Using Asset Prices to Measure the Cost of Business Cycles

Fernando Alvarez

University of Chicago and National Bureau of Economic Research

Urban J. Jermann

University of Pennsylvania and National Bureau of Economic Research

We measure the cost of consumption fluctuations using an approach that does not require the specification of preferences and instead uses asset prices. We measure the *marginal cost of consumption fluctuations*, the per unit benefit of a marginal reduction in consumption fluctuations expressed as a percentage of lifetime consumption. We find that the gains from eliminating all consumption uncertainty are very large. However, for consumption fluctuations corresponding to business cycle frequencies, we estimate the marginal cost to be between 0.08 percent and 0.49 percent of lifetime consumption.

In a seminal contribution, Lucas (1987) proposes a measure of the welfare cost of economic fluctuations. His measure is defined as the compensation required to make the representative agent indifferent between consumption plans with and without business cycle fluctua-

tions. With this measure, Lucas finds a very small cost of business cycles. Subsequently, several studies have proposed estimates of this cost of business cycles under alternative assumptions on preferences and consumption processes. As a function of these assumptions, estimates vary widely across studies.[1] In our paper, we measure the welfare cost of business cycles through an approach that does not require the specification of consumer preferences; instead, we directly use financial market data.

We define the *marginal cost of consumption fluctuations* as the per unit benefit of a marginal reduction in consumption fluctuations. Because it is marginal, we can relate this cost directly to asset prices. In particular, we show the marginal cost to be equal to the ratio of the prices of two long-lived securities: one representing a claim to stabilized consumption, the other a claim to actual consumption. Measuring the cost of economic fluctuations then becomes a task in asset pricing.

The literature has in general focused on the potential benefits of eliminating all consumption uncertainty, that is, replacing the actual consumption process by its expected path. We take this as a starting point of our analysis, but we also focus specifically on the welfare gain of eliminating business cycle fluctuations without eliminating all consumption risk. We believe that this difference is important because a large part of consumption fluctuations may not be directly related to business cycles and as such to policies related to business cycle stabilization. On the basis of no-arbitrage principles, we derive simple expressions for the marginal benefit of eliminating all uncertainty and for the benefit of eliminating business cycle fluctuations. These expressions are simple functions of an interest rate, the average growth rate of consumption, a consumption risk premium, and the moving average coefficients that define the process for stabilized consumption.

Estimating the marginal cost presents two challenges. First, we need to price a nontraded security, an equity claim to consumption. To do this, we use an extension of the method proposed by Cochrane and Saa-Requejo (2000) that is based on no-arbitrage restrictions when existing assets do not completely span the payoff of the asset to be priced. A second issue concerns the measurement of the business cycle components of consumption. We use a frequency domain approach following the work of Baxter and King (1998, 1999). This application is complicated because our requirement that the stabilized consumption be

[1] See, e.g., Atkeson and Phelan (1994), Obstfeld (1994), Campbell and Cochrane (1995), Dolmas (1998), Hansen, Sargent, and Tallarini (1999), Krusell and Smith (1999), Tallarini (2000), Otrok (2001), and Lucas (2003) for a recent survey of this literature; for the related literature on the welfare gains from international integration, see Lewis (1996) and van Wincoop (1999).

defined as the dividend of a security precludes the use of the standard two-sided moving average representation.

We have two sets of quantitative results. First, our estimate of the cost of all consumption uncertainty, while noisy, is extremely high. Essentially, offering agents a perpetual bond whose coupons are growing at the average growth rate of the economy would be extremely valuable. On the other hand, the cost of business cycle fluctuations is found to be small. We find that the costs of business cycle fluctuations are between 0.08 percent and 0.49 percent of consumption. This finding is robust to, among other things, the set of reference security returns used for pricing consumption risk, the specifications of the stochastic processes of consumption and returns, the possible imperfections of the frequency domain filters we use, and the introduction of durable goods consumption.

We organize the paper as follows. In Section I we define the marginal cost and present characterizations in terms of yields and growth rates. Sections II, III, and IV contain the detailed empirical analysis. Section V presents analytical results about the marginal cost and its relationship to Lucas's approach of measuring the cost of business cycles.

I. The Marginal Cost of Consumption Fluctuations

We start this section by defining the marginal cost of consumption fluctuations. We characterize this cost for two definitions of consumption fluctuations. The first includes all consumption uncertainty, and the second covers business cycle fluctuations. In both cases we derive expressions for the marginal cost as functions of three variables: an interest rate, the average growth rate of consumption, and a consumption risk premium. We then quantify the marginal costs using the values of these variables estimated in Sections II and III of the paper.

A. Defining the Marginal Cost of Consumption Fluctuations

Assume that $\{x\}$ is a stochastic process for payoffs, that is, a stream of random payoffs for all dates $t \geq 1$, and that $V_0[\{x\}]$ is the time 0 price of a security that pays $\{x\}$. Consider the processes $\{c\}$ that represent aggregate consumption and $\{C\}$ a more stable version of aggregate consumption, which we call trend. We define the marginal cost of consumption fluctuations ω_0 as the ratio of the values of two securities: a claim to the consumption trend, $V_0[\{C\}]$, and a claim to aggregate consumption, $V_0[\{c\}]$:

$$\omega_0 \equiv \frac{V_0[\{C\}]}{V_0[\{c\}]} - 1. \tag{1}$$

If an agent can trade these two securities, the difference in prices $V_0[\{C\}] - V_0[\{c\}]$ measures the benefit of removing the business cycle fluctuations from this agent's consumption. This is achieved by selling the aggregate consumption process $\{c\}$ and buying the consumption trend $\{C\}$. In equation (1), ω_0 expresses this cost in terms of $V_0[\{c\}]$, the value of aggregate consumption $\{c\}$.

Estimating the marginal cost ω_0 in (1) presents two challenges, which occupy most of the body of the paper. We need to develop a workable definition of $\{C\}$, and we need to measure the prices $V_0[\{C\}]$ and $V_0[\{c\}]$, which may not be directly observable.

We provide here an interpretation of the marginal cost ω_0 for the particular case of a representative agent economy.[2] Assume that in each period t, the economy experiences one of finitely many events $z_t \in Z$, and denote by $z^t = (z_0, z_1, \ldots, z_t)$ the history of events up through and including period t. We index commodities by histories, so we write $x : \mathbf{Z} \to R_+$, where $\mathbf{Z} \equiv \prod_{t \geq 1} Z^t$, or simply $\{x\} = \{x_t(z^t) : \forall t \geq 1, z^t \in Z^t\}$. Let $U(\cdot)$ be a utility function, mapping consumption processes into R. We define the total cost of consumption fluctuations function $\Omega(\alpha)$ as the solution of

$$U([1 + \Omega(\alpha)]\{c\}) = U((1 - \alpha)\{c\} + \alpha\{C\}), \qquad (2)$$

where $\alpha \in [0, 1]$, $c : \mathbf{Z} \to R_+$, and $C : \mathbf{Z} \to R_+$. Without writing it explicitly, we assume that $c_0(z^0)$ enters the utility function in (2) in such a way as not to be multiplied by $1 + \Omega(\alpha)$ and that $c_0(z^0) = C_0(z^0)$. The scalar α measures the fraction of consumption $\{c\}$ that has been replaced by the less risky trend consumption $\{C\}$. The total cost function gives the total benefit from reducing consumption fluctuations as a function of the fraction of the reduction in fluctuations. It is straightforward to see that $\Omega(0) = 0$, since no reduction in fluctuations generates no benefit. Thus $\Omega'(0)$ is the first-order approximation of $\Omega(1)$ around $\alpha = 0$.[3] We find $\Omega'(0)$ a useful approximation of $\Omega(1)$ because we can estimate $\Omega'(0)$ using asset prices; indeed $\Omega'(0) = \omega_0$. To see this, assuming that U is differentiable with respect to each $c_t(z^t)$ for all t and z^t and denoting the partial derivatives by $U_{z^t}(\{c\}) \equiv \partial U(\{c\})/\partial c_t(z^t)$, we obtain

$$\Omega'(0) = \frac{\sum_{t=1}^{\infty} \sum_{z^t \in Z^t} U_{z^t}(\{c\}) \cdot [C_t(z^t) - c_t(z^t)]}{\sum_{t=1}^{\infty} \sum_{z^t \in Z^t} U_{z^t}(\{c\}) \cdot c_t(z^t)}. \qquad (3)$$

[2] We present a nonrepresentative agent interpretation in Sec. V below.

