

IS IGNORANCE BLISS? THE COST OF BUSINESS-CYCLE UNCERTAINTY

FRODE BREVIK

VU University Amsterdam

STEFANO D'ADDONA

University of Rome 3

We investigate the cost of business-cycle uncertainty (lack of firm knowledge about the prevailing state of the economy) in a setup where the economy switches between booms and recessions at random intervals. Calibrating an exchange economy model to match the properties of the postwar U.S. data, we find that giving consumers additional information beyond that already contained in the endowment growth rates yields only moderate gains. In a second stage, we investigate the effect of nonperfect information processing in this setting. Surprisingly, we find that opting for slow learning might yield large utility gains, especially for consumers with a strong preference for early resolution of uncertainty.

Keywords: Business-Cycle Uncertainty, Epstein–Zin Preferences, Regime Switching, Learning

1. INTRODUCTION

News announcements that can be used to gauge the state of the economy are amongst the most keenly surveyed by financial market participants, and it is well known that such announcements move prices in both fixed-income and equity markets [Andersen et al. (2007); Kliesen and Schmid (2006)]. This is not restricted to forward-looking measures such as consumer confidence surveys, but also holds for data that pertain to past equilibria, such as employment numbers, inflation, or GDP estimates.

This points to considerable uncertainty, not only on future economic developments, but also on the prevailing state of the economy. It seems reasonable to ask what the aggregate cost of this uncertainty is. Or, what is the aggregate welfare gain from providing research that reduces such uncertainty? As far as we are aware, this question has not yet found a place in economic research.

We thank seminar participants at the University of St. Gallen and at the 2008 Annual Meeting of the Society for Non-linear Dynamics and Econometrics for helpful discussions. We especially thank William A. Barnett (the Editor) and two anonymous referees for their comments and suggestions. Address correspondence to: Stefano d'Addona, Department of International Studies, University of Rome 3, Via G. Chiabrera, 199, I-00145 Rome, Italy; e-mail: daddona@uniroma3.it.

The related topic of the cost of business-cycle fluctuations per se has received considerable attention. In an influential contribution, Lucas (1987) measured the cost of the business cycle by computing the equivalent reduction in consumption that a representative consumer would accept to eliminate any deviation from the trend. Such measurements can be used to evaluate the benefits of policies designed to reduce macroeconomic fluctuations. Similarly, the calculations we provide can help evaluate the benefits of generating better information on the state of the economy. In particular, in the benchmark endowment economy representation of the United States we use here, the costs of business cycle uncertainty are about 8 basis points of annual consumption. With an aggregate consumption of US\$10 trillion, reducing aggregate consumption by up to US\$8 billion to fund research would be welfare-improving, provided it removed all uncertainty about the prevailing state of the economy.

Lucas' estimate was based on a linear trend consumption growth rate and a representative consumer with a power utility function. To evaluate the cost of business cycle *uncertainty*, we need to relax both assumptions. Our main question cannot be properly addressed using a power utility function, because the expected utility theory effectively maintains that agents are indifferent to the timing of resolution of risk. Instead, we use the generalized utility function proposed by Epstein and Zin (1989) and Weil (1989), more commonly known as Epstein–Zin preferences. The extension to non-expected utility functions is a natural one and has been widely adopted in the cost of business cycle literature [e.g., in Obstfeld (1994); Pemberton (1996); Dolmas (1998); and Tallarini (2000)]. Also, it cannot be addressed within the standard linear models used in most business-cycle research. Instead, we use a regime-switching model as pioneered by Hamilton (1989).

Our concern is different from those of the rich literature on the *value of information* [see, e.g., Gollier (2001, Chs. 24–26)]. This literature deals with the value of receiving information that can be acted on to make a profit. From the perspective of firms or consumers, better information on the state of the economy can translate into tangible gains through better-informed decision making or other channels. We will not look at such effects, but focus on the subjective costs incurred by consumers from late resolution of uncertainty.

The subjective benefit of knowing the current state of the economy is that it enables consumers to better forecast their future consumption levels. If, as both microeconomic studies and financial market data indicate,¹ consumers have a preference for early resolution of uncertainty, then better forecasts translate into lower uncertainty.

Using a model where agents have a strong preference for early resolution of uncertainty, we find that the benefits from removing any doubt about the state of the economy are quite modest. To some extent this seem to explain why the general public, unlike financial market professionals, seems not to be very well informed about the state of the economy, even to the extent of making flawed economic decisions. As Chauvet and Hamilton (2006) note,

the widespread belief by the American public that the U.S. was still in recession in 2003 may have played a role in tax cuts approved by the U.S. Congress, the outcome of a special election for the governor of California, and a host of other policy and planning decisions by government bodies, private firms, and individual households.

Throughout the main parts of our paper, we assume that consumers are Bayesian learners, but the low benefits we find from providing them with better information raise the question of what would be the result if agents would employed a different learning mechanism. In the last part of our paper, we attempt a first look at this question by computing the utility level of a consumer who chooses an extremely low speed of learning: never updating his state beliefs at all. As it turns out, such consumers could actually achieve significantly higher utility levels than Bayesian learners. This effect is stronger the higher their preference for early resolution of uncertainty. This is quite striking, because such slow learners are facing the greatest amount of short-term consumption uncertainty. We link this result to the cyclical properties of the stochastic discount factor in our model.

The paper is organized as follows: Section 2 introduces the general model and explains consumers learning about the economy, Section 3 provides an overview of the data and estimation of the model, and in Section 4 we proceed to choose the preference parameters. Results for the standard model are presented and discussed in Section 5, whereas in Section 6 we look at slow learning. Section 7 concludes.

