconditionally, i.e. satisfies the unconditional Euler equations producing zero unconditional pricing errors. The extension of the methodology considered in this section recovers an I-SDF that satisfies the more stringent conditional Euler equation restrictions, thereby producing zero conditional pricing errors. The recovered SDF, therefore, must also price assets unconditionally. As described below, the SEL and SET estimators rely on the same principles as the EL and ET estimators, respectively, but incorporate additional constraints

[8] Appendix C shows that the compounded I-SDF is extremely similar to the I-SDF estimated by varying the sampling frequency of the data, yielding a correlation with the latter of 99.7% at the five year horizon.

through conditional moment restrictions. We focus on the SEL estimator in the main body of the paper and report very similar results for the SET in Appendix III.3.

Absence of arbitrage opportunities implies the following conditional pricing restrictions:

$$\mathbb{E}^{\mathbb{P}_t}\left[M_{t+1}\mathbf{R}^e_{t+1}|\mathcal{F}_t\right] = \mathbb{E}^{\mathbb{P}_t}\left[(\Delta C_{t+1})^{-\gamma}\,\psi_{t+1}\mathbf{R}^e_{t+1}|\mathcal{F}_t\right] = \mathbf{0}, \tag{24}$$

where $\mathcal{F}_t$ denotes the information set up to time t, and the first equality follows from the assumed multiplicative decomposition of the SDF. Under weak regularity conditions, we have

$$\mathbb{E}^{\mathbb{P}_t}\left[(\Delta C_{t+1})^{-\gamma}\frac{\psi_{t+1}}{\mathbb{E}^{\mathbb{P}_t}[\psi_{t+1}|\mathcal{F}_t]}\mathbf{R}^e_{t+1}|\mathcal{F}_t\right] = \mathbb{E}^{\mathbb{F}_t}\left[(\Delta C_{t+1})^{-\gamma}\,\mathbf{R}^e_{t+1}|\mathcal{F}_t\right] = \mathbf{0}, \tag{25}$$

where $\frac{d\mathbb{F}_t}{d\mathbb{P}_t} = \frac{\psi_{t+1}}{\mathbb{E}^{\mathbb{P}_t}[\psi_{t+1}|\mathcal{F}_t]}$ is the Radon-Nikodym derivative of $\mathbb{F}$ with respect to $\mathbb{P}$.

We assume that the time-t information set of the investors, $\mathcal{F}_t$, can be summarized by a finite state vector, that we denote by $X_t \in \mathbb{R}^m$. Suppose that the historical realizations of consumption growth, excess returns, and the conditioning variables are given by $(\Delta c_t, \mathbf{r}^e_t, x_t)_{t=1}^T$, and that these realizations characterize the possible states of the world. Let $f_{i,j}$ denote the conditional probability (under the measure $\mathbb{F}$) of observing the joint outcome $(\Delta c_j, \mathbf{r}^e_j, x_j)$ at time $t+1$, i.e. the probability of state j being realized at time $t+1$, given that state i was realized at time t.

The SEL estimator of the transition matrix $\{f_{i,j}; i,j = 1,\ldots,T\}$ is such that it belongs to the simplex:

$$\Delta := \cup_{i=1}^T \Delta_i = \cup_{i=1}^T \left\{(f_{i,1},...,f_{i,T}):\ \sum_{j=1}^T f_{i,j} = 1,\ f_{i,j} \geq 0\right\}$$

and that: $\forall i \in \{1,\ldots,T\}, \quad \forall \gamma \in \Gamma,$

$$\left\{\widehat{f}^{SEL}_{i,\cdot}(\gamma)\right\} = \underset{(f_{i,\cdot})\in\Delta_i}{\arg\min}\ \sum_{j=1}^T \log\left(\frac{\omega_{i,j}}{f_{i,j}}\right)\omega_{i,j} \quad \text{s.t.} \quad \sum_{j=1}^T f_{i,j}\times(\Delta c_j)^{-\gamma}\,\mathbf{r}^e_j = \mathbf{0}. \tag{26}$$

where $f_{i,\cdot}$ denotes the T-dimensional vector $(f_{i,1},...,f_{i,T})$, Γ is the set of all admissible parameters γ, and $\omega_{i,j}$ are nonparametric kernel density weights:

$$\omega_{i,j} = \frac{\mathcal{K}\left(\dfrac{x_i - x_j}{b_T}\right)}{\displaystyle\sum_{t=1}^T \mathcal{K}\left(\dfrac{x_i - x_t}{b_T}\right)}, \tag{27}$$

where $\mathcal{K}$ is a kernel function belonging to the class of second order product kernels,[9] and the bandwidth b_T is a smoothing parameter.[10]

The objective function in Equation (26) is the KLIC divergence between the measure $\mathbb{F}_t \equiv \{f_{t,j}\}_{j=1}^{T}$ that is consistent with asset prices, i.e. satisfies the conditional Euler equations for the test assets, and the physical measure proxied by the nonparametric kernel density weights, $\mathbb{P}_t \equiv \{\omega_{t,j}\}_{j=1}^{T}$. And, $\frac{f_{t,j}}{\omega_{t,j}} = \frac{\psi_{t,j}}{\mathbb{E}^{\mathbb{P}_t}(\psi_{t,j}|\mathcal{F}_t)}$ is the Radon-Nikodym derivative of $\mathbb{F}$ with respect to $\mathbb{P}$. Suppose that the consumption growth component of the pricing kernel, $(\Delta C)^{-\gamma}$, is sufficient to price assets perfectly. Then, we have that $\forall t = 1, 2, ..., T$, the second component of the pricing kernel $\psi_{t,j} \equiv 1, \forall j = 1, 2, ..., T$, implying that $f_{t,j} = \omega_{t,j}, \forall j = 1, 2, ..., T$, the latter being the physical measure. However, if the consumption growth component is not sufficient to price assets, the estimated measure $\mathbb{F}_t$ is distorted relative to the physical measure $\mathbb{P}_t$. And, the SEL estimator searches for a measure $\mathbb{F}_t$ that is as close as possible to the physical measure $\mathbb{P}_t$. In other words, the approach distorts the physical probabilities as little as possible in order to satisfy the *conditional* Euler equation restrictions.

The solution to Equation (26) is analytical and given by:

$\forall i, j \in \{1, \ldots, T\}$,

$$\widehat{f}_{i,j}^{SEL}(\gamma) = \frac{\omega_{i,j}}{1 + (\Delta c_j)^{-\gamma}\, \widehat{\theta}_i(\gamma)'\, \mathbf{r}_j^e}\,, \tag{28}$$

where $\widehat{\theta}_i(\gamma) \in \mathbb{R}^N : i = \{1, \ldots, T\}$ are the Lagrange multipliers associated with the conditional Euler equation constraints, and solve the following unconstrained problem:

$$\widehat{\theta}_i(\gamma) = \underset{\theta_i \in \mathbb{R}^N}{\arg\max} \sum_{j=1}^{T} \omega_{i,j} \log\left[1 + (\Delta c_j)^{-\gamma}\, \theta_i'\, \mathbf{r}_j^e\right]. \tag{29}$$

Equations (28) and (29) show that the SEL procedure delivers a $(T \times T)$ matrix of probabilities $\left\{\widehat{f}_{i,j}^{SEL}(\gamma)\right\}$ for each value of the parameter γ. Each row $i : i = \{1, 2, ..., T\}$ contains the probabilities of transitioning to each of the T possible states $j : \{j = 1, 2, ..., T\}$ in the subsequent period, conditional on state i having been realized in the current period. Therefore, the approach recovers the *entire conditional distribution* of the data, under the measure $\mathbb{F}$, that is consistent with observed asset prices, i.e. that satisfies the conditional Euler equations. As shown in Appendix B.2 with a cross-sectional out of sample pricing exercise, the ψ component recovered with this method greatly improves pricing ability relative to the CRRA benchmark.

[9] $\mathcal{K}$ should satisfy the following. For $X = (X^{(1)}, X^{(2)}, ..., X^{(m)})$, let $\mathcal{K} = \prod_{i=1}^{m} k(X^{(i)})$. Here $k : \mathbb{R} \to \mathbb{R}_+$ is a continuously differentiable p.d.f. with support $[-1, 1]$. k is symmetric about the origin, and for some $a \in (0, 1)$ is bounded away from zero on $[-a, a]$.

[10] In theory, b_T is a null sequence of positive numbers such that $Tb_T \to \infty$.

Using the SEL-estimated conditional distribution, the cost of all one-period consumption fluctuations at each date (or state) t can be calculated as:

$$\frac{\frac{V_t\left(C_{t+1}^{stab}\right)}{C_t}}{\frac{V_t(C_{t+1})}{C_t}} - 1 = \frac{\mathbb{E}^{\mathbb{F}_t}\left[(\Delta C_{t+1})^{-\gamma}(1+\mu_c)\,|\mathcal{F}_t\right]}{\mathbb{E}^{\mathbb{F}_t}\left[(\Delta C_{t+1})^{-\gamma}(\Delta C_{t+1})\,|\mathcal{F}_t\right]} - 1 = \frac{(1+\mu_c)\sum_{j=1}^{T}\widehat{f}_{t,j}^{SEL}\times(\Delta c_j)^{-\gamma}}{\sum_{j=1}^{T}\widehat{f}_{t,j}^{SEL}\times(\Delta c_j)^{1-\gamma}} - 1. \quad (30)$$

Finally, note that the question naturally arises as to the economic interpretation of the recovered ψ-component of the kernel. For instance, it could capture a misspecification of investors' risk preferences relative to the power utility SDF. Alternatively, the ψ-component could capture investors' subjective beliefs about future macroeconomic and financial outcomes. While our approach does not need to take a stance on the identity of this component, Ghosh and Roussellet (2019) present evidence in favour of the latter interpretation. Specifically, they show that the beliefs about the stock market co-move positively with both Robert Shiller's survey data on institutional investors' confidence in the stock market as well as the Livingston Survey, and the beliefs about inflation are strongly correlated with inflation forecasts contained in the Survey of Professional Forecasters.

II.2.3 Properties of the Estimators

The asymptotic properties of the EL/ET and SEL/SET estimators of a finite-dimensional parameter vector (the SDF parameter γ in our setting) have been studied in the case of a *correctly specified model* – see, e.g., Qin and Lawless (1994) and Kitamura and Stutzer (1997) for the EL and ET estimators, respectively, and Kitamura, Tripathi, and Ahn (2004) for the SEL estimator. In this context, a correctly specified model refers to the setting in which the ψ-component of the SDF in Equation (1) is degenerate, i.e. $\psi \equiv 1$, and, therefore, the physical measure $\mathbb{P}$ equals the distorted measure $\mathbb{F}$ needed for the SDF to price assets.

Our framework differs from the above literature in two important respects. First, our starting premise is that the standard power utility model is insufficient to price assets, as evidenced by a large existing literature, and, therefore, this model-implied SDF needs to be augmented by an additional ψ-component, i.e. $\mathbb{P} \neq \mathbb{F}$. In other words, model misspecification emerges naturally in our set-up and this alters the properties of the above estimators. Second, instead of estimates of the SDF parameter γ, we are more interested in the missing ψ-component of the SDF. This involves an important extension of the econometric results in the above papers.

Ghosh and Otsu (2020) show that, in this set up, under mild regularity conditions, the estimated ψ-distribution converges in probability to its pseudo true value – the distribution that is minimally distorted with respect to the physical measure $\mathbb{P}$, within the class of models parametrized by a known SDF (the power utility kernel in our setup). More formally, under suitable regularity conditions, such as the ones in Komunjer and Ragusa (2016), the probability limit of the estimated distribution function can be interpreted as the information projection by the relative entropy divergence from the data generating distribution function $\mathbb{P}$ to the set of distribution functions satisfying the moment restrictions given by the Euler equations (see the constraints in Equations (16) and (17)). While our approach is nonparametric, not relying on a fully specified SDF, the above property of the estimator parallels that of parametric maximum likelihood estimators for misspecified models (see, e.g., White (1982), Vuong (1989)).

In the absence of knowledge of the true SDF, this methodology offers a more robust approach to measuring the cost of fluctuations. At the least, it offers a valuable alternative relative to fully structural approaches, having well-behaved asymptotics and, as we show in the following section, finite-sample behaviour. Moreover, simulation evidence, presented in Appendix D, suggests that this approach accurately recovers the cost of aggregate fluctuations for empirically realistic sample sizes. Specifically, our results suggest that the latter conclusion holds in both correctly specified settings where the econometrician has knowledge of the true SDF as well as in misspecified settings where, in the absence of knowledge of the true SDF, the econometrician erroneously uses the power utility SDF when recovering the ψ-component of the kernel and using it to measure the cost of fluctuations.

III The Market Value of Aggregate Uncertainty

In this section, we use the I-SDF, extracted using the information-theoretic procedure outlined in Section II, to obtain the cost of aggregate consumption fluctuations.

Before presenting the empirical results, we turn to a discussion of the SDF parameter γ that enters the welfare cost calculations (see, e.g., Equations (22)–(23) for the expected cost and Equation (30) for the time-variation in the cost). As highlighted in Section II, virtually all representative agent consumption-based models proposed in the literature imply the multiplicative form for the SDF assumed in this paper, $M_t = \left(\frac{C_t}{C_{t-1}}\right)^{-\gamma} \psi_t$. Different models use different calibrations for the utility curvature parameter γ.

For example, in the time and state separable power utility model, γ is the CRRA of the representative agent and an upper bound of 10 is generally considered plausible for it. However, much higher levels of risk aversion are needed for the model to explain several observed features of financial market data.

In models with Epstein and Zin (1989) recursive preferences,

$$M_t = \delta^\eta \left(\frac{C_t}{C_{t-1}}\right)^{-\frac{\eta}{\rho}} R_{c,t}^{\eta-1} = \delta^\eta \left(\frac{C_t}{C_{t-1}}\right)^{-\gamma} \underbrace{\left(\frac{\frac{P_{c,t}}{C_t}+1}{\frac{P_{c,t-1}}{C_{t-1}}}\right)^{\eta-1}}_{\psi_t},$$

where ρ is the elasticity of intertemporal substitution, $\eta = \frac{1-\gamma}{1-\frac{1}{\rho}}$, R_c (the return on total wealth) denotes the latent gross return on an asset that delivers aggregate consumption as its dividend, P_c/C is its price-to-consumption ratio, and the second equality follows from factorizing out consumption growth from the return on total wealth.[11] These models typically calibrate $\gamma = 10$ (see, e.g., Bansal and Yaron (2004)). Some models with recursive preferences calibrate γ to much larger values (e.g., Piazzesi and Schneider (2007)).

In models with external habit formation (see, e.g., Campbell and Cochrane (1999)), $M_t = \delta \left(\frac{C_t}{C_{t-1}}\right)^{-\gamma} \left(\frac{S_t}{S_{t-1}}\right)^{-\gamma}$, where S_t is the surplus consumption ratio and γ the utility curvature parameter that determines the time-varying risk aversion $\frac{\gamma}{S_t}$. Campbell and Cochrane (1999) calibrate $\gamma = 2$. However, Ghosh, Julliard, and Taylor (2017) show that the model needs a higher γ (typically in excess of 7) to satisfy entropy bounds for admissible SDFs, that are tighter than the seminal variance bounds of Hansen and Jagannathan (1991). In models with complementarities in consumption, (see e.g., Piazzesi, Schneider, and Tuzel (2007)), $M_t = \delta \left(\frac{C_t}{C_{t-1}}\right)^{-\gamma} \left(\frac{A_t}{A_{t-1}}\right)^{\frac{\gamma\zeta-1}{\zeta-1}}$, where A_t is the expenditure share on non-housing consumption, γ^{-1} is the intertemporal elasticity of substitution, and ζ is the intratemporal elasticity of substitution between housing services and non-housing consumption. The authors' consider two alternative calibrations of $\gamma = 5$ and $\gamma = 16$. However, Ghosh, Julliard, and Taylor (2017) show that the model needs a higher γ (typically in excess of 20) to satisfy entropy bounds for admissible SDFs. In models with rare disasters, γ is typically calibrated to values between 3 and 4 (see, e.g., Barro (2006), Wachter (2013)). However, Julliard and Ghosh (2012) show that these models need much higher levels of risk aversion, typically in excess of 20, to explain the equity premium puzzle. To summarize, most models in the literature either calibrate the SDF parameter γ to 10 or higher values and/or require such high values of the parameter to explain asset prices.

Also, in addition to recovering the ψ-component of the SDF, our information-theoretic procedure offers a way to estimate γ. For instance, the EL estimator of γ is defined as (see

[11] This re-encoding of the SDF in terms of consumption growth and the price-consumption ratio is also used, e.g., in Bansal and Yaron (2004, p. 1487).

Kitamura (2006)):

$$\widehat{\gamma}^{EL} = \min_{\gamma} \min_{\mathbb{F}} \int \log\left(\frac{d\mathbb{P}}{d\mathbb{F}}\right) d\mathbb{P} \equiv \int_{\mathbf{z}} \log\left(\frac{p(\mathbf{z})}{f(\mathbf{z})}\right) p(\mathbf{z}) d\mathbf{z}, \text{ s.t. } \mathbf{0} = \int_{\mathbf{z}} \mathbf{R}^e(\mathbf{z}) \left(\Delta C(\mathbf{z})\right)^{-\gamma} f(\mathbf{z}) d\mathbf{z}. \tag{31}$$

In other words, the EL approach searches for a value of γ in the admissible parameter space that minimizes the KLIC divergence between the $\mathbb{F}$ and $\mathbb{P}$ measures, subject to the Euler equation constraints. The ET estimator of γ is similarly defined, albeit swapping the roles of $\mathbb{P}$ and $\mathbb{F}$.

We estimate γ in our baseline 1929–2015 sample, using total consumption expenditures as the measure of aggregate consumption and the excess returns on the market as the sole test asset (see Appendix E for a description of the data). Figure 1 plots the EL objective function in Equation (31) as a function of γ. The point estimate of γ is 22.1 (red dotted line). A similar point estimate of 26.2 is obtained when nondurables and services consumption is used as the measure of aggregate consumption expenditures. Identical point estimates are obtained with the ET approach.

