# The Market Cost of Business Cycle Fluctuations *

*By* Anisha Ghosh, Christian Julliard, and Michael J. Stutzer†

Draft: August 15, 2022

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.

*JEL: E3, E2, G12, C5.*
*Keywords: Aggregate Uncertainty, Business Cycle Risk, Pricing Kernel, Empirical Likelihood, Smoothed Empirical Likelihood.*

* 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.
† Anisha Ghosh: Desautel Faculty of Management, McGill University, anisha.ghosh@mcgill.ca. Christian Julliard: Department of Finance, London School of Economics, c.julliard@lse.ac.uk. Michael J. Stutzer: Leeds School of Business, The 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., Lars Peter Hansen and Kenneth J. Singleton (1983), Martin Lettau and Sydney Ludvigson (2001), Jonathan A. Parker and Christian Julliard (2005), Christian Julliard and Anisha 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

[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 Rajnish Mehra and Edward C. Prescott (1985).

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 Robert E. Lucas (1987) original setting is nested within this family, corresponding to the case in which $\psi_t$ is a positive constant. Prominent examples of models in this class are: habit formation models (see, e.g., John Y. Campbell and John H. Cochrane (1999), Lior Menzly, Tano Santos and Pietro Veronesi (2004)); long run risks models (e.g., Ravi Bansal and Amir Yaron (2004)); models with complementarities in consumption (e.g., Monika Piazzesi, Martin Schneider and Selale Tuzel (2007), Motohiro Yogo (2006)); models in which $\psi_t$ captures aggregation over heterogeneous agents who face uninsurable idiosyncratic shocks to their labor income (e.g. George M. Constantinides and Darrell Duffie (1996), George M. Constantinides and Anisha Ghosh (2017)), as well as solvency constraints (e.g. Hanno N. Lustig and Stijn G. Van 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 $\psi_t$ captures deviations from rational expectations (e.g. Suleyman Basak and Hongjun Yan (2010), Lars Peter Hansen and Thomas J. Sargent (2010)).

Estimates of the cost of business cycles vary widely across these model specifications (see, e.g., Gadi 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, we reflective of agents' attitude towards risk. Anisha Ghosh, Christian Julliard and Alex 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, our approach does not require investors' beliefs to be rational. The approach 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 $\psi$-component of the pricing kernel $M$ such that the resultant $M$ satisfies the unconditional Euler equations for the assets, i.e. successfully prices broad cross-sections 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. Fernando Alvarez and Urban J. Jermann (2004) show that, in the context of a 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 Art B. Owen (2001) and the Exponential Tilting (ET) estimator of Y. Kitamura and M. 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

[3]See also Anisha Ghosh, Christian Julliard and Alex 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.

Empirical Likelihood (SEL) estimator of Yuichi Kitamura, Gautam Tripathi and Hyungtaik Ahn (2004) and the Smoothed Exponential Tilting (SET) estimator (see, e.g., Anisha Ghosh, Taisuke Otsu and Guillaume Roussellet (2020)) that extend the EL and 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 somewhat sensitive to the assumed value of the utility curvature parameter ($\gamma$ 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.

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 Michael Stutzer (1995, 1996) and Kitamura and Stutzer (1997) (see Yuichi 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., Caio Almeida and René Garcia (2012), David Backus, Mikhail Chernov and Stanley E. Zin (2013), Caio Almeida and René Garcia (2016)), construct bounds on the SDF and its components and recover the missing component from a candidate SDF (see, e.g., Jaroslav Borovicka, Lars P. Hansen and Jose A. Scheinkman (2016), Ghosh, Julliard and Taylor (2017), Mirela Sandulescu, Fabio Trojani and Andrea Vedolin (2018)), performance evaluation of funds (see, e.g., Caio Almeida, Kym Ardison and René Garcia (2019)), and recover investors' beliefs from observed asset prices (see, e.g., Lars Peter Hansen (2014), Anisha Ghosh and Guillaume Roussellet (2019), Xiaohong Chen, Lars P. Hansen and Peter G. 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), Ayse Imrohoroglu (1989), Andrew Atkeson and Christopher Phelan (1994), Maurice Obstfeld (1994), James Pemberton (1996), Jim Dolmas (1998), Thomas Tallarini (2000), Paul Beaudry and Carmen Pages (2001), Christopher Otrok (2001), Kjetil Storesletten, Chris I. Telmer and Amir Yaron (2001), Alvarez and Jermann (2004), Tom Krebs (2007), Ian Martin (2008), Robert J. Barro (2009), Per Krusell and Anthony A. Smith (2009), Larry G. Epstein, Emmanuel Farhi and Tomasz Strzalecki (2014), Hanno N. Lustig, Stijn G. Van Nieuwerburgh and Adrien 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. 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 Hang Bai and Lu Zhang (2020) that develops a general equilibrium model with recursive utility, search frictions, and capital accumulation. Jessie Davis and Gill 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 V 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.A, we follow Alvarez and Jermann (2004) and define the marginal cost of aggregate consumption fluctuations, for two alternative definitions of fluctuations. In Subsection II.B, 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.

### *A. The Cost of Aggregate Fluctuations*

The cost (or, the market price) of consumption fluctuations, $\omega_0$, 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 $\omega_0$ 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) 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 $\mu_c$ 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. Calibrating $\mu_c = 2.3\%$, $y_0 = 3.0\%$ and $r_0 - y_0 \geq 0.2\%$, they obtain a very high estimate of the 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 $\mu_c$. In fact, in Alvarez and Jermann (2004), the estimate of the cost varies in the range 28.0%–1535.7% based on different calibrations of the parameters. Specifically, as $y_0 \to \mu_c$, we have $\omega_0 \to \infty$, and the approach breaks down. Olivier J. Blanchard (2019) points out that, at the current time, the nominal rate on a 10-year government bond is 2.7%, while the expected nominal growth rate is 4.0%, causing $y_0 - \mu_c$ to be negative, thereby negating 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 1980$s$. This reveals the fragility of the results obtained using Equation (3).

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. 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. Gadi 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) underlying 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.

Note that 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, 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 data generating process.

### *B. Measuring the Cost of Aggregate Fluctuations*

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

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

where $M_t$ is the SDF, $\underline{\mathcal{F}_t} = \{\mathcal{F}_t, \mathcal{F}_{t-1}, \ldots\}$ denotes the investors' information set at time $t$, and $\mathbb{E}^{\mathbb{P}}\left[.|\underline{\mathcal{F}_t}\right]$ 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

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

$\tilde{pc}_{1,t}$ 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

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

implying that

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

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 $\gamma$ denotes the CRRA and $\delta$ the subjective discount factor. Moreover, assuming log-normality of the aggregate consumption growth as in Lucas (1987):

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

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

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

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 data generating process. 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 $\psi$-component.

### 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:

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

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 $\frac{d\mathbb{F}(\mathbf{z})}{d\mathbb{P}(\mathbf{z})} = \frac{\psi(\mathbf{z})}{\mathbb{E}^{\mathbb{P}}(\psi(\mathbf{z}))}$.

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:

(14)
$$\min_{\mathbb{F}} \int \log\left(\frac{d\mathbb{P}}{d\mathbb{F}}\right) d\mathbb{P} \equiv \int_{\mathbf{z}} \log\left(\frac{p(\mathbf{z})}{f(\mathbf{z})}\right) p(\mathbf{z}) d\mathbf{z}, \text{ s.t. } \mathbf{0} = \int_{\mathbf{z}} \mathbf{R}^e(\mathbf{z}) \left(\Delta C(\mathbf{z})\right)^{-\gamma} f(\mathbf{z}) d\mathbf{z}.$$

Since relative entropy is not symmetric, we can reverse the roles of $\mathbb{P}$ and $\mathbb{F}$ in Equation (14) to obtain an alternative divergence criterion between these two measures, which can be minimized to recover an alternative estimate of the measure, $\mathbb{F}$:

(15)
$$\min_{\mathbb{F}} \int \mathbf{log}\left(\frac{d\mathbb{F}}{d\mathbb{P}}\right) d\mathbb{F} = \int \mathbf{log}\left(\frac{f(\mathbf{z})}{p(\mathbf{z})}\right) f(\mathbf{z}) d\mathbf{z} \text{ s.t. } \mathbf{0} = \int \mathbf{R}^e(\mathbf{z}) \left(\Delta C(\mathbf{z})\right)^{-\gamma} f(\mathbf{z}) d\mathbf{z},$$

Equations (14) and (15) are, respectively, the Empirical Likelihood (EL) estimator of Owen (2001) and the Exponentially-Tilted (ET) estimator of Kitamura and Stutzer (1997) (see also Susanne M. 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, $\psi$, 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.

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 (14) and (15), naturally imposes the non-negativity of the pricing kernel.

Second, our approach to recover the $\psi_t$ component satisfies the Occam's razor, or law of parsimony, since it adds the *minimum amount of information* needed for the pricing kernel to price assets. 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 (14) (the same holds for Equation (15)). 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 $\psi$-component of the kernel or the physical distribution $\mathbb{P}$.