[3] In Sec. V below, we present a more detailed analysis of $\Omega(\cdot)$ and a comparison of ω_0 with the cost used in Lucas (1987).

Furthermore, notice that the shadow price of a security with payoff $\{x\}$ for the agent with consumption $\{c\}$ must be

$$V_0[\{x\}] \equiv \frac{1}{U_{z^0}(\{c\})} \sum_{t=1}^{\infty} \sum_{z^t \in Z^t} U_{z^t}(\{c\}) \cdot x_t(z^t).$$

Combining this expression with (3), we obtain $\omega_0 = \Omega'(0)$.

B.   Cost of All Uncertainty

Consider a definition of C_t that implies the elimination of all consumption uncertainty, namely,

$$C_t = E_0 c_t. \tag{D1}$$

Assume that the unconditional expectation of consumption growth does not depend on calendar time:

$$E\left[\frac{c_{t+1}}{c_t}\right] = 1 + g. \tag{A1}$$

Hence, using the definition in equation (1), we have the marginal cost of all uncertainty:

$$\omega_0 = \frac{r_0 - g}{y_0 - g} - 1,$$

where we define y_0 as the yield to maturity that corresponds to the price $V_0(\{C_t\})$, and likewise r_0 for $V_0(\{c_t\})$, implicitly by

$$\frac{V_0(\{C\})}{c_0} = \frac{1 + g}{y_0 - g} \tag{D2}$$

and

$$\frac{V_0(\{c\})}{c_0} = \frac{1 + g}{r_0 - g}, \tag{D3}$$

which implies that $y_0 > g$ and $r_0 > g$.[4]

The yields to maturity y_0 and r_0 are convenient transformations of the prices obtained by setting the expected growth rates of consumption for each period equal to its unconditional expectation g. Consistent with the standard properties of yields to maturity, if consumption growth

---

[4] Clearly,

$$\frac{V_0(\{C\})}{c_0} = \frac{1 + g}{y_0 - g} = \sum_{t=1}^{\infty} \left(\frac{1 + g}{1 + y_0}\right)^t = \sum_{t=1}^{\infty} \left[\frac{E(c_t/c_{t-1})}{1 + y_0}\right]^t,$$

and similarly for $V_0(\{c\})$.

were independently and identically distributed (IID) and if one-period interest rates were constant, then y_0 would be equal to the one-period interest rate. Moreover, if consumption growth were IID and if dividend-price ratios were constant, then r_0 would be the expected one-period return to consumption equity.

As shown in table 1 below, for the period 1954–2001, the average per capita growth rate of consumption g is 2.3 percent, and the average yield after inflation for long-term government bonds is 3.0 percent. As we shall discuss in the next section, we estimate the consumption risk premium, $r_0 - y_0$, to have a mean of at least 0.2 percent. Combining these numbers gives us an estimate of the marginal cost of all uncertainty of at least

$$\omega_0 = \frac{r_0 - g}{y_0 - g} - 1 = \frac{(0.030 + 0.002) - 0.023}{0.030 - 0.023} - 1 = 28.6\%.$$

As we show below, substantially larger numbers can be obtained under reasonable alternative assumptions. This finding highlights the fact that security markets implicitly attach a very high value to a perpetual bond whose coupons are growing at the average growth rate of per capita consumption. Note that, as the yield y_0 gets close to the growth rate g, this value tends to infinity. It is also clear that the formula for the cost of all uncertainty is very sensitive to potential measurement errors in r_0, y_0, and g.

C.   Cost of Business Cycles

To consider business cycle fluctuations, we define the trend as a one-sided moving average of consumption,

$$C_t = a_0 c_t + a_1(1 + g)c_{t-1} + a_2(1 + g)^2 c_{t-2} + \cdots + a_K(1 + g)^K c_{t-K} \quad \text{(D4)}$$

for a vector of weights $\boldsymbol{a} = (a_0, \ldots, a_K)$ satisfying

$$\sum_{k=0}^{K} a_k = 1. \quad \text{(A2)}$$

Definition D4 and assumptions A1 and A2 imply that

$$E\left(\frac{C_t}{c_0}\right) = (1 + g)^t,$$

so that, in expectation, the trend tracks consumption. We further assume

that interest rates are constant and equal to y (A3) and that the following initial conditions hold:

$$\frac{c_0}{c_{-1}} = \frac{c_{-1}}{c_{-2}} = \cdots = \frac{c_{-K+1}}{c_{-K}} = 1 + g. \tag{A4}$$

The next proposition derives an expression for the marginal cost of business cycles ω_0, as a function of r_0, y, g, and $\boldsymbol{a}$.

PROPOSITION 1. Assume that discount bonds for all maturities and a consumption equity claim are traded. Then, ruling out arbitrage opportunities, and under assumptions A1, A2, A3, and A4, we have

$$\omega_0 = \sum_{t=1}^{\infty} w_{0,t} \sum_{k=0}^{K} a_k \left(\frac{1 + r_0}{1 + y}\right)^{\min\{t,k\}} - 1, \tag{4}$$

where the weights $w_{0,t}$ are defined as

$$w_{0,t} \equiv \frac{r_0 - g}{1 + g}\left(\frac{1 + g}{1 + r_0}\right)^t. \tag{5}$$

The essence of the proof consists of a replication argument like the ones used to price a derivative security, which in our case is the consumption trend. To this effect, we design portfolio strategies, one for each time t, with payoffs that exactly replicate the realizations of the consumption trend C_t. To exactly replicate the payoffs, we use the linearity of the trend consumption and the assumption of constant interest rates, so that portfolios of bonds can be rolled over into the future at known interest rates. The details of the proof are in Appendix A. In this argument, the assumption of constant interest rates can be replaced, with no loss of generality, by the requirement that interest rates are known in advance. Finally, we use the yield to maturity for the consumption equity r_0 and the unconditional growth rate of consumption g to state the formula for the marginal cost ω_0, but we do not assume that either the returns of the consumption equity or the consumption growth rates are IID in this proposition.