Figure 1: Profile Likelihood

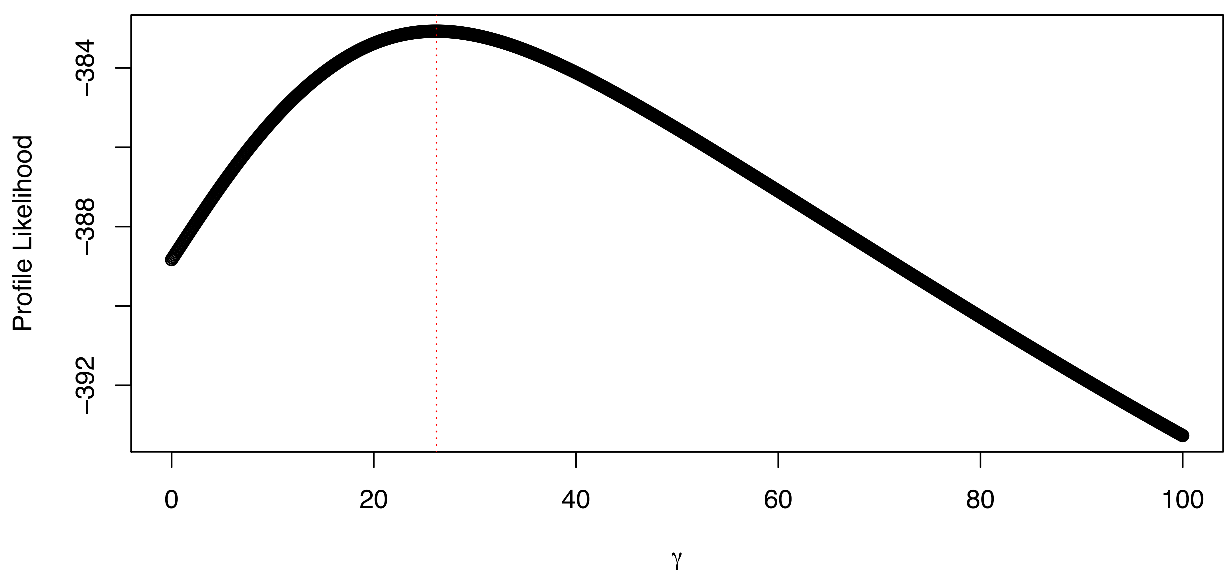

Note: The figure plots the EL objective function as a function of the SDF parameter γ. The dotted vertical line denotes the point estimate of γ. Consumption denotes the real personal total consumption expenditure (includes durables, nondurables, and services). The excess return on the market portfolio is the sole test asset. The sample is annual, covering the period 1929-2015.

Motivated by the observations that most theoretical models calibrate γ to 10 or higher values and that the point estimate obtained in the historical sample is much higher, we set

$\gamma = 10$ in our baseline results as a conservative benchmark. Note that higher values of γ serve to further increase the marginal utility of the representative agent in bad states with low realizations of the consumption growth rate and, therefore, would further increase the estimates of the cost of consumption fluctuations. We also assess the sensitivity of our results to alternative choices of γ.

We next proceed to estimate the cost of fluctuations. Section III.1 presents the term structure of the expected (or, average) cost of consumption fluctuations. Section III.3 reports the nature of time-variation in the cost. Finally, in Appendix F, we present a host of robustness checks, including alternative definitions of relative entropy, an alternative longer sample period going back as far as 1890, and an alternative larger cross section of base assets used in the recovery of the I-SDF.

III.1 The Average Cost of All Consumption Uncertainty

Recall that, rather than estimating the cost of aggregate consumption fluctuations over an infinite time horizon, we focus on the term structure of the cost for finite time periods. Specifically, we estimate the (cumulative) cost for one- to ten-year time horizons.

Equation (13) defines the expected cost of one-period fluctuations. The cost is percentage difference between the prices of two hypothetical securities: a claim to a stabilized consumption in the next period, $\tilde{pc}_1^{stab}$, and a claim to the actual consumption next period, $\tilde{pc}_1$. When measuring the cost of all fluctuations, stabilized consumption refers to a consumption path from which all fluctuations have been removed, i.e. consumption growth in each period is replaced with its unconditional mean. When measuring the cost of business cycle fluctuations, on the other hand, stabilized consumption refers to the residual after the business cycle component has been removed from the aggregate consumption series. We compute the stabilized consumption series using the widely used Hodrick-Prescott filter. Since our empirical analysis uses annual data, we use a smoothing parameter of 6.25 in the application of the Hodrick-Prescott filter, following the suggestions in Ravn and Uhlig (2002).

Equations (22)–(23) reveal that the prices of these two securities and, therefore, the cost of one-period consumption fluctuations, depend on the SDF. We use the I-SDF, recovered using the EL approach, to measure this cost.[12]

The costs of multi-year fluctuations are obtained with a multi-period SDF constructed by compounding the I-SDF, as explained in Section II.2. Note that the recovered I-SDF depends on the particular measure of the aggregate consumption expenditures as well as on the set of base assets used for its construction (see Equations (19)-(20)). To ensure robustness, we estimate the I-SDF using two different measures of consumption expenditures and two

[12] Very similar results are obtained with the ET approach and are reported in Appendix F.

alternative sets of assets. The latter includes (a) the market return and (b) the returns on the 6 Fama-French size and book-to-market equity sorted portfolios. Our choice of base assets is motivated by two arguments. First, existing empirical evidence shows that size and book-to-market equity sorted portfolios have significant predictive power for consumption and GDP growth rates, not only in the United States but also in ten developed markets (see, e.g., Liew and Vassalou (2000), Parker and Julliard (2005)). They, therefore, constitute appropriate base assets to infer agents' preferences toward consumption risk. Second, the equity premium is, perhaps, the most robust feature of stock market data for over a century. Over 300 risk factors have been proposed in the literature, and many of the associated risk premia have disappeared or greatly diminished over time including ones that were once believed to be the most robust (e.g., the size and value premia). This makes the market return the most natural choice of test asset.[13]

The estimates of the cost of consumption fluctuation are presented in Table 1. Panel A presents results when consumption refers to the expenditures on nondurables and services, while Panel B does the same for total consumption expenditures (including durables). Consider first Panel A. In Row 1, only the market portfolio is used as base asset for recovering the I-SDF. Column 2 shows that in this case the estimated cost of all one-period consumption fluctuations is 1.53%. In Row 2, the six size and book-to-market-equity sorted portfolios of Fama-French are used to recover the I-SDF. Column 2 shows that the estimated cost of all one-period consumption fluctuations remains quite similar at 1.29%. Row 3 shows that the one-year cost, estimated using the pricing kernel implied by power utility preferences with a constant CRRA (hereafter referred to as the CRRA kernel), is smaller at .93%. And, Row 4 shows that, if the assumption of lognormal consumption growth is imposed on the CRRA kernel – this corresponds to Lucas' original specification – the cost of one-period consumption fluctuations further reduces to .75%.

Note that the above results pertain to the cost of one-period fluctuations alone. Columns 3, 4, 5, and 6 of Panel A present the costs of all consumption fluctuations over two, three, four, and five year horizons, respectively. Row 1 shows that, when the market portfolio alone is used to recover the I-SDF, the costs of consumption fluctuations over two, three, four, and five years increase to 5.2%, 11.8%, 14.3%, and 14.4%, respectively. That is, the cost of consumption fluctuations over two years is more than three times higher than the cost of fluctuations over one year alone (5.2% versus 1.5%). Similarly, the cost of consumption fluctuations over a three-year period is more than seven times higher than the cost over one

[13] Our focus on the term structure of the cost of fluctuations might make it seem natural to use dividend strips as base assets. However, data on dividend strips are only available over relatively short time periods (typically the early 2000s). This feature limits their usage in our setting that relies on longer time periods so as to have a greater coverage of the different possible economic states.

Table 1: Cumulative Cost of All Consumption Fluctuations

	1 Yr	2 Yr	3 Yr	4 Yr	5 Yr
Panel A: Nondurables & Services Consumption					
I-SDF (Mkt)	1.53	5.15	11.75	14.28	14.44
	[.41,2.24]	[.74,6.61]	[.96,13.74]	[1.12,18.60]	[1.24,21.87]
	[1.28,2.50]	[3.35,7.25]	[5.03,15.98]	[5.72,21.07]	[6.06,25.01]
I-SDF (FF6)	1.29	3.52	6.65	10.63	11.20
	[.10,1.62]	[.29,4.20]	[.33,8.37]	[.30,13.85]	[.24,17.75]
	[.38,2.02]	[.81,5.61]	[1.01,11.81]	[1.07,19.73]	[1.13,23.97]
CRRA Kernel	.93	2.08	3.73	4.87	5.03
Lucas	.75	1.09	1.40	1.68	1.94
$\left(\frac{C_t}{C_{t-1}}\right)^{-\gamma_1} R_{M,t}^{-\gamma_2}$ (Mkt)	1.23	3.30	6.67	7.87	7.93
$\left(\frac{C_t}{C_{t-1}}\right)^{-\gamma_1} R_{M,t}^{-\gamma_2}$ (FF6)	1.31	3.68	7.57	8.70	8.70
Panel B: Total Consumption					
I-SDF (Mkt)	2.15	6.77	16.13	19.65	19.73
	[.67,3.12]	[.1.17,8.67]	[1.48,18.67]	[1.68,24.86]	[1.87,28.79]
	[1.82,3.42]	[4.59,9.46]	[6.77,21.62]	[7.39,27.99]	[8.09,32.74]
I-SDF (FF6)	1.88	4.89	9.46	15.00	15.57
	[.27,2.26]	[.68,5.64]	[.85,11.41]	[.84,18.73]	[.83,23.64]
	[.66,2.75]	[1.35,7.17]	[1.70,15.30]	[1.82,25.32]	[1.92,30.73]
CRRA Kernel	1.42	3.08	5.80	7.63	7.77
Lucas	1.15	1.68	2.16	2.61	3.03
$M_t = \left(\frac{C_t}{C_{t-1}}\right)^{-\gamma_1} R_{M,t}^{-\gamma_2}$ (Mkt)	1.76	4.61	9.80	11.67	11.69
$M_t = \left(\frac{C_t}{C_{t-1}}\right)^{-\gamma_1} R_{M,t}^{-\gamma_2}$ (FF6)	1.86	5.13	11.09	12.84	12.81

The table reports the (cumulative) costs of *all* aggregate consumption fluctuations over one-to five-year horizons. Panel A presents results when consumption denotes the real personal consumption expenditure of nondurables and services, while Panel B does the same for total personal consumption expenditure (that includes durables). In each panel, the costs are calculated using the I-SDF recovered from the market portfolio alone (Row 1), the I-SDF recovered from the six size and book-to-market-equity sorted portfolios of Fama and French (Row 2), the kernel implied by power utility preferences with a constant CRRA (Row 3), and Lucas' original specification that involves power utility preferences and i.i.d. lognormal aggregate consumption growth dynamics (Row 4), a kernel log-linear in aggregate consumption growth and the market return with the coefficients estimated with the GMM approach using as test assets the market portfolio alone (Row 5) and the six size and book-to-market-equity sorted portfolios of Fama and French (Row 6). Under the I-SDF based estimates we report two sets of bootstrapped 90% confidence intervals, with the bottom one computed conditional on the simulated samples having at least one observation from the Great Depression period. The data sample is annual covering the period 1929-2015.

year alone (11.8% versus 1.5%); and the costs over four- and five-year periods are each almost ten times higher than the cost over one year (14.3% and 14.4%, respectively, versus 1.5%). Similarly, row 2 shows that, even with a broader set of base assets, the cost of consumption fluctuation grows more than proportionally with the time horizon, with e.g. the five-year cost being almost ten times larger than the one-year cost.

Computing confidence bands for the estimated costs of fluctuations is challenging since the annual data sample contains only 86 independent observations, hence asymptotic theory would hardly be applicable, especially for the long horizon estimates. Consequently, we resort to a conservative bootstrap approach, and we report centered 90% confidence region for the I-SDF based estimates. Given the small number of observations, we compute these unconditionally (across all bootstrapped samples, top numbers in brackets) and conditionally across only those bootstrapped samples containing at least one Great Depression type episode. As the table shows, bootstrapped samples with no Great Depression observations cause a large downward bias in the estimates, as expected from the rare events based asset pricing literature (see, e.g., Barro (2005), Wachter (2013)). Overall, these confidence bands are wide, not surprisingly more so for long horizons (since these are computed via block bootstrap) and when we use more base assets (hence having a larger degree of parameter uncertainty). Nevertheless, when we avoid the rare events induced bias (and in particular when focusing on the market as the sole base asset to reduce sampling uncertainty), the confidence regions still indicate large costs of consumption fluctuation with very high probability.

Row 3 shows that the CRRA kernel implies much smaller costs of two, three, four, and five year consumption fluctuations of 2.1%, 3.7%, 4.9%, and 5.0%, respectively. In fact, the costs are an order of magnitude smaller than the costs implied by the I-SDF (with the exception of the two-year fluctuations that is also less than half of that implied by the I-SDF). Lucas' kernel in Row 4 implies even smaller costs of 1.1%, 1.4%, 1.7%, and 1.9% at two-, three-, four-, and five-year horizons, respectively.

Next, we present evidence that our non-linear adjustment to the pricing kernel (the ψ-component that we recover is a highly nonlinear function of consumption growth and the base assets) has certain desirable properties relative to alternative linear (or log-linear) adjustments that have been proposed in the literature, and also produces markedly different results.

First (not reported in the table), we have constructed an SDF based on the Hansen and Jagannathan (1997) linear adjustment to the CRRA kernel SDF to enable it to successfully price assets.[14] This approach, due to the linear dependency of the SDF on returns, often yields a negative SDF (from 1.2% to 5.2% of the times depending on the horizon), and implausible costs of consumption fluctuations ranging from 1.76% at the one-year horizon to 43.7% at the 5-year horizon to 304.0% at the ten-year horizon, with extremely wide confidence bands. This highlights the instability arising from the use of (compounded) linear SDFs based on tradable factors (see, e.g., Chernov, Lochstoer, and Lundeby (2021)).

Second, to preserve the non-negativity of the pricing kernel, we consider an SDF that is log-linear in the aggregate consumption growth rate and the market return. This specification corresponds to the SDF implied by Epstein and Zin (1989) recursive preferences with the

[14]This yields the SDF

$$M_{t+1} = (C_{t+1}/C_t)^{-\gamma} - \frac{\sum_{t=1}^{T}\left((C_{t+1}/C_t)^{-\gamma}\right)(R_{M,t+1} - R_{F,t})}{\sum_{t=1}^{T}(R_{M,t+1} - R_{F,t})^2}(R_{M,t+1} - R_{F,t}). \quad (32)$$

where, to make the results comparable with those obtained using the I-SDF, we set $\gamma = 10$.

stock market return used as a proxy for the return on the total wealth portfolio:

$$M_{t+1} = \exp\{(C_{t+1}/C_t)^{-\gamma_1}\, R_{M,t+1}^{\gamma_2}\}. \tag{33}$$

To make the results comparable with those obtained using the I-SDF, we set $\gamma_1 = \gamma = 10$ and estimate γ_2 using GMM on the Euler equation for the excess market return (row 5) or the six Fama-French portfolios used to construct their celebrated SMB and HML factors (row 6). These SDFs yield substantially smaller estimates of the cost of fluctuations compared to the I-SDF. Specifically, the cost of all fluctuations at the one- to five-year horizons are 1.2%, 3.3%, 6.7%, 7.9%, and 7.9%, respectively. With the exception of the one year horizon, these estimates are about half of those obtained with I-SDF, suggesting that the later captures priced consumption risk that is missed by the simple Epstein-Zin pricing kernel.

But why does the ψ component of the I-SDF increase the cost of consumption fluctuations so much more than a (log-linear) pricing kernel that is based on the same base assets? The reason lies in the differential effects that the added term, relative to CRRA, has on the covariance of consumption growth with the resulting SDF. For instance, the Epstein-Zin formulation (in row 5) increases (in absolute terms) the covariance between consumption growth and the SDF at the one (five) year horizon by about 31% (22%), while the I-SDF does so by about 64% (164%), hence implying a much riskier consumption process. As discussed in Section IV below, where we discuss the sensitivity of our results to the curvature parameter γ, the I-SDF always yields an SDF that has strong negative correlation with consumption growth – even when we exclude consumption from the pricing kernel (when $\gamma = 0$, this correlation is about -35%).

Figure 2 sheds light on how the estimated ψ component adjusts the amount of consumption risk in the pricing kernel to best fit the data. As outlined in Panel A, as γ varies from $[1, 20]$ – the economically plausible range that contains the calibrated values from most asset pricing models – the estimated ψ strongly negatively correlates with consumption growth. That is, the I-SDF uses ψ to add more consumption risk to the pricing kernel. In other words, the ψ-component of the SDF helps to magnify the increased marginal utility during bad states characterized by low realizations of consumption growth, thereby implying a higher cost of consumption fluctuations relative to that arising from the CRRA kernel alone.

Instead, as the γ coefficient becomes too large (in excess of its point estimate of 26.2), Panel A shows that our method detects too much consumption risk in the pricing kernel to be consistent with the base asset returns. Consequently, the estimated correlation of ψ and consumption growth over this range of the curvature parameter γ becomes *positive*. That is, the relative entropy finds the consumption risk in the CRRA pricing kernel for such high γ to be excessive, and uses ψ to reduce it. In other words, for these high values of γ, the

Figure 2: Correlation $\left(\frac{C_t}{C_{t-1}}, \psi_t\right)$ and time series of ψ for Alternative values of γ

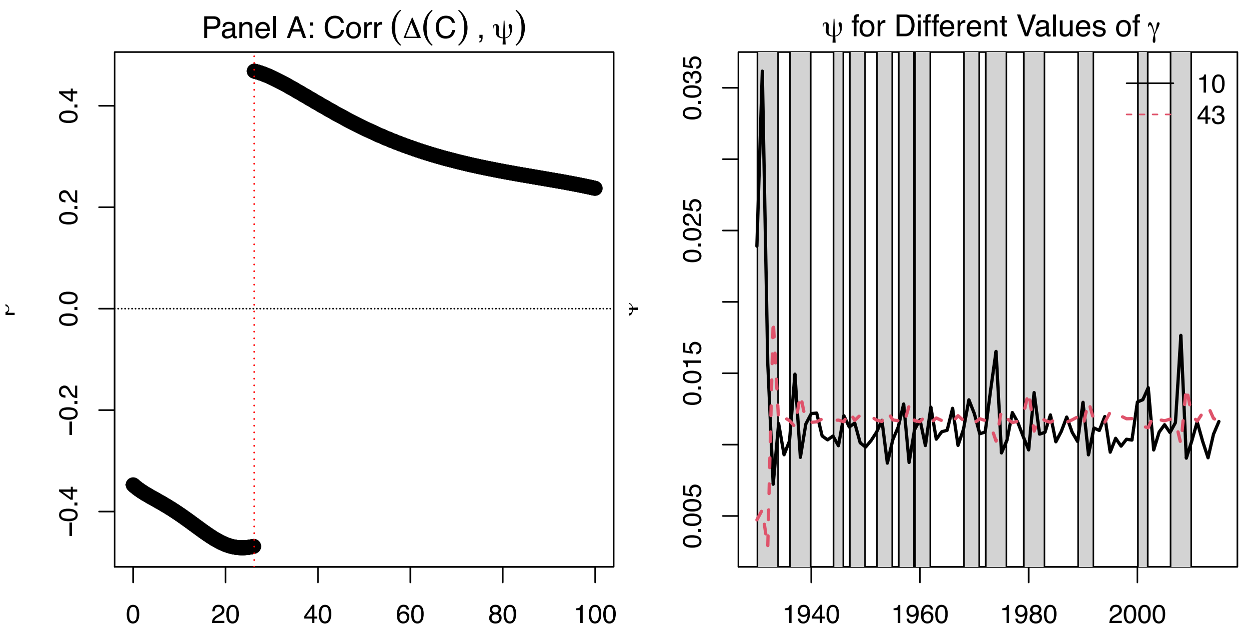

Notes: Panel A plots the correlation between the two multiplicative components of the pricing kernel, namely the recovered ψ and consumption growth, for alternative values of the SDF parameter γ. Panel B plots the time series of the recovered ψ for two different values of γ, namely those corresponding to our baseline value of 10 (black solid line) and a value of 43 that is symmetrically above its point estimate of 26.2 (red-dashed line). Consumption denotes the real personal consumption expenditure of nondurables and services and the excess return on the market portfolio is the sole test asset used in the recovery of the ψ component using the EL approach. The sample is annual covering the period 1929-2015.