Third, as implied by the work of Donald E. Brown and Robert L. Smith (1990), the use of entropy is desirable if we think that tail events are an important component of the risk measure.[4]

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 (14) can be made

[4]Brown and Smith (1990) develop what they call "a Weak Law of Large Numbers for rare events;" that is, they show that the empirical distribution observed in a very large sample converges to the distribution that minimizes the relative entropy.

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 $\psi$-component of the SDF:[5]

$$\text{(16)} \qquad \arg\max_{\{\psi_t\}_{t=1}^T} \frac{1}{T}\sum_{t=1}^{T} \log \psi_t \quad \text{s.t.} \quad \frac{1}{T}\sum_{t=1}^{T} (\Delta c_t)^{-\gamma}\, \psi_t \mathbf{r}_t^e = \mathbf{0}.$$

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

$$\text{(17)} \qquad \hat{\psi}_t = \frac{1}{T(1+\hat{\theta}(\gamma)'\mathbf{r}_t^e\,(\Delta c_t)^{-\gamma})} \quad \forall t,$$

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

$$\text{(18)} \qquad \hat{\theta}(\gamma) = \arg\min_{\theta} -\sum_{t=1}^{T} \log(1+\theta'\mathbf{r}_t^e\,(\Delta c_t)^{-\gamma}).$$

The solution to Equation (15) is similarly simple to implement and is presented in Appendix A.A1.

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 (14)). 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 more recent Fama-French 3 and 5 factor models. Appendix A.A5 reproduces a table from Ghosh, Julliard and Taylor (2022) on the out-of-sample performance of the I-SDF vis a vis other popular factor models. This suggests that the I-SDF is more successful at capturing the relevant sources of

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

priced risk and, therefore, offers a more reliable candidate kernel with which to measure the cost of aggregate economic fluctuations.

Finally, we show via simulations in Appendix A.A2 that the methodology is quite successful in estimating the cost of fluctuations in hypothetical economies for empirically realistic sample sizes.

With the recovered $\psi$-component, the I-SDF is obtained (up to a positive scale factor, $\kappa$) as

$$\left\{\widehat{M}_t\right\}_{t=1}^{T} = \left\{\kappa\left(\Delta c_t\right)^{-\gamma}\widehat{\psi}_t\right\}_{t=1}^{T}. \tag{19}$$

The proportionality constant, $\kappa$, 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:

$$\widehat{p\tilde{c}_1^{stab}/p\tilde{c}_1} - 1 = \frac{\sum_{t=1}^{T}\widehat{M}_t\left(1+\mu_c\right)}{\sum_{t=1}^{T}\widehat{M}_t\Delta c_t} - 1. \tag{20}$$

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

$$\widehat{p\tilde{c}_1^{stab}/p\tilde{c}_1} - 1 = \frac{\sum_{t=1}^{T}\widehat{M}_t\Delta c_t^{stab}}{\sum_{t=1}^{T}\widehat{M}_t\Delta c_t} - 1, \tag{21}$$

where $\Delta c_t^{stab}$ 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.14 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 (20) and (21) 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 (19), can be compounded to recover the $j$-period discount factor:

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

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 (20)-(21).

### 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, Otsu and Roussellet (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 data generating process.

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 through conditional moment restrictions.

We illustrate the methodology using the SEL estimator. The absence of arbitrage opportunities implies the following conditional pricing restrictions:

$$\mathbb{E}^{\mathbb{P}_t}\left[M_{t+1}\mathbf{R}^e_{t+1}|\underline{\mathcal{F}_t}\right] = \mathbb{E}^{\mathbb{P}_t}\left[(\Delta C_{t+1})^{-\gamma}\,\psi_{t+1}\mathbf{R}^e_{t+1}|\underline{\mathcal{F}_t}\right] = \mathbf{0}, \tag{22}$$

where the first equality follows from the assumed multiplicative decomposition of the SDF. Under weak regularity conditions, we have

$$\mathbb{E}^{\mathbb{P}_t}\left[(\Delta C_{t+1})^{-\gamma}\,\frac{\psi_{t+1}}{\mathbb{E}^{\mathbb{P}_t}(\psi_{t+1}|\underline{\mathcal{F}_t})}\mathbf{R}^e_{t+1}|\underline{\mathcal{F}_t}\right] = \mathbb{E}^{\mathbb{F}_t}\left[(\Delta C_{t+1})^{-\gamma}\,\mathbf{R}^e_{t+1}|\underline{\mathcal{F}_t}\right] = \mathbf{0}, \tag{23}$$

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

We assume that the time-$t$ information set of the investors, $\underline{\mathcal{F}_t}$, 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: $\forall i \in \{1, \ldots, T\}, \quad \forall \gamma \in \Gamma$,

$$\text{(24)}\quad \left\{\widehat{f}_{i,\cdot}^{SEL}(\gamma)\right\} = \underset{(f_{i,\cdot})\in\Delta_i}{\arg\min} \sum_{j=1}^{T} \log\left(\frac{\omega_{i,j}}{f_{i,j}}\right)\omega_{i,j} \quad \text{s.t.} \quad \sum_{j=1}^{T} f_{i,j}\times(\Delta c_j)^{-\gamma}\,\mathbf{r}_j^e = \mathbf{0}.$$

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

$$\text{(25)}\qquad \omega_{i,j} = \frac{\mathcal{K}\left(\dfrac{x_i - x_j}{b_T}\right)}{\sum\limits_{t=1}^{T}\mathcal{K}\left(\dfrac{x_i - x_t}{b_T}\right)},$$

where $\mathcal{K}$ is a kernel function belonging to the class of second order product kernels,[6] and the bandwidth $b_T$ is a smoothing parameter.[7]

The objective function in Equation (24) 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, $\frac{f_{t,j}}{\omega_{t,j}} = \frac{\psi_{t,j}}{\mathbb{E}^{\mathbb{P}_t}(\psi_{t,j}|\underline{\mathcal{F}_t})}$ 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 (24) is analytical and given by:

[6] $\mathcal{K}$ 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 $[-a, a]$.

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

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

$$\widehat{f}_{i,j}^{SEL}(\gamma) = \frac{\omega_{i,j}}{1 + (\Delta c_j)^{-\gamma}\, \widehat{\theta}_i(\gamma)'\, \mathbf{r}_j^e}\,, \tag{26}$$

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:

$$\widehat{\theta}_i(\gamma) = \arg\max_{\theta_i \in \mathbb{R}^N} \sum_{j=1}^{T} \omega_{i,j} \log\left[1 + (\Delta c_j)^{-\gamma}\, \theta_i'\, \mathbf{r}_j^e\right]. \tag{27}$$

Equations (26) and (27) 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 $\gamma$. 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.

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

$$\frac{\frac{V_t\left(C_{t+1}^{stab}\right)}{C_t}}{\frac{V_t(C_{t+1})}{C_t}} - 1 = \frac{\mathbb{E}^{\mathbb{F}_t}\left[(\Delta C_{t+1})^{-\gamma}(1+\mu_c)\,|\underline{\mathcal{F}_t}\right]}{\mathbb{E}^{\mathbb{F}_t}\left[(\Delta C_{t+1})^{-\gamma}(\Delta C_{t+1})\,|\underline{\mathcal{F}_t}\right]} - 1 = \frac{(1+\mu_c)\sum_{j=1}^{T}\widehat{f}_{t,j}^{SEL} \times (\Delta c_j)^{-\gamma}}{\sum_{j=1}^{T}\widehat{f}_{t,j}^{SEL} \times (\Delta c_j)^{1-\gamma}} - 1. \tag{28}$$

Finally, note that the question naturally arises as to the economic interpretation of the recovered $\psi$-component of the kernel. For instance, it could capture a misspecification of investors' risk preferences relative to the power utility SDF. Alternatively, the $\psi$-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 recovered component is quite similar across a range of preference specifications. Moreover, consistent with the interpretation of $\psi$ as capturing investors' beliefs, they show that the recovered beliefs about consumption growth have strong forecasting power for consumption growth, 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.

## Properties of the Estimators

The asymptotic properties of the EL/ET and SEL/SET estimators of a finite-dimensional parameter vector (the SDF parameter $\gamma$ in our setting) have been studied in the case of a *correctly specified model* – see, e.g., J. Qin and J. 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 $\psi$-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 $\psi$-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 $\gamma$, we are more interested in the missing $\psi$-component of the SDF. This involves an important extension of the econometric results in the above papers.

Ghosh, Otsu and Roussellet (2020) show that, in this set up, under mild regularity conditions, the estimated $\psi$-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 Yuichi Kitamura (2003) and Ivana Komunjer and Giuseppe 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 (14) and (15)). 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., H. White (1982), Quang H. 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 A.A2, suggests that this approach accurately recovers the cost of aggregate fluctuations for empirically realistic sample sizes. Specifically, the 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 $\psi$-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 $\gamma$ that enters the welfare cost calculations (see, e.g., Equations (20)–(21) for the expected cost and Equation (28) 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 $\gamma$.