Since the expression for (4) is complex, we introduce an approximation for the marginal cost:

$$\omega_0 \cong (r_0 - y) \cdot \sum_{k=0}^{K} a_k k, \tag{6}$$

which is accurate for deviations from trend corresponding to business cycle fluctuations; see Appendix B for a derivation and Section III below for an illustration. Thus the marginal cost of business cycles is approximately equal to the consumption risk premium, a measure of the market price of risk, times a constant that depends on the moving average coefficients, a measure of the volatility of the deviations from trend. For

instance, let us compare the marginal costs ω_0 and ω_0' for two moving average coefficient vectors $\boldsymbol{a} \geq 0$ and $\boldsymbol{a}' \geq 0$, respectively, and assume that $\boldsymbol{a}'$ puts more weight on higher k's, or formally that $\boldsymbol{a}'$ first-order stochastically dominates $\boldsymbol{a}$. If, furthermore, $r_0 > y$, then $\omega_0' > \omega_0$ (this comparative static result holds for the exact expression [4]). The intuition for this result is obvious for the extreme case in which $a_0 = 1$, so that the deviations from trend will be identically zero, and hence $\omega_0 = 0$. Finally, the following limiting case relates the marginal cost of business cycles to the marginal cost of all uncertainty.

PROPOSITION 2. Setting $a_0 = a_1 = \cdots = a_{K-1} = 0$ and $a_K = 1$ and letting K go to infinity, under the assumptions A1–A4, we obtain that

$$\omega_0 = \frac{r_0 - g}{y - g} - 1;$$

that is, the marginal cost of business cycles equals the marginal cost of all uncertainty.

Consider selecting the moving average coefficients $\boldsymbol{a}$ so that the deviations from trend correspond to the conventional view that business cycles last no more than eight years. As described later in the paper, this results in a value of $\sum_{k=0}^{K} a_k k$ of 0.387. On the basis of the estimates presented in the next section for the 1954–2001 period, we conclude that the mean of the consumption risk premium $r_0 - y$ is between 0.2 percent and 1.3 percent. Thus, using equation (6), we estimate the mean of the marginal cost of business cycles ω_0 to be between 0.08 percent and 0.49 percent.

II.   Valuing Consumption Equity

In this section, we present our estimates of the value of a security with payoffs equal to aggregate consumption. We have shown that under the assumption of constant interest rates y, we can compute the marginal cost of business cycles as a simple function of the consumption growth rate g, and the moving average weights defining business cycle fluctuations $\boldsymbol{a}$, once we know the value of consumption equity, with implicit yield to maturity r_0. Valuing consumption equity is nontrivial because this is not a traded security. We use as much as possible a preference-free asset pricing approach to value consumption equity as a function of other asset prices under the assumption of no arbitrage. However, because consumption cannot be completely replicated by existing assets, additional assumptions are needed. The first two estimates for $r_0 - y$ are obtained by adapting the method developed by Cochrane and Saa-Requejo (2000) for the computation of bounds on the price of a security whose payoffs cannot be perfectly replicated by existing assets. The key

of their method is to use the prices of observed portfolios as a reference, together with a restriction on the highest possible Sharpe ratio to infer plausible prices for the unobserved security. In addition to this, we also present estimates based on a parametric model for the stochastic discount factor.

We are interested in finding the price, V_t, of a claim to an infinite sequence of payoffs $\{c_{t+k}\}_{k=1}^{\infty}$. To save on notation and to focus on the main ideas, we start by assuming that the growth rates of the payoffs are IID and that the price-dividend ratios $v_t \equiv V_t/c_t$ are constant; we relax these assumptions later. In the IID case, we focus on the (constant) price of a security with a single payoff $c'/c \equiv c_{t+1}/c_t$, denoted by q_t. The price-dividend ratio for the security that has payoff $\{c_{t+k}/c_t\}_{k=1}^{\infty}$ is then given by $v = q/(1 - q)$.[5] Overall, we shall present three different estimates for q.

We assume that there is an observed set of $J + 1$ reference portfolios with current price vector $\boldsymbol{p}$ and with the payoffs to be received next period given by vector $\boldsymbol{x}$. We assume that there is a risk-free asset among these $J + 1$ reference portfolios. Our first estimate of the price q is denoted by q^*, and it is given by the price of the part of the consumption payoff that is spanned by the reference portfolio $\boldsymbol{x}$. That is, q^* is the price of a claim to $\boldsymbol{b}^{\top}\boldsymbol{x}$, where $c'/c = \boldsymbol{b}^{\top}\boldsymbol{x} + u$ and u is orthogonal to $\boldsymbol{x}$, so it satisfies $E[u\boldsymbol{x}] = 0$. Thus $\boldsymbol{b}^{\top}\boldsymbol{x}$ has the interpretation of the payoff of a portfolio $\boldsymbol{b}$ of the reference assets, and hence its value equals $\boldsymbol{b}^{\top}\boldsymbol{p}$. We assume that the component u is priced as if it were a risk-free asset; that is, it has no risk premium. Since $\boldsymbol{x}$ includes a risk-free asset, it must be that $E[u] = 0$, and hence we have $q^* = \boldsymbol{b}^{\top} \cdot \boldsymbol{p}$.

Now we describe our second estimate of the price q, denoted by $\underline{q}$, which we take to be a lower bound of the price of the consumption strip. For this, we find it useful to introduce the concept of a stochastic discount factor. As is well known, no arbitrage guarantees the existence of a stochastic discount factor $m_{t+1} \geq 0$ that satisfies $p_t = E_t[m_{t+1}x_{t+1}]$ for all prices and payoffs p_t and x_{t+1}. An example of a valid stochastic discount factor in our setup is

$$m_{t+1}(z^{t+1}) = \frac{U_{z^{t+1}}/U_{z^t}}{P(z_{t+1}|z^t)},$$

where P is the probability measure on histories z^t, and the U_{z^t} are the derivatives of U with respect to $c_t(z^t)$. Recall that the stochastic discount

factor m_{t+1} is unique if and only if markets are complete. We define $\underline{q} = E[mc'/c]$, where the discount factor m has been suitably restricted. In particular, we follow Cochrane and Saa-Requejo by restricting the set of stochastic discount factors to be consistent with the prices of the reference payoffs and impose an upper bound on its volatility. Specifically, $\underline{q}$ solves

$$\underline{q} = \min_{m \geq 0} E\left[m\frac{c'}{c}\right]$$

subject to (i) $p = E[m\boldsymbol{x}]$, (ii) $m \geq 0$, and (iii) $\sigma(m)/E(m) \leq h$. Let R and $1 + y$ be any gross return and the gross risk-free rate. Then condition iii limits the Sharpe ratio of any gross return R, defined as $|E[R - (1 + y)]|/\sigma(R)$, to be lower than h. To see this, notice that $E[m[R - (1 + y)]] = 0$, and hence

$$\frac{|E_t[R - (1 + y)]|}{\sigma_t(R)} \leq \frac{\sigma_t(m)}{E_t(m)},$$

with $E(m) = 1/(1 + y)$. Thus $\sigma_t(m)/E_t(m)$ provides an upper bound to the *market price of risk*, that is, the expected excess returns that one can trade off at market prices per unit of risk, as measured by the standard deviation of the returns. In the language of Cochrane and Saa-Requejo, portfolios with large Sharpe ratios are good deals, and hence restriction iii on the discount factors is interpreted to mean that there should be no deals that are "too good."

Cochrane and Saa-Requejo show how the prices q^* and $\underline{q}$ are related. In particular, when it is assumed that the nonnegativity constraint ii is not binding,

$$\underline{q} = q^* - \frac{1}{1 + y}\sqrt{(h^2 - \tilde{h}^2)}\sqrt{(1 - R^2)}\sigma\left(\frac{c'}{c}\right),$$

where R^2 is the R^2 from the regression of c'/c on $\boldsymbol{x}$ and $\tilde{h}$ is the highest Sharpe ratio that can be obtained with the reference assets. Clearly, $\underline{q} \leq q^*$. The difference between q^* and $\bar{q}$ depends on how well c'/c is fitted by the reference assets $\boldsymbol{x}$, as measured by the R^2, and on how far the highest allowable Sharpe ratio h is from the highest Sharpe ratio that is achievable with the reference portfolios $\tilde{h}$. This formula shows that condition iii limits the size of the risk premium that is attributed to u, the part of the payoff c'/c not spanned by $\boldsymbol{x}$. We estimate $\underline{q}$ and q^* by replacing the population moments in the expression by their sample analogues.