2. BASIC FRAMEWORK

We adopt a variation of Lucas's (1978) exchange economy. The log growth rate of the completely perishable endowment good c is given by

$$\Delta \log c_t = \mu_{s_t} + \sigma \epsilon_t, \quad (1)$$

where the mean growth rate μ fluctuates with the state of the economy s_t ; σ denotes its volatility and ϵ_t is an i.i.d. standard normal innovation. The state of the economy is not directly observable by agents, but the realized growth rates provide some information on it. We restrict the economy to be in either state 1, a high-growth (boom) state, or state 2, a low-growth (recession) state. The time-independent transition probabilities between the two states are given by

$$\theta_{ij} = \Pr \{s_{t+1} = j \mid s_t = i\}, \quad (2)$$

with $i, j \in \{1, 2\}$ and $\sum_i \theta_{ij} = 1$. The economy is populated by a continuum of agents whose utility can be represented by Epstein–Zin preferences

$$V_t = \left[(1 - \beta)c_t^{1-1/\psi} + \beta \mathcal{R}_t(V_{t+1})^{1-1/\psi} \right]^{-1/\psi}, \quad (3)$$




FIGURE 1. Preferences for resolution of uncertainty: (left) early resolution and (right) late resolution.

where

$$\mathcal{R}_t(V_{t+1}) = E_t[V_{t+1}^{1-\gamma}]^{1/(1-\gamma)}$$

The parameter γ is the Arrow–Pratt coefficient of constant relative risk aversion (CRRA); ψ is the elasticity of intertemporal substitution (EIS); and β measures the subjective time-discount rate under certainty. The function reduces to a monotone transformation of the standard power utility function for $\psi = \gamma^{-1}$.

As Kocherlakota (1990) shows, consumers whose EIS is larger than the inverse of their CRRA will prefer to have uncertainty resolved as early as possible, whereas the opposite holds for those whose EIS is smaller than the inverse of their CRRA. This is illustrated in Figure 1: both trees depicted in the figure have two alternative paths, both with a consumption of c_0 and c_1 in the first two periods and then a consumption of either \bar{c}_2 or \tilde{c}_2 in the last period. The only difference between the two trees is that in the one on the left-hand side, consumers learn their period-2 consumption already in period 1. A consumer with a relatively high EIS will prefer the tree on the left-hand side, because uncertainty about final-period consumption is resolved earlier.

The way we model consumers’ learning in this setting is standard, and we will provide only a short account here. For a more thorough treatment, we refer to Hamilton (1994, Ch. 22).

Let $\hat{\xi}_{t|t}$ be the vector of inferred (unsmoothed) posterior probabilities of being in each state conditional on all the data available up to time t , given complete knowledge about the population parameters, so that the j element of the vector is given by

$$\hat{\xi}_{t|t} = \Pr \{s_t = j \mid \mathbf{Y}_t, \Gamma\}. \tag{4}$$

\mathbf{Y}_t is a vector of all data up to time t and Γ contains all the model parameters. When the transition probabilities θ_j are collected in the matrix Θ , Bayes’s rule implies that an optimal forecast and inference for each date t can be found from

the equations

$$\hat{\xi}_{t|t-1} = \Theta \hat{\xi}_{t-1|t-1}, \tag{5}$$

$$\hat{\xi}_{t|t} = \frac{\hat{\xi}_{t|t-1} \odot \eta_t}{\mathbf{1}'(\hat{\xi}_{t|t-1} \odot \eta_t)}, \tag{6}$$

where \odot denotes element-by-element multiplication, $\mathbf{1}$ is a (2×1) vector of ones, and η_t is the likelihood of observing the realized consumption growth rate (and other signals) in each of the two states. When no external signals are available, the j element of η_t is given by

$$\eta_t(j) \propto e^{-(\Delta \log c_t - \mu_j)/(2\sigma)}. \tag{7}$$

Besides the information embedded in the endowment growth rates, consumers may have available additional sources of information for inferring the current state of the economy. We choose to model all this other information as an independent noisy signal. For convenience, we let this signal take the form

$$y_t = 1_{\{s_t=1\}}\mu_{y,1}(h) + \epsilon_{y,t}, \tag{8}$$

where $\epsilon_{y,t}$ is an i.i.d. white noise term, and $1_{\{s_t=1\}}$ is an indicator function that equals one if we are in the first state and zero otherwise. The strength of the signal is determined by $h \in [0, 1]$. An h of zero implies that the signal contains no information, whereas an h of one implies that the signal is strong enough to reveal the state of the economy with certainty. Normalizing the mean of the signal in the recession state to 0, and assuming without loss of generality that the signal has a positive mean in the boom state, the mean that generates an $h\%$ reduction of the probability of making a type I error is given by

$$\mu_{y,1}(h) = -2F^{-1}\left(\frac{1-h}{2}\right), \tag{9}$$

where F denotes the cumulative distribution function of a standard normal.

The final state beliefs are obtained by computing the joint likelihood of observing both the realized endowment growth rate and the realized external signal in both states:

$$\eta_t(j) \propto e^{-(\Delta \log c_t - \mu_j)/2\sigma} e^{-(y_t - \mu_y)/2}. \tag{10}$$

As with equation (7), the likelihood vector is passed through the filter (i.e., equation (6)) to generate inferred state probabilities.