ψ-component serves to counteract the effect of the consumption growth component, i.e., it serves to mitigate the extreme effects of low realizations of consumption growth on the marginal utility. To stress this point further, Panel B of Figure 2 plots the time series of the recovered ψ for two alternative values of γ – our baseline value of 10 (black solid line), which is below its EL point estimate, and a value of 43 (red-dashed line) that is symmetrically above the point estimate. The figure shows clearly the strong negative co-movement between the two time series – the correlation between them is -74.8% – consistent with the findings in Panel A that the estimated ψ adds consumption risk to the pricing kernel for low values of γ, while it reduces it for too high values.

As Panel A of the figure also highlights, our calibrated baseline value of $\gamma = 10$ is conservative, since higher values (not exceeding 26.2), would imply much higher costs of consumption fluctuations as reported and discussed in section IV below. In particular, as discussed therein, when we estimate, rather than calibrate, γ, the cost of consumption fluctuations roughly doubles.

Figure 3: Marginal Cost of All Consumption Fluctuations, 1929-2015

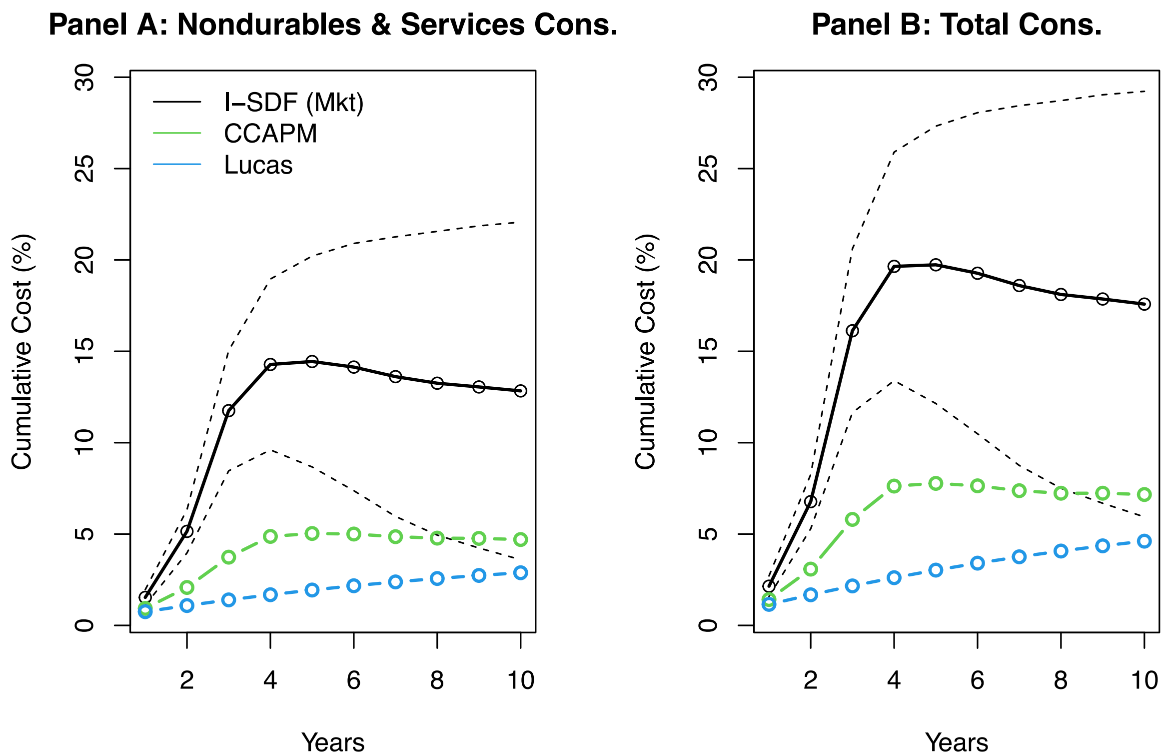

Notes: The figure plots the cumulative costs of *all* aggregate consumption fluctuations over one- to ten-year horizons, for different choices of the pricing kernel and measures of consumption. Panel A presents results when consumption refers to the real personal consumption expenditure of nondurables and services, while Panel B does the same when consumption denotes the total personal consumption expenditure. The costs are presented for the I-SDF extracted with the EL approach using the excess return on the market portfolio as the sole test asset (black line) along with bootstrapped two standard deviation bands (black dashed lines), the pricing kernel implied by power utility preferences with a constant CRRA (green line), and Lucas' original specification that involves power utility preferences and i.i.d. lognormal aggregate consumption growth dynamics (blue line).

But does the cost of consumption fluctuation keep on growing after the first 5 year horizon considered in Table 1? Figure 3 plots the term structure of the cost of fluctuations over one- to ten-year horizons. The black solid line corresponds to the costs obtained with the I-SDF recovered with the market portfolio as the sole test asset. The green and blue lines denote, respectively, the costs estimated with the CRRA kernel and Lucas' specification. The figure highlights that the costs of consumption fluctuations implied by the I-SDF are much higher, at any horizon, relative to those obtained with Lucas' and CRRA specifications. Furthermore, the figure shows that the I-SDF implied costs flatten out, at about 20% (14%) for total (nondurables and services) consumption, at the business cycle frequency i.e. about 4–5 year horizon. That is, consumption fluctuations beyond this horizon do not seem to be perceived as particularly costly. Nevertheless, since we have only 86 observation to (block) bootstrap from, the confidence regions become particularly wide at longer horizons. Interestingly, this term structure behaviour matches the finding in Bryzgalova, Huang, and Julliard (2024b) that, in estimating the risk premia demanded by an asset that delivers a return equal to the aggregate consumption growth rate, using a very different method, data, and a very large cross-section of base assets, finds a steeply increasing term structure of consumption risk that flattens out exactly at the same horizon.

III.2 The Average Cost of Business Cycle Fluctuation

We now turn to the estimation of the costs of only business cycle consumption fluctuations.

Table 2 Panels A and B report, respectively, estimates for nondurables and services, and total consumption. Panel A, Row 1 shows that, using the I-SDF extracted from the market portfolio alone the cost of business cycle fluctuations in nondurables and services consumption over a one-year time horizon is estimated to be 0.56%. This cost grows more than proportionally with the time horizon and is seven times as large at the 4 year horizon. Row 2 shows similar results when using the six size and book-to-market-equity sorted portfolios as base assets. Row 3 shows that, for the CRRA kernel, while the cost of business cycle fluctuations over a one-year period is similar to that obtained with the I-SDF (0.5% versus 0.5%–0.6%), the cost increases little for multi-year horizons. For instance, the cost of five-year fluctuations is only 1.4% – less than half of the cost implied by the I-SDF in Rows 1 and 2. Interestingly, even the Epstein-Zin kernels (in Rows 4–5) show increasing costs of business cycle fluctuations when moving from the one- to the five-year horizon. Nevertheless, the point estimates with these kernels are always smaller than those obtained with the I-SDF.[15] Panel B shows similar patterns, but with I-SDF implied costs of business cycle fluctuation being even larger (4.1%-5.1%) at the five-year horizon.

An important point to note is that, while the estimates of the costs of business cycle fluctuations in Table 2 are smaller than the corresponding costs of all consumption uncertainty in Table 1, the former nonetheless represents a substantial fraction of the latter. For instance, Panel A, Row 1 of both Tables shows that the cost of business cycle fluctuations constitutes 36.3% of the cost of all consumption fluctuations over a one-year horizon. Similarly, the costs of business cycle fluctuations over two, three, four, and five years account for 28.7%, 28.9%, 27.3%, and 24.7%, respectively, of the cost of all fluctuations at these horizons. Similarly, when the 6 FF portfolios are used for the recovery of the I-SDF in Row

[15] Note that we do not report the pure business cycle costs in the Lucas formulation since, given the i.i.d. assumption in that setting, these are identical to the costs of all consumption fluctuations previously discussed.

Table 2: Cumulative Cost of Business Cycle Fluctuations

	1 Yr	2 Yr	3 Yr	4 Yr	5 Yr
Panel A: Nondurables & Services Consumption					
I-SDF (Mkt)	.56	1.48	3.39	3.90	3.57
	[.15, 1.11]	[.29, 2.17]	[.35, 3.88]	[.36,5.11]	[.37, 5.96]
	[.43, 1.24]	[1.27, 2.32]	[2.13, 4.31]	[2.32, 5.70]	[2.37,7.06]
I-SDF (FF6)	.46	1.03	2.07	3.03	2.90
	[.01, .68]	[.11, 1.42]	[.16, 2.59]	[.16, 4.01]	[.15, 5.07]
	[.16, .76]	[.38, 1.53]	[.55, 3.26]	[.62, 5.31]	[.67, 6.64]
CRRA Kernel	.46	.85	1.32	1.52	1.40
$\left(\frac{C_t}{C_{t-1}}\right)^{-\gamma_1} R_{M,t}^{-\gamma_2}$ (Mkt)	.53	1.15	2.15	2.38	2.20
$\left(\frac{C_t}{C_{t-1}}\right)^{-\gamma_1} R_{M,t}^{-\gamma_2}$ (FF6)	.54	1.23	2.38	2.60	2.41
Panel B: Total Consumption					
I-SDF (Mkt)	.90	2.09	4.85	5.55	5.12
	[.32, 1.61]	[.56, 2.98]	[.65, 5.48]	[.68, 7.10]	[.69, 8.21]
	[.75, 1.75]	[1.85, 3.12]	[2.93, 6.15]	[3.06, 7.97]	[3.12, 9.69]
I-SDF (FF6)	.77	1.60	3.05	4.35	4.14
	[.11, 1.02]	[.27, 2.11]	[.32, 3.67]	[.31, 5.61]	[.29, 7.88]
	[.32, 1.14]	[.67, 2.18]	[.91, 4.53]	[1.00, 7.18]	[1.08, 8.87]
CRRA Kernel	.76	1.32	2.08	2.40	2.21
$M_t = \left(\frac{C_t}{C_{t-1}}\right)^{-\gamma_1} R_{M,t}^{-\gamma_2}$ (Mkt)	.84	1.68	3.20	3.54	3.27
$M_t = \left(\frac{C_t}{C_{t-1}}\right)^{-\gamma_1} R_{M,t}^{-\gamma_2}$ (FF6)	.86	1.79	3.55	3.87	3.58

The table reports the (cumulative) costs of business cycle fluctuations in consumption, over one-to five-year horizons. Panel A presents results when consumption denotes the real personal consumption expenditure of nondurables and services, while Panel B does the same for total personal consumption expenditure (that includes durables). In each panel, the costs are calculated using the I-SDF recovered from the market portfolio alone (Row 1), the I-SDF recovered from the six size and book-to-market-equity sorted portfolios of Fama and French (Row 2), the kernel implied by power utility preferences with a constant CRRA (Row 3), a kernel log-linear in aggregate consumption growth and the market return with the coefficients estimated with the GMM approach using as test assets the market portfolio alone (Row 4) and the six size and book-to-market-equity sorted portfolios of Fama and French (Row 5). Under the I-SDF based estimates we report two sets of bootstrapped 90% confidence intervals, with the bottom one computed conditional on the simulated samples having at least one observation from the Great Depression period. The sample is annual covering the period 1929-2015.

2, the costs of business cycle fluctuations over one to five years account for 35.8%, 29.3%, 31.1%, 31.0%, and 25.9%, respectively, of the costs of all consumption fluctuations over these time horizons.

Figure 4, Panel A plots the term structure of the cost of all consumption fluctuations (solid line) and business cycle fluctuations in consumption (dashed line) at one to ten year horizons, when the market portfolio alone is as base asset to recover the I-SDF. The fairly large ratio of the cost of business cycle fluctuations to the cost of all consumption fluctuations, at all time horizons, is evident from the figure. Moreover, as with the cost of all fluctuations, the cost of business cycles seems to stabilize after 4–5 years, thereby suggesting well-behaved asymptotics of our approach.

Overall, two salient conclusions emerge from the results of the last two subsections. First,

Figure 4: Marginal Cost of All versus Business Cycle Consumption Fluctuations, 1929-2015

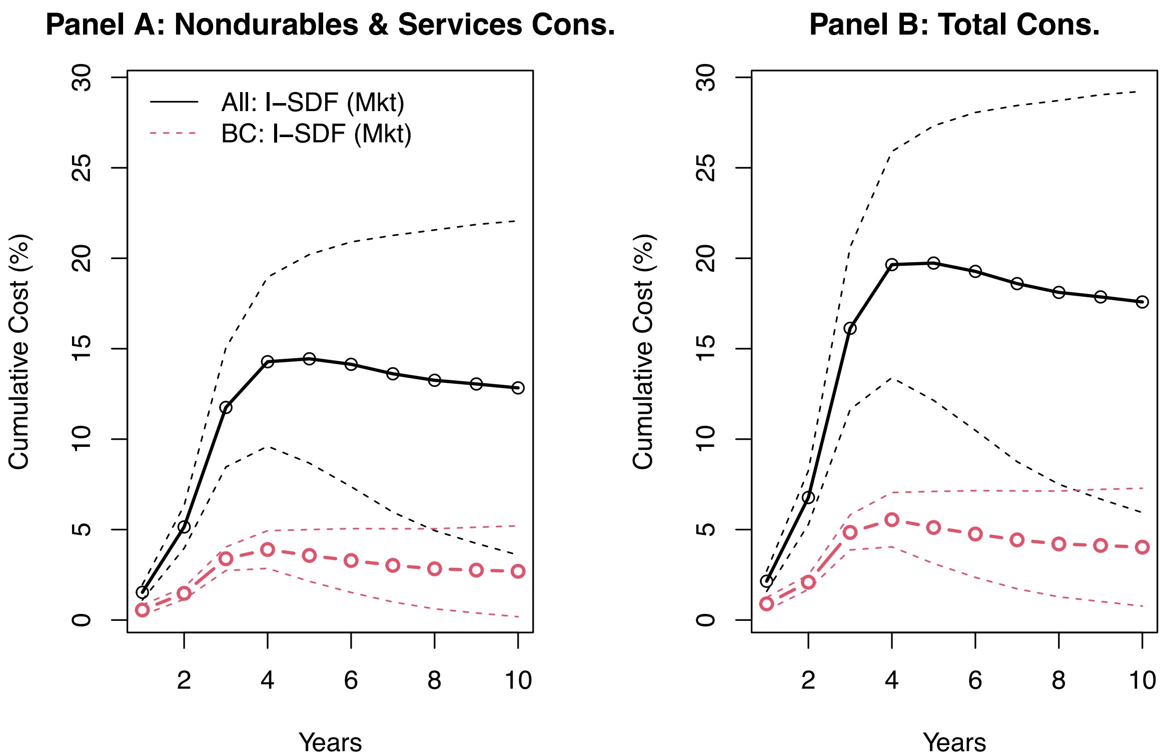

Notes: The figure plots the term structure of the (cumulative) cost of *all* aggregate consumption fluctuations (black solid line) and business cycle fluctuations in consumption (red line), over 1-10 years, obtained using the I-SDF. The black dashed lines and red dashed lines correspond to the bootstrapped two standard deviation confidence intervals for the cost of all and business cycle fluctuations, respectively. Panel A presents results when consumption refers to the real personal consumption expenditure of nondurables and services, while in Panel B consumption denotes total personal consumption expenditure. The I-SDF is extracted using the excess return on the market portfolio as the sole test asset with the EL approach. The sample is annual covering the period 1929-2015.

economic agents perceive the cost of aggregate economic uncertainty to be quite substantial. For instance, our point estimates of the cost of all consumption fluctuations over a five-year horizon vary from 11.2% to 19.7%, depending on the measure of aggregate consumption and the set of base assets used to recover the I-SDF. The cost is substantially higher than that originally obtained by Lucas as well as CRRA and Epstein-Zin preferences. Second, we find that the costs of business cycle fluctuations are large and constitute between a quarter to a third of the cost of all consumption fluctuations.

The second result is in sharp contrast with the finding in Alvarez and Jermann (2004) who argue that, while the cost of all consumption fluctuations is very high, the cost of business cycle fluctuations in consumption is miniscule, ranging from 0.1% to 0.5%. Our estimates of the cost of business cycle fluctuations over a cumulative five-year period are as high as 5.1%

– between ten and fifty times higher than the estimates in Alvarez and Jermann (2004). The question, therefore, naturally arises as to what drives this difference.

This discrepancy is driven, at least in part, by the choice of the smoothing filter used to remove business cycle variation from the historical consumption series in Alvarez and Jermann (2004).[16] We use the widely used Hodrick-Prescott (HP) two-sided filter to obtain a long run trend consumption series from which fluctuations corresponding to business cycle frequencies (fluctuations lasting less than eight years) have been removed. Alvarez and Jermann (2004) (AJ), on the other hand, use a one-sided filter, whereby trend consumption at time-t is expressed as a weighted average of $K(= 20)$ lags, with the coefficients chosen so as to represent a low-pass filter that lets pass frequencies that correspond to cycles of eight years and more. Figure 5 presents a comparison of the HP and AJ filters. The figure plots the historical consumption growth (blue line), the trend consumption growth obtained using the HP filter (red line), and the trend consumption growth obtained using the AJ filter (black line). Consumption refers to the total personal consumption expenditures. The figure shows that the HP filter delivers a smoother trend consumption growth relative to the AJ filter, and the trend component of the AJ filter instead closely tracks the yearly frequencies fluctuations in consumption.

In fact, it is known that a one-sided filter of the AJ type, with coefficients chosen to let pass frequencies that correspond to cycles of at least a given length, cannot fully eliminate higher frequency fluctuations. In other words, it also lets pass some fluctuations corresponding to higher frequencies. Consequently, in the context of the present application, the computed trend contains a non-negligible amount of business cycle variability (as evident in the figure). Moreover, the trend consumption growth is markedly different between the HP and AJ filters over our sample period. Specifically, the historical real consumption growth has a volatility of 3.4%, while the trend consumption growth obtained with the HP and AJ filters have volatilities of 1.9% and 2.8%, respectively. Thus, the trend growth obtained with the AJ filter has 46% higher volatility than that obtained with the HP filter. Therefore, not surprisingly, the cost of business cycles obtained with the AJ filter are lower than those obtained with the HP filter. As an illustration, when the I-SDF is recovered with the market portfolio as base asset, the cost of business cycle fluctuations over one- to five-year horizons takes values 0.9%, 2.1%, 4.8%, 5.5%, and 5.1%, respectively, with the HP filter. The corresponding costs obtained using the AJ filter are 0.5%, 1.2%, 2.6%, 2.2%, and 1.5%, respectively – still higher than the values reported in Alvarez and Jermann (2004) but smaller than those obtained with the HP filter.

Of course, a two-sided filter like the HP filter also has an undesirable feature, namely

[16]We thank Jaroslav Borovicka for pointing this out.