For example, in the time and state separable power utility model, $\gamma$ 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 Larry G. Epstein and Stanley E. Zin (1989) recursive preferences, $M_t = \delta^{\eta}\left(\frac{C_t}{C_{t-1}}\right)^{-\frac{\eta}{\rho}} R_{c,t}^{\eta-1} = \delta^{\eta}\left(\frac{C_t}{C_{t-1}}\right)^{-\gamma}\left(\frac{\frac{P_{c,t}}{C_t}+1}{\frac{P_{c,t-1}}{C_{t-1}}}\right)^{\eta-1}$, where $\rho$ is the elasticity of intertemporal substitution and $\eta = \frac{1-\gamma}{1-\frac{1}{\rho}}$ (the second equality follows from factorizing out consumption growth from the return on total wealth). These models typically calibrate $\gamma = 10$ (see, e.g., Bansal and Yaron (2004)). Some models with recursive preferences calibrate $\gamma$ to much larger values (e.g., Monika Piazzesi and Martin 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 $\gamma$ 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 $\gamma$ (typically in excess of 7) to satisfy entropy bounds for admissible SDFs, that are tighter than the seminal variance bounds of Lars Peter Hansen and Ravi 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, $\gamma^{-1}$ is the intertemporal elasticity of substitution, and $\zeta$ 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 $\gamma$ (typically in excess of 20) to satisfy entropy bounds for admissible SDFs. In models with rare disasters, $\gamma$ is typically calibrated to values between 3 and 4 (see, e.g., Robert J. Barro (2006), Jessica 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 $\gamma$ to 10 or higher values and/or require such high values of the parameter to explain asset prices.

Also, in addition to recovering the $\psi$-component of the SDF, our information-theoretic procedure offers a way to estimate $\gamma$. For instance, the EL estimator of $\gamma$ is defined as (see Kitamura (2006)):

(29)

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

In other words, the EL approach searches for a value of $\gamma$ 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 $\gamma$ is similarly defined, albeit swapping the roles of $\mathbb{P}$ and $\mathbb{F}$.

We estimate $\gamma$ 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 A.A3 for a description of the data). Figure 1 plots the EL objective function in Equation (29) as a function of $\gamma$. The point estimate of $\gamma$ 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.

Motivated by the observations that most theoretical models calibrate $\gamma$ 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 $\gamma$ 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 $\gamma$.

Figure 1. Profile Likelihood

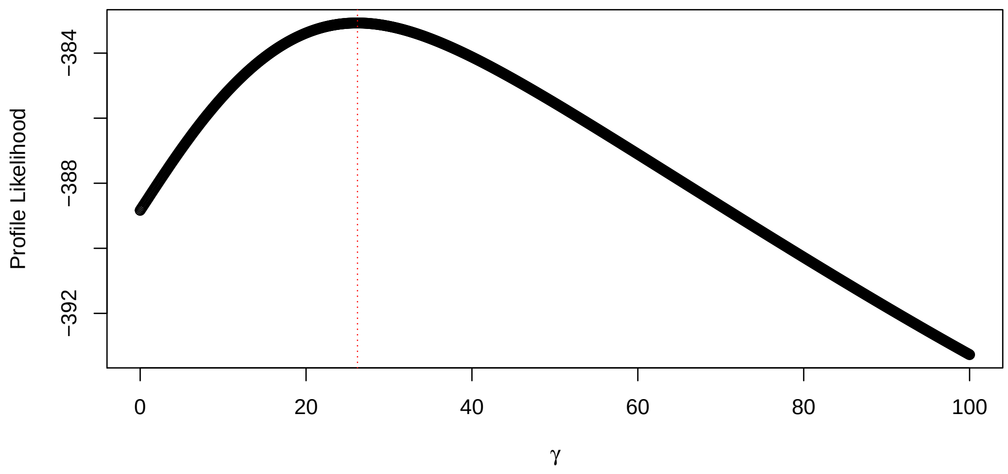


*Note: The figure plots the EL objective function as a function of the SDF parameter* $\gamma$*. The dotted vertical line denotes the point estimate of* $\gamma$*. Consumption denotes the real personal total consumption expenditure (includes durables, nondurables, and services). The excess return on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

We next proceed to estimate the cost of fluctuations. Section III.A presents the term structure of the expected (or, average) cost of consumption fluctuations. Section III.B reports the nature of time-variation in the cost. Finally, in Appendix A.A4, we present a host of robustness checks, including alternative definitions of relative entropy and an alternative longer sample period going back as far as 1890.

### *A. The Average Cost of 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 the ratio of 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 Morten O. Ravn and Harald Uhlig (2002).

Equations (20)–(21) 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 and ET approaches, to measure this cost. The costs of multi-year fluctuations are obtained by compounding the I-SDF, as explained in Section II.B. Note that the recovered I-SDF depends on the particular measure of the aggregate consumption expenditures as well as on the set of assets used (see Equations (17)-(18)). To ensure robustness, we estimate the I-SDF using two different measures of consumption expenditures and two 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 choices of test assets are 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., Jimmy Liew and Maria Vassalou (2000), Parker and Julliard (2005)). They, therefore, constitute appropriate test 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.[8]

The results 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 Rows 1–2, the market portfolio is the sole test asset and the EL and ET approaches, respectively, are used in the extraction of the I-SDF. Row 1, Column 2 shows that the estimated cost of all one-period consumption fluctuations is 1.53% using the EL estimator. Row 2, Column 2 shows that the ET estimator implies a very similar cost at 1.47%. In Rows 3–4, the six size and book-to-market-equity sorted portfolios of Fama-French are used to recover the I-SDF, using the EL and ET approaches, respectively. Column 2 of Rows 3–4 show that the estimated cost of all one-period consumption fluctuations remains quite similar at 1.29% and 1.28%, respectively, with the EL and ET approaches. Row 5 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 6 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 (Row 2) 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 (4.5%, 10.3%, 11.9%, and 11.9%, respectively) using the EL (ET) approach. For the EL approach, 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

[8]Our focus on the term structure of the cost of fluctuations might make it seem natural to use dividend strips as test 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 states.

of consumption fluctuations over a three-year period is more than seven times higher than the cost over one 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%). Similar results are obtained with the ET estimator.

**Table 1: Cumulative Cost of Consumption Fluctuations**

| | 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: Nondurables & Services Consumption | | | | | | | | | | |
| I-SDF (Mkt, EL) | 1.53 | 5.15 | 11.75 | 14.28 | 14.44 | .556 | 1.48 | 3.39 | 3.90 | 3.57 |
| I-SDF (Mkt, ET) | 1.47 | 4.49 | 10.32 | 11.94 | 11.94 | .574 | 1.37 | 3.04 | 3.35 | 3.07 |
| I-SDF (FF6, EL) | 1.29 | 3.52 | 6.65 | 10.63 | 11.20 | .462 | 1.03 | 2.07 | 3.03 | 2.90 |
| I-SDF (FF6, ET) | 1.28 | 3.50 | 6.17 | 9.93 | 10.64 | .411 | .926 | 1.87 | 2.79 | 2.69 |
| CRRA Kernel | .933 | 2.08 | 3.73 | 4.87 | 5.03 | .457 | .854 | 1.32 | 1.52 | 1.40 |
| Lucas | .751 | 1.09 | 1.40 | 1.68 | 1.94 | - | - | - | - | - |
| Panel B: Total Consumption | | | | | | | | | | |
| I-SDF (Mkt, EL) | 2.15 | 6.77 | 16.13 | 19.65 | 19.73 | .896 | 2.09 | 4.85 | 5.55 | 5.12 |
| I-SDF (Mkt, ET) | 2.07 | 6.13 | 14.68 | 17.32 | 17.31 | .904 | 1.98 | 4.47 | 4.98 | 4.59 |
| I-SDF (FF6, EL) | 1.88 | 4.89 | 9.46 | 15.00 | 15.57 | .770 | 1.60 | 3.05 | 4.35 | 4.14 |
| I-SDF (FF6, ET) | 1.83 | 4.81 | 8.72 | 14.09 | 14.75 | .691 | 1.43 | 2.73 | 4.01 | 3.83 |
| CRRA Kernel | 1.42 | 3.08 | 5.80 | 7.63 | 7.77 | .761 | 1.32 | 2.08 | 2.40 | 2.21 |
| Lucas | 1.15 | 1.68 | 2.16 | 2.61 | 3.03 | - | - | - | - | - |

The table reports the (cumulative) costs of *all* aggregate consumption fluctuations (Columns 2-6) and the costs of business cycle fluctuations in consumption (Columns 7-11), 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 (Rows 1–2), the I-SDF recovered from the six size and book-to-market-equity sorted portfolios of Fama and French (Rows 3–4), the kernel implied by power utility preferences with a constant CRRA (Row 5), and Lucas' original specification that involves power utility preferences and i.i.d. lognormal aggregate consumption growth dynamics (Row 6). The sample is annual covering the period 1929-2015.