3. DATA AND ESTIMATION

We use quarterly U.S. data spanning 1952:I–2005:IV to calibrate the model. The variables used in the estimation are mainly from the NIPA tables published on the Bureau of Economic Analysis’ website (<http://bea.gov/>). All series are

TABLE 1. Estimated regime parameters

State	$\mu_c(s)$	$\sigma_c(s)$	θ_{ij}
Panel A (PCE)			
Boom ($s = b$)	0.074 (0.0005)	0.0065	0.0618 (0.0243)
Recession ($s = r$)	-0.0007 (0.0010)	(0.0006)	0.2326 (0.0781)
Panel B (NDS)			
Boom ($s = b$)	0.0067 (0.0003)	0.0045	0.0618 (0.0243)
Recession ($s = r$)	0.0019 (0.0007)	(0.0002)	0.2326 (0.0781)

Note: This table reports the estimated parameters of the regime switching model for the US postwar data based on an MCMC algorithm from Kim and Nelson (1999). Panel A reports estimates using real quarterly per capita PCE, whereas Panel B reports estimates using real quarterly per capita consumption of NDS (Q1:1952–Q4:2005; *Source:* BEA). Standard errors are reported in parentheses.

expressed in real per capita terms. We use two consumption measures: personal consumption expenditures (PCE) and consumption of services and nondurables (NDS). In the construction of the second consumption series, care was taken to avoid the problems related the addition of chain-weighted series [Whelan (2002)].

The nominal risk-free rate was imputed from the end of quarter average of bid and ask quotes for 3-month treasury bills in the secondary market as reported on the monthly CRSP data base. To arrive at real interest rates, we estimated an AR(1) process for the inflation rate (taken to be the change in the log PCE deflator) and used it to compute expected inflation rates. These were then used to compute the expected real rates from the observed nominal rates. As a cross check we used realized real interest rates (which is equivalent to assuming perfect foresight) and found almost identical results.

Parameter estimates for the regime switching model are obtained in two steps. In the first step we use a Markov-chain Monte Carlo (MCMC) procedure on the real per capita GDP series from NIPA. In this, we closely followed the algorithm described in Sect. 9.1 of Kim and Nelson (1999). From this estimation, we obtain the transition matrix Θ and smoothed state probabilities for each quarter. In a second step, we estimate the consumption process parameters for each consumption series by maximum likelihood, taking the smoothed state probabilities as given. We chose this two-step procedure over estimating the process directly by applying MCMC to the consumption series for two reasons: First, it ensures that we get the same state transition probabilities for both series. Second, GDP has a stronger signal-to-noise ratio than the two consumption series. The reported standard errors were computed from the derivatives of the inverse of the Fisher information matrix. The resulting estimates are given in Table 1.

Some preliminary intuition on the model economy can be inferred from the regime-switching estimation: an important variable for our analysis is the high persistence of both states, especially the boom state. The probabilities of switching from the two states are 6.18% and 23.26%, respectively. These probabilities imply an average duration of 16.2 quarters for booms and 4.3 quarters for recessions. The high persistence of the states, coupled with the higher mean growth rates, implies that the conditional consumption growth rate is higher in booms than in recessions.

PCE is more volatile than NDS, and the spread of the mean growth rate between the two states is also wider for PCE. This is due to the high cyclicality of durables expenditures. Durables yield a stream of consumption over their lifetime, so durables expenditures do not fit perfectly with our theoretical consumption aggregate. Rather than trying to impute to true durables consumption series [as in, e.g., Dunn and Singleton (1986); Eichenbaum and Hansen (1990); Ogaki and Reinhart (1998); or Yogo (2006)], we choose to report the raw numbers. Our estimated consumption volatilities hence form an upper bound of the true underlying volatility.

On the other hand, the reported numbers for nondurables and services are likely to be on the lower end of the true consumption series because of time aggregation issues, and hence form a lower bound on the true volatility of the underlying series [Breedon et al. (1989)].

4. PARAMETERIZATION

The main factor determining the cost of business-cycle uncertainty is the relative size of the EIS ψ and the coefficient of relative risk aversion γ .

For our purposes, we choose not to rely on standard calibration techniques from business-cycle research in the tradition of Kydland and Prescott (1982), where the model parameters are chosen to match long-run aggregate ratios: A standard calibration for the coefficient γ would be problematic. As Tallarini (2000) forcefully demonstrates, the business cycle predictions are mostly determined by the EIS.

Given the controversies regarding calibration,² we use parameter values from the asset-pricing literature for the risk aversion coefficient γ and confirm the estimates found in recent micro studies for the EIS ψ with empirical financial market data. [See also the discussion in Lucas (2003).]

For γ , we use a benchmark of 25. Although this is high in comparison to what is classically used in business-cycle research, it is in the middle of the range of what is employed in the asset-pricing literature [see, e.g., Campbell and Cochrane (1999) or Lettau et al. (2008)]. The time discount factor β is set to 0.9925 throughout the paper. This is standard level in the asset pricing literature [see, e.g., Bansal and Yaron (2004) or Lettau et al. (2008)].

As recent research has shown, the intertemporal elasticity parameter is best estimated using disaggregate data. These studies typically find EIS parameters around or above 1. [See Beaudry and van Wincoop (1996); Vissing-Jørgensen

TABLE 2. Implied vs. empirical risk-free spread

γ (CRRA)	ψ (EIS)				
	0.5	1.15	1.3	1.75	2
0.5	4.77 (4.90)	2.08 (2.47)	1.84 (2.24)	1.37 (1.77)	1.20 (1.67)
5	4.81 (4.73)	2.12 (2.33)	1.88 (2.19)	1.41 (4.55)	1.24 (5.34)
15	4.88 (4.34)	2.19 (2.00)	1.95 (1.86)	1.48 (6.76)	1.31 (9.46)
25	4.91 (3.94)	2.19 (1.64)	1.95 (1.46)	1.47 (3.99)	1.30 (7.46)
30	4.91 (3.73)	2.16 (1.46)	1.92 (1.26)	1.44 (2.24)	1.26 (4.76)

Note: This table reports the spreads between the boom and recession risk-free rates predicted by the model letting both the EIS and the CRRA vary. The values reported are for the parameters estimated using PCE as a consumption measure. The gray area highlights the range of plausible EIS according to Vissing-Jørgensen and Attanasio (2003). The unconditional risk-free is reported in parentheses.