Figure 5: Comparison of HP and AJ Filters

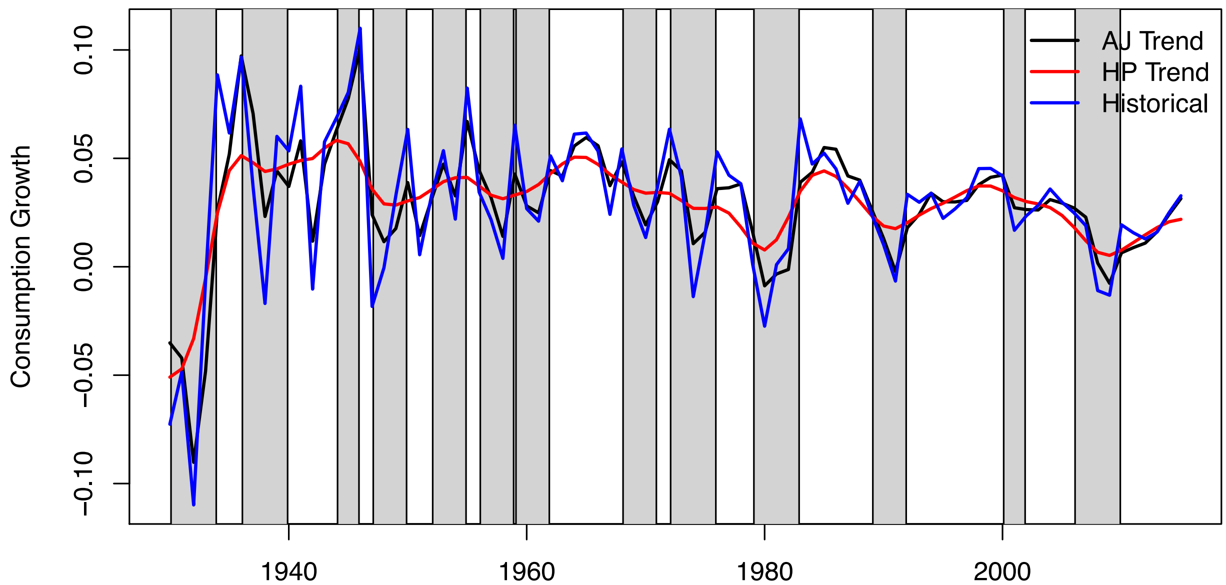

Notes: The figure plots the historical consumption growth (blue line), the trend consumption growth obtained using the HP filter (red line), and the trend consumption growth obtained using the AJ filter (black line). Consumption refers to the total personal consumption expenditures. The sample is annual over 1929–2015.

that the current filtered consumption is contaminated by future consumption realizations. In other words, the approach contaminates the current information set with future information, which may affect the covariance between consumption growth and the SDF, and, therefore, the pricing of consumption claims. To provide further evidence on the costs of business cycle fluctuations, we try a third approach to filtering recently suggested by Hamilton (2018) to overcome the shortcomings of the two-sided HP filter. This involves performing a least squares regression of C_{t+h} on a constant and the p most recent values of C as of date t. The fitted values from this regression provide the trend component of the consumption process, while the residuals identify the transient component. Since we are interested in removing business cycle fluctuations from consumption, we follow Hamilton (2018) and set $h = 8$ quarters and $p = 4$. The estimates of the cost of business cycles obtained using the Hamilton filter are even higher than those obtained using the HP filter – 1.4% versus 0.9% at the one-year horizon, 3.9% versus 2.1% at the two-year horizon, 7.0% versus 4.8% at the three-year horizon, 7.9% versus 5.5% at the four-year horizon, and 6.0% versus 5.1% at the five-year horizon. Table 3 summarizes the results obtained with the three smoothing filters.

Separately, in Section III.3 below, we present further evidence supporting the high cost of business cycle fluctuations using an approach that does not involve a smoothing filter.

Table 3: Cost of Business Cycle Fluctuations with Alternative Smoothing Filters

		1 Yr	2 Yr	3 Yr	4 Yr	5 Yr
Panel A: Nondurables & Services Consumption						
I-SDF (Mkt)	AJ	0.35	0.95	1.95	1.68	1.16
	HP	0.56	1.48	3.39	3.90	3.57
	Hamilton	0.95	2.86	5.00	5.44	3.86
I-SDF (FF6)	AJ	0.47	0.96	1.33	1.35	0.99
	HP	0.46	1.03	2.07	3.03	2.90
	Hamilton	1.00	2.88	3.69	4.73	3.49
Panel B: Total Consumption						
I-SDF (Mkt)	AJ	0.52	1.21	2.57	2.16	1.50
	HP	0.90	2.09	4.85	5.55	5.12
	Hamilton	1.41	3.94	7.01	7.86	5.98
I-SDF (FF6)	AJ	0.67	1.31	1.82	1.74	1.28
	HP	0.77	1.60	3.05	4.35	4.14
	Hamilton	1.74	4.30	5.47	7.07	5.62

The table reports the (cumulative) costs of business cycle fluctuations in consumption over one-to five-year horizons, for alternative smoothing filters: the HP filter, the AJ filter, and the Hamilton filter. Panel A presents results when consumption denotes the real personal consumption expenditure of nondurables and services, while Panel B does the same for total personal consumption expenditure (that includes durables). In each panel, the costs are calculated using the I-SDF recovered from the market portfolio alone (Rows 1–3) and the I-SDF recovered from the six size and book-to-market-equity sorted portfolios of Fama and French (Row 4–6). The sample is annual covering the period 1929-2015.

III.3 Time-Variation in Cost of Fluctuations

We now proceed to use our methodology to estimate the cost of aggregate consumption fluctuations in different states (or, times). We focus on the cost of all one-period consumption fluctuations, given by Equation (30). We present the results obtained with the SEL estimator. The alternative SET estimator produces very similar results, reported in Appendix G.

We first estimate the time series of the cost in our baseline sample covering the period 1930-2015. Each year corresponds to a particular state and the SEL approach estimates the welfare benefits of eliminating all consumption uncertainty in the subsequent year. In our implementation, we use nondurables and services consumption as the measure of the aggregate consumption expenditures and the excess return on the market portfolio as the base asset.

Note that the SEL procedure requires the specification of the investors' conditioning set. But what constitutes a good instrument to capture the relevant conditioning set? Within the consumption-based asset pricing frameworks (e.g., long run risks and habit models), consumption and its history are sufficient statistics for the current state, albeit through different channels (via the conditional consumption mean in long run risks models, and via the habit level in habit models). For this reason, in our baseline results, we use an exponentially-weighted moving average of lagged consumption growth as the conditioning variable. Nevertheless, in the real world, more than just consumption might be informative about the overall state of the economy. Consequently, to capture overall economic conditions, we also complement the (exponentially-weighted) moving average of past consumption with an exponentially-weighted moving average of the first principal component extracted from a broad cross section of macro variables: the panel of 106 macroeconomic variables from Sydney Ludvigson's web site, based on the Global Insights Basic Economics Database and The Conference Board's Indicators Database. These variables cover six broad categories of macroeconomic data: output, labor market, housing sector, orders and inventories, money and credit, and price levels. We transform each variable to make it stationary and then extract the first principal component from their cross section.[17]

Figure 6 presents the time series of the cost.[18]. Several features are immediately evident from the figure. First, the cost is strongly time-varying – it varies from 0.15% to 8.0% a year, with an average of 0.75%. Second, the cost is strongly countercyclical, rising sharply during (most) recessionary episodes. The average of the cost over the subsample that corresponds to recession years, where a year is classified as a recession year if there is an NBER-designated recession in any of its quarters, is 1.17%. The estimated costs are particularly high during the period of the Great Depression 1930-1933, with a mean of 5.8% and a maximum as high as 8.0%. In contrast, the average cost over the subsample comprised of expansionary episodes alone is less than half of that during recessions at 0.53%. The correlation between the cost and a dummy variable that takes the value 1 in a given year if there is an NBER-designated recession in any of its quarters and 0 otherwise is 36.1%. Finally, the estimates of the cost are economically large, given that they represent the welfare benefits of eliminating all consumption uncertainty for one period alone.[19]

Next, to assess the sensitivity of the results, we present the time series of the cost for alternative choices of the sample period and conditioning set. First, note that our baseline results were obtained for the 1929-2015 sample period. This raises the potential concern that our results may be largely driven by the volatile prewar period, that includes the episodes

[17]We refer the reader to Ludvigson's website for a detailed description of these variables and the transformations applied to make them stationary.

[18]Note that the SEL delivers overall very small pricing errors in sample: the mean (median) pricing error is -6.2×10^{-10} (2.5×10^{-10}), while the minimum is -5.2×10^{-8}, and the maximum is 3.6×10^{-8}

[19]Note that the average of the conditional estimates is not equal to the unconditional one year cost estimate reported in Table 1. This is due to two reasons. First, in the conditional approach, we lose, mechanically, the first data point (the worst year of the Great Depression) and that observation raises the unconditional estimate. Second, even without the loss of a data point, there is a Jensen's inequality term that should be accounted for in comparing the unconditional cost estimate with the average of the conditional ones.

of the Great Depression and the aftermath of World War II (the only two macroeconomic disaster episodes in the US identified in Barro (2006) over this period). To mitigate this concern, Figure 7 presents the time series of the annualized cost (red line) using quarterly data over the postwar period 1947:$Q1$–2015:$Q4$.

Figure 6: Time-Varying Cost of One-Period Consumption Fluctuations, 1929-2015

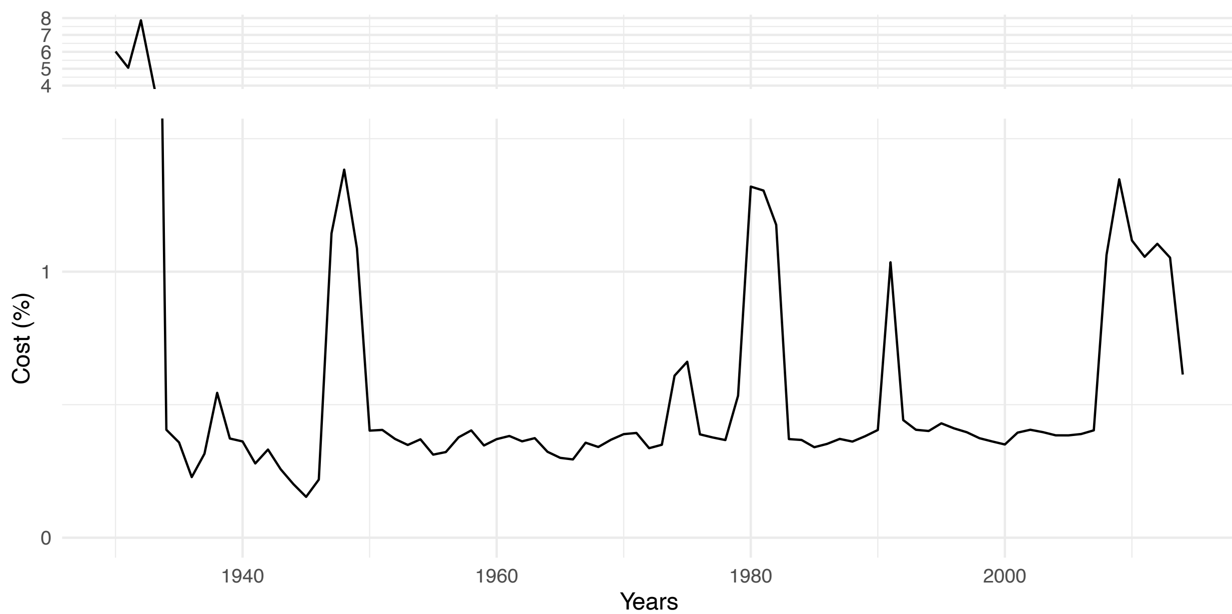

Notes: The figure plots the time series of the cost of one-period consumption uncertainty. The cost is estimated using the SEL approach, using nondurables and services consumption as the measure of the consumption expenditures, the excess return on the market portfolio as the sole test asset, and an exponentially-weighted moving average of lagged consumption growth as the conditioning variable. The sample is annual, covering the period 1930-2015.

The strong countercyclical variation in the cost is immediately evident from the figure. In fact, the countercyclicality is even more pronounced in the postwar period, compared to the longer 1929–2015 sample – the correlation with the recession dummy is 49.2% over the former period compared with 36.1% in the latter longer sample. Also, the magnitudes of the costs over the postwar subperiod are similar, regardless of whether the full 1929-2015 sample or the postwar period alone is used in the estimation of these costs. For instance, over the two years of the Great Recession, 2008–2009, the cost of removing one-year fluctuations is estimated to be 1.20% on average using the longer sample, and this is quite similar to the average cost for the Great Recession period obtained using the postwar sample data only (0.86%). Finally, as outlined in Figure A.3 of the Appendix, the strong countercyclical pattern of the estimates is independent of the calibration of the γ parameter.

As a second robustness check, we present results for an expanded conditioning set. Note that our baseline results were obtained using a weighted average of past consumption growth

Figure 7: Time-Varying Cost: Robustness to Sample Period and Conditioning Set

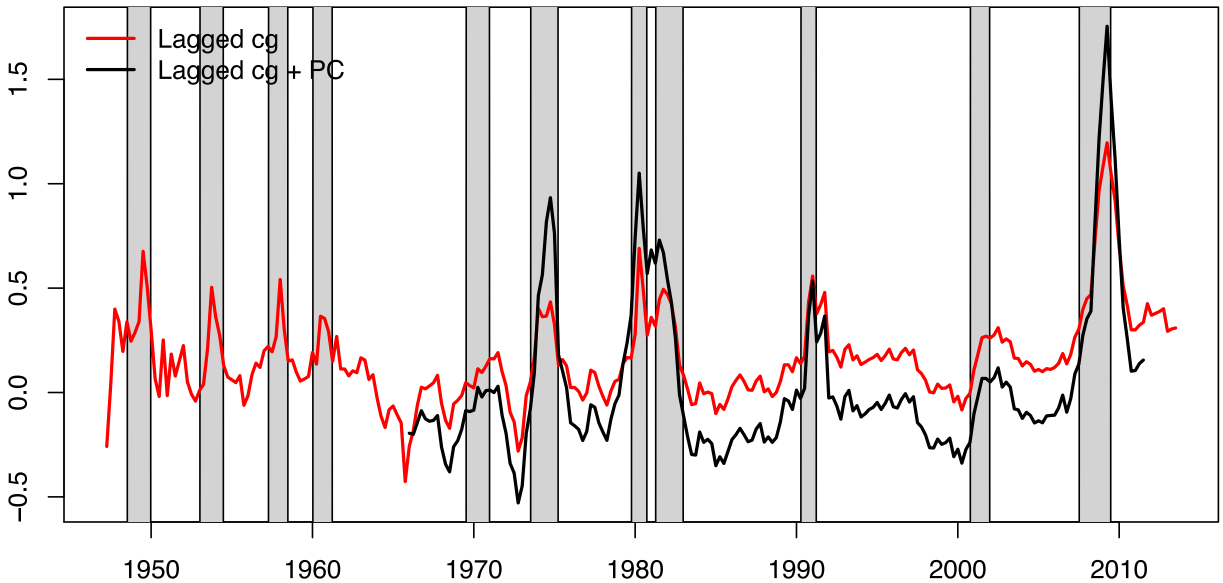

Notes: The figure plots the time series of the cost of one-period consumption uncertainty. The cost is estimated using the SEL approach, using nondurables and services consumption as the measure of the consumption expenditures and the excess return on the market portfolio as the sole test asset. The conditioning set consists of an exponentially-weighted moving average of lagged consumption growth (red line) and lagged consumption growth and a principal component extracted from a broad cross section of 106 macro variables (black line). The sample is quarterly, covering the period 1947:$Q$1–2015:$Q$4 (red line) or 1966:$Q$1–2015:$Q$4 (black line).

as the sole conditioning variable. This may potentially raise concerns about the robustness of the findings. Therefore, we estimate the time series of the cost when the conditioning set includes not only an exponentially-weighted average of past consumption growth, but also an exponentially-weighted average of a principal component extracted from a cross section of 106 macroeconomic variables that cover six broad categories of macroeconomic data: output, labor market, housing sector, orders and inventories, money and credit, and price levels. The time series of the cost is presented in Figure 7 (black line). Since data on these macro variables are only available from the mid-sixties, the cost estimates start from 1966:$Q$1. The figure shows that the recovered time series of the cost seems quite robust to the choice of the conditioning set.

Overall, our results suggest that the cost of consumption fluctuations is strongly countercyclical and this offers, at least a partial, explanation of the high costs of business cycle fluctuations that we estimate in Section III.1.

IV Sensitivity to Alternative Choices of γ

We assess the sensitivity of the magnitude of the cost of fluctuations to alternative choices of γ. So far, we have reported estimates of the cost of consumption fluctuations using a baseline value of $\gamma = 10$. As discussed in Section III, this baseline value corresponds to the minimum value of this parameter typically required by a broad class of consumption-based models to explain asset prices.

Figure 8, as well as Table A.VI of the Appendix, report the estimated cost of all consumption fluctuations (Panel A), as well as the share of the cost ascribable to business cycle fluctuations (Panel B), for alternative values of γ. In addition to $\gamma = 10$, we consider: $\gamma = 2$, the canonical calibration in habit type models; $\gamma = 3.5$, the value commonly used in models with rare disasters; $\gamma = 16$, the value assumed in models with complementarities in consumption; $\gamma = 26.2$, the point estimate obtained with the EL method; and $\gamma = 0$, that corresponds to the case of no observable consumption in the pricing kernel. Three salient points emerge from this analysis.

Figure 8: Cost of fluctuations for alternative γ values.

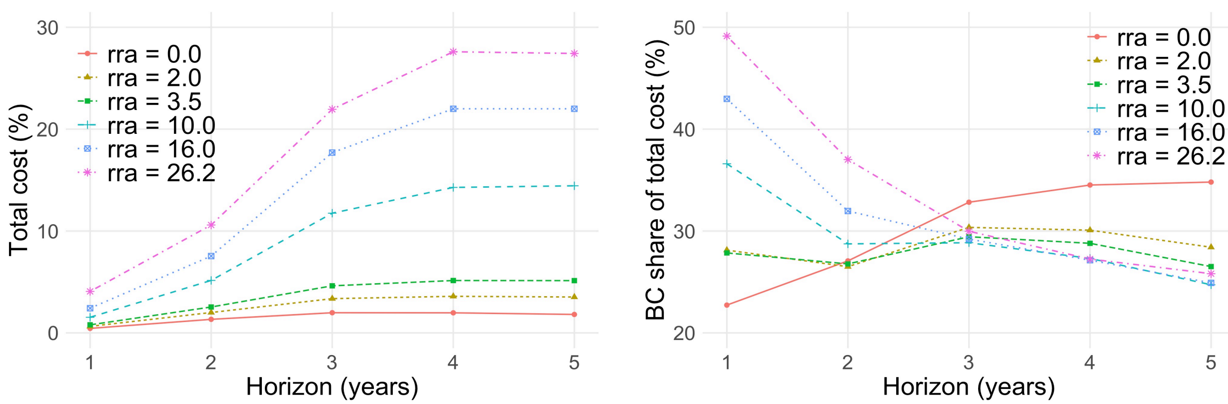

Panel A. Total cost of fluctuations

Panel B. Business Cycle share of total cost

Notes: The figure plots the (cumulative) costs of *all* aggregate consumption fluctuations (Panel A) and the share of the cost ascribable to business cycle fluctuations in consumption (Panel B), over one-to five-year horizons, for different values of γ. The consumption measure is the real personal consumption expenditure of nondurables and services. The costs are calculated using the I-SDF recovered from the market portfolio alone with the EL approach. The sample is annual covering the period 1929-2015.