Row 5 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 6 implies even smaller costs of 1.1%, 1.4%, 1.7%, and 1.9% at two-, three-, four-, and five-year horizons, respectively. Finally, very similar results are obtained when the 6 FF portfolios are used to recover the I-SDF. Rows 3–4 show that the costs of fluctuations for two-, three-, four-, and five-year periods are substantially higher for the I-SDF compared to the CRRA kernel – for instance, 3.5% versus 2.1% for two years, 6.7% versus 3.7% for three years, 10.6% versus 4.9% for four years, and 11.2% versus 5.0% for five years using the EL estimator (Row 3). And the costs are even higher when compared to Lucas' specification.

Figure 2, Panel A plots the term structure of the cost of fluctuations over one- to ten-year horizons.[9] The black solid line corresponds to the costs obtained with the I-SDF recovered with the market portfolio as the sole test asset. The black-dashed line, on the other hand, denotes the costs implied by the I-SDF extracted from the 6 FF portfolios. The green and blue lines denote the costs estimated with the CRRA kernel and Lucas' specification, respectively. The figure highlights the higher costs implied by the I-SDF relative to those obtained with Lucas' specification.

FIGURE 2. MARGINAL COST OF ALL CONSUMPTION FLUCTUATIONS, 1929-2015

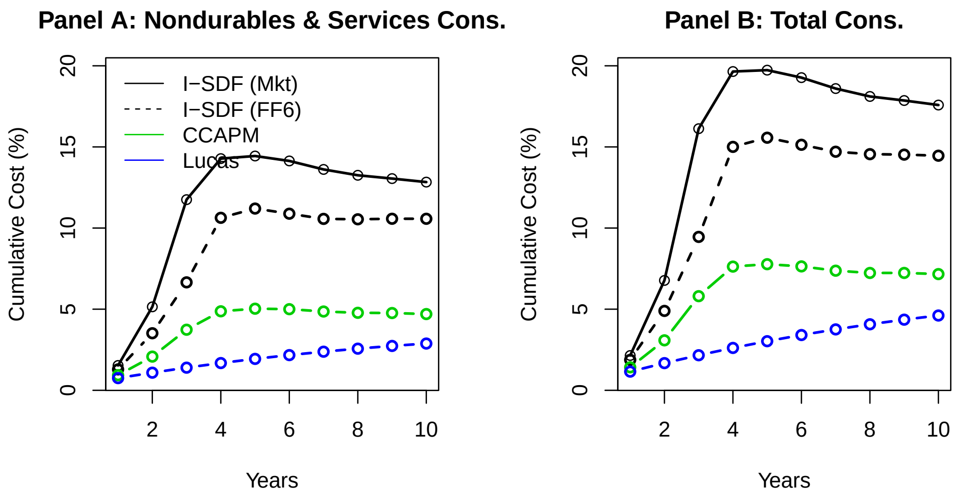


[9]We present the term structure of the cost of fluctuations from one to ten years in Figure 2 but only from one to five years in Table 3 to avoid cluttering of numbers in the table.

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 using the excess return on the market portfolio as the sole test asset (black line) or the excess returns on the 6 FF portfolios as test assets (black dashed line) with the EL approach, 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).

The last five columns of Table 1, Panel A present our estimates of the cost of business cycle fluctuations in consumption. Row 1 shows that, using the I-SDF extracted from the market portfolio alone with the EL approach, the cost of business cycle fluctuations in consumption over a one-year time horizon is estimated to be 0.6%. The costs of business cycle fluctuations over two, three, four, and five year horizons increase to 1.5%, 3.4%, 3.9%, and 3.7%, respectively. Similar results are obtained in Row 2 with the ET estimator – the costs of business cycle fluctuations over two, three, four, and five year horizons increase to 1.4%, 3.0%, 3.4%, and 3.1%, respectively. The results also remain similar in Rows 3–4 when the six size and book-to-market-equity sorted portfolios are used in the recovery of the I-SDF with the EL and ET approaches, respectively – the costs of business cycle fluctuations increase from 0.5% at the one-year horizon to 2.9% for a five-year time period using the EL and from 0.4% to 2.7% using the ET.

Row 5 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 in the case of the former. 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 1–4.

An important point to note is that while the estimates of the costs of business cycle fluctuations are smaller than the corresponding costs of all consumption uncertainty, the former, nonetheless, represents a substantial fraction of the latter. For instance, Panel A, Row 1 shows that, when the market portfolio is used in the extraction of the I-SDF with the EL approach, the cost of business cycle fluctuations constitutes 36.3% of the cost of all consumption fluctuations over a one-year horizon. 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 consumption fluctuations over these time horizons. Similarly, when the 6 FF portfolios are used for the recovery of the I-SDF in Row 3, 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. Similar results are obtained in Rows 2 and 4 that use the ET estimator.

Figure 3, Panel A plots the term structure of the cost of all consumption fluctuations (solid line) and business cycle fluctuations in consumption (dashed line) over 1-10 years. The black lines present the estimates obtained when the market portfolio alone is used as the test asset to recover the I-SDF. The red lines, on the other hand, are based on the estimates obtained when the I-SDF is recovered from the 6 FF portfolios. 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 with increase in the time horizon, thereby suggesting well-defined asymptotics of our approach.

The results in Table 1, Panel A were obtained using personal consumption expenditures on nondurables and services as the measure of consumption. Panel B, that uses the total consumption expenditures (including durables) as the measure of consumption, produces results similar to those in Panel A. Note that, not surprisingly, the costs of fluctuations are bigger with total consumption compared to those obtained with nondurables and services consumption (see also Figure 2, Panel B and Figure 3, Panel B).

Overall, two salient conclusions emerge from the results of this section. First, that economic agents perceive the cost of aggregate economic uncertainty to be quite substantial. For instance, our estimates of the cost of all consumption fluctuations over a five-year horizon vary from 10.6%-19.7%, depending on the measure of aggregate consumption expenditure or the set of assets or the precise definition of relative entropy used to recover the I-SDF. The cost is substantially higher than that originally obtained by Lucas.

Figure 3. Marginal Cost of All versus Business Cycle Consumption Fluctuations, 1929-2015

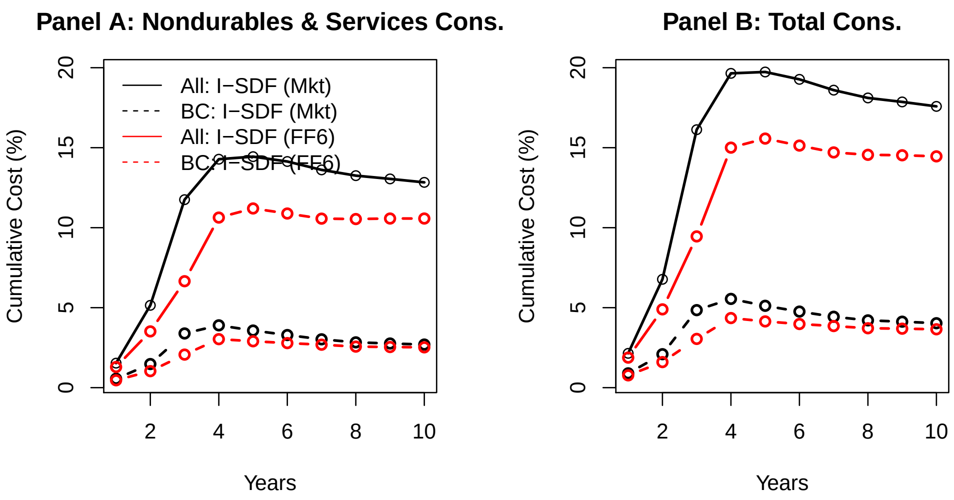


*Notes: The figure plots the term structure of the (cumulative) cost of all aggregate consumption fluctuations (solid line) and business cycle fluctuations in consumption (dashed line), over* 1-10 *years, obtained using the I-SDF. Panel A presents results when consumption refers to the real personal 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 (black lines) and the 6 FF portfolios (red lines) with the EL approach. The sample is annual covering the period* 1929-2015.

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. Our results are in contrast to those 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, varying 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). Therefore, the question naturally arises as to what drives this difference.

We show that the 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.[10] 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 4 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.

Figure 4. Comparison of HP and AJ Filters

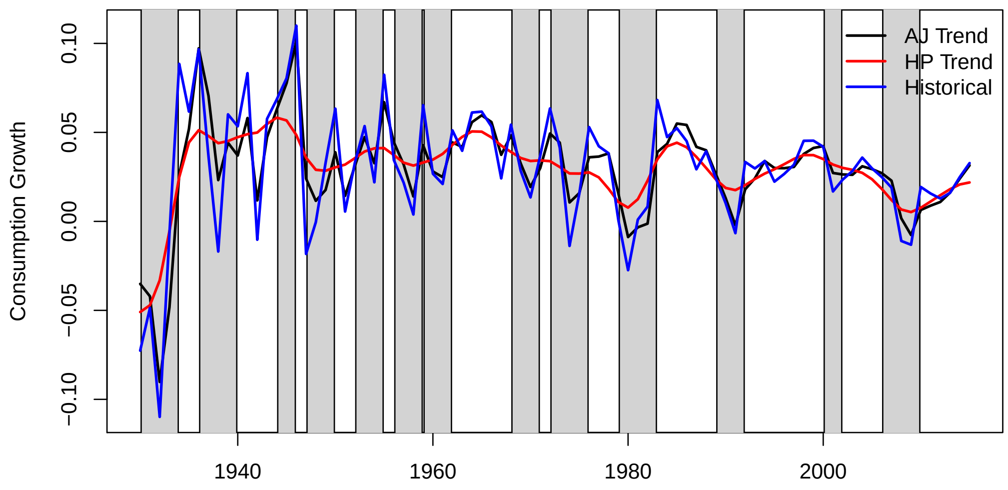


Notes: The figure plots the historical consumption growth (blue line), the trend consumption growth obtained using the HP filter (red line), and the trend consumption growth obtained using the AJ filter

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

(black line). Consumption refers to the total personal consumption expenditures. The sample is annual over 1929–2015.