(2002)]. The most relevant study for our purposes is by Vissing-Jørgensen and Attanasio (2003), who employ the same Epstein–Zin framework that we do. They find values in the range of 1.17 to 1.75 for the EIS. In the tables that are cited hereafter, we will mark this range with a gray backdrop.

To further pin down the parameter ψ , we match the predicted fluctuations of the risk-free rate over the business cycle with those in the U.S. data.³ Regressing our estimated one-period ahead risk-free rate on the boom probabilities and a constant, we find a regression coefficient of 1.9. That is, moving from a situation where agents know for sure that they were in a recession to one where they would know for sure that they were in a boom entails a increase in the annualized 3M risk-free rate of 1.9 percentage points.

We compute the model's predicted interest rate using the method described in the Appendix. Table 2 collects the predicted spread between the boom and recession interest rates. The average predicted interest rate is given in parentheses. From the table, it is clear that the spread is determined mainly by the EIS parameter (ψ), with only limited variation with the coefficient of risk aversion (γ).⁴ The empirical spread over the business cycle is well matched for $\psi = 1.3$ and we choose this value for our baseline calibration.

5. RESULTS

5.1. Reduction in Consumption Uncertainty

Knowledge of the state of the economy affect utility by reducing uncertainty about future consumption levels. Figure 2 illustrates how knowledge of the current state




FIGURE 2. Forecast precision and state uncertainty. This figure shows three-quarters ahead forecasts of consumption levels measured as the increase in log consumption relative to its current level. The left-hand panel shows the case of an agent being in a boom state, whereas the right-hand panel shows the case of an agent being in a recession state. In both panels forecasts given state certainty are indicated by the bold dashed lines, with the one-standard deviation bound given by the thin dashed lines. The solid lines in both subplots give the expectation of a consumer who bases her forecast on the ergodic state probabilities. The shaded area gives a one-standard deviation bound.

translates into better forecasts. Plotted are expected future consumption levels, measured as the increase in log consumption relative to the current level, for the next three quarters. The solid lines in both subplots give the expectation of a consumer who has no particular information on the current state of the economy and bases her forecast on the ergodic state probabilities. The shaded area gives a one-standard deviation bound for this forecast.

Forecasts given state certainty are indicated by the bold dashed lines in each subplot. Being in a boom translates into a higher expected consumption level for all future periods. The difference between this forecast and that made by the ignorant consumer is increasing with the horizon, but at a decreasing rate as the state forecast converges with the horizon to the ergodic state probabilities. The one-standard deviation bound on the forecast is given by the thin dashed lines. Conversely, being in a recession translates into sharply lower expected consumption levels for

all future periods. The difference in expectations from the unconditional case is more pronounced, reflecting the lower incidence of recession periods.

As the figure suggests, knowing the state of the economy can yield significant forecasting gains, especially if this state happens to be a recession. Moreover, these gains are increasing with the forecasting horizon. The average root-mean-square error (RMSE) of the forecast under full information is 10% lower than the one when no information is available. (At the longer horizons of 1 and 2 years ahead, it is 12% and 18% lower, respectively.) Most of this reduction comes from the benefit of being able to identify recessions correctly. When only recession states are considered, the fall in the RMSE is almost twice as high.

At the lower limit (in our setup), consumers will have available the endowment growth rate for estimating the state of the economy. Hence, the relevant question for measuring the cost of business-cycle uncertainty is how much additional uncertainty the lack of firm knowledge introduces for such consumers. We estimated this by running a 100,000-period Monte Carlo simulation of our model economy. In each period, consumers' beliefs were updated according to the filter presented in Section 2. As it turns out, such business-cycle uncertainty does not translate in a large predictability loss. For most horizons, the average RMSE is roughly only 3% higher than under certainty. Most of this loss occurs in recessions, because they are on average shorter and hence harder to detect. The relatively low level of uncertainty left after the information in the consumption series is processed contributes to the relatively low costs of business cycle uncertainty reported in the next section.

5.2. Implied Utility Gains

Figure 3 shows how much consumption an agent would be willing to give up to learn the current state of the economy, for different levels of uncertainty, using both consumption measures. The left-hand panel gives the values when we use PCE as our consumption measure; the right-hand panel when we use NDS. Two measures of cost are provided. The solid lines show the amount of consumption drop that the agents would be willing to incur to always have a perfect signal on the state of the economy available. This is not to be confused with a one-time payment, because a consumption drop would have a permanent effect on future consumption growth rates through our process assumptions. An alternative measure is how much consumption agents would be willing to give up to learn the current state of the economy with certainty; this cost is marked with dashed lines. If they have very strong priors about the state, i.e., if they are on either on the ends of the horizontal axes, such a signal would be superfluous and they would not be willing to pay for it. This is not true for the first measure: if business-cycle uncertainty is permanently removed, consumers benefit from the knowledge that there will also be no uncertainty in future periods.