First, as one would expect, a larger value for the curvature parameter results in overall larger estimates for the cost of consumption fluctuations. For instance, for $\gamma = 26.2$, the cost of all fluctuations is as high as 27% of one period consumption at the 5 year horizon, while it is reduced to only 1.8% at the same horizon for $\gamma = 0$. However, as shown by the (bootstrapped) p-values in Table A.VI, values of $\gamma < 10$ are extremely unlikely given the data: their probability being less than 1.5%. That is, the data strongly favour the inclusion of consumption growth in the pricing kernel, possibly with a larger curvature parameter than in our baseline calibration, and the p-value of $\gamma = 0$ is at most 1.3%. Hence, our baseline estimates, obtained with $\gamma = 10$, can be viewed as conservative.

Note also that the effect of γ on the estimated costs is non linear. This is due to the fact that, as Table A.VI in the Appendix shows, a change in γ changes both the predictability of future consumption using the lagged SDF, as well as the predictability of the SDF itself using lagged consumption. To see why this change in predictability is important, consider the (log) per period risk premium on an asset that gives a return equal to the log consumption growth over the next $1+S$ periods. Let $\bar{r}_c^{1+S}$ denote the (per period) risk premium of this asset. Let further denote the multi-period log consumption growth as $\Delta c_{t-1\to t+S} \equiv \ln C_{t+S} - \ln C_{t-1}$ and the log SDF between $t-1$ and $t+S$ as $m_{t-1\to t+S}$. Log-linearizing the fundamental asset pricing equation for this asset, it follows that

$$\begin{aligned}\bar{r}_c^{1+S} \;\; = & \;\; -\frac{1}{1+S} cov\left(m_{t-1\to t+S}; \Delta c_{t-1\to t+S}\right) = \underbrace{-cov\left(m_{t-1,t}, \Delta c_{t-1,t}\right)}_{\bar{r}_c^1} \\ & -\frac{1}{1+S}\sum_{\tau=1}^{S}\sum_{s=1}^{\tau}\left[cov\left(m_{t-1,t}; \Delta c_{t-1+s\to t+s}\right) + cov\left(\Delta c_{t-1,t}; \Delta m_{t-1+s\to t+s}\right)\right]\end{aligned}$$

where $\bar{r}_c^1$ is the risk premium for the one period investment. The above makes clear that the pricing of consumption growth risk can be horizon dependent if, and only if, either the SDF predicts future consumption or consumption predicts the future SDF.

In the above equation, an increase (decrease) in risk aversion, is expected to increase (decrease) $\bar{r}_c^1$ but also, as Table A.VII of the Appendix shows, the predictability of future consumption growth with the lagged SDF as well as the predictability of the future SDF with lagged consumption growth. But why can't the estimated ψ-component restore the "right" degree of predictability even when $\gamma = 0$? The results is intuitive: the dual solution for ψ in Equation (19) is a function of $(\Delta c_t)^\gamma$ and asset returns. Hence, when $\gamma = 0$, the I-SDF cannot use the information in the consumption process to learn about, and encode predictability in, the SDF. Instead, when $\gamma > 0$ the future SDF is very predictable with current consumption. For instance, the R^2 of the predictive regression of the SDF 5 years in the future with current consumption growth changes from about 3% to 10% moving from $\gamma = 0$ to $\gamma = 10$. This is due to the fact that when $\gamma = 0$ the I-SDF is purely a (non-linear) function of asset returns, and consequently it exhibits (like asset returns themselves) little predictability.

Second, regardless of the calibrated value of γ, the estimated cost of all consumption fluctuations increases sharply during the first 3-4 years, but flattens thereafter.

Third, as Panel B of Figure 8 highlights, the share of the costs of consumption fluctuations ascribable to the business cycle is both substantial (ranging from 23% to 49%), and relatively stable at around 30% for horizons of two or more years, regardless of the calibrated value of γ. This is in contrast to Alvarez and Jermann (2004) that estimates minuscule share of the cost of consumption fluctuations as being generated by business cycle variations.

Overall, our results suggest that the magnitude of the total cost of fluctuations is sensitive to the choice of the parameter γ, with larger values implying larger costs. However, for the values of this parameter that are more likely given the data (as large, or larger, than 10 in Table A.VI), the total costs are very large. Furthermore, the conclusion that the cost of business cycles is substantial, accounting for at least a quarter of the total cost of fluctuations, is relatively stable across alternative calibrated values of γ. Finally, as Figure A.3 in the Appendix shows, the finding that the cost of fluctuations is strongly countercyclical is robust to the choice of the curvature parameter γ.

V Implications for Macro Finance Models

As we are about to show, our estimates of the term structure of the cost of consumption fluctuations are informative not only for structural models commonly considered in the literature, but also about the pricing of "consumption strips."

First, Figure 9 compares our estimated costs of consumption fluctuations, based on $\gamma = 10$, with the ones implied by the long run risks model of Bansal and Yaron (2004) (BY) and the habit model of Campbell and Cochrane (1999) (CC). To generate the model-implied costs, we use exactly the same calibrations of the parameters as in the original papers. Panels A and B show that, regardless of whether we use nondurables and services consumption (Panel A) or total consumption (Panel B) as the empirical measure of consumption expenditures, our estimates are quite different from the model-implied ones. At the one year horizon, our estimates of the cost are very similar to the CC one, but they increase sharply afterwards and stabilize (at value of about 14-19%) after about four years. Instead, the CC implied costs grow almost linearly over the first ten years, and reaches a value of about 10% only at the 10 year horizon. The BY model-implied cost is much smaller than our I-SDF based one for any horizon between one and ten years (ranging from 0.6% to 3%). Panel C also highlights the very long run (up to 50 years) behaviour of the model implied costs. The CC-implied cost flattens out at about 20 years, to a level similar to what we estimate at the 10 year horizon. Instead, the BY-implied cost increases almost linearly with the horizon and, only at horizons exceeding 30 years does it reach values similar to our estimates at the 10 year horizon, but overshoots our estimates thereafter and the cost keeps growing steeply for more than a century (not reported). This is due to the very strong persistence in the calibrated conditional consumption mean process of Bansal and Yaron (2004). Our findings, instead, are more in line with the estimates in Bryzgalova, Huang, and Julliard (2024a), that documents strong persistence in the consumption process that dies off after the business cycle horizon, but generates nevertheless enough consumption risk to rationalize the equity premium puzzle with a low level of relative risk aversion.

Second, our method for estimating the cost of consumption fluctuations has also direct implications for the costs of "consumption strips." The literature (see, e.g., van Binsbergen and Koijen (2017), van Binsbergen, Brandt, and Koijen (2012), and Bansal, Miller, Song, and Yaron (2021)) has primarily focused on the term structure of dividend strips,[20] since consumption strips are not traded assets.

Our method, instead, can deliver the term structure of consumption strips using the exact same framework that we have employed to evaluate consumption fluctuations. To see this, notice that the yield-to-maturity of consumption in s periods can be expressed as

$$R^c_{s,t+s} = \frac{C_{t+s}}{P_{s,t}} = \frac{C_{t+s}/C_t}{P_{s,t}/C_t}.$$

We can then express the (hold-to-maturity) conditional expected return of a long horizon consumption strip in excess of the one-period strip as

$$\mathbb{E}_t\left[\frac{1}{s}\ln R^c_{s,t+s} - \ln R^c_{1,t+1}\right] = \mathbb{E}_t\left[\frac{1}{s}\ln\frac{C_{t+s}}{C_t} - \ln\frac{C_{t+1}}{C_t}\right] - \frac{1}{s}\ln\frac{P_{s,t}}{C_t} + \ln\frac{P_{1,t}}{C_t}.$$

The last two terms in the above expression can be estimated, at every t, using the conditional method in Section III.3. Hence, the (unconditional) expected term structure of yield spreads for the consumption strips is

$$\mathbb{E}\left[\frac{1}{s}\ln\frac{C_{t+s}}{C_t} - \ln\frac{C_{t+1}}{C_t} - \frac{1}{s}\ln\tilde{pc}_{s,t} + \ln\tilde{pc}_{1,t}\right], \tag{34}$$

and can be estimated by replacing the unconditional expectation with the sample average operator.

Figure 10 plots the estimated yield spreads of consumption strips using Equation (34). The (unconditional) term structure appears to be significantly downward sloping after one quarter, but essentially flat at longer horizons. The very short horizon behaviour of our estimates matches the downward sloping term structure of dividend strip yields found in van Binsbergen and Koijen (2017), that studies the dividend strip term structure from six

[20] Giglio, Kelly, and Kozak (2023) do not rely on dividend strips data and, instead, estimate the term structure of equity yields using an affine model for returns and dividends.

Figure 9: Comparison with macrofinance models

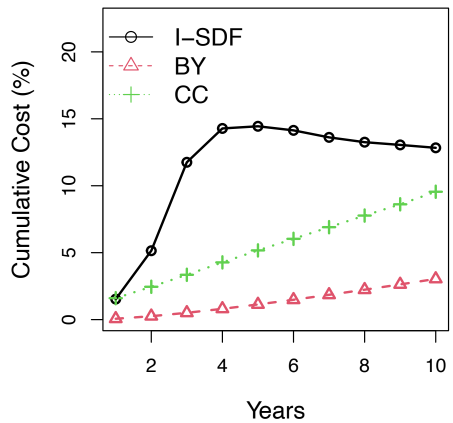

Panel A. Nondurable and service consumption

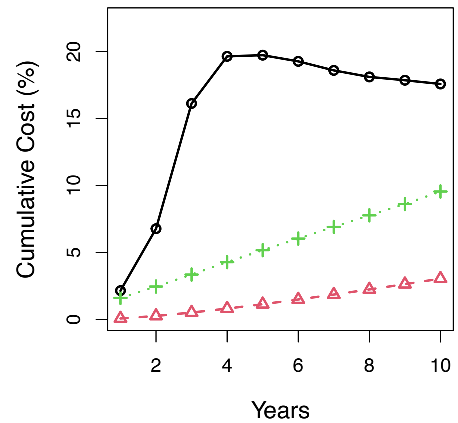

Panel B. Total consumption

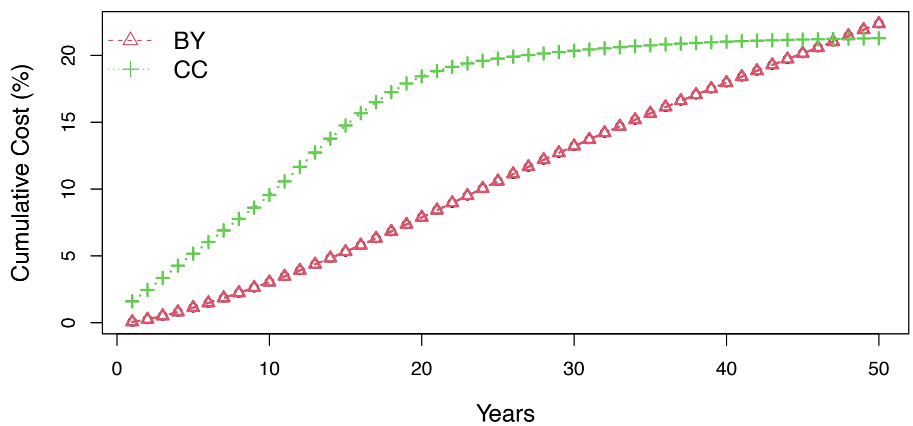

Panel C. Long run costs of fluctuations in Campbell and Cochrane (1999) and Bansal and Yaron (2004)

Notes: The top two panels plot the term structure of the (cumulative) cost of *all* aggregate consumption fluctuations, over 1-10 years, obtained using the I-SDF (black solid line with circles), implied by the long run risks model (red-dashed line with triangles), and implied by the external habit model (green-dotted line with crosses). Panel A presents results when consumption refers to the real personal consumption expenditure of nondurables and services, while in Panel B consumption denotes total personal consumption expenditure. The I-SDF is extracted using the excess return on the market portfolio as the sole test asset with the EL approach. The sample is annual covering the period 1929-2015. In Panel C we extend the horizon to 50 years for the long run risks model (red-dashed line with triangles) and implied by the external habit model (green-dotted line with crosses).

months to two years horizons. Instead, Bansal, Miller, Song, and Yaron (2021) studies the term structure at the one to five year horizons, and find an essentially flat unconditional term structure for the dividend yield: their estimates imply a weakly upward sloping (unconditional) term structure for the S&P500, and a marginally downward sloping one for the Eurostoxx 50 and Nikkei 225 indexes (but the precise slope measurement is significantly affected, as they show, by the large bid-ask spread in dividend strips data). Our estimates, like their structural long run risk model with regime switching, imply an essentially flat unconditional term structure at the one to five year horizons.

Figure 10: Yield Spread on Consumption Strips

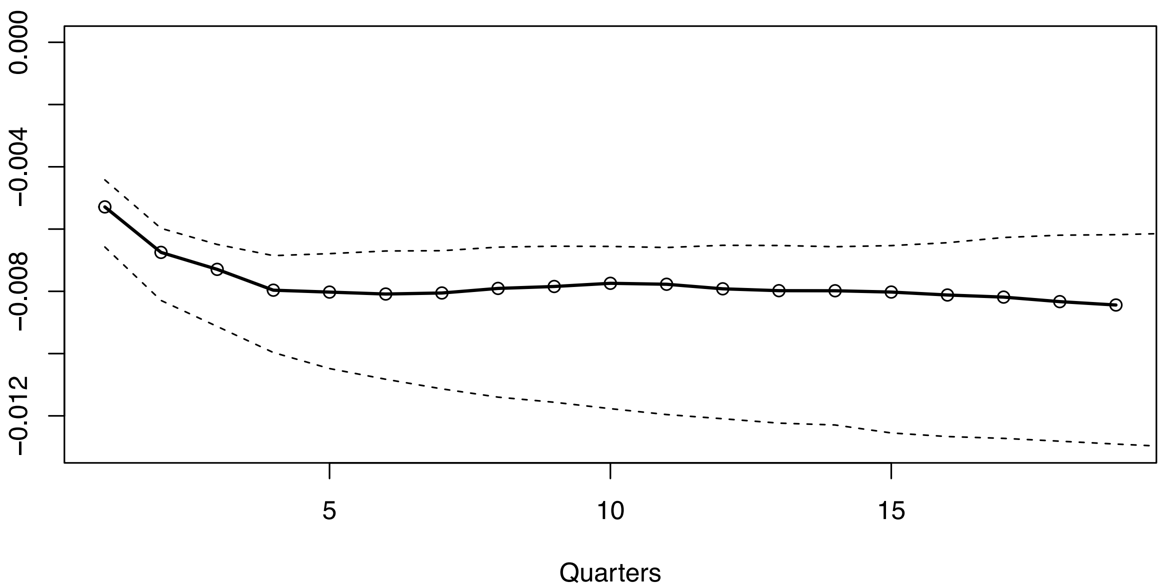

Notes: The figure plots the yield spreads on consumption strips up to 20 quarters (5 years) estimated as sample average of equation (34). The consumption measure is the real personal consumption expenditure of nondurables and services. The consumption yields are calculated using the conditional I-SDF recovered from the market portfolio alone with the EL approach. The sample is quarterly covering the period 1947:$Q1$–2015:$Q4$. Dashed lines denote bootstrap 95% confidence interval.

VI Conclusion

We propose a novel approach to measure the welfare costs of aggregate economic fluctuations. Our methodology does not require a full specification of the preferences of consumers or any assumptions about the dynamics of the data generating process. Instead, using data on consumption growth and returns on a chosen set of assets, we rely on an information-theoretic (or relative entropy minimization) approach to estimate the pricing kernel. We refer to the resulting kernel as the *information kernel*, or the I-SDF, because of the information-theoretic approach used in its recovery. Unlike the CRRA kernel, or Lucas' original specification that imposes the additional assumption of i.i.d. lognormality of consumption growth on the CRRA model, the I-SDF accurately prices a broad set of assets – unconditionally as well as conditionally, in-sample as well as out-of-sample – thereby successfully capturing the relevant sources of priced risk in the economy.

Using the I-SDF, we show that the welfare benefits from the elimination of *all* consumption uncertainty are large – typically, an order of magnitude bigger than those implied by Lucas' specification. Furthermore, our estimated cost of consumption fluctuation at the short and medium horizon are much larger than those implied by the original calibrations of the external habit (Campbell and Cochrane (1999)) and long-run risk (Bansal and Yaron (2004)) models, and the latter implies counterfactually high costs at the very long horizon.

We find that the costs of *business cycle* fluctuations in consumption constitute a substantial proportion – typically between a quarter to a third – of the costs of all consumption uncertainty. Moreover, using an extension of our information-theoretic methodology, we present evidence that the welfare benefits of eliminating aggregate consumption fluctuations are strongly time-varying and countercyclical.

The difference in the results from earlier literature can be attributed, at least in part, to two factors. First, the I-SDF correctly prices broad cross sections of assets, and thereby identifies the relevant sources of priced risk more accurately than existing models. Second, the I-SDF has a strong business cycle component, suggesting that business cycle risk is an important source of priced risk. Also, the non-requirement of a fully specified utility function characterizing consumers' preferences and assumptions about the dynamics of the data generating process makes the I-SDF, and therefore the resulting estimates of the costs of fluctuations, more robust to misspecification.

Our results indicate that the cost of business-cycle fluctuations may be much higher than previously thought. Nevertheless, our estimates do not incorporate the possibility that government policies effective in curbing fluctuations may alter the trend growth in consumption. Barlevy (2004), for instance, shows that in an endogenous growth framework, shutting down aggregate uncertainty increases annual consumption growth by .35-.40%. This would serve to further increase the cost of fluctuations.

Finally, the present paper focuses on estimating the welfare costs of aggregate consumption uncertainty. However, our methodology is considerably general and may also be applied to obtain the costs of uninsurable idiosyncratic risk, such as labor income risk. This is left for future research.

References

Almeida, C., K. Ardison, and R. Garcia (2019): “Nonparametric Assessment of Hedge Fund Performance,” *Journal of Econometrics (forthcoming).*

Almeida, C., and R. Garcia (2012): “Assessing misspecified asset pricing models with empirical likelihood estimators,” *Journal of Econometrics*, 170(2), 519–537.

——— (2016): “Economic Implications of Nonlinear Pricing Kernels,” *Management Science*, 63(10), 3147–3529.

Alvarez, F., and U. J. Jermann (2004): “Using Asset Prices to Measure the Cost of Business Cycles,” *Journal of Political Economy*, 112(6), pp. 1223–1256.

Atkeson, A., and C. Phelan (1994): “Reconsidering the costs of business cycles with incomplete markets,” *NBER Macroeconomics Annual*, pp. 187–207.

Backus, D., M. Chernov, and S. E. Zin (2013): “Sources of Entropy in Representative Agent Models,” *Journal of Finance*, forthcoming.

Bai, H., and L. Zhang (2020): “Searching for the Equity Premium,” NBER Working Paper 28001.

Bansal, R., S. Miller, D. Song, and A. Yaron (2021): “The term structure of equity risk premia,” *Journal of Financial Economics*, 142(3), 1209–1228.

Bansal, R., and A. Yaron (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 59(4), 1481–1509.

Barlevy, G. (2004): “The Cost of Business Cycles under Endogenous Growth,” *The American Economic Review*, 94(4), pp. 964–990.

——— (2005): “The cost of business cycles and the benefits of stabilization,” *Economic Perspectives, Federal Reserve Bank of Chicago*, Q(I), pp. 32–49.

Barro, R. J. (2005): “Rare Events and the Equity Premium,” NBER Working Papers 11310, National Bureau of Economic Research, Inc.