The figure shows that the HP filter delivers a smoother trend consumption growth relative to the AJ filter. 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. In fact, 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 from the market portfolio with the EL approach, 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 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 James D. 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. Separately, in Section III.B, we present further evidence supporting the high cost of business cycle fluctuations, using an approach that does not involve a smoothing filter.

Finally, we present evidence that our non-linear adjustment to the pricing kernel – note that the $\psi$-component recovered using the EL/ET approaches are highly nonlinear functions of consumption growth and the test asset returns – has certain desirable properties relative to alternative linear (or log-linear) adjustments that have been proposed in the literature and also produce markedly different results compared to the latter. To demonstrate this, we first consider an SDF that is log-linear in the aggregate consumption growth rate and the market return (note that this specification corresponds to the SDF implied by Epstein and Zin (1989) recursive preferences with the stock market return used as a proxy for the return on the total wealth portfolio):

$$M_{t+1} = \exp\{(C_{t+1}/C_t)^{-\gamma_1} \, R_{M,t+1}^{\gamma_2}\}. \tag{30}$$

To make the results comparable with those obtained using the EL/ET approaches, we set $\gamma_1 = \gamma = 10$ and estimate $\gamma_2$ using the GMM approach to match the Euler equation for the excess market return. Substantially smaller estimates of the cost of fluctuations are obtained relative to those with EL and ET. Specifically, using total consumption as the measure of the aggregate consumption expenditures, the cost of all fluctuations at the one- to five-year horizons are 1.7%, 4.6%, 9.80%, 11.7%, and 11.7%, respectively. In contrast, the corresponding cost estimates obtained using the EL (Table 1, Panel B, Row 1) are almost double at each horizon at 2.1%, 6.8%, 16.1%, 19.6%, and 19.7%, respectively.

Second, we present results for the linear adjustment to the SDF proposed in Lars P. Hansen and Ravi Jagannathan (1997) to enable it to successfully price assets. When the excess market return is the sole test asset, this implies:

(31)
$$M_{t+1} = (C_{t+1}/C_t)^{-\gamma} - \frac{\sum_{t=1}^{T} \left((C_{t+1}/C_t)^{-\gamma}\right)(R_{M,t+1} - R_{F,t})}{\sum_{t=1}^{T} (R_{M,t+1} - R_{F,t})^2} (R_{M,t+1} - R_{F,t}) .$$

Once again, to make the results comparable with those obtained using the EL/ET approaches, we set $\gamma = 10$. The resulting estimates of the costs are economically implausible and unstable, varying from 1.76% at the one-year horizon to 43.7% at the 5-year horizon to 304.0% at the ten-year horizon. The absence of well-behaved asymptotic behaviour in this case, unlike those obtained with EL/ET, highlights the numerical instability of linear adjustments and, therefore, the desirability of the EL and ET approaches in such settings.

### B. *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 (28). We present the results obtained with the SEL estimator. The alternative SET estimator produces similar results, reported in Appendix A.A6.

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 test asset. Note that the SEL procedure requires the specification of the investors' conditioning set. In our baseline results, we use an exponentially-weighted moving average of lagged consumption growth as the conditioning variable – a natural candidate in the context of consumption-based models.

Figure 5 presents the time series of the cost. 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 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%.[11] 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.[12]

Figure 5. Time-Varying Cost of One-Period Consumption Fluctuations, 1929-2015

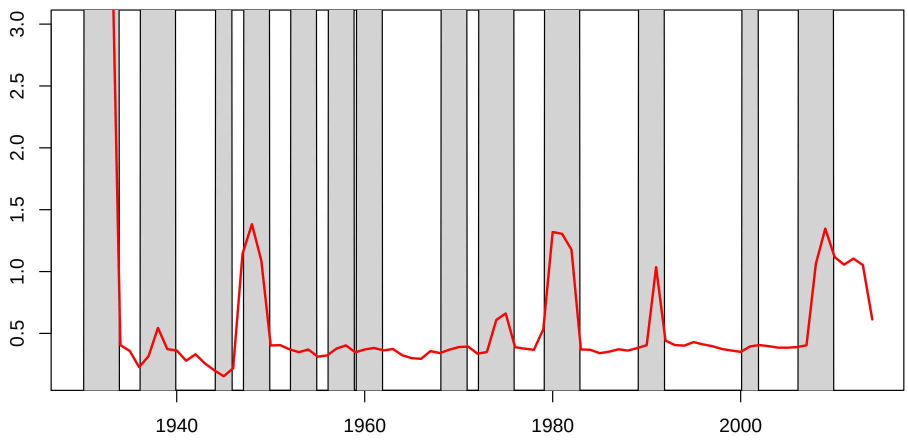


[11]The y-axis in Figure 5 is capped at 3% despite the maximum cost being as high as 8% so as to render the remaining (smaller) cost estimates readable.

[12]Note that, because of Jensen's inequality, the average of the time series of the cost does not exactly equal the average cost of all one-year fluctuations computed in the previous section.

*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.

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 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 6 presents the time series of the annualized cost (red line) using quarterly data over the postwar period 1947:$Q1$–2015:$Q4$.

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, similar to the average cost of 0.86% obtained using the postwar sample alone.

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 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 broad cross section of over a hundred macro variables. Specifically, we obtain panel data on 106 macroeconomic variables from Sydney Ludvigson's web site, based on the Global Insights Basic Economics Database and The Conference Board's Indicators Database. The 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 a principal component from the cross section of transformed variables.[13] The time series of the cost is presented in Figure 6 (black line). Since data on the broad cross section are only available from the mid-sixties, the cost estimates start from 1966:$Q1$. The figure shows that the recovered time series of the cost seems quite robust to the choice of the conditioning set.

Figure 6. Time-Varying Cost: Robustness to Sample Period and Conditioning Set

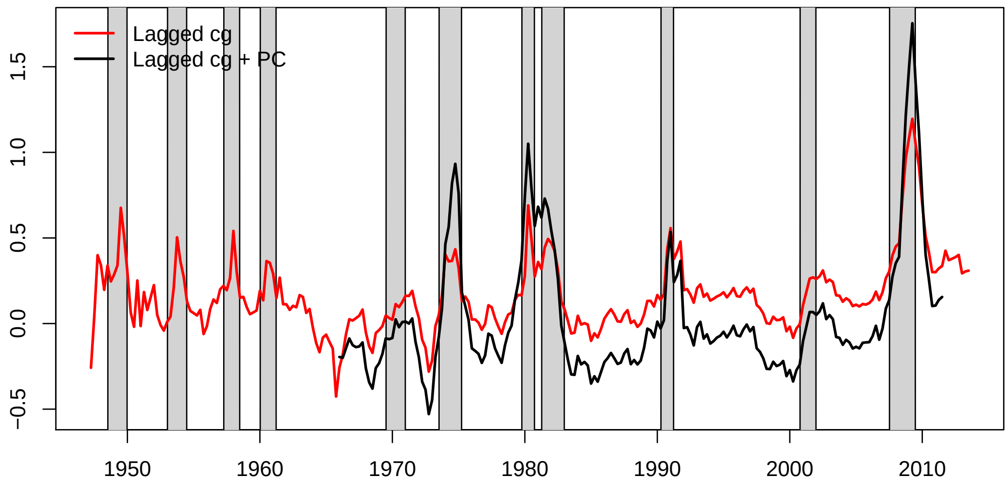


*Notes: The figure plots the time series of the cost of one-period consumption uncertainty. The cost is estimated using the SEL approach, using nondurables and services consumption as the measure of the consumption expenditures and the excess return on the market portfolio as the sole test asset. The 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:$Q1$–2015:$Q4$ *(red*

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

*line) or* 1966*:Q*1–2015*:Q*4 *(black line).*

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.A.

## IV. Interpretation and Sensitivity of Results to Alternative Choices of $\gamma$

Note that our methodology takes the stance that the true (unknown) SDF has the multiplicative form in Equation ((1)), with a component that is a function of consumption growth as in the power utility model and a second (unknown) $\psi$-component. In this section, we first present evidence supporting the presence of the consumption growth component in the SDF, thereby justifying our choice. Specifically, we show that the recovered $\psi$-component of the I-SDF has similar properties for all values of $\gamma$ in the economically plausible range, including $\gamma = 0$ in which case the SDF is not constrained to depend on consumption growth at all apriori. We use nondurables and services consumption as the measure of the aggregate consumption expenditures, the excess return on the market as the sole test asset, and the EL approach to recover the I-SDF.