We relax the assumptions of IID growth rates for the payoffs and constant price-dividend ratios by considering a setup with a Markov regime-switching process. In particular, we let $z_t = (s_t, \epsilon_t)$ be as follows:

let s_t be a Markov chain with $s \in \{1, 2, \ldots, n\} = S$ and transition function $\pi(s'|s)$, and let $\epsilon_t \in E$ be independent of the history ϵ^{t-1} and with a cumulative distribution function $F(\epsilon|s) = \Pr\{\epsilon_t \leq \epsilon | s_t = s\}$. We let consumption growth rates $c_{t+1}/c_t = 1 + g(z_{t+1})$ and reference payoffs $x_{t+1} = x(z_{t+1})$ be functions of z_{t+1}, whereas the vector of prices of the $J + 1$ reference assets $p_t = p(s_t)$ and the price-dividend ratio $V_t/c_t = v(s_t)$ are functions of s_t. In Appendix C, we define operators whose fixed points give the prices V_t^*/c_t and $\underline{V}_t/c_t$, corresponding, respectively, to the parts of consumption equity spanned by the reference assets and the lower bound of the value of consumption equity. For empirical implementation we consider two non-IID specifications: a two-state regime-switching process and a bivariate vector autoregression (VAR), which we further describe below.

Our third estimate for the price q is based on a parametric model for the stochastic discount factor m_{t+1}. We let $\log m_{t+1}$ be a linear function of aggregate consumption and the market return. This specification is motivated by the Lucas asset pricing model for a utility function with constant relative risk aversion (CRRA), where $\log m_t$ is linear in the log of consumption growth, as well as by the generalization of Epstein and Zin (1991), which allows for a constant intertemporal elasticity of substitution different from the reciprocal of the coefficient of relative risk aversion, where $\log m_{t+1}$ is linear in the log of consumption growth and in the log return of consumption equity. In particular, we assume that m_{t+1} is given by

$$m_{t+1} = \delta \exp(\lambda^\top n_{t+1}), \tag{7}$$

where n_{t+1} is a vector of "factors" with "loading" vector λ and constant δ. Using reference payouts x_{t+1} with prices p_t, we estimate the factor loadings using generalized method of moments on

$$0 = E\left[\exp(\lambda^\top n_{t+1}) \cdot \left[\frac{x_{t+1}}{p_t} - (1 + y)\right]\right].$$

Then, under the assumption that the factors n_{t+1} and the returns x_{t+1}/p_t are IID, we estimate q through the sample analogue to

$$E\left[\exp(\lambda^\top n_{t+1}) \cdot \left[\frac{c_{t+1}/c_t}{q} - (1 + y)\right]\right] = 0.$$

Tables 1–3 contain our estimates of the value of consumption equity for different specifications. Following Cochrane and Saa-Requejo, we have assumed that the highest admissible Sharpe ratio is one in annual terms. As they point out, this is a rather large number, since the observed Sharpe ratio of a market portfolio is about 0.5, and this value is regarded

TABLE 1
MARGINAL COST OF CONSUMPTION FLUCTUATIONS, ω, IID CASE, FOR SELECTED
REFERENCE PORTFOLIOS

| | $r - y$ | | | ω | | | |
| | | | | Business Cycles | | All Uncertainty | |
	% (1)	Upper Bound (%) (2)	R^2 (3)	% (4)	Upper Bound (%) (5)	% (6)	Upper Bound (%) (7)
	A. 1954–2001 ($y=2.96$, $g=2.27$)						
R(Market)	.19	1.17	.17	.07	.44	28.00	169.88
R(10dec)	.27	1.00	.38	.10	.38	38.91	145.08
R(17ind)	.20	.54	.48	.07	.19	28.74	78.13
	B. 1889–2001 ($y=2.15$, $g=1.96$)						
R(Market)	.56	3.30	.24	.21	1.25	219.00	1,722.40
	C. 1927–2001 ($y=2.16$, $g=1.93$)						
R(10dec)	.56	2.22	.44	.21	.84	243.50	958.77
R(17ind)	.44	1.67	.47	.17	.64	190.61	723.77

NOTE.—$r - y$ stands for the consumption risk premium; R^2 comes from the regression of consumption growth on returns. R(Market) stands for the CRSP value-weighted return covering NYSE and AMEX. R(10dec) stands for the returns of the 10 CRSP size-decile portfolios. R(17ind) stands for the returns of the 17 industry portfolios from French's (2002) data. All returns are real.

by the equity premium literature as puzzlingly high. To facilitate the use of the formulas derived in Section I, we express the value of consumption equity in yields to maturity in excess of the risk-free rate, which we call the consumption risk premium; that is, $r_0 - y = (1 + g)/v_0 + g - y$, for both v_0^* and $\underline{v}_0$. Since $\underline{v}_0 \leq v_0^*$, the yield spread attributable to $\underline{v}_0$ determines the *upper* bound of the consumption risk premium.

Table 1 contains estimates of the consumption risk premium under the assumptions of IID consumption growth and returns. We consider three sets of reference portfolios. In addition to a risk-free rate, we use either the Center for Research in Security Prices (CRSP) value-weighted portfolio return covering the New York Stock Exchange (NYSE) and the American Stock Exchange (AMEX), 10 size-decile CRSP portfolios, or 17 industry portfolios constructed by French (2002). Consumption is defined as consumption expenditures on nondurables and services. For the postwar period we find that the consumption risk premium of the spanned part is between 0.19 percent and 0.27 percent, with upper bounds between 0.54 percent and 1.17 percent, depending on the reference portfolios used.[6] The best replication is achieved through the 17 industry portfolios, with an R^2 of .48 for the regression of c'/c onto

---

[6] In computing the lower bound of the price, we do not explicitly impose nonnegativity constraints on the stochastic discount factor. Imposing such constraints would tighten the bound closer toward the price of the spanned component.

TABLE 2
MARGINAL COST OF CONSUMPTION FLUCTUATIONS, ω, NON-IID CASES, WITH THE
PORTFOLIO BEST FITTING CONSUMPTION: 1954–2001

| | $r - y$ | | | ω | | | |
| | | | | Business Cycles | | All Uncertainty | |
	% (1)	Upper Bound (%) (2)	R^2 (3)	% (4)	Upper Bound (%) (5)	% (6)	Upper Bound (%) (7)
	A. VAR(1) for Consumption Growth and Returns						
R(Market)	.11	1.77	.02	.04	.67	16.08	256.84
R(10dec)	.14	1.72	.06	.05	.65	20.94	249.67
R(17ind)	.28	1.47	.48	.11	.56	40.43	213.67
	B. Regime-Switching Process for Consumption Growth and Returns						
R(Market)	.21	1.64	.16	.08	.62	30.93	238.25
R(10dec)	.27	1.35	.38	.10	.51	39.38	195.35
R(17ind)	.18	1.14	.47	.07	.43	26.62	165.82

NOTE.—$r - y$ is the unconditional mean of the consumption risk premium obtained from the model. R(Market) stands for the CRSP value-weighted return covering the NYSE and AMEX. R(10dec) stands for the return of the combination of the 10 CRSP size-decile portfolios that best replicate consumption in a least-squares sense. R(17ind) stands for the return of the combination of 17 industry portfolios from French's (2002) data that best replicate consumption in a least-squares sense. R^2 is the unconditional mean of the one-step-ahead R^2 obtained from the model in a regression of consumption on returns. All returns are real.

real returns. Considering longer sample periods increases the estimated consumption risk premium by about two to three times.