——— (2006): “Rare Disasters and Asset Markets in the Twentieth Century,” *Quarterly Journal of Economics*, 121(3), 823–866.

——— (2009): “Rare Disasters, Asset Prices, and Welfare Costs,” *American Economic Review*, 99(1), 243–264.

Basak, S., and H. Yan (2010): *The Review of Economic Studies*77(3), 914–936.

Beaudry, P., and C. Pages (2001): “The cost of business cycles and the stabilization value of unemployment insurance,” *European Economic Review*, 45, 1545–1572.

Blanchard, O. J. (2019): “Public Debt and Low Interest Rates,” AFA Lecture.

Borovicka, J., L. P. Hansen, and J. A. Scheinkman (2016): "Misspecified Recovery," *Journal of Finance*, 71(6), 2493–2544.

Brown, D. E., and R. L. Smith (1990): "A Correspondence Principle for Relative Entropy Minimization," *Naval Research Logistics*, 37(2).

Bryzgalova, S., J. Huang, and C. Julliard (2024a): "Consumption in Asset Returns," SSRN working paper.

——— (2024b): "Macro Strikes Back: Term Structure of Risk Premia and Market Segmentation," SSRN working paper.

Campbell, J. Y., and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205–51.

Chen, X., L. P. Hansen, and P. G. Hansen (2020): "Robust Identification of Investor Beliefs," *Proceedings of the National Academy of Science*, 117, 33130–33140.

Chernov, M., L. A. Lochstoer, and S. R. H. Lundeby (2021): "Conditional Dynamics and the Multihorizon Risk-Return Trade-Off," *The Review of Financial Studies*, 35(3), 1310–1347.

Constantinides, G. M., and D. Duffie (1996): "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104(2), 219–40.

Constantinides, G. M., and A. Ghosh (2011): "Asset Pricing Tests with Long Run Risks in Consumption Growth," *The Review of Asset Pricing Studies*, 1, 96–136.

——— (2017): "Asset Pricing with Countercyclical Household Consumption Risk," *Journal of Finance*, 72(1), 415–460.

Csiszár, I. (1975): "I-Divergence Geometry of Probability Distributions and Minimization Problems," *Annals of Probability*, 3, 146–158.

Davis, J., and G. Segal (2020): "Trendy Business Cycles and Asset Prices," Working Paper.

Dolmas, J. (1998): "Risk preferences and the welfare cost of business cycles," *Review of Economic Dynamics*, 1, 646–676.

Epstein, L. G., E. Farhi, and T. Strzalecki (2014): "How Much Would You Pay to Resolve Long-Run Risk?," *AER*, 104(9), 2680–2697.

Epstein, L. G., and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–968.

Fama, E. F., and K. R. French (1993): "Common Risk Factors in the Returns on Stocks and Bonds," *The Journal of Financial Economics*, 33, 3–56.

——— (2015): "A five-factor asset pricing model," *Journal of Financial Economics*, 116(1), 1–22.

Ghosh, A., C. Julliard, and A. Taylor (2017): “What is the Consumption-CAPM missing? An information-Theoretic Framework for the Analysis of Asset Pricing Models,” *Review of Financial Studies*, 30(2), 442?504.

——— (2022): “An Information-Based One-Factor Asset Pricing Model,” Working paper.

Ghosh, A., and T. Otsu (2020): “Properties of Subjective Beliefs Estimators,” Working paper, london school of economics and mcgill university.

Ghosh, A., and G. Roussellet (2019): “Identifying Beliefs from Asset Prices,” Working paper, mcgill university.

Giglio, S., B. T. Kelly, and S. Kozak (2023): “Equity Term Structures without Dividend Strips Data,” *Journal of Finance (forthcoming)*.

Hamilton, J. D. (2018): “Why You Should Never Use the Hodrick-Prescott Filter,” *The Review of Economics and Statistics*, 100.

Hansen, L., T. Sargent, and T. Tallarini (1999): “Robust Permanent Income and Pricing,” *The Review of Economic Studies*, 66(4), 873–907.

Hansen, L. P. (2014): “Nobel Lecture: Uncertainty Outside and Inside Economic Models,” *Journal of Political Economy*, 122(5), 945 – 987.

Hansen, L. P., and R. Jagannathan (1991): “Implications of Security Market Data for Models of Dynamic Economies,” *Journal of Political Economy*, 99(2), 225–62.

Hansen, L. P., and R. Jagannathan (1997): “Asssessing Specification Errors in Stochastic Discount Factor Models,” *The Journal of Finance*, 52, 557–590.

Hansen, L. P., and T. J. Sargent (2010): “Fragile Beliefs and the Price of Uncertainty,” *Quantitative Economics*, 1(1), 129–162.

Hansen, L. P., and K. J. Singleton (1983): “Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns,” *Journal of Political Economy*, 91, 249–68.

Hou, K., C. Xue, and L. Zhang (2014): “Digesting Anomalies: An Investment Approach,” *The Review of Financial Studies*, 28(3), 650–705.

Imrohoroglu, A. (1989): “Cost of business cycles with indivisibilities and liquidity constraints,” *Journal of Political Economy*, 97(6), 1364–1383.

Julliard, C., and A. Ghosh (2012): “Can Rare Events Explain the Equity Premium Puzzle?,” *Review of Financial Studies*, 25(10), 3037–3076.

Kitamura, Y. (2006): “Empirical Likelihood Methods in Econometrics: Theory and Practice,” Cowles Foundation Discussion Papers 1569, Cowles Foundation, Yale University.

Kitamura, Y., and M. Stutzer (1997): “An Information-Theoretic Alternative To Generalized Method Of Moments Estimation,” *Econometrica*, 65(4), 861–874.

Kitamura, Y., G. Tripathi, and H. Ahn (2004): "Empirical Likelihood-Based Inference in Conditional Moment Restriction Models," *Econometrica*, 72(6), 1667–1714.

Komunjer, I., and G. Ragusa (2016): "Existence and Characterization of Conditional Density Projections," *Econometric Theory*, 32, 947–987.

Krebs, T. (2007): "Job Displacement Risk and the Cost of Business Cycles," *AER*, 97, 664–686.

Krusell, P., and A. A. Smith (2009): "Revisiting the welfare effects of eliminating business cycles," *Review of Economic Dynamics*, 12(3), 393–404.

Lettau, M., and S. Ludvigson (2001): "Consumption, Aggregate Wealth, and Expected Stock Returns," *Journal of Finance*, 56(3), 815–49.

Liew, J., and M. Vassalou (2000): "Can book-to-market, size and momentum be risk factors that predict economic growth?," *The Journal of Financial Economics*, 57, 221–245.

Lucas, R. E. (1987): *Models of Business Cycles*. Blackwell, Oxford.

Lustig, H. N., and S. G. V. Nieuwerburgh (2005): "Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective," *The Journal of Finance*, 60(3), pp. 1167–1219.

Lustig, H. N., S. G. V. Nieuwerburgh, and A. Verdelhan (2013): "The Wealth-Consumption Ratio," *The Review of Asset Pricing Studies*, 3(1), 38–94.

Martin, I. (2008): "Disasters and the Welfare Cost of Uncertainty," *aerpp*, 98(2).

Mehra, R., and E. C. Prescott (1985): "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15(2), 145–61.

Menzly, L., T. Santos, and P. Veronesi (2004): "Understanding Predictability," *Journal of Political Economy*, 112(1), 1–47.

Obstfeld, M. (1994): "Evaluating risky consumption paths: The role of intertemporal substitutability," *European Economic Review*, 38(7), 1471–1486.

Otrok, C. (2001): "On measuring the welfare cost of business cycles," *Journal of Monetary Economics*, 47(1), 61–92.

Owen, A. B. (2001): *Empirical Likelihood*. Chapman and Hall.

Parker, J. A. (2001): "The consumption risk of the stock market," *Brookings Papers on Economic Activity*, 2001(2), 279–333.

Parker, J. A., and C. Julliard (2003): "Consumption Risk and Cross-Sectional Returns," NBER Working Papers 9538, National Bureau of Economic Research, Inc.

——— (2005): "Consumption Risk and the Cross-Section of Expected Returns," *Journal of Political Economy*, 113(1).

Pemberton, J. (1996): "Growth trends, cyclical fluctuations, and welfare with non-expected utility preferences," *Economics Letters*, 50(3), 387–392.

Piazzesi, M., and M. Schneider (2007): "Equilibrium Yield Curves.," in *N.B.E.R. Macroeconomics Annual 2006*, vol. 21. Cambridge: MIT Press.

Piazzesi, M., M. Schneider, and S. Tuzel (2007): "Housing, consumption and asset pricing," *Journal of Financial Economics*, 83(3), 531–569.

Qin, J., and J. Lawless (1994): "Empirical Likelihood and General Estimating Equations," *Annals of Statistics*, 22, 300–325.

Ravn, M. O., and H. Uhlig (2002): "On adjusting the Hodrick-Prescott filter for the frequency of observations," *The Review of Economics and Statistics*, 84(2), 371–375.

Sandulescu, M., F. Trojani, and A. Vedolin (2018): "Model-Free International Stochastic Discount Factors," Working paper, boston university.

Schennach, S. M. (2005): "Bayesian Exponentially Tilted Empirical Likelihood," *Biometrika*, 92(1), 31–46.

Storesletten, K., C. I. Telmer, and A. Yaron (2001): "The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk," *European Economic Review*, 45(7), 1311–1339.

Stutzer, M. (1995): "A Bayesian approach to diagnosis of asset pricing models," *Journal of Econometrics*, 68(2), 367 – 397.

——— (1996): "A simple Nonparametric Approach to Derivative Security Valuation," *Journal of Finance*, LI(5), 1633–1652.

Tallarini, T. (2000): "Risk sensitive real business cycles," *Journal of Monetary Economics*, 45(3), 507–532.

van Binsbergen, J. H., M. Brandt, and R. Koijen (2012): "On the timing and pricing of dividends," *American Economic Review*, 102(4), 1596–1618.

van Binsbergen, J. H., and R. S. Koijen (2017): "The term structure of returns: Facts and theory," *Journal of Financial Economics*, 124(1), 1–21.

Vuong, Q. H. (1989): "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica*, pp. 307–333.

Wachter, J. (2013): "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?," *Journal of Finance*, 68(3), 987–1035.

White, H. (1982): "Maximum Likelihood of Misspecified Models," *Econometrica*, 50, 1–25.

Yogo, M. (2006): "A Consumption-Based Explanation of Expected Stock Returns," *Journal of Finance*, 61(2), 539–580.

Appendix

A Solution for the ET Estimator

As with the EL estimator, the ET estimator is also numerically simple to implement. Specifically, the ψ-component is estimated (up to a positive constant scale factor) as:

$$\hat{\psi}_t = \frac{e^{\hat{\theta}(\gamma)'\mathbf{r}_t^e(\Delta c_t)^{-\gamma}}}{\frac{1}{T}\sum_{t=1}^{T} e^{\hat{\theta}(\gamma)'\mathbf{r}_t^e(\Delta c_t)^{-\gamma}}} \quad \forall t, \tag{A.1}$$

where $\hat{\theta}(\gamma) \in \mathbb{R}^N$ is the vector of Lagrange multipliers that solves the unconstrained dual problem:

$$\hat{\theta}(\gamma) = \arg\min_{\theta} \left[\log \left(\frac{1}{T} \sum_{t=1}^{T} e^{\theta' \mathbf{r}_t^e (\Delta c_t)^{-\gamma}} \right) \right]. \tag{A.2}$$

B Out-of-Sample Performance of the I-SDF

B.1 Unconditional I-SDF

This section compares the out-of-sample performance of the I-SDF vis a vis other popular factor models. Following Ghosh, Julliard, and Taylor (2022), we construct the out-of-sample I-SDF in a rolling fashion. In particular, for a given cross section of asset returns, we divide the time series of returns and consumption growth into rolling subsamples of length $\overline{T}$ and final date T_i, $i = 1, 2, 3, ...$, and constant $s := T_{i+1} - T_i$. In subsample i, we estimate the vector of Lagrange multipliers $\widehat{\theta}_{T_i}$ by solving the minimization in Equations (16)–(18). Using the estimates of the Lagrange multipliers, $\widehat{\theta}_{T_i}$, the out-of-sample I-SDF, $\widehat{M}_{T_i}$ is obtained for the subsequent s periods (i.e. for t such that $T_i + 1 \leq t \leq T_{i+1}$) using equation (19). This process is repeated for each subsample to obtain the time series of the estimated kernel over the out-of-sample evaluation period. We then use this out-of-sample I-SDF as the single factor to price different cross-sections of test assets.

To place the widely used multifactor models – (Fama and French, 1993, 2015) 3- and 5-factor models (FF3 and FF5, respectively) or the Hou, Xue, and Zhang (2014) q-factor model (HXZ) – on an equal footing with the one-factor I-SDF, we present the empirical performance of the multifactor models when a multi-factor model-implied SDF is constructed as a linear function of the risk factors, with the coefficients estimated in a rolling out-of-sample fashion using only past returns data on the same cross section of portfolios used to recover the I-SDF.

For instance, we define the FF3 model-implied SDF as:

$$M_t^{FF3} = b_0 + \sum_{j=1}^{3} b_j f_{j,t}, \tag{A.3}$$

where $\{f_{j,t}\}_{j=1}^3 = \{R_{M,t} - R_{F,t}, R_{SMB,t}, R_{HML,t}\}$ and the coefficients b_j, $j = 0, 1, 2, 3$, are estimated in a rolling out-of-sample fashion using only past returns data on the cross-section of portfolios, so as to satisfy the Euler equation restrictions for these portfolios:

$$0 = E\left[(R_{i,t} - R_{f,t})(b_0 + \sum_{j=1}^{3} b_j f_{j,t})\right]. \tag{A.4}$$

The resulting M_t^{FF3} is then used as the single risk factor in standard Fama-Macbeth cross sectional regressions for different sets of test assets to assess its empirical performance. Similarly, the FF5 model-implied SDF is defined as $M_t^{FF5} = b_0 + \sum_{j=1}^{5} b_j f_{j,t}$.

In practice, we set the size of the rolling window $\overline{T} = 30$ years, i.e. 360 months, and $s = 12$. Since the FF5 and HXZ factors are only available from the mid sixties and given the 30-years rolling window, our out-of-sample evaluation period covers 1993:07–2017:06. The results, presented in Table A.I, show the superior performance of the I-SDF relative to the CAPM, FF3, FF5, and HXZ models. In particular, for all three sets of test assets the I-SDF yields: smaller intercept, higher OLS/GLS adjusted R^2, higher degree of statistical significance (albeit this is low for all models due to the short OOS evaluation period), and better T^2 and q statistics.

B.2 Conditional I-SDF

Next we turn to the assessment of the out of sample performance of the conditional I-SDF. By construction, the conditional estimate cannot be use directly for an OOS analysis in the time series dimension, hence we turn to a cross-sectional OSS evaluation. That is, we estimated the I-SDF in the full time series sample using a set of base assets, and then evaluate its ability to explain a different cross-section of test asset during the same sample period.

As base asset to construct the conditional I-SDF we simply use the excess-return on the market, and as test assets we employ the same cross-sections as in Table A.I: FF25, 55 Decile portfolios, and a composite cross-section containing the FF25, 30 Industry portfolios, and 10 Momentum portfolios. We evaluate this OOS performance based on the Euler equation errors produced by the conditional I-SDF. As a benchmark, we also report the Euler equation errors obtained with the CRRA pricing kernel and the kernel density estimate of the conditional

Table A.I: Out of Sample Performance of Model-Implied SDFs, 1993:07–2017:06

Row	Assets	*const.*(%)	λ_{sdf}	$\overline{R}^2_{OLS}$ (%)	$\overline{R}^2_{GLS}$ (%)	T^2	q
	Panel A: I-SDF (0% < 0)						
(1)	FF 25	.66 (.03)	−.33 (.04)	42.6	43.8	43.1 (.030)	.19
(2)	55 *Decile Portfolios* (*ME*, *BE/ME*, *Mom*, *STR*, *LTR*, 5 *Ind.*)	.71 (.00)	−.10 (.08)	43.9	25.7	63.1 [.06]	.22
(3)	25 FF + 30 Ind + 10 Mom	.72 (.00)	−.10 (.08)	24.3	25.5	150.1 (.00)	.53
	Panel B: FF3-SDF (.35% < 0)						
(1)	FF 25	.81 (.03)	.000 (.49)	−4.3	39.5	56.7 (.00)	.20
(2)	55 *Decile Portfolios* (*ME*, *BE/ME*, *Mom*, *STR*, *LTR*, 5 *Ind.*)	.96 (.00)	.02 (.16)	16.7	22.2	66.8 (.00)	.23
(3)	25 FF + 30 Ind + 10 Mom	.88 (.00)	.01 (.26)	4.2	23.2	156.9 (.00)	.55
	Panel C: FF5-SDF (11.5% < 0)						
(1)	FF 25	.81 (.00)	.02 (.49)	−4.34	38.7	58.1 (.00)	.20
(2)	55 *Decile Portfolios* (*ME*, *BE/ME*, *Mom*, *STR*, *LTR*, 5 *Ind.*)	.81 (.00)	1.11 (.14)	22.6	22.7	64.7 (.13)	.23
(3)	25 FF + 30 Ind + 10 Mom	.80 (.00)	.67 (.23)	7.01	26.9	147.8 (.00)	.52
	Panel D: HXZ-SDF (5.56% < 0)						
(1)	FF 25	.82 (.00)	−.04 (.25)	10.7	36.9	59.9 (.00)	.21
(2)	55 *Decile Portfolios* (*ME*, *BE/ME*, *Mom*, *STR*, *LTR*, 5 *Ind.*)	.74 (.00)	−.03 (.32)	3.73	24.0	65.7 (.11)	.23
(3)	25 FF + 30 Ind + 10 Mom	.75 (.00)	−.03 (.34)	2.21	23.9	156.7 (.00)	.55

Cross-sectional regressions of average excess returns listed in column 2 on the estimated factor loadings for different model-implied SDFs. For each model, the model-implied SDF is recovered in a rolling out-of-sample fashion using only past returns data on a cross section of 15 equity portfolios: 5 industry portfolios and the smallest and largest deciles of size, BM, momentum, short term reversal, and long term reversal sorted portfolios. Panels A–D report results for the I-SDF, FF3, FF5, and HXZ models, respectively. For each set of test assets, the table presents the intercept and slopes, along with p-values (Shanken corrected) in parentheses, the OLS and GLS adjusted R^2, Shanken's (1985) cross-sectional T^2 statistic along with its asymptotic p-value in parentheses; and the q statistic that measures how far the factor-mimicking portfolios are from the mean–variance frontier. The last column reports the average absolute Euler Equation error for each cross section of test assets, along with the cross-sectional root mean squared error in parentheses.

distribution. Note that in this cross-sectional OOS exercise no parameters are estimated at the OOS evaluation stage to maximize fit.