If we set $\gamma = 0$, i.e. when no structure is imposed on the pricing kernel, the correlation between the recovered $\psi$ and consumption growth is $-0.35$. Thus, even without including a consumption growth component in the pricing kernel, the recovered I-SDF is strongly negatively correlated with consumption growth. As $\gamma$ is varied from $[1, 20]$ – the economically plausible range that contains the calibrated values of this parameter from most asset pricing models – Figure 7, Panel A shows that the correlation between the recovered $\psi$ and consumption growth remains very similar both in magnitude and sign. This suggests that the additional $\psi$-component of the SDF helps to magnify the increased marginal utility during bad states characterized by low realizations of consumption growth and, therefore, implies a higher cost of consumption fluctuations relative to that implied by the power utility kernel that is a function of consumption growth alone.

Highlighting further the information content of the consumption growth component, the figure panel also shows that, for values of $\gamma$ higher than its EL point estimate of 26.2 (red-dashed line), the correlation between consumption growth and the recovered $\psi$ switches sign, becoming strongly positive. In other words, for these values of $\gamma$, the $\psi$-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. Figure 7, Panel B plots the time series of the recovered $\psi$ for two alternative values of $\gamma$ – a value of 10, our baseline value, which is below its EL point estimate (black solid line) and a value of 43 that is symmetrically above the point estimate (red-dashed line). 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.

Overall, the results suggest that the consumption growth component of the SDF, that is relied upon in a large class of consumption-based macro models, is not merely a misspecification but rather an important ingredient of the true underlying SDF. This helps further justify the inclusion of this component in our multiplicative specification of the SDF (see Equation (1)).

FIGURE 7. CORRELATION $\left(\frac{C_t}{C_{t-1}}, \psi_t\right)$ AND TIME SERIES OF $\psi$ FOR ALTERNATIVE VALUES OF $\gamma$

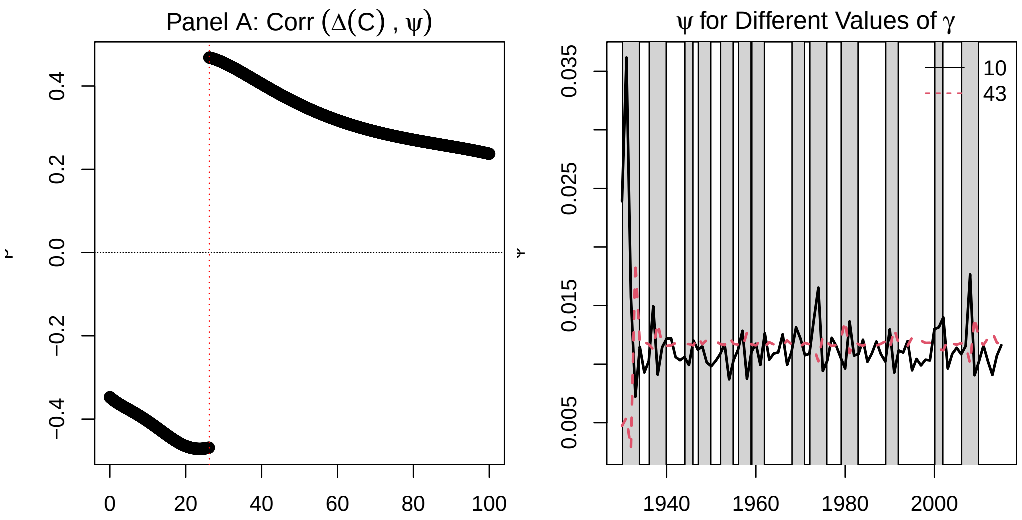


*Notes: Panel A plots the correlation between the two multiplicative components of the pricing kernel, namely the recovered $\psi$ and consumption growth, for alternative values of the SDF parameter $\gamma$. Panel*

*B plots the time series of the recovered $\psi$ for two different values of $\gamma$, 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 $\psi$ component using the EL approach. The sample is annual covering the period* 1929-2015.

Next we assess the sensitivity of the magnitudes of the cost of fluctuations to alternative choices of $\gamma$. 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.

Table 2 presents the one- to five-year term structure of the cost of all (Columns 2-6) and business cycle (Columns 7-11) fluctuations in consumption, for several different choices of $\gamma$. To facilitate comparison, in Row 1, we repeat the cost estimates for the baseline value of $\gamma = 10$ from Table 1, Panel $A$, Row 1. This corresponds to the value of $\gamma$ typically used in the long run risks paradigm.

**Table 2: Cumulative Cost of Consumption Fluctuations for Alternative $\gamma$**

| | 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: Nondurables & Services Consumption | | | | | | | | | | |
| $\gamma$= 10.0 | 1.53 | 5.15 | 11.75 | 14.28 | 14.44 | .56 | 1.48 | 3.39 | 3.90 | 3.57 |
| $\gamma$= 2.0 | 0.64 | 2.00 | 3.36 | 3.59 | 3.52 | .18 | 0.53 | 1.02 | 1.08 | 1.00 |
| $\gamma$= 3.5 | 0.79 | 2.54 | 4.62 | 5.14 | 5.13 | .22 | .68 | 1.36 | 1.48 | 1.36 |
| $\gamma$= 16.0 | 2.42 | 7.54 | 17.7 | 22.0 | 22.0 | 1.04 | 2.41 | 5.17 | 5.97 | 5.48 |

The table reports the (cumulative) costs of *all* aggregate consumption fluctuations (Columns 2-6) and the costs of business cycle fluctuations in consumption (Columns 7-11), over one-to five-year horizons. Rows 1-4 present the results for alternative values of the SDF parameter $\gamma$. 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.

Row 2 presents the cost estimates for $\gamma = 2$, the value commonly used in the external habit paradigm. The magnitudes of the cost estimates are smaller for

$\gamma = 2$ compared to those obtained for $\gamma = 10$. Specifically, the cost of all one- to five-year fluctuations varies from 1.5%–14.4% for the latter value of $\gamma$ compared to 0.6%–3.5% for the former value. Similar patterns obtain for the cost of business cycle fluctuations – the cost of one- to five-year fluctuations varies from 0.6%–3.6% for $\gamma = 10$ compared to 0.2%–1.0% for $\gamma = 2$. Note, however, that while the magnitudes of the costs reduce when moving to a lower value of $\gamma$, the cost of business cycles continues to constitute a substantial fraction (a quarter to a third) of the cost of all fluctuations – 28.1%, 26.5%, 30.4%, 30.1%, and 28.4% at the one- to five-year horizons, respectively.

Very similar conclusions obtain in Row 3, that presents the cost estimates for $\gamma = 3.5$, the value commonly used in the rare disasters paradigm. The only difference with respect to Row 2 is that the cost estimates are higher in Row 3 compared to those obtained in Row 2 owing to the higher value of $\gamma$.

We next present the results for $\gamma = 16$, the value assumed in models with complementarities in consumption. This value is higher than our baseline value of 10 and Row 4 shows that the estimates of the cost are much higher than those obtained in the baseline case. In particular, the cost of one- to five-year fluctuations varies from 2.4%–22.0% in the case of all fluctuations and from 1.0%–5.5% for business cycle fluctuations. Like with all the other values of $\gamma$, the costs of business cycles constitute between a quarter to a third of the costs of all fluctuations.

Overall, the results suggest that the precise magnitude of the cost of fluctuations is somewhat sensitive to the assumed value of the utility curvature parameter $\gamma$, with larger values implying larger costs. However, the findings that business cycle fluctuations constitute a substantial proportion – between a quarter to a third – of the cost of all consumption fluctuations is robust to the value of this parameter.

## V. 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. Moreover, 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. Finally, 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.

Note that, our results indicate that the cost of business-cycle fluctuations may be much higher than previously thought. 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, Caio, and René Garcia.** 2012. "Assessing misspecified asset pricing models with empirical likelihood estimators." *Journal of Econometrics*, 170(2): 519–537.

**Almeida, Caio, and René Garcia.** 2016. "Economic Implications of Nonlinear Pricing Kernels." *Management Science*, 63(10): 3147–3529.

**Almeida, Caio, Kym Ardison, and René Garcia.** 2019. "Nonparametric Assessment of Hedge Fund Performance." *Journal of Econometrics (forthcoming).*

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

**Atkeson, Andrew, and Christopher Phelan.** 1994. "Reconsidering the costs of business cycles with incomplete markets." *NBER Macroeconomics Annual*, 187–207.

**Backus, David, Mikhail Chernov, and Stanley E. Zin.** 2013. "Sources of Entropy in Representative Agent Models." *Journal of Finance.* forthcoming.

**Bai, Hang, and Lu Zhang.** 2020. "Searching for the Equity Premium." NBER Working Paper 28001.

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

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

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

**Barro, Robert J.** 2006. "Rare Disasters and Asset Markets in the Twentieth Century." *Quarterly Journal of Economics*, 121(3): 823–866.

**Barro, Robert J.** 2009. "Rare Disasters, Asset Prices, and Welfare Costs." *American Economic Review*, 99(1): 243–264.