Table 2 reports results when we allow for departures from the IID case. In panel A, we use a Markov chain approximation of a bivariate VAR process with normal innovations consisting of the consumption growth rate and one excess return. We consider bivariate VARs and hence include only one excess return, given the cost to numerically solve for prices q^* and q. We consider three different specifications for the excess returns, which correspond to the three cases considered in table 1. For the two cases that cover several portfolios, that is, the 10 size-decile portfolios and the 17 industry portfolios, we use the combination of these returns that has the highest correlation with consumption. In panel B, we use a two-state Markov regime, where, conditional on the state, consumption growth and the excess return are IID. We consider the same three specifications for the excess returns as in the VAR(1) case. Regimes are assumed to be observable and to be determined by splitting the sample into high and low growth rates of consumption. The cutoff is set at 0.5 percent below the mean annual growth rates in the sample, with the aim to capture the difference between recessions and expansions. We also explored alternative choices for regimes based on the NBER chronology. These results are not reported since they resulted in little quantitative differences. We find that, based on the spanned part, the consumption risk premium is between

TABLE 3
Marginal Cost of Consumption Fluctuations, ω, with Consumption as a Factor

			ω	
Factors	Returns (1)	$r - y$ (%) (2)	Business Cycles (%) (3)	All Uncertainty (%) (4)
		A. 1954–2001		
ΔC	Market	1.11	.42	160.94
ΔC, R(Market)	Market, $R(10) - R(1)$	.21	.08	31.05
		B. 1927–2001		
ΔC	Market	1.60	.61	702.28
ΔC, R(Market)	Market, $R(10) - R(1)$	3.49	1.33	1,535.70

Note.—$r - y$ is the consumption risk premium. R(Market) stands for the CRSP value-weighted return covering the NYSE and AMEX. $R(10)$ and $R(1)$ are the largest and smallest of the 10 CRSP size-decile portfolios. All returns are real. Logarithms are taken of the variables used as factors.

0.11 percent and 0.28 percent and the upper bound is between 1.14 percent and 1.77 percent, depending on whether the VAR or the two-state regime-switching process is used and depending on which excess return is used.[7]

As a summary statistic of our main findings, we average the estimates in tables 1 and 2 for the postwar period, thus obtaining a risk premium of consumption equity of 0.2 percent for the part of consumption spanned by existing assets with an upper good-deal bound of 1.3 percent. While the value of the spanned part of consumption does not correspond to a lower bound according to the good-deal methodology, it seems reasonable to take this estimate as a lower bound because our prior beliefs would not be to attribute a negative risk premium to the part of consumption that is not spanned by the returns in our sample. On the other hand, we consider the upper good-deal bound of 1.3 percent truly as an upper bound for the risk premium of consumption equity. Indeed, while it might be possible to come up with portfolios with large average excess returns that are more strongly correlated with consumption, our choice of a largest admissible Sharpe ratio of one seems generous enough, given that this is about twice what is implied by historical returns of a value-weighted market portfolio. Moreover, explicitly imposing nonnegativity constraints would also tighten the bounds for annual data frequencies.

Table 3 contains estimates of the consumption equity premium under the parametric specification of the stochastic discount factor in (7). We present results for two specifications. In the first row, we use the log

[7] Table 2 does not report results for the longer sample period covering 1927–2001, since this does not result in any significant changes compared to the corresponding IID cases in table 1.

consumption growth rate as the only factor in (7), following the Lucas asset pricing model, and we choose the loading vector λ to fit the excess return of the market portfolio. In the second row, we consider a specification with two factors, the log consumption growth rate and the log market return, and we choose the vector λ to fit the market return and the difference in return between the smallest and largest CRSP size-decile portfolios. Column 2 shows that the consumption risk premium is estimated to be 1.11 percent for the one-factor case and 0.21 percent for the two-factor case. These values are in between the ones estimated by the methods reported in columns 1 and 2 in tables 1 and 2.

A comparison of the results in table 3 to those in table 1 for the case in which the aggregate stock market is used as the reference return provides further insights. In table 3, with consumption growth as the factor, the consumption risk premium, $r - y$, is estimated at 1.11 percent for the postwar period. In table 1, the corresponding estimate of $r - y$ is 0.19 percent. These two estimates are closely related. For the results reported in table 3, we have

$$r - y \cong \frac{1}{\beta_{R^M, c'/c}} \cdot (r^M - y), \tag{8}$$

where $\beta_{R^M, c'/c}$ is the slope coefficient in the regression of R^M on c'/c, $r^M \equiv E(R^M) - 1$ is the expected market return, and $r^M - y$ is the market equity premium. With consumption growth as the factor, the stochastic discount factor given by equation (7) corresponds to the one implied by CRRA utility. Thus, to the extent that the covariance of consumption and the market return is small, as is well known from the equity premium literature, the risk aversion required to explain the equity premium is large, $\beta_{R^M, c'/c}$ is small, and the consumption risk premium $r - y$ is relatively large.

For the corresponding case in table 1, we have

$$r - y \cong \beta_{c'/c, R^M} \cdot (r^M - y), \tag{9}$$

which is just the capital asset pricing model relationship. It is easy to see that unless R^M and c'/c are perfectly correlated, we have $1/\beta_{R^M, c'/c} > \beta_{c'/c, R^M}$, and the ratio of the first to the second is given by $[1/\text{corr}(R^M, c'/c)]^2$. In the case here, $\text{corr}(R^M, c'/c)$ is roughly equal to 0.4, so that $[1/\text{corr}(R^M, c'/c)]^2 \cong 6$, which indeed corresponds approximately to the ratio of the consumption risk premiums of 1.11/0.19. See Appendix D for a derivation of (8) and (9).

We have further explored the sensitivity of our results to five sets of auxiliary assumptions without reporting them here in detail. First, the exact value of the risk-free rate used to estimate the consumption equity premium $r_0 - y$ turns out not to be important. To a first approximation,

our methods just estimate covariance risk. Second, we have considered an alternative timing convention for combining consumption growth rates and returns. For the benchmark case reported here, we have paired consumption growth from year t to $t + 1$ with returns from the first to the last day of year t. Alternatively, we have considered returns from the last day of June in t until the last day of June in $t + 1$. The findings are barely distinguishable across the two cases. Third, we have considered quarterly data for the postwar period 1954–2001. In general, consumption risk premia are somewhat smaller (after annualization) than for the annual results reported here. The robustness of our estimates across specifications and return sets that we have reported for annual data also holds for the quarterly period. Fourth, we have included the return spread between long-term corporate bonds and government bonds from Ibbotson Associates (2002) and found that the results were not sensitive to the addition of these portfolios. Fifth, in the NBER working paper version of this paper (Alvarez and Jermann 2000), we have considered richer specifications of the stochastic discount factor (7), allowing for non-IID returns—including variable interest rates—and consumption growth rates in a multivariate VAR context; results were similar.