Results are summarized in Table A.II. Two observations are in order. First, the mean absolute (conditional) pricing errors of the conditional I-SDF are quite small in this out of sample evaluation, ranging from 1.9% to 2.7%. Second, the *psi* component of the conditional I-SDF greatly improves the pricing ability relative to the CRRA pricing kernel: the latter yields pricing errors that are larger than the ones achievable with the I-SDF by a factor of

Table A.II: Conditional Cross-sectional out of Sample Performance

Row	Assets	Mean	Max	Min
	Panel A: conditional I-SDF Euler equation errors			
(1)	FF 25	.027	.060	.003
(2)	55 *Decile Portfolios* (*ME*, *BE/ME*, *Mom*, *STR*, *LTR*, 5 *Ind.*)	.019	.074	.003
(3)	25 FF + 30 Ind + 10 Mom	.022	.074	.003
	Panel B: CRRA SDF conditional Euler equation errors			
(1)	FF 25	.084	.128	.055
(2)	55 *Decile Portfolios* (*ME*, *BE/ME*, *Mom*, *STR*, *LTR*, 5 *Ind.*)	.070	.117	.013
(3)	25 FF + 30 Ind + 10 Mom	.074	.128	.013

Euler equation errors of conditional SDFs in the cross-section of asset returns reported in the second column. The conditional SDF is estimated using the market excess return as base asset. Panel A reports results of the conditional I-SDF while Panel B reports results for the conditional pricing kernel with consumption growth as the only factor. The third column, report the mean average absolute pricing error across time and assets, while the fourth and fifth column reports the max and min values across assets.

about 3.

C Pricing Multi-Period returns Using I-SDF

Multi-period I-SDFs can be obtained in two ways. First, the 1-year I-SDF can be compounded to obtain the I-SDF at lower (e.g., 2-year, 3-year, and so on) frequencies. Second, the lower frequency I-SDFs can be directly recovered using the EL approach using corresponding lower frequency returns. Figure A.1 plots the time series of the 5-year I-SDF recovered using the two approaches. Clearly, the estimates are highly correlated at all frequencies – the correlation coefficients are 93.0%, 95.2%, 99.7%, and 99.7%, at the 2-, 3-, 4-, and 5-year frequencies, respectively (note that the correlation is 100%, by construction, at the annual frequency). Furthermore, the Euler equation errors are quite small in both cases. The mean absolute Euler equation errors for the excess market return obtained using the former approach are 4.0×10^{-7}, 0.10, 0.15, 0.07, and 0.03 at the 1-year, 2-year, 3-year, 4-year, and 5-year frequencies, respectively. The corresponding Euler equation errors obtained using the latter approach should be theoretically zero, and indeed are numerically very small (4.0×10^{-7}, 0.8×10^{-5}, 1.2×10^{-9}, 5.0×10^{-6}, and 9.8×10^{-7}, respectively).

Figure A.1: Multi-Period Stochastic Discount Factor, 1929-2015

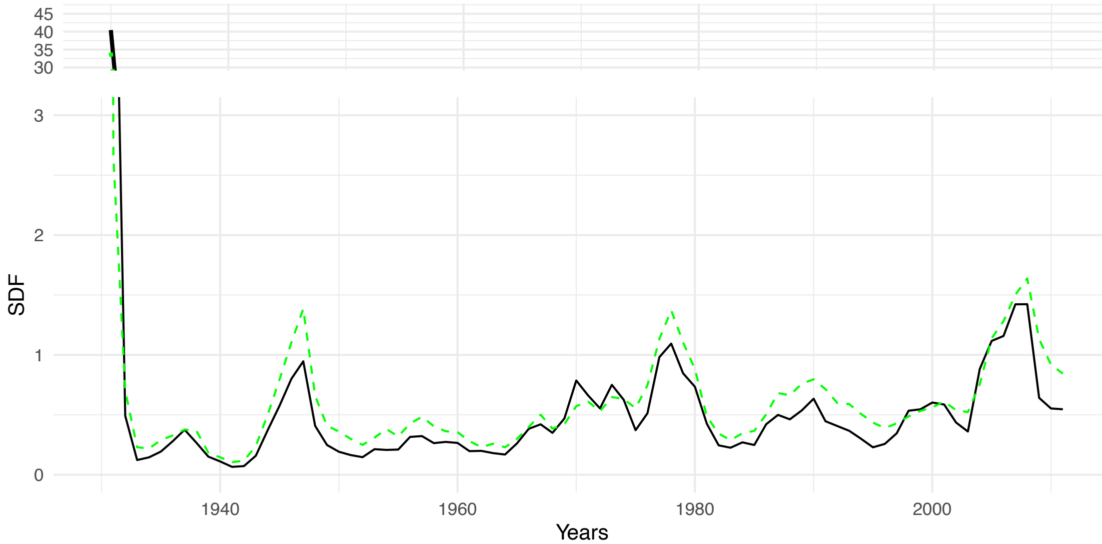

Notes: The figure plots the time series of the 5-year I-SDF. The 5-year I-SDF is obtained by compounding the annual I-SDF (solid black line) as well as by recovering the 5-year I-SDF directly from overlapping 5-year returns using the EL approach (green-dotted line). The I-SDFs are recovered using nondurables and services consumption as the measure of the consumption expenditures and the excess return on the market portfolio as the sole test asset. The sample is annual, covering the period 1930-2015.

D Performance of the Estimator: Simulation Evidence

In this section, we provide simulation evidence on the performance of the EL and ET estimators in measuring the cost of aggregate fluctuations, in both correctly specified and misspecified settings. In our first example (hereafter referred to as Economy I), we consider a hypothetical exchange economy in which the representative agent has power utility preferences with a constant CRRA and consumption growth is *i.i.d.* log-normal. Note that, in this correctly specified setup, $\psi_t \equiv 1$. In our second example (hereafter referred to as Economy II), we consider a standard long run risks economy where the representative investor has Kreps-Porteus recursive preferences and the aggregate consumption growth rate has a persistent predictable component and fluctuating volatility. With Economy II, we consider two scenarios – the first corresponds to when the econometrician correctly uses the SDF implied by recursive preferences and the second to when she erroneously uses the power utility preferences when recovering the ψ-component with the EL/ET approach. Note that, in the former scenario, the true $\psi_t \equiv 1$ whereas in the latter scenario the ψ-component captures the return on the investor's total wealth portfolio. We assess whether the estimators successfully recover the cost of fluctuations in these settings for empirically realistic sample sizes.

Consider first Economy I. The aggregate consumption growth evolves according to:

$\log\left(\Delta C_t\right) \overset{\mathbb{P}}{\sim} \mathcal{N}\left(\mu_c, \sigma_c^2\right)$. The following Euler equation holds in equilibrium:

$$0 = \mathbb{E}^{\mathbb{P}}\left[\left(\Delta C_{t+1}\right)^{-\gamma}\left(R_{m,t+1} - R_{f,t+1}\right)\right], \tag{A.5}$$

where $R_{m,t}$ and $R_{f,t}$ denote the market return and the risk free rate, respectively, at time t. Note that, Equation (A.5) may be rewritten as

$$0 = \mathbb{E}^{\mathbb{P}}\left[\left(\Delta C_{t+1}\right)^{-\gamma}\psi_{t+1}\left(R_{m,t+1} - R_{f,t}\right)\right], \tag{A.6}$$

where $\psi_t \equiv 1$.

This example economy fits into the framework described in Section II of the paper. Therefore, given time series data on consumption growth, the market return, and risk free rate, the EL/ET approaches can be used to estimate (up to a strictly positive constant scale factor) the ψ-component of the kernel:

$$\left\{\widehat{\psi}_t^{EL}\right\}_{t=1}^{T} = \underset{\{\psi_t\}_{t=1}^{T}}{\arg\max} \sum_{t=1}^{T}\log(\psi_t) \quad \text{s.t.} \quad \frac{1}{T}\sum_{t=1}^{T}\left(\Delta c_{t+1}\right)^{-\gamma}\psi_{t+1}(R_{m,t+1} - R_{f,t+1}) = 0. \tag{A.7}$$

$$\left\{\widehat{\psi}_t^{ET}\right\}_{t=1}^{T} = \underset{\{\psi_t\}_{t=1}^{T}}{\arg\min} \sum_{t=1}^{T}\psi_t\log(\psi_t) \quad \text{s.t.} \quad \frac{1}{T}\sum_{t=1}^{T}\left(\Delta c_{t+1}\right)^{-\gamma}\psi_{t+1}(R_{m,t+1} - R_{f,t+1}) = 0. \tag{A.8}$$

Using the recovered ψ, the term structure of the cost of fluctuations may be computed as described in Section II.B of the paper.

We show, via simulations, that the EL and ET approaches successfully identify the cost of fluctuations. Note that, in this economy, the equilibrium price-dividend ratio is $\frac{P_t}{D_t} = \nu$, a constant, where

$$\nu = \frac{\exp\left[\log(\delta) + (1-\gamma)\mu_c + \frac{(1-\gamma)^2\sigma_c^2}{2}\right]}{1 - \exp\left[\log(\delta) + (1-\gamma)\mu_c + \frac{(1-\gamma)^2\sigma_c^2}{2}\right]}, \tag{A.9}$$

and the equilibrium risk free rate is also constant at:

$$R_f = \frac{1}{\exp\left(\log(\delta) - \gamma\mu_c + \frac{\gamma^2\sigma_c^2}{2}\right)}. \tag{A.10}$$

To perform our simulation exercise, we calibrate μ_c and σ_c to the sample mean (2.8%) and volatility, (3.4%) respectively, of (log) consumption growth in our data (real per capita total consumption over 1929-2015). The preference parameters are calibrated at $\delta = 0.99$ and $\gamma = 10$. We simulate a time series of consumption growth. Using the simulated consumption growth, we obtain the market return as $R_{m,t+1} = \frac{\frac{P_{t+1}}{C_{t+1}}+1}{\frac{P_t}{C_t}} \cdot \frac{C_{t+1}}{C_t} = \frac{\nu+1}{\nu} \cdot \frac{C_{t+1}}{C_t}$, where ν is defined in Equation (A.9). The risk free rate is simply a constant, given by Equation (A.10).

Using the above time series, we recover $\{\psi_t\}_{t=1}^{T}$ using the EL and ET approaches in Equations (A.7) and (A.8), respectively. Armed with the ψ-component, we obtain the term structure of the cost of all consumption fluctuations. We repeat the above exercise for $10,000$ simulated samples. We report the averages and 90% confidence intervals of the costs of fluctuations across these simulations and compare them to their true population values. We present results for a sample size of $T = 87$, corresponding to the length of the historical sample at the annual frequency that we use in our empirical analysis.

The results are reported in Table A.III. Panel A reports the true values of the cost of one- to five-year fluctuations. Panel B presents the results across the simulated samples. Each cell in Panel B has four entries – the left (right) two are obtained with the EL (ET) approach, with the top number denoting the average of the cost across the simulated samples and the bottom numbers in square brackets its 90% confidence interval. The results show that the EL method is successful at accurately estimating the cost of fluctuations. Specifically, the EL-implied mean costs of fluctuations at the one- to five-year horizons across the $10,000$ simulations are essentially identical to the corresponding true costs. The 90% confidence intervals are fairly tight. The results obtained with the ET approach are virtually identical to those obtained with EL.

Table A.III: Simulation Results for Economy I, Cost of All Fluctuations (%)

	1 Yr	2 Yr	3 Yr	4 Yr	5 Yr
	Panel A: True Values				
	1.15	1.68	2.16	2.62	3.03
	Panel B: Simulated Values				
T=87	1.15 / 1.15 [1.15,1.16] [1.15,1.16]	1.66 / 1.66 [1.42,1.66] [1.43,1.89]	2.11 / 2.12 [1.66,2.57] [1.68,2.58]	2.51 / 2.53 [1.84,3.23] [1.87,3.24]	2.88 / 2.89 [1.98,3.86] [2.02,3.85]

The table reports the (cumulative) costs of *all* aggregate consumption fluctuations, over one-to five-year horizons in a hypothetical economy. The samples are simulated from a hypothetical endowment economy in which a representative agent has power utility preferences and the aggregate consumption growth is *i.i.d.* log-normal. Panel A presents the true values of these costs of fluctuations. Panel B presents the average of the costs, along with the 90% confidence intervals (in square brackets below), computed from $10,000$ simulated samples of size corresponding to the length of the historical data (T=87). To obtain the costs in Panel B, the ψ-component of the SDF is recovered using the EL (left entries in each cell) and ET approach (right entries in each cell).

Consider next Economy II. This is the standard Bansal and Yaron (2004) long run risks economy, i.e. the representative investor has recursive preferences and aggregate consumption and dividend growth rates have a small persistent predictable component and fluctuating volatility that captures time-varying economic uncertainty. Therefore, the SDF takes the form $M_t = \left(\frac{C_t}{C_{t-1}}\right)^{-\frac{\eta}{\rho}} \underbrace{\delta^\eta R_{c,t}^{\eta-1}}_{\psi_t}$, where γ denotes the CRRA, ρ the elasticity of intertemporal substitution, $\eta = \frac{1-\gamma}{1-\frac{1}{\rho}}$, and $R_{c,t} = \frac{P_{c,t}+C_t}{P_{c,t-1}}$ denotes the unobservable return on total wealth.

For this economy, we first consider the scenario when the econometrician uses the correct SDF when recovering the ψ-component of the pricing kernel using the EL and ET approaches. For the EL approach, for instance, this involves:

$$\left\{\widehat{\psi}_t^{EL}\right\}_{t=1}^T = \underset{\{\psi_t\}_{t=1}^T}{\arg\max} \sum_{t=1}^T \log(\psi_t) \quad \text{s.t.} \quad \frac{1}{T}\sum_{t=1}^T \left(\frac{c_{t+1}}{c_t}\right)^{-\frac{\eta}{\rho}} r_{c,t}^{\eta-1}\psi_{t+1}(r_{m,t+1} - r_{f,t+1}) = 0. \tag{A.11}$$

We also consider the scenario when the econometrician, in the absence of knowledge of the true SDF, incorrectly uses the power utility SDF when recovering its ψ-component:

$$\left\{\widehat{\psi}_t^{EL}\right\}_{t=1}^T = \underset{\{\psi_t\}_{t=1}^T}{\arg\max} \sum_{t=1}^T \log(\psi_t) \quad \text{s.t.} \quad \frac{1}{T}\sum_{t=1}^T \left(\frac{c_{t+1}}{c_t}\right)^{-\gamma} \psi_{t+1}(r_{m,t+1} - r_{f,t+1}) = 0. \tag{A.12}$$

As with Economy I, we recover ψ in both the above scenarios using the EL and ET approaches and use it to compute the term structure of the cost of all consumption fluctuations. The results are presented in Table A.IV. Panel A presents the true costs whereas Panels B and C report the mean and 90% confidence intervals of the costs obtained from 10,000 simulations for the correctly specified and misspecified scenarios, respectively. To simulate the model, we use the annual parameter estimates from Constantinides and Ghosh (2011).

Panel B shows that, for the correctly specified setup, the EL and ET approaches identify the term structure of the cost of one- to five-year fluctuations almost perfectly – the mean costs across the simulated samples are very close to the corresponding true costs in Panel A. The confidence intervals are tight for empirically realistic sample sizes. Panel C shows that, even when the SDF is misspecified, the EL and ET approaches identify the term structure of the cost of one- to five-year fluctuations fairly accurately. The mean cost estimates across the simulations are a bit higher than the true values – as an example, for $T = 87$, the EL approach produces an estimate of 0.39% at the 1-year horizon versus the true value of 0.14%, 0.75% versus 0.32% at the 2-year horizon, and 2.28% versus 0.98% at the 5-year horizon. However, note that the difference is economically small in all cases.

Overall, the simulation results suggest that the EL and ET estimators perform reasonably

Table A.IV: Simulation Results for Economy II, Cost of All Fluctuations (%)

	1 Yr	2 Yr	3 Yr	4 Yr	5 Yr
Panel A: True Values					
	.14	.32	.54	.76	.98
Panel B: Simulated Values, Correctly Specified SDF					
T=87	.12 / .12	.30 / .30	.54 / .54	.82 / .81	1.14 / 1.12
	[−.01,.26] [−.01,.26]	[.07,.61] [.06,.60]	[.16,1.09] [.15,1.09]	[.26,1.67] [.24,1.66]	[.35,2.34] [.33,2.32]
Panel C: Simulated Values, Misspecified SDF					
T=87	.39 / .42	.75 / .80	1.19 / 1.26	1.71 / 1.80	2.28 / 2.39
	[.23,.58] [.25,.65]	[.42,1.21] [.44,1.29]	[.62,2.02] [.65,2.13]	[.83,3.03] [.87,3.16]	[1.04,4.20] [1.10,4.37]

The table reports the (cumulative) costs of *all* aggregate consumption fluctuations, over one-to five-year horizons in a hypothetical economy. The samples are simulated from a hypothetical endowment economy in which a representative agent has Epstein-Zin recursive preferences and the aggregate consumption growth rate has a persistent component and fluctuating volatility. Panel A presents the true values of these costs of fluctuations. Panels B and C present the average of the costs, along with the 90% confidence intervals (in square brackets below), computed from $10,000$ simulated samples of size corresponding to the length of the historical annual data (T=87). To obtain the costs in Panels B–C, the ψ-component of the SDF is recovered using the EL (left entries in each cell) and ET approach (right entries in each cell). In Panel B, the econometrician uses the correct SDF implied by recursive preferences whereas in Panel C, she erroneously uses the power utility SDF when recovering the ψ-component of the SDF using the EL/ET approaches.

well at recovering the cost of fluctuations for empirically realistic sample sizes, in both correctly specified and misspecified settings. This lends further support for its use in the recovery of the pricing kernel for welfare cost calculations.

E Data Description

The extraction of the I-SDF for use in welfare cost calculations requires data on the aggregate consumption expenditures and returns on a set of traded assets. Ideally, we would like to use the longest available time series of these variables in the estimation to mitigate concerns that certain possible states may not have been realized in the sample. At the same time, to assess the robustness of our key results, we would like to repeat our analysis for different measures of consumption expenditures as well as different sets of assets. While data on total consumption is available from 1890 onwards, disaggregated expenditures on different consumption categories (e.g., durables, nondurables, and services) are only available from 1929 onwards. Moreover, data on broad cross sections of asset returns are also not available prior to the late 1920s. Therefore, we focus on a baseline data sample starting at the onset of the Great Depression (1929-2015).

For the 1929-2015 data sample, we consider two alternative measures of consumption: (i) the personal consumption expenditure on nondurables and services, and (ii) the personal consumption expenditure on durables, nondurables and services. The consumption data are obtained from the Bureau of Economic Analysis. Nominal consumption is converted to real using the Consumer Price Index (CPI).

We use different sets of assets to extract the I-SDF: (i) the market portfolio, proxied by the Center for Research in Security Prices (CRSP) value-weighted index of all stocks on the NYSE, AMEX, and NASDAQ, and (ii) the 6 equity portfolios formed from the intersection of two size and three book-to-market-equity groups. The proxy for the risk-free rate is the one-month Treasury Bill rate. The returns on all the above assets are obtained from Kenneth French's data library. Annual returns for the assets are computed by compounding monthly returns within each year and converted to real using the CPI. Excess returns on the portfolios are then computed by subtracting the risk free rate.

To further assess the robustness of our results, we also repeat our analysis using two alternative data sets: (i) total personal consumption expenditure over the 1890-2015 sample and the excess return on the $S\&P$ 500 as the sole asset, and (ii) the personal consumption expenditure on nondurables and services along with the excess return on the CRSP value-weighted market portfolio, over the entire available quarterly sample 1947:Q1-2015:Q4.