**Basak, Suleyman, and Hongjun Yan.** 2010. "Equilibrium Asset Prices and Investor Behaviour in the Presence of Money Illusion." *The Review of Economic Studies*, 77(3): pp. 914–936.

**Beaudry, Paul, and Carmen Pages.** 2001. "The cost of business cycles and the stabilization value of unemployment insurance." *European Economic Review*, 45: 1545–1572.

**Blanchard, Olivier J.** 2019. "Public Debt and Low Interest Rates." AFA Lecture.

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

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

**Campbell, John Y., and John 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, Xiaohong, Lars P. Hansen, and Peter G. Hansen.** 2020. "Robust Identification of Investor Beliefs."

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

**Constantinides, George M., and Anisha Ghosh.** 2017. "Asset Pricing Tests with Long Run Risks in Consumption Growth." *Journal of Finance*, 72(1): 415–460.

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

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

**Davis, Jessie, and Gill Segal.** 2020. "Trendy Business Cycles and Asset Prices." Working Paper.

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

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

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

**Ghosh, Anisha, and Guillaume Roussellet.** 2019. "Identifying Beliefs from Asset Prices." Working Paper, McGill University.

**Ghosh, Anisha, Christian Julliard, and Alex 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.

**Ghosh, Anisha, Christian Julliard, and Alex Taylor.** 2022. "An Information-Based One-Factor Asset Pricing Model." Working paper.

**Ghosh, Anisha, Taisuke Otsu, and Guillaume Roussellet.** 2020. "Properties of Subjective Beliefs Estimates." Working Paper, London School of Economics and McGill University.

**Hamilton, James D.** 2018. "Why You Should Never Use the Hodrick-Prescott Filter." *The Review of Economics and Statistics*, 100.

**Hansen, Lars P., and Ravi Jagannathan.** 1997. "Asssessing Specification Errors in Stochastic Discount Factor Models." *The Journal of Finance*, 52: 557–590.

**Hansen, Lars Peter.** 2014. "Nobel Lecture: Uncertainty Outside and Inside Economic Models." *Journal of Political Economy*, 122(5): 945 – 987.

**Hansen, Lars Peter, and Kenneth J. Singleton.** 1983. "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *Journal of Political Economy*, 91: 249–68.

**Hansen, Lars Peter, and Ravi Jagannathan.** 1991. "Implications of Security Market Data for Models of Dynamic Economies." *Journal of Political Economy*, 99(2): 225–62.

**Hansen, Lars Peter, and Thomas J. Sargent.** 2010. "Fragile Beliefs and the Price of Uncertainty." *Quantitative Economics*, 1(1): 129–162.

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

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

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

**Kitamura, Yuichi.** 2003. "A Likelihood-Based Approach to the Analysis of a Class of Nested and Non-Nested Models." Working Paper.

**Kitamura, Yuichi.** 2006. "Empirical Likelihood Methods in Econometrics: Theory and Practice." Cowles Foundation, Yale University Cowles Foundation Discussion Papers 1569.

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

**Komunjer, Ivana, and Giuseppe Ragusa.** 2016. "Existence and Characterization of Conditional Density Projections." *Econometric Theory*, 32: 947–987.

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

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

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

**Liew, Jimmy, and Maria 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, Robert E.** 1987. *Models of Business Cycles.* Oxford:Blackwell.

**Lustig, Hanno N., and Stijn G. Van Nieuwerburgh.** 2005. "Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective."

*The Journal of Finance*, 60(3): pp. 1167–1219.

**Lustig, Hanno N., Stijn G. Van Nieuwerburgh, and Adrien Verdelhan.** 2013. "The Wealth-Consumption Ratio." *The Review of Asset Pricing Studies*, 3(1): 38–94.

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

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

**Menzly, Lior, Tano Santos, and Pietro Veronesi.** 2004. "Understanding Predictability." *Journal of Political Economy*, 112(1): 1–47.

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

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

**Owen, Art B.** 2001. *Empirical Likelihood.* Chapman and Hall.

**Parker, Jonathan A., and Christian Julliard.** 2005. "Consumption Risk and the Cross-Section of Expected Returns." *Journal of Political Economy*, 113(1).

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

**Piazzesi, Monika, and Martin Schneider.** 2007. "Equilibrium Yield Curves." Vol. 21. Cambridge: MIT Press.

**Piazzesi, Monika, Martin Schneider, and Selale 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, Morten O., and Harald Uhlig.** 2002. "On adjusting the Hodrick-Prescott filter for the frequency of observations." *The Review of Economics and Statistics*, 84(2): 371–375.

**Sandulescu, Mirela, Fabio Trojani, and Andrea Vedolin.** 2018. "Model-Free International Stochastic Discount Factors." Working Paper, Boston University.

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

**Storesletten, Kjetil, Chris I. Telmer, and Amir 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, Michael.** 1995. "A Bayesian approach to diagnosis of asset pricing models." *Journal of Econometrics*, 68(2): 367 – 397.

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

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

**Vuong, Quang H.** 1989. "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses." *Econometrica*, 307–333.

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

## Online Appendix

### *A1. Solution for the ET Estimator*

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

$$\hat{\psi}_t = \frac{e^{\hat{\theta}(\gamma)'\mathbf{r}_t^e(\Delta c_t)^{-\gamma}}}{\frac{1}{T}\sum_{t=1}^{T} e^{\hat{\theta}(\gamma)'\mathbf{r}_t^e(\Delta c_t)^{-\gamma}}} \quad \forall t, \tag{A1}$$

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

$$\hat{\theta}(\gamma) = \arg\min_{\theta} \left[\log\left(\frac{1}{T}\sum_{t=1}^{T} e^{\theta'\mathbf{r}_t^e(\Delta c_t)^{-\gamma}}\right)\right]. \tag{A2}$$

### *A2. 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 $\psi$-component with the EL/ET approach. Note that, in the former scenario, the true $\psi_t \equiv 1$ whereas in the latter scenario the $\psi$-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: $\log(\Delta C_t) \overset{\mathbb{P}}{\sim} \mathcal{N}\left(\mu_c, \sigma_c^2\right)$. The following Euler equation holds in equilibrium:

$$0 = \mathbb{E}^{\mathbb{P}}\left[(\Delta C_{t+1})^{-\gamma}\left(R_{m,t+1} - R_{f,t+1}\right)\right], \tag{A3}$$

where $R_{m,t}$ and $R_{f,t}$ denote the market return and the risk free rate, respectively, at time $t$. Note that, Equation (A3) may be rewritten as

$$0 = \mathbb{E}^{\mathbb{P}}\left[(\Delta C_{t+1})^{-\gamma}\,\psi_{t+1}\,(R_{m,t+1} - R_{f,t})\right], \tag{A4}$$

where $\psi_t \equiv 1$.

This example economy fits into the framework described in Section II. 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 $\psi$-component of the kernel:

(A5)

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

(A6)

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

Using the recovered $\psi$, the term structure of the cost of fluctuations may be computed as described in Section II.B.

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

$$\nu = \frac{\exp\left[\log(\delta) + (1-\gamma)\mu_c + \dfrac{(1-\gamma)^2\sigma_c^2}{2}\right]}{1 - \exp\left[\log(\delta) + (1-\gamma)\mu_c + \dfrac{(1-\gamma)^2\sigma_c^2}{2}\right]}, \tag{A7}$$

and the equilibrium risk free rate is also constant at:

$$R_f = \frac{1}{\exp\left(\log(\delta) - \gamma\mu_c + \dfrac{\gamma^2\sigma_c^2}{2}\right)}. \tag{A8}$$

To perform our simulation exercise, we calibrate $\mu_c$ and $\sigma_c$ 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 $\nu$ is defined in Equation (A7). The risk free rate is simply a constant, given by Equation (A8).

Using the above time series, we recover $\{\psi_t\}_{t=1}^{T}$ using the EL and ET approaches in Equations (A5) and (A6), respectively. Armed with the $\psi$-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.1. 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.1: 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 | 1.66 | 1.66 | 2.11 | 2.12 | 2.51 | 2.53 | 2.88 | 2.89 |
| | [1.15,1.16] | [1.15,1.16] | [1.42,1.66] | [1.43,1.89] | [1.66,2.57] | [1.68,2.58] | [1.84,3.23] | [1.87,3.24] | [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 $\psi$-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 $\gamma$ denotes the CRRA, $\rho$ 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 $\psi$-component of the pricing kernel using the EL and ET approaches. For the EL approach, for instance, this involves:

(A9)
$$\left\{\widehat{\psi}_t^{EL}\right\}_{t=1}^T = \arg\max_{\{\psi_t\}_{t=1}^T} \sum_{t=1}^T \log(\psi_t) \quad \text{s.t.} \quad \frac{1}{T}\sum_{t=1}^T \left(\frac{c_{t+1}}{c_t}\right)^{-\frac{\eta}{\rho}} r_{c,t}^{\eta-1}\psi_{t+1}(r_{m,t+1}-r_{f,t+1}) = 0.$$

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 $\psi$-component:
(A10)

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

As with Economy I, we recover $\psi$ 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.2. 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 George M. Constantinides and Anisha 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.