III.   Measuring Business Cycles

In this section, we describe the choice of the moving average coefficients $\{a_k\}$ that determine the consumption trend $\{C\}$, as defined in D4 and A2. We define the trend $\{C\}$, so that the deviations of consumption from trend, $c_t - C_t$, are fluctuations that last eight years or less. Thus the trend $\{C\}$ contains fluctuations that last more than eight years. Our definition of business cycles as fluctuations that last up to eight years is consistent with the definition of Burns and Mitchell (1946) and also corresponds approximately to the definition of business cycles implied by the widely used Hodrick-Prescott filter for quarterly data with a smoothing parameter of 1,600.

We choose the moving average coefficients $\{a_k\}$ so as to represent a low-pass filter that lets pass frequencies that correspond to cycles of eight years and more. Low-pass filters are represented in the time domain by infinite-order two-sided moving averages. However, a requirement of our analysis is to have trend consumption in time t be a function of information available at time t; thus our choice of a one-sided moving average. To do this, we follow the approach presented by Baxter and King (1998, 1999). Let $\beta(v)$ be the frequency response function of the desired low-pass filter, which in our case is equal to one for frequencies lower than eight years and zero otherwise. Let $\alpha_K(v)$ be the frequency response function associated with a set of moving average coefficients

$\{a_k\}_{k=0}^{K}$. We select the moving average coefficients $\{a_k\}_{k=0}^{K}$ so that α_K approximates β. In particular, our choice of $\{a_k\}$ minimizes

$$\int_{-\pi}^{\pi} |\beta(v) - \alpha_K(v)|^2 f(v) dv, \tag{10}$$

where $f(v)$ is a weighting function representing (an approximation to) the spectral density of the series to be filtered. In this minimization, we impose the condition $\alpha_K(0) = 1$, which implies that $\sum_{k=0}^{K} a_k = 1$.

We use the spectral density of an AR(1) with an autoregressive coefficient of one as the weighting function f, because this matches approximately the spectral density of consumption. See also Alvarez and Jermann (2002) for another view about how consumption fluctuations are largely permanent. We set the number of lags $K = 20$. In our case, using more coefficients does not significantly affect quantitative results; with fewer coefficients, results are slightly different. The coefficients are given in Appendix E.

The costs of business cycles corresponding to the estimates of consumption risk premiums that we discussed above are presented in tables 1–3. Take, for instance, table 2, the regime-switching case, labeled $R(17\text{ind})$. In this case, the cost of business cycles is 0.07 percent based on the spanned part of consumption as displayed in column 4, with 0.43 percent as an upper-bound estimate, as displayed in column 5.

All results reported in the tables are based on the exact formula derived in proposition 1. We illustrate here the accuracy of the approximation given by equation (6). For instance, for the case just discussed, table 2 shows the consumption risk premium based on the spanned part as 0.18 percent and that based on the good-deal upper bound as 1.14 percent. For $K = 20$, with the optimal filter weights, $\sum_{k=0}^{K} a_k k = 0.387$, so that the approximate cost of business cycles is 0.07 percent based on the spanned part with an approximate upper bound of 0.44 percent.

Following our discussion in the previous section, we summarize the main quantitative results by averaging the estimates of the marginal cost ω_0 based on postwar data presented in tables 1 and 2. We find the cost of business cycles to be between 0.08 percent based on the spanned part of consumption and 0.49 percent based on the upper good-deal bound. As we discuss further below, these conclusions are quite robust to alternative filters and the introduction of durable goods consumption.

A.   *Discussion of One-Sided Filters*

We provide here some discussion about the extent to which our results are robust to the particular filter choice. As a specific requirement of
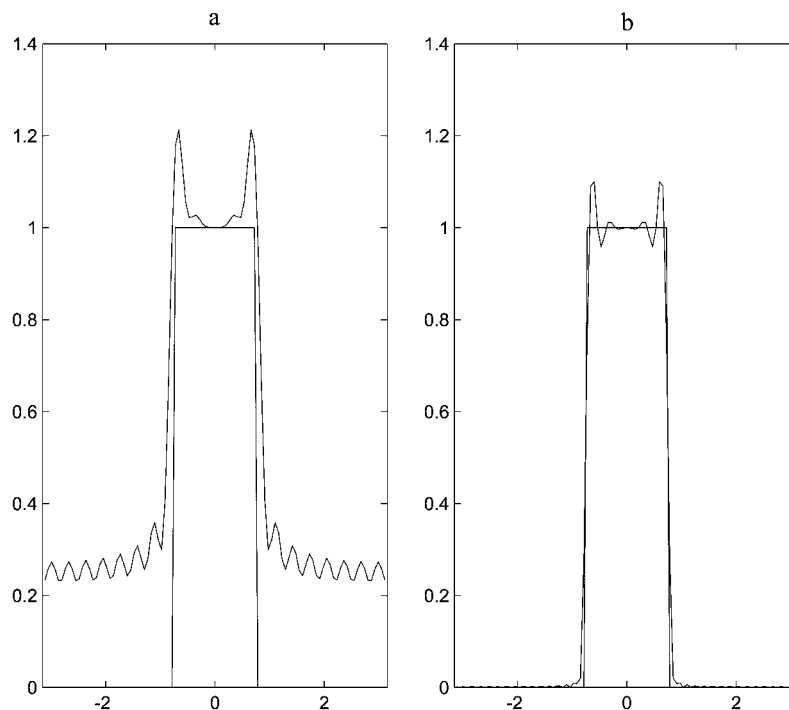
Fig. 1.—Transfer function for low-pass filter with an eight-year cutoff: *a*, one-sided; *b*, two-sided.

our analysis we need a one-sided filter. However, since this filter is one-sided, it cannot avoid introducing a phase shift. As a result, the trend lags the original series. In particular, the objective function displayed in equation (10) can be written as the integral of the square of the differences of the gains of the filters, $[|\beta(v)| - |\alpha_K(v)|]^2$, plus a term that depends on the phase shift. This second term is zero if the filter has no phase shift. Figure 1*a* illustrates this issue by plotting the transfer function (the squared gain) of this filter. The transfer function should be one in between the desired frequencies and zero for higher frequencies. Instead, it tends to let pass up to 30 percent of the variance at higher frequencies, so that the computed trend contains a nonnegligible amount of cyclical variability. As shown in figure 1*b*, and as is well known, two-sided band-pass filters fit the ideal filter's step function much more closely—remember that a symmetric two-sided filter does not introduce a phase shift. The corresponding time domain represen-

tation is in figure 2.[8] Specifically, deviations from trend scaled by a growth factor $(c_t - C_t)/c_0(1 + g)^t$ are shown for one-sided and two-sided filters. Clearly, the one-sided filter generates cyclical movements that are less volatile than those from the corresponding two-sided filter.

On the basis of this comparison, we can consider an ad hoc adjustment to the one-sided filter so as to replicate the amount of business cycle volatility obtained from the more accurate two-sided filter. As shown in figure 2, the series generated by the one-sided filter is strongly correlated with the series from the two-sided filter, but the series generated by the one-sided filter is less volatile. In particular, for the postwar period 1954–2001, the plotted deviations from trend, $(c_t - C_t)/c_0(1 + g)^t$, have standard deviations of 0.55 and 0.65 for the one- and two-sided filters, respectively. We can scale up the volatility of business cycles by multiplying the cyclical deviations by a constant $\theta > 1$, so that the cyclical component is adjusted to become $\theta(c_t - C_t)$.[9] Specifically, with $\theta = 1.2$, the standard deviation of the scaled one-sided filter is about equal to the one from the two-sided filter. A little algebra shows that with this adjustment the approximate cost of business cycles defined in equation (6) is just multiplied by θ, becoming $\theta(r_0 - y) \times \sum_{k=0}^{K} a_k k$. Thus, to the extent that adjusting business cycles obtained from a one-sided filter requires an increase in standard deviation of 20 percent, the cost of business cycles is also increased by a factor of 0.2.