F Robustness: Alternative Definitions of Relative Entropy and Data Sample

In this section, we perform a number of checks to establish the robustness of our estimates of the cost of all consumption uncertainty as well as the cost of business cycle fluctuations in consumption reported in Section III of the paper. For all the robustness tests, consumption refers to the total personal consumption expenditure.[1]

Our first robustness check uses yet another definition of relative entropy (a third alternative to the EL and ET approaches). Specifically, we recover the risk-neutral measure $\mathbb{Q}$ such that:

$$\widehat{\mathbb{Q}} = \min_{\mathbb{Q}} \int \mathbf{log}\left(\frac{d\mathbb{Q}}{d\mathbb{Q}^m}\right) d\mathbb{Q} = \int \mathbf{log}\left(\frac{q(\mathbf{z})}{q^m(\mathbf{z})}\right) q(\mathbf{z}) d\mathbf{z} \text{ s.t. } \mathbf{0} = \int \mathbf{R}^e(\mathbf{z}) q(\mathbf{z}) d\mathbf{z}, \quad \text{(A.13)}$$

where $\frac{d\mathbb{Q}^m}{d\mathbb{P}} = \frac{(\Delta C)^{-\gamma}}{E\left[(\Delta C)^{-\gamma}\right]}$. In other words, $\mathbb{Q}^m$ is the risk neutral measure implied by the power utility model with a constant CRRA. Thus, Equation (A.13) recovers the risk neutral measure $\mathbb{Q}$ that is minimally distorted relative to the CRRA model implied risk neutral measure $\mathbb{Q}^m$, while also successfully pricing the set of test assets used in the estimation. Note that the main difference between Equation (A.13) and the EL and ET estimators defined in Equations (14) and (15), respectively, is that while the latter two minimize the

[1] Very similar results are obtained using nondurables and services consumption and are omitted for brevity.

relative entropy (or distance) between the recovered measure and the physical measure, the former minimizes the distance between the recovered risk neutral measure and the risk neutral measured implied by a candidate model SDF.

The solution to Equation (A.13) is obtained as:

$$\hat{q}_t = \frac{e^{\hat{\theta}(\gamma)'\mathbf{r}_t^e} (\Delta c_t)^{-\gamma}}{\frac{1}{T}\sum_{t=1}^{T} e^{\hat{\theta}(\gamma)'\mathbf{r}_t^e} (\Delta c_t)^{-\gamma}} \quad \forall t, \tag{A.14}$$

where $\hat{\theta}(\gamma) \in \mathbb{R}^N$ is the vector of Lagrange multipliers that solves the dual problem:

$$\hat{\theta}(\gamma) = \arg\min_{\theta} \left[\log \left(\frac{1}{T} \sum_{t=1}^{T} e^{\theta' \mathbf{r}_t^e} (\Delta c_t)^{-\gamma} \right) \right]. \tag{A.15}$$

We use the recovered risk neutral measure $\hat{q}_t$ to calculate the cost of consumption fluctuations. The results, reported in Table A.3, Row 1 of Panels A and B, for the scenarios when the test assets consist of the market portfolio alone and the six Fama-french portfolios, respectively, are very similar to those obtained with the EL and ET (Table 1, Panel B, Rows 1-4) approaches.

Table A.V: Cumulative Cost of Total Consumption Fluctuations, Robustness Checks

	All Fluctuations					B. C. Fluctuations				
	1 Yr	2 Yr	3 Yr	4 Yr	5 Yr	1 Yr	2 Yr	3 Yr	4 Yr	5 Yr
Panel A: Market Portfolio										
$I\text{-}SDF^{ET}$	2.07	6.13	14.68	17.32	17.31	0.90	1.98	4.47	4.98	4.59
$I\text{-}SDF^{Alt}$	1.83	4.92	10.87	12.75	12.70	.851	1.75	3.48	3.82	3.52
1890-2015	1.38	2.69	4.85	6.84	8.24	.931	1.39	2.08	2.54	2.69
Panel B: FF 6 Portfolios										
$I\text{-}SDF^{ET}$	1.83	4.81	8.72	14.09	14.75	.691	1.43	2.73	4.01	3.83
$I\text{-}SDF^{Alt}$	1.76	4.46	8.96	14.12	14.67	.764	1.61	3.02	4.20	4.00
Panel C: Mkt, S, B, G, V, W, L										
$I\text{-}SDF$	3.51	4.57	9.83	9.69	9.23	1.75	2.60	4.08	3.61	3.34

The table reports the (cumulative) cost of *all* aggregate consumption fluctuations (Columns 2-6) and the cost of business cycle fluctuations in consumption (Columns 7-11), for one- to five-year time horizons. Consumption denotes the real *total* personal consumption expenditure (includes durables, nondurables, and services). Panel A presents results when the excess return on the market portfolio is the sole asset used to recover the I-SDF. In Panel B, on the other hand, the I-SDF is estimated using the 6 Fama-French size and book-to-market-equity sorted portfolios. In each of Panels A–B, the costs are calculated using the I-SDF extracted using the exponential tilting (ET) approach (Row 1), the risk-neutral measure recovered by minimizing the distance from the CRRA model-implied risk-neutral measure while satisfying the pricing restrictions (Row 2) and the I-SDF extracted with the EL approach over the longer 1890-2015 sample (Row 3). Panel C presents the costs obtained with the I-SDF extracted using the EL approach for a larger cross section of test assets, namely the excess returns on the market portfolio and the top and bottom deciles of portfolios formed by sorting stocks on the basis of size, book-to-market-equity ratio, and momentum. The sample is annual covering the period 1929-2015, except for Row 3 in Panel A where it extends over 1890-2015.

Second, we present the costs of fluctuations using the EL approach with data going back as far as 1890. The excess return on the market is the sole test asset, with the return on the S&P composite index used as a proxy for the market return and the prime commercial paper rate as a proxy for the risk free rate. The data are obtained from Robert Shiller's website. The costs of all and business cycle fluctuations in consumption, presented in Row 2 of Panel A, are smaller than those obtained using the baseline 1929-2015 sample (see Table 1, Panel B, Rows 1 and 3 for the EL and ET, respectively). The smaller estimates of the cost obtained in this longer data sample can be accounted for, at least partly, by the usage of the commercial paper rate as a proxy for the risk free rate, thereby leading to an underestimation of the magnitude of the equity premium in this sample. Specifically, the average level of the equity premium is 7.9% in the baseline sample, more than double the value of 3.1% in the longer 1890 onwards sample. Moreover, just as with the baseline sample, the cost of business cycle fluctuations still accounts for a substantial fraction (more than a third) of the cost of all consumption fluctuations for all the horizons considered.[2]

Overall, our results suggest that the estimates of the cost of aggregate economic fluctuations are fairly robust to the measure of consumption expenditures, the set of test assets used to recover the I-SDF, the choice of sample period, as well as the definition of relative entropy. This lends further support to the quantitative estimates in the paper.

G Time-Varying Cost of One-Period Consumption Fluctuations Using SET

The Smoothed Exponential Tilting (SET) estimator is defined as:

$\forall i \in \{1, \dots, T\}, \quad \forall \gamma \in \Theta,$

$$\left\{\widehat{f}_{i,\cdot}^{SET}(\gamma)\right\} = \underset{(f_{i,\cdot}) \in \Delta_i}{\arg\min} \sum_{j=1}^{T} \log\left(\frac{f_{i,j}}{\omega_{i,j}}\right) f_{i,j} \quad \text{s.t.} \quad \sum_{j=1}^{T} f_{i,j} \times (\Delta c_j)^{-\gamma} \, \mathbf{r}_j^e = \mathbf{0}. \tag{A.16}$$

Figure A.2 plots the time series of the cost of one-period consumption uncertainty obtained with the SET approach (blue-dashed line). For the sake of comparison, we also plot the time series of the cost obtained with the SEL approach (solid red line). The correlation between the two time series is 99.94%.

[2]Since the size and book-to-market-equity sorted portfolios are not available prior to the late 1920s, we cannot recover the I-SDF using these portfolios over the 1890-2015 sample.

H Additional Tables

Table A.VI: Cumulative Cost of Consumption Fluctuations for Alternative γ

	1 Yr	2 Yr	3 Yr	4 Yr	5 Yr
Panel A: Cost of All Fluctuations					
$\gamma = 0.0$ (.013/.000)	.44	1.33	1.98	1.97	1.81
$\gamma = 2.0$ (.014/.000)	.64	2.00	3.36	3.59	3.52
$\gamma = 3.5$ (.014/.000)	.79	2.54	4.62	5.14	5.13
$\gamma = 10.0$ (.015/.000)	1.53	5.15	11.75	14.28	14.44
$\gamma = 16.0$ (.022/.016)	2.42	7.54	17.7	22.0	22.0
$\gamma = 26.2$ (.161/.453)	4.07	10.59	21.95	27.60	27.43
Panel B: Share of Costs due to Business Cycle Fluctuations					
$\gamma = 0.0$ (.013/.000)	22.73%	27.07%	32.83%	34.52%	34.81%
$\gamma = 2.0$ (.014/.000)	28.13%	26.50%	30.36%	30.08%	28.41%
$\gamma = 3.5$ (.014/.000)	27.85%	26.77%	29.44%	28.79%	26.51%
$\gamma = 10.0$ (.015/.000)	36.60%	28.74%	28.85%	27.31%	24.72%
$\gamma = 16.0$ (.022/.016)	42.98%	31.96%	29.21%	27.14%	24.91%
$\gamma = 26.2$ (.161/.453)	49.14%	37.02%	29.98%	27.25%	25.81%

The table reports the (cumulative) costs of *all* aggregate consumption fluctuations (Panel A) and the costs of business cycle fluctuations in consumption (Panel B), over one-to five-year horizons, using. Rows 1-6 present the results for alternative values of the SDF parameter γ. Under the calibrated value of γ we report bootstrapped p-values for the true value being smaller than, or equal to, the reported value. The first p-value calculated is across all simulated sample, while the second is conditional on having at least one bootstrapped observation corresponding to the Great Depression period. Consumption denotes the real personal consumption expenditure of nondurables and services. The costs are calculated using the I-SDF recovered from the market portfolio alone with the EL approach. The sample is annual covering the period 1929-2015.

Table A.VII: Predictive Regressions of Consumption and the SDF

	$\gamma=0$		$\gamma=2$		$\gamma=10$	
	Coeff.	$R^2(\%)$	Coeff.	$R^2(\%)$	Coeff.	$R^2(\%)$
	Panel A: $m_{t+1\to t+s}=a+b\Delta c_{t+1}+\epsilon_{t+s}$					
$s=2$	-3.88	4.73	-6.89	13.84	-18.30	53.24
	(-2.03)		(-3.65)		(-9.72)	
$s=3$	-3.08	2.61	-6.55	10.37	-19.77	41.04
	(-1.48)		(-3.08)		(-7.55)	
$s=4$	-0.85	0.18	-4.41	4.58	-17.77	31.30
	(-0.39)		(-1.97)		(-6.07)	
$s=5$	-0.10	0.00	-3.38	2.52	-15.50	24.17
	(-0.04)		(-1.44)		(-5.05)	
	Panel B: $\Delta c_{t+1\to t+s}=a+bm_{t+1}+\epsilon_{t+s}$					
$s=2$	-0.08	30.05	-0.10	46.29	$-.11$	89.95
	(-6.09)		(-8.46)		(-21.64)	
$s=3$	-0.09	24.27	-0.11	36.88	$-.12$	67.14
	(-5.13)		(-6.92)		(-12.94)	
$s=4$	-0.09	18.51	-0.11	28.29	$-.12$	51.53
	(-4.29)		(-5.65)		(-9.28)	
$s=5$	-0.07	12.89	-0.09	20.12	$-.10$	37.44
	(-3.44)		(-4.49)		(-6.92)	

The Table reports univariate regressions of the multi-period log SDF (m) on the lagged log consumption growth (Δ), in Panel A, and univariate regressions of multi-period consumption growth on the lagged one-period log SDF, Panel B. The sample is annual, covering the period 1930-2015. t-statics are reported in parenthesis under the estimated slope coefficient.

I Additional Figures

Figure A.2: Time-Varying Cost of One-Period Consumption Fluctuations: SEL vs SET

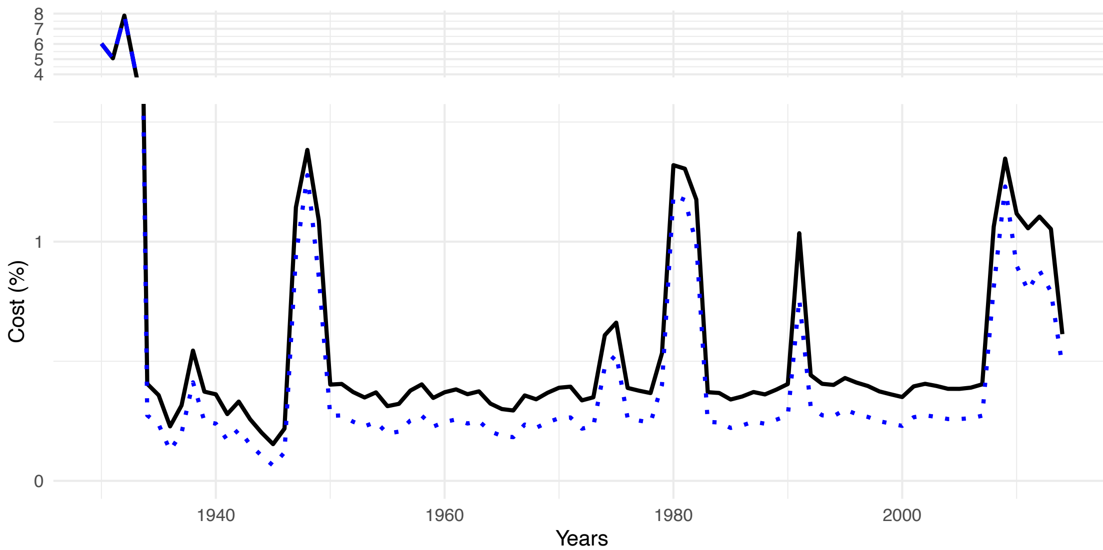

Notes: The figure plots the time series of the cost of one-period consumption uncertainty. The cost is estimated using the SEL (solid black line) and SET (blue-dotted line) approaches, using nondurables and services consumption as the measure of the consumption expenditures, the excess return on the market portfolio as the sole test asset, and an exponentially-weighted moving average of lagged consumption growth as the conditioning variable. The sample is annual, covering the period 1930-2015.

Figure A.3: Time-Varying Cost: Robustness to SDF Parameter

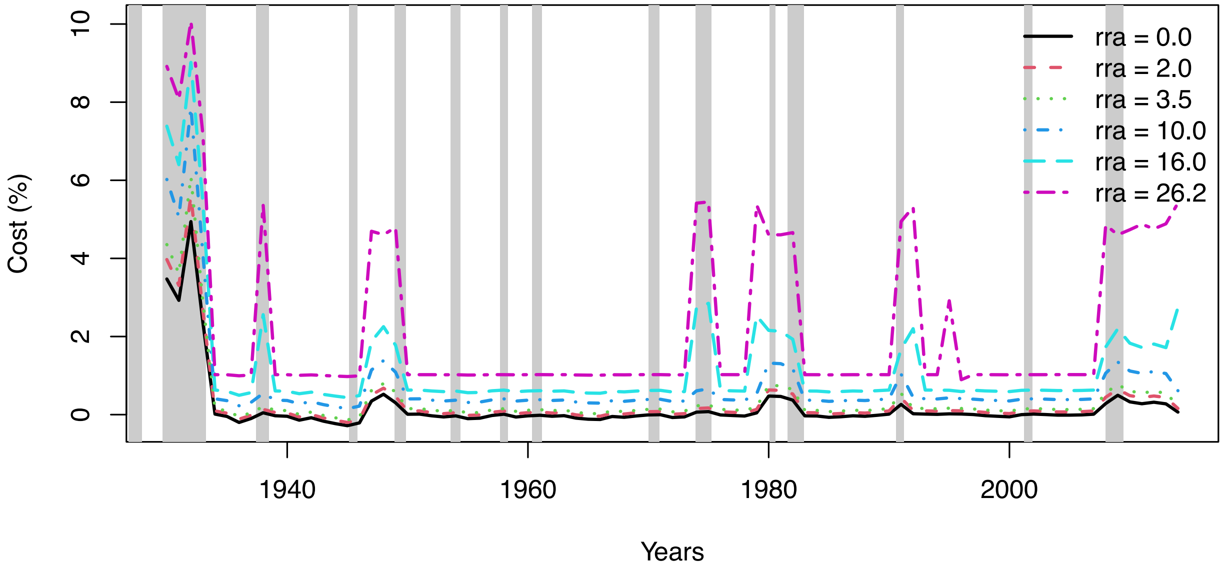

Panel A: Annual data, 1929-2015

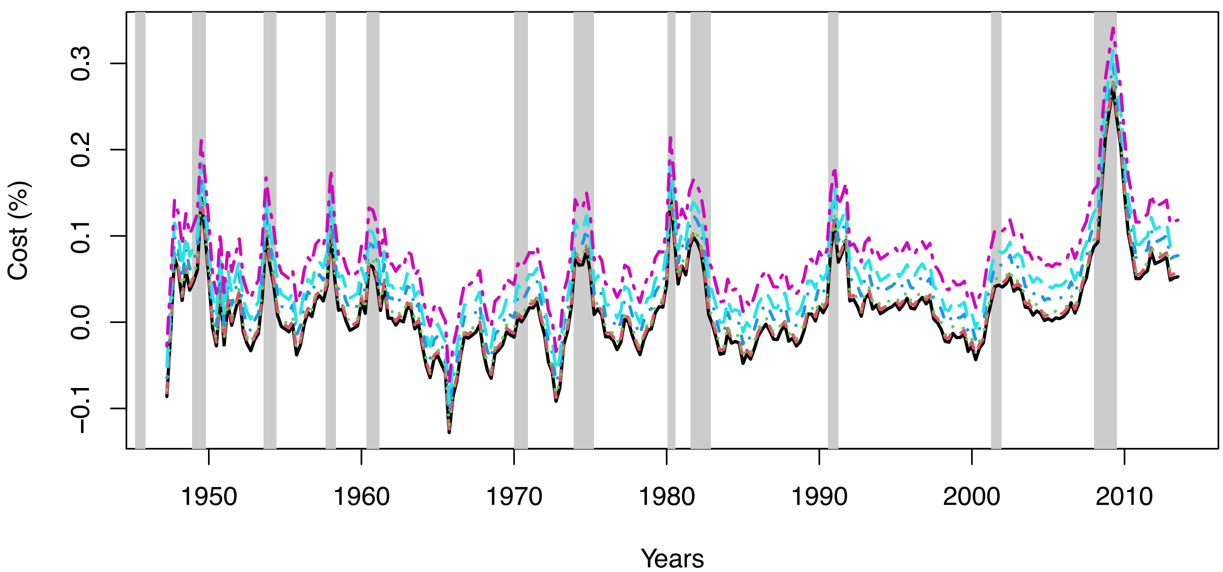

Panel B: Quarterly data, 1947:Q1-2015:Q4

Notes: The figure plots the time series of the cost of one-period consumption uncertainty, for different values of the SDF parameter γ. The cost is estimated using the SEL approach, using nondurables and services consumption as the measure of the consumption expenditures, the excess return on the market portfolio as the test asset, and an exponentially-weighted moving average of lagged consumption growth as the conditioning variable. The sample is annual, covering the period 1930-2015 (Panel A) or quarterly, covering the period 1947:$Q$1–2015:$Q$4 (Panel B).