**Table A.2: 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<br>[−.01,.26] [−.01,.26] | .30 / .30<br>[.07,.61] [.06,.60] | .54 / .54<br>[.16,1.09] [.15,1.09] | .82 / .81<br>[.26,1.67] [.24,1.66] | 1.14 / 1.12<br>[.35,2.34] [.33,2.32] |
| | | | Panel C: Simulated Values, Misspecified SDF | | |
| T=87 | .39 / .42<br>[.23,.58] [.25,.65] | .75 / .80<br>[.42,1.21] [.44,1.29] | 1.19 / 1.26<br>[.62,2.02] [.65,2.13] | 1.71 / 1.80<br>[.83,3.03] [.87,3.16] | 2.28 / 2.39<br>[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 $\psi$-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 $\psi$-component of the SDF using the EL/ET approaches.

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

### *A3. 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.

### A4. 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.A. For all the robustness tests, consumption refers to the total personal consumption expenditure.[14]

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:

(A11)

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

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 (A11) 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 (A11) and the EL and ET estimators defined in Equations (14) and (15), respectively, is that while the latter two minimize the 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 (A11) is obtained as:

$$\hat{q}_t = \frac{e^{\hat{\theta}(\gamma)'\mathbf{r}_t^e}\,(\Delta c_t)^{-\gamma}}{\frac{1}{T}\sum_{t=1}^{T} e^{\hat{\theta}(\gamma)'\mathbf{r}_t^e}\,(\Delta c_t)^{-\gamma}} \quad \forall t, \tag{A12}$$

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

$$\hat{\theta}(\gamma) = \arg\min_{\theta} \left[\log\left(\frac{1}{T}\sum_{t=1}^{T} e^{\theta'\mathbf{r}_t^e}\,(\Delta c_t)^{-\gamma}\right)\right]. \tag{A13}$$

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.

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

**Table A.3: 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-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-SDF*$^{Alt}$ | 1.76 | 4.46 | 8.96 | 14.12 | 14.67 | .764 | 1.61 | 3.02 | 4.20 | 4.00 |
| 1890-2015 | - | - | - | - | - | - | - | - | - | - |

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). The costs are calculated using the I-SDF extracted the risk-neutral measure recovered by minimizing the distance from the CRRA model-implied risk-neutral measure while satisfying the pricing restrictions (Row 1) and the I-SDF extracted with the EL approach over the longer 1890-2015 sample (Row 2). 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. The sample is annual covering the period 1929-2015, except for Row 3 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.[15]

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.

### *A5. Out-of-Sample Performance of the I-SDF*

This section reproduces a table from Ghosh, Julliard and Taylor (2022) on the out-of-sample performance of the I-SDF vis a vis other popular factor models.

[15]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.

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 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 (17)–(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 multi factor models (e.g., the Fama-French three- and five-factor models) on an equal footing with the one-factor I-SDF, we present the empirical performance of the multi factor 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:

$$M_t^{FF3} = c_0 + \sum_{j=1}^{3} c_j f_{j,t}, \tag{A14}$$

where $\{f_{j,t}\}_{j=1}^3 = \{R_{M,t}, R_{SMB,t}, R_{HML,t}\}$ and the coefficients $c_j$, $j = 0, 1, 2, 3$, 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:

$$0 = E[(R_{i,t} - R_{f,t})(c_0 + \sum_{j=1}^{3} c_j f_{j,t})]. \tag{A15}$$

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 $M_t^{FF5} = c_0 + \sum_{j=1}^{5} c_j f_{j,t}$.

In practice, we set the size of the rolling window $\overline{T} = 30$ years, i.e. 360 months, and $s = 12$. The results, presented in Table A.4, show the superior performance of the I-SDF relative to the CAPM, the FF3, and the FF5 models – the estimated intercept is smaller and the OLS adjusted $R^2$ larger for the I-SDF relative to the other models for all three sets of test assets.

**Table A.4: Out-of-Sample Performance of I-SDF**

| Row | *const.* | $\lambda_{sdf}$ | $\overline{R}^2_{OLS}$ (%) | $\overline{R}^2_{GLS}$ (%) | $T^2$ | $q$ |
|---|---|---|---|---|---|---|
| Panel A: 25 FF Portfolios | | | | | | |
| I-SDF | .003 (4.83) | −1.37 (−4.95) | 49.4 | 27.6 | 45.0 (.004) | .119 |
| CAPM | .012 (3.79) | −.004 (−1.41) | 3.95 | 30.3 | 74.5 (.000) | .116 |
| FF3 | .008 (15.6) | .743 (2.46) | 17.4 | 33.9 | 53.7 (.000) | .109 |
| FF5 | .009 (9.16) | −.005 (−1.38) | 1.20 | 36.6 | 54.2 (.000) | .210 |
| Panel B: 55 Test Assets | | | | | | |
| I-SDF | .004 (21.3) | −.64 (−8.88) | 59.0 | 39.2 | 74.4 (.028) | .133 |
| CAPM | .007 (3.56) | −.001 (−.44) | −1.51 | 5.88 | 132.5 (.000) | .205 |
| FF3 | .006 (18.6) | 1.46 (0.57) | −1.26 | 6.40 | 130.7 (.000) | .203 |
| FF5 | .007 (46.8) | −.479 (−4.78) | 28.8 | 24.0 | 63.7 (.149) | .229 |
| Panel C: 25 FF + 30 Industry + 10 Momentum | | | | | | |
| I-SDF | .005 (17.6) | −.66 (−7.06) | 43.3 | 27.4 | 151.1 (.000) | .272 |
| CAPM | .009 (5.49) | −.002 (−1.40) | 1.50 | 19.6 | 194.4 (.000) | .301 |
| FF3 | .005 (7.64) | −.718 (−2.30) | 6.28 | 25.7 | 167.3 (.000) | .283 |
| FF5 | .007 (30.3) | −4.83 (−1.21) | .70 | 23.9 | 156.6 (.000) | .544 |

Cross-sectional regressions of average excess returns of different sets of test assets on the estimated factor loadings for different asset pricing models. Panel A presents the results when the test assets consist of the 25 size and BM sorted portfolios of FF. Panel B presents results when the test assets consist of 10 size-sorted, 10 BM-sorted, 10 momentum-sorted, 10 short term reversal sorted, 10 long term reversal sorted, and 5 industry-sorted portfolios. Panel C presents results when the test assets consist of the 25 FF, 30 Industry, and 10 Momentum-sorted portfolios. In each panel, the first row presents the results when the factor is the I-SDF. This I-SDF is constructed from 15 portfolios – consisting of the smallest and largest deciles of size, BM, momentum, short term reversal, and long term reversal sorted portfolios and 5 industry portfolios – using a relative entropy minimizing procedure, in a rolling out-of-sample fashion starting in 1963:07. Rows 2–4 present the results for the SDFs implied by the CAPM, the FF3, and the FF5 models, respectively. For each model, the table presents the intercept and slope, along with $t$-statistics in parentheses. It also presents the OLS adjusted $R^2$ and the GLS adjusted $R^2$. The last two columns present, respectively, 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.

### *A6. Time-Varying Cost of One-Period Consumption Fluctuations Using SET*

The Smoothed Exponential Tilting (SET) estimator is defined as:
$\forall i \in \{1, \ldots, T\}, \quad \forall \gamma \in \Theta,$

(A16)

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

Figure A1 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%.

Figure A1. Time-Varying Cost of One-Period Consumption Fluctuations, 1929-2015

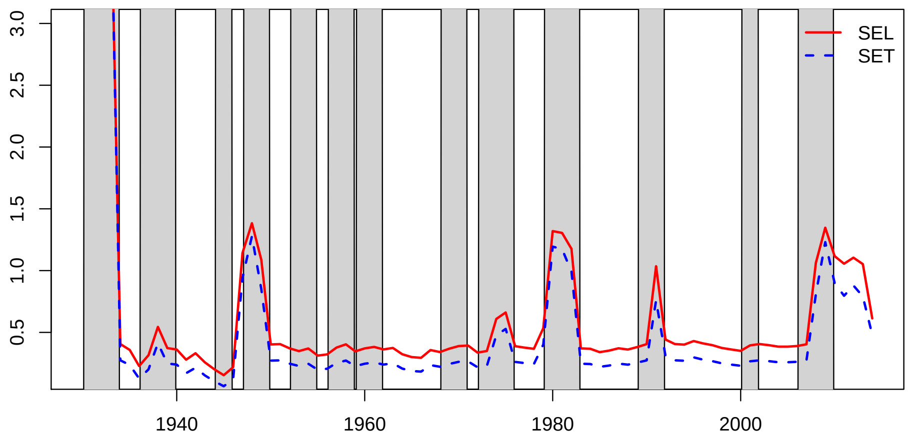


*Notes: The figure plots the time series of the cost of one-period consumption uncertainty. The cost is estimated using the SEL (solid red line) and SET (blue-dashed line) approaches, using nondurables and services consumption as the measure of the consumption expenditures, the excess return on the market 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.