As we see from the figure, the cost of business-cycle uncertainty is higher when we rely on PCE as our consumption measure than when we use NDS. This reflects

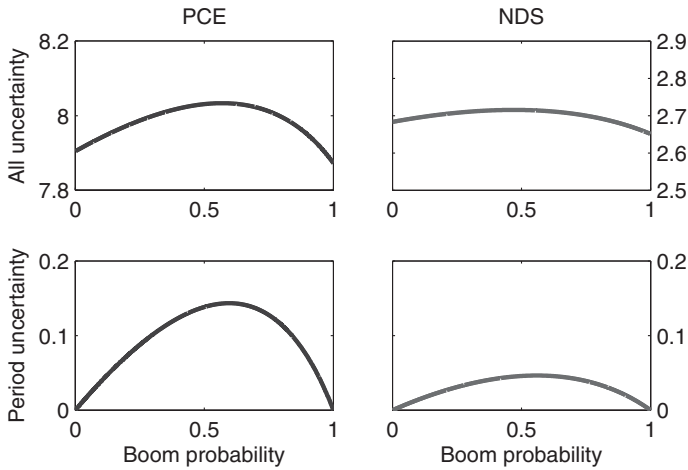


FIGURE 3. Degree of uncertainty and its period cost. This figure shows the cost for an agent to learn the current state of the economy. The left-hand panel depicts the values using per capita PCE; the right-hand panel is obtained using per capita consumption of NDS. The solid lines show the payment, measured with a consumption drop, that the agent would make to always have a perfect signal on the state of the economy. The dashed lines depict the current consumption the agent would be willing to give up to learn the current state of the economy with certainty.

the larger difference between the mean boom and recession growth rates of the two aggregates.

The average cost of business cycle uncertainty depends, of course, on how much uncertainty there is on average. Figure 4 shows the ergodic distribution of state beliefs given only the information available in each of the two consumption measures. Both densities are right-skewed, reflecting the higher incidence of booms in the population. The larger difference between the mean boom and recession growth rates of PCE makes it a better measure for detecting the state of the economy. This advantage is balanced somewhat by its larger volatility. In sum, the higher degree of uncertainty under NDS is dominated by its lower cost.

Table 3 reports the results for the cost of business cycle uncertainty for a range of alternative parameters. All numbers are reported for a quarterly time-discount factor (β) of 0.9925 and are given in basis points. The first number reported is the average amount of consumption that an agent would be willing to permanently give up to move to an economy where she would always know the state of the economy. Even for our baseline calibration, where agents have a strong preference for early resolution of uncertainty, it does not exceed 0.1% of current consumption.⁵ The second set of numbers, given in parentheses, gives the average amount of consumption a household would be willing to give up to learn only the current state of the economy. As expected, the numbers are increasing as we move along any path from the northwest corner of the table (where $\gamma < 1/\psi$)




FIGURE 4. Densities of state beliefs. This figure shows the ergodic distribution of state beliefs. The solid line is plotted using the information available in the per capita PCE, whereas the dashed line is plotted using the information available in the per capita consumption of NDS.

to the southeast corner (where $\gamma > 1/\psi$). This reflects a shift to an ever greater preference for early resolution of uncertainty.

6. BLISS IN IGNORANCE AFTER ALL?

The preceding results indicate that only modest utility gains can be achieved from trying to tweak out information from other sources than the endowment growth rates per se. If information processing is costly, it might not be worthwhile for a consumer to actually invest a lot of effort in this.

Given this finding, it seems natural to ask what is the benefit of processing even the information that is given in the endowment growth rates. As it turns out, there might be significant utility gains from not processing this information. That is, there are cases where the consumer prefers always to be kept ignorant about the state of the economy.

We look at the case of a consumer who, instead of updating his beliefs according to whatever news is available, chooses to use the unconditional boom and recession probabilities for forming his expectations about future consumption paths—in other words, a consumer who knows the underlying structure of the economy, but

TABLE 3. Equivalent consumption reductions

γ (CRRA)	ψ (EIS)					
	0.5	1.15	1.3	1.5	1.75	2
	PCE					
0.5	-0.470 (-0.002)	-0.451 (0.002)	-0.467 (0.003)	-0.491 (0.003)	-0.522 (0.003)	-0.552 (0.003)
5	0.445 (0.015)	0.554 (0.012)	0.549 (0.012)	0.538 (0.011)	0.522 (0.011)	0.504 (0.010)
15	3.367 (0.062)	3.722 (0.040)	3.746 (0.038)	3.768 (0.036)	3.785 (0.034)	3.795 (0.032)
25	7.197 (0.116)	7.761 (0.077)	7.809 (0.073)	7.857 (0.069)	7.900 (0.066)	7.932 (0.063)
30	9.115 (0.143)	9.724 (0.097)	9.777 (0.093)	9.831 (0.088)	9.880 (0.084)	9.917 (0.080)
	NDS					
0.5	-0.139 (-0.001)	-0.089 (0.001)	-0.087 (0.001)	-0.087 (0.001)	-0.087 (0.001)	-0.089 (0.001)
5	0.202 (0.006)	0.273 (0.004)	0.277 (0.004)	0.280 (0.004)	0.282 (0.004)	0.282 (0.004)
15	1.165 (0.024)	1.299 (0.015)	1.309 (0.014)	1.318 (0.013)	1.326 (0.012)	1.331 (0.011)
25	2.412 (0.045)	2.621 (0.027)	2.638 (0.026)	2.656 (0.024)	2.671 (0.022)	2.682 (0.021)
30	3.125 (0.057)	3.372 (0.034)	3.393 (0.032)	3.415 (0.030)	3.434 (0.028)	3.448 (0.026)

Note: This table reports for the two consumption measures the size of the permanent drop in consumption which agents would be willing to incur to always have full information on the state of the business cycle. The numbers in parenthesis is the corresponding amount that agents would be willing to give up to learn the *current* state of the economy. Both quantities are measured in basis points of current consumption.

who does not use any conditioning information when forming expectations about future consumption growth rates. We will refer to such a consumer as ignorant.