An alternative one-sided filter can be obtained from the two-sided filter by forecasting future values on the basis of available information at the time of the payout. Under the assumption that consumption follows a random walk, this would imply that the sum of all the leading coefficients would be added to a_0, without changing the coefficients corresponding to lagged values of consumption. As can be shown, for our case with $f(\omega)$ the pseudo spectrum of a random walk, this one-sided filter equals the one used in this paper.

Overall, we conclude that possible adjustments to the one-sided filter used in this paper are not likely to result in considerable changes in the cost of business cycles, as long as the definition of business cycles is based on the idea of cyclical movements lasting no more than eight years.

IV.   Durable Goods

In this section we examine the impact of expanding the definition of consumption to include durables in addition to nondurables and ser-

---

[8] Note that for this figure and the corresponding calculations, we use filters with $K = 5$, so as not to lose too many observations. For the period of overlap, the case with $K = 20$ (not shown) results in very similar time-series realizations.

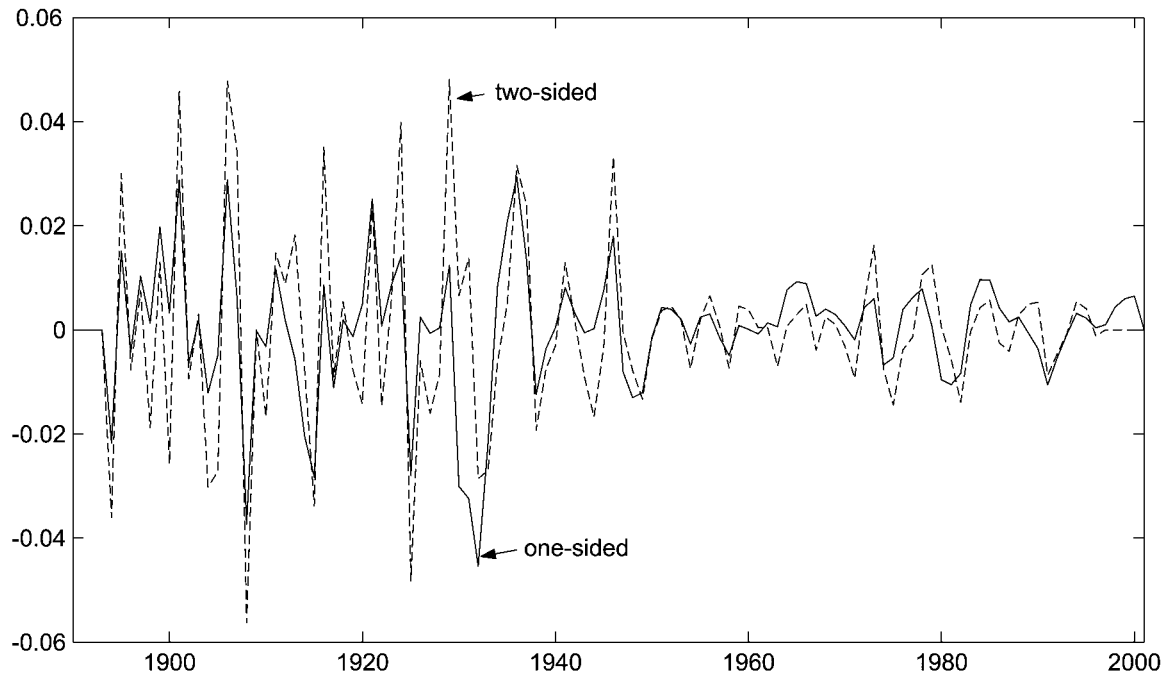[9] Note, in this case, that the trend is given by $(1 - \theta)c_t + \theta C_t$.

Fɪɢ. 2.—Deviations from trend with one- and two-sided filters

vices. We find that stabilizing durable goods consumption creates a sizable gain when measured in percentage terms of this type of consumption goods. However, because the value of the lifetime consumption of durables is so much smaller than for nondurables and services, the overall effect on the marginal cost of business cycles is small.

We derive an expression for the marginal cost of fluctuations that includes both durable consumption goods and nondurable consumption goods and services. We assume that the utility function has nondurables and services, c^{ns}, and durables, c^d, and define the cost of fluctuations Ω as before:

$$U([1 + \Omega(\alpha)]\{c^{ns}\}, \ [1 + \Omega(\alpha)]\{c^d\}) =$$

$$U((1 - \alpha)\{c^{ns}\} + \alpha\{C^{ns}\}, \ (1 - \alpha)\{c^d\} + \alpha\{C^d\}), \tag{11}$$

where C^{ns} and C^d are the trends in consumption of nondurables and services and consumption of durables, respectively. As in the previously discussed case with one type of goods, the marginal cost is obtained by differentiating (11) with respect to α:

$$\Omega'(0) \equiv \bar{\omega}_0 = \frac{\sum_{t \geq 1} \sum_{z^t \in Z^t} \left[\dfrac{\partial U}{\partial c_t^{ns}(z^t)} C_t^{ns}(z^t) + \dfrac{\partial U}{\partial c_t^d(z^t)} C_t^d(z^t) \right]}{\sum_{t \geq 1} \sum_{z^t \in Z^t} \left[\dfrac{\partial U}{\partial c_t^{ns}(z^t)} c_t^{ns}(z^t) + \dfrac{\partial U}{\partial c_t^d(z^t)} c_t^d(z^t) \right]} - 1.$$

This can be written here as

$$\bar{\omega}_0 = \frac{V_0^{ns}(\{C^{ns}\}) + P_0 V_0^d(\{C^d\})}{V_0^{ns}(\{c^{ns}\}) + P_0 V_0^d(\{c^d\})} - 1,$$

where P_0 is the time 0 spot price of durables in terms of nondurables, and V_0^{ns} and V_0^d are the prices to streams of nondurables and services and to durable consumption goods, each in terms of their own time 0 goods' units, respectively, defined as

$$V_0^i(\{x^i\}) = \frac{1}{\partial U/\partial c_0^i} \sum_{t \geq 1} \sum_{z^t \in Z^t} \frac{\partial U}{\partial c_t^i(z^t)} x_t^i(z^t)$$

for $i \in (ns, d)$, $x \in (c, C)$, and

$$P_0 = \frac{\partial U/\partial c_0^d}{\partial U/\partial c_0^{ns}},$$

where the utility function U is evaluated at $\{c^{ns}\}$ and $\{c^d\}$. The expression

for the aggregate marginal cost of fluctuations can be written more compactly as

$$\bar{\omega}_0 = (1 - s_0)\omega_0^{ns} + s_0\omega_0^d, \qquad (12)$$

where $\omega_0^i \equiv V_0^i(\{C^i\})/V_0^i(\{c^i\}) - 1$ for $i \in (ns, d)$, and s_0 denotes the share of the value of the durable consumption equity in aggregate consumption equity, that is,

$$s_0 = \frac{P_0 V_0^d(\{c^d\})}{V_0^{ns}(\{c^{ns}\}) + P_0 V_0^d(\{c^d\})}.$$

In our previous sections we have estimated ω_0^{ns}. Thus our remaining tasks in order to estimate $\bar{\omega}_0$ are to obtain empirical counterparts of ω_0^d and s_0.