Defining the continuation value scaled with the current consumption level as $v_t = V_t/C_t$, we can represent the agents’ preferences from (3) as

$$v_t = \{ (1 - \beta) + \beta [\mathcal{R}_t(v_{t+1}G_{t+1})]^{1-1/\psi} \}^{1/(1-1/\psi)}, \tag{11}$$

where G_{t+1} denotes the gross growth rate of consumption between t and $t + 1$:

$$G_{t+1} = C_{t+1}/C_t.$$

The information set of the ignorant consumer is time-invariant. We signify this by using the unconditional version of the risk-adjustment operator \mathcal{R} , which is given

TABLE 4. Gains from not learning

γ (CRRA)	ψ (EIS)					
	0.5	1.15	1.3	1.5	1.75	2
	PCE					
0.5	-0.093	-0.192	-0.210	-0.232	-0.257	-0.280
5	0.808	1.599	1.746	1.925	2.126	2.305
15	3.356	6.376	6.906	7.543	8.244	8.859
25	6.702	12.111	13.009	14.070	15.215	16.199
30	8.643	15.191	16.247	17.484	18.806	19.932
	NDS					
0.5	-0.032	-0.067	-0.073	-0.081	-0.089	-0.097
5	0.271	0.542	0.593	0.655	0.725	0.788
15	1.046	2.061	2.247	2.473	2.726	2.951
25	1.967	3.810	4.140	4.539	4.982	5.373
30	2.483	4.764	5.168	5.655	6.194	6.666

Note: This table reports the increase in utility possible from not updating the state probabilities at all. This is measured as the percentage increase in consumption necessary to make the consumer indifferent between always observing the state of the economy with certainty and committing himself to never updating his prior subjective state probabilities.

by

$$\mathcal{R}(x_{t+1}) = E[x_{t+1}^{1-\gamma}]^{\frac{1}{1-\gamma}}.$$

By guess and verify, we can confirm that the scaled continuation value of the ignorant consumer is given by the fixed point of

$$\tilde{v} = \{ (1 - \beta) + \beta[\mathcal{R}(\tilde{v}G_{t+1})]^{1-1/\psi} \}^{1/(1-1/\psi)}. \tag{12}$$

Table 4 gives the utility gains (or losses) that not learning yields for the consumer. Analogously to the last section, we define utility gains as the increase in consumption necessary to make the consumer indifferent between (1) learning the state with certainty (and enjoying higher consumption) and (2) not updating the state beliefs at all (but staying at the current consumption level). Formally, they are given by the value x that solves

$$\tilde{v} = e^x \mathcal{R}(v_t). \tag{13}$$

The left-hand side of the equation gives the scaled continuation value under ignorance. v_t is the continuation value at time t after uncertainty is resolved. The term $\mathcal{R}(v_t)$ gives the certainty equivalent of the continuation value under the ignorant agent’s information set.

The results in the table show that there are significant utility gains from never updating the state beliefs, especially compared to the modest gains that

could be made by eliminating business-cycle uncertainty, whereas for the benchmark calibration, the gains from eliminating business-cycle uncertainty for the two consumption measures were only about three and eight basis points. (See Table 3.) The gains from not learning at all reported in Table 4 are about 4 and 13 percentage points—an increase by more than two orders of magnitude. These gains increase as we move toward the lower right corner of the table. This might be quite surprising, because a consumer whose preferences can be represented by parameter constellations such as those in the lower right corner has an extreme preference for early resolution of uncertainty. Moreover, in the preceding section, we demonstrated that these very consumers gain from more precise knowledge about the state of the economy.

To explain this seemingly contradictory finding, we rely on a covariance decomposition of the certainty equivalents in equations (11) and (12). By the computational formula for the covariance to express certainty equivalent as

$$\begin{aligned} \mathcal{R}_t(v_{t+1}G_{t+1}) &= E_t \left[v_{t+1}^{1-\gamma} G_{t+1}^{1-\gamma} \right]^{\frac{1}{1-\gamma}} \\ &= \left\{ E_t \left[v_{t+1}^{1-\gamma} \right] E_t \left[G_{t+1}^{1-\gamma} \right] + \text{cov}_t \left(v_{t+1}^{1-\gamma}, G_{t+1}^{1-\gamma} \right) \right\}^{\frac{1}{1-\gamma}}, \end{aligned} \tag{14}$$

the magnitude of the covariance term depends on whether agents learn or not. In particular,

$$\text{cov}_t \left(v_{t+1}^{1-\gamma}, G_{t+1}^{1-\gamma} \right) \begin{cases} = 0, & \text{without learning} \\ > 0, & \text{with learning.} \end{cases}$$

With learning, the covariance term is positive, because a higher-than expected consumption growth rate leads to an upward revision of the probability that the agent assigns to the boom state and thus an increase of the continuation value v_{t+1} .

With ignorance, the agents do not update their beliefs about the growth state at $t + 1$. Although the continuation value V_{t+1} depends on the realized consumption level at $t + 1$, the scaled continuation value v_{t+1} is constant and given by $v_{t+1} = \tilde{v}$, so the covariance must be zero.

Applying the chain rule to equation (14), we see that

$$\frac{\partial \mathcal{R}_t(v_{t+1}G_{t+1})}{\partial \text{cov}_t \left(v_{t+1}^{1-\gamma}, G_{t+1}^{1-\gamma} \right)} \begin{cases} > 0 & \text{if } \gamma < 1 \\ = 0 & \text{if } \gamma = 1 \\ < 0 & \text{if } \gamma > 1. \end{cases} \tag{15}$$

Because the conditional covariance is positive only with learning, equation (14) shows that $\gamma = 1$ is a cutoff for when, for a given value function, learning increases current utility. For levels of risk aversion above 1, future learning decreases the certainty equivalent.

The covariance decomposition also explains why the utility effect of ignorance is stronger, the higher the EIS. The higher ψ , the more the scaled continuation value reacts to changes in expected growth rates. This implies that the magnitude of the covariance term in equation (14) increases in ψ .

One caveat to the discussion in this section is that we have treated the information processing of the agents as separate from the economic outcome. This is innocuous in the exchange economy setting we use here, but would not necessarily hold in a richer model with a production side. In such a model, at least a subset of agents could choose not to process economic information without affecting the path of the general economy.

7. CONCLUSION

In this paper, we investigate and quantify the cost of business-cycle uncertainty in a simple setup where the economy switches between booms and recessions at random intervals. Firm knowledge about the prevailing state of the economy enables better forecasts of future consumption levels and hence reduces consumption uncertainty, a boon for consumers who prefer to have uncertainty resolved as early as possible.

We found two apparently contradictory results: (1) consumers experience a (very modest) utility gain when they are provided with information beyond what is already incorporated into their endowment process; (2) there are large utility gains from either committing to ignoring macroeconomic information or assuming a simplified model for the economy. These results were reconciled by noting the different mechanism behind each of the results: In the second case, ignorance removes some of the positive covariance between realized consumption growth and expected future consumption. For preferences such as those of our baseline calibration, this leads to a utility gain. In the first case, the information is orthogonal to the current consumption, and adding more of it does not change the covariance between realized consumption and expected utility. Here the normal results obtain: consumers with a preference for early resolution of uncertainty will also benefit from more information.

As we see it, ours is but a first attempt to quantify the cost of business-cycle uncertainty. One natural extension would be to look at the costs of uncertainty in economies with incomplete insurance markets, such as those studied by Imrohoroglu (1989) and Atkeson and Phelan (1994).

NOTES

1. For a review of this literature see Section 4.
2. See, e.g., Altug (1989) and Christiano and Eichenbaum (1992) for critical views, or Kydland and Prescott (1991) and Hoover (1995) for supporting arguments.
3. The results reported in this section are obtained using PCE as a consumption measure. Because of the smaller difference between the mean growth rates of NDS consumption over the business cycle, the model-implied spreads are generally lower for this measure.

4. The relation can be explored by log-linearizing equation (A.2) in the Appendix. The resulting expression for the continuously compounded interest rate is

$$r_t^f = -\log \beta + \frac{1}{\psi} E_t [g_{t+1}] - \frac{\kappa}{2} \left[\frac{1}{\psi^2} \sigma_g^2 + \left(\frac{1}{\kappa} - 1 \right) \text{var}_t (r_{t+1}^w) \right],$$

where $\kappa = (1 - \gamma)/(1 - 1/\psi)$, $\text{var}_t (g_{t+1})$ is the conditional standard deviation of the growth rate of log consumption, and $\text{var}_t (r_{t+1}^w)$ is the conditional variance of the return to the aggregate wealth portfolio. [See Brevik and d'Addona (2010)]. ψ has a first-order effect on the spread between boom and recession interest rates through the term $(1/\psi)E_t [g_{t+1}]$. The risk aversion parameter γ influences the spread through the variance terms and, as can be seen from the columns in Table 2, the effect of increasing γ is not always of the same sign. The reason is that the variance of returns to the aggregate wealth portfolio is endogenous and changes with γ .

5. This might appear quite high given Lucas's estimate of 0.1% for the overall cost, but we are assuming a level of risk aversion much higher than the value of 2 chosen in his contribution.

REFERENCES

- Altug, Sumru (1989) Time-to-build and aggregate fluctuations: Some new evidence. *International Economic Review* 30, 889–920.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega (2007) Real-time price discovery in global stock, bond and foreign exchange markets. *Journal of International Economics* 73, 251–277.
- Atkeson, Andrew and Christopher Phelan (1994) Reconsidering the costs of business cycles with incomplete markets. In Stanley Fischer and Julio J. Rotemberg (eds.), *NBER Macroeconomics Annual 1994*, vol. 9, pp. 187–218. Cambridge, MA: MIT Press.
- Bansal, Ravi and Amir Yaron (2004) Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59, 1481–1509.
- Beaudry, Paul and Eric van Wincoop (1996) The intertemporal elasticity of substitution: An exploration using a U.S. panel of state data. *Economica* 63, 495–512.
- Breedeen, Douglas T., Michael R. Gibbons, and Robert H. Litzenberger (1989) Empirical test of the consumption-oriented CAPM. *Journal of Finance* 44, 231–262.
- Brevik, Frode and Stefano d'Addona (2010) Information quality and stock returns revisited. *Journal of Financial and Quantitative Analysis* 45, 1419–1446.
- 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, 205–251.
- Chauvet, Marcelle and, James D. Hamilton (2006) Dating business cycle turning points. In Costas Milas, Philip Rothman, and Dick van Dijk (eds.), *Nonlinear Time Series Analysis of Business Cycles*, pp. 1–54. Amsterdam: Elsevier.
- Christiano, Lawrence J. and Martin Eichenbaum (1992) Current real-business-cycle theories and aggregate labor-market fluctuations. *American Economic Review* 82, 430–450.
- Dolmas, Jim (1998) Risk preferences and the welfare cost of business cycles. *Review of Economic Dynamics* 1, 646–676.
- Dunn, Kenneth B. and Kenneth J. Singleton (1986) Modelling the term structure of interest rates under habit formation and durability. *Journal of Financial Economics* 17, 27–55.
- Eichenbaum, Martin and Lars Peter Hansen (1990) Estimating models with intertemporal substitution using aggregate time series data. *Journal of Business and Economic Statistics* 8, 53–69.
- Epstein, Larry G. and Stanley Zin (1989) Substitution, risk aversion, and the temporal behavior of consumption growth and asset returns-I. A theoretical framework. *Econometrica* 57, 937–969.
- Gollier, Christian (2001) *The Economics of Risk and Time*. Cambridge, MA: MIT Press.
- Hamilton, James D. (1989) A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57, 357–384.
- Hamilton, James D. (1994). *Time Series Analysis*. Princeton, NJ: Princeton University Press.

- Hoover, Kevin D. (1995) Facts and artifacts: Calibration and the empirical assessment of real-business-cycle models. *Oxford Economic Papers* 47, 24–44.
- Imrohorglu, Ayse (1989) Cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political Economy* 97, 1364–1383.
- Judd, Kenneth L. (1998). *Numerical Methods in Economics*. Cambridge, MA: MIT Press.
- Kim, Chang-Jin and Charles R. Nelson (1999) *State-Space Models with Regime Switching*. Cambridge, MA: MIT Press.
- Kliesen, Kevin L. and Frank A. Schmid (2006) Macroeconomic news and real interest rates. *Federal Reserve Bank of St. Louis Review* 88, 133–144.
- Kocherlakota, Narayana R. (1990) Disentangling the coefficient of relative risk aversion from the elasticity of intertemporal substitution: An irrelevance result. *Journal of Finance* 45, 175–90.
- Kydland, Finn E. and Edward C. Prescott (1982) Time to build and aggregate fluctuations. *Econometrica* 50, 1345–1370.
- Kydland, Finn E. and Edward C. Prescott (1991) The econometrics of the general equilibrium approach to business cycles. *Scandinavian Journal of Economics* 93, 161–178.
- Lettau, Martin, Sydney C. Ludvigson, and Jessica A. Wachter (2008) The declining equity premium: What role does macroeconomic risk play? *Review of Financial Studies* 21, 1653–1687.
- Lucas, Robert E. (1987). *Models of Business Cycles*. Oxford, UK: Basil Blackwell.
- Lucas, Robert E. (2003) Macroeconomic priorities. *American Economic Review* 93, 1–14.
- Lucas, Robert E., Jr. (1978) Asset prices in an exchange economy. *Econometrica* 46, 1429–1445.
- Obstfeld, Maurice (1994) Evaluating risky consumption paths: The role of intertemporal substitutability. *European Economic Review* 38, 1471–1486.
- Ogaki, Masao and Carmen M. Reinhart (1998) Measuring intertemporal substitution: The role of durable goods. *Journal of Political Economy* 106, 1078–1098.
- Pemberton, James (1996) Growth trends, cyclical fluctuations, and welfare with non-expected utility preferences. *Economics Letters* 50, 387–392.
- Tallarini, Thomas D. (2000) Risk-sensitive real business cycles. *Journal of Monetary Economics* 45, 507–532.
- Vissing-Jørgensen, Annette (2002) Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy* 110, 825–853.
- Vissing-Jørgensen, Annette and Orazio P. Attanasio (2003) Stock-market participation, intertemporal substitution, and risk-aversion. *American Economic Review* 93, 383–391.
- Weil, Philippe (1989) The equity premium puzzle and the risk free rate puzzle. *Journal of Monetary Economics* 24, 401–421.
- Whelan, Karl (2002) A guide to U.S. chain aggregated NIPA data. *Review of Income and Wealth* 48, 217–233.
- Yogo, Motohiro (2006) A consumption-based explanation of expected stock returns. *Journal of Finance* 61, 539–580.

APPENDIX: COMPUTING THE IMPLIED RISK-FREE RATE

To solve for the risk-free rate, we rely on one of the key results of Epstein and Zin (1989), which is that the stochastic discount factor can be expressed as

$$M_{t+1} = \beta^\kappa \left(\frac{c_{t+1}}{c_t} \right)^{-\frac{\kappa}{\psi}} (R_{t+1}^w)^{\kappa-1}, \quad (\mathbf{A.1})$$

where R_{t+1}^w is the equilibrium gross return to aggregate wealth between t and $t + 1$. As usual, the gross risk-free rate is given by the inverse of the expected value of the stochastic

discount factor, or

$$R_t^f = E_t \left[\beta^\kappa \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma} \left(\frac{1 + w_{t+1}}{w_t} \right)^{\kappa-1} \right]^{-1}, \quad (\text{A.2})$$

where w_t denotes the wealth–consumption ratio (i.e., the price in units of current consumption of a claim to aggregate consumption). Thus, using the Epstein–Zin specification, the interest rate will fluctuate not only with the expected growth rate of consumption, but also with the expected changes in the price–consumption ratio (w_t).

Because all processes are Gaussian, the only variable that is hard to compute is the wealth–consumption ratio. In a similar setting, we have shown that the wealth–consumption ratio is a nonlinear combination of expected price wealth ratios in boom and recession states [Brevik and d'Addona (2010)]; hence we cannot rely on the standard linear-algebra closed form solutions. To solve for the state prices, we relied on numerical integration using Gauss–Hermite quadrature. Using a large number of nodes, we made sure that the computed prices were arbitrarily close to the true values. [See, e.g., Judd (1998, Ch. 7)].

Given the estimated price–consumption ratios, we compute the expected value of the stochastic discount factor. Taking its inverse yields the implied risk-free rate.