We start by describing our estimation of the cost of fluctuations of durable consumption ω_0^d. We distinguish between expenditure on durables and durable consumption. Specifically, we assume that consumption services are provided by the stock of durables, which is assumed to depreciate at a constant rate and to increase by current-period durable expenditures. Then durable consumption, c_t^d, can be represented as a one-sided moving average of current and past expenditures, e_{t-j}, on consumer durables $c_t^d = \sum_{j=0} d_j e_{t-j}$.[10] The value of a claim to lifetime durable consumption is computed in two steps. First, we estimate the value of lifetime durable expenditure the way we did in Section II for the consumption of nondurables and services. Second, following the derivations in proposition 1, we can write the value of lifetime durable consumption as a linear function of the value of lifetime durable expenditure, with the linear coefficient functions of $\{d_j\}$, y, and g. Indeed, this is possible because durable consumption is specified as a one-sided moving average of expenditure, just as the consumption trend has been specified as a one-sided moving average of consumption.

Table 4 reports the estimated price of a claim to durable consumption in terms of durable consumption by using the corresponding yields, $r_0^d - y$, as in tables 1 and 2. The estimated risk premium for durable consumption goods is between 0.45 percent and 1.48 percent based on the spanned part, with upper good-deal bounds between 5.77 percent and 6.49 percent. These values are more than three and seven times higher than the risk premiums estimated for consumption of nondurables and services. The main reason for the increase is the higher volatility of the growth rates of durable expenditure, which have an annual standard deviation of 6.7 percent compared to only 1.16 percent for nondurables and services, for the sample covering 1954–2001. Col-

---

[10] We end up truncating the lags at 10 years for the computations. We found that the truncation lag was not quantitatively important.

TABLE 4
Marginal Cost of Fluctuations in Durable Good Consumption, ω (1954–2001):
$y = 5.08$, $g = 4.34$

| | $r - y$ | | | ω | | | |
| | | | | Business Cycles | | All Uncertainty | |
	% (1)	Upper Bound (%) (2)	R^2 (3)	% (4)	Upper Bound (%) (5)	% (6)	Upper Bound (%) (7)
IID case	1.43	5.77	.31	.53	2.13	194.96	785.42
VAR(1)	.45	6.58	.03	.17	2.43	61.92	896.47
Regime switching	1.48	6.49	.31	.55	2.40	200.90	883.80

Note.—$r - y$ is the consumption risk premium; for the non-IID cases it is the unconditional mean obtained from the model. The return used is the CRSP value-weighted return covering the NYSE and AMEX. R^2 is taken from the regression of growth rates of durable consumption expenditure on returns for the IID case; otherwise, it is the unconditional mean of the one-step-ahead R^2 obtained from the model. All returns are real.

umns 4 and 5 of table 4 display estimates of the business cycle cost ω_0^d using the same weights $\{a_k\}$ as in tables 1 and 2.

We estimate the average of the value share of durable consumption equity in total consumption equity s_0 to be 6 percent and 4.3 percent corresponding, respectively, to the spanned part and the upper-bound estimates from the IID cases in tables 1 and 4. These shares are smaller than the average expenditure share for durable consumption, which for the postwar period is about 13 percent of total consumption expenditure. The reason is that the price-consumption ratios for durables $V_0(\{c^d\})/c_0^d$ are smaller than $V_0(\{c^{ns}\})/c_0^{ns}$, the counterparts for nondurables and services. See Appendix F for more details about the calculation of s_0.

Finally, combining the estimates of ω_0^{ns}, ω_0^d, and s_0 as in equation (12), we can compute an estimate for the aggregate cost of fluctuations including both durable and nondurable consumption goods. For the IID case, we estimate the aggregate cost $\bar{\omega}_0$ to be 0.10 percent based on the prices for the spanned parts; this is higher than the corresponding estimate of $\omega_0^{ns} = 0.07$ percent for nondurables and services in table 1. When the estimates based on the upper bound of $r_0 - y$ are used, the aggregate cost is $\bar{\omega}_0 = 0.51$ percent, compared to the corresponding $\omega_0^{ns} = 0.44$ percent for nondurables and services in table 1. We conclude that adding durable consumption goods does not significantly change our estimates.

V. Comparing Marginal Cost and Total Cost of Consumption Fluctuations

In this section, we present some results about the properties of the marginal cost function that allow us to link our approach more closely

to the large literature that has focused on computing total costs in the line of Lucas (1987). Our main result is a set of conditions under which the marginal cost is an upper bound for the total cost. We also present an example for the cost of all uncertainty with expected, time-separable utility. In this case, we show that the marginal cost equals twice the total cost up to a second-order approximation.[11]

We start this section by comparing our approach to that of Lucas (1987). For that purpose, we define the *total cost of consumption fluctuations* as $\Omega(1)$, that is, $U([1 + \Omega(1)]\{c\}) = U(\{C\})$. Defining the trend consumption to be $\{C\} = \{E_0(c)\}$, that is, where $C(z^t) = E_0(c_t)$ for all t and z^t, we obtain

$$U([1 + \Omega(1)]\{c\}) = U(\{E_0(c)\}), \qquad (13)$$

which is Lucas's definition of the cost of business cycles. Thus Lucas's definition can be seen as the total benefit associated with eliminating all the consumption fluctuations, that is, $\alpha = 1$, and consumption fluctuations are defined as consumption uncertainty, that is, resulting in the exchange of consumption for its expected path.

Note that the specification in equation (13) differs slightly from Lucas's and the literature's standard specification because we have chosen to begin compensation as of $t = 1$; the standard has been to start compensation at $t = 0$. We choose this departure because our definition is more consistent with the idea of ex-dividend security prices, and some of our qualitative results present themselves more tractably with our definition. In any case, the quantitative difference between Lucas's definition and ours should be insignificant.

We provide here also an alternative interpretation of our marginal cost ω_0 that is valid with incomplete markets. For that purpose, assume that for individual agents indexed by i, consumption is given as $c^i = c + \epsilon^i$, where ϵ^i is the idiosyncratic component and $c = C + d$, so that d stands for the deviation from the (aggregate) trend. To save on notation, we omit time subscripts. If we then define Ω as compensating only the aggregate component $\{c\}$, so that

$$U^i(\{[1 + \Omega_i(\alpha)]c + \epsilon^i\}) = U^i((1 - \alpha)\{c^i\} + \alpha\{C + \epsilon^i\}),$$

and if we assume that all agents i have access to claims paying $\{c\}$ and $\{C\}$, we have that

$$\Omega_i'(0) = \frac{V_0[\{C\}]}{V_0[\{c\}]} - 1 = \omega_0.$$

Indeed, under the stated assumptions, even with agents subject to pos-

---

[11] Additional results, for instance, about consumption externalities, are available in our working paper (Alvarez and Jermann 2000).

sibly uninsurable idiosyncratic risk, they would end up equalizing their valuations for $\{c\}$ and $\{C\}$.

A.   *Homothetic Preferences and Scale-Free Cost Functions*

To analyze the marginal cost function, we make the following initial assumptions: $U(\{c\})$ is increasing and concave in $\{c\}$. We also assume that the process $\{C\}$ is preferred to $\{c\}$, that is, $U(\{C\}) > U(\{c\})$. If we require that the cost of fluctuations $\Omega(\alpha)$ be the same for the processes $\{c\}$ and $\{C\}$ as for the processes $\{\lambda c\}$ and $\{\lambda C\}$, where λ is any positive scalar, then we must impose some additional restrictions on the utility function U. This requirement implies that the cost of consumption fluctuations will not differ merely because economies are rich and poor. Specifically, we require U to be homothetic; that is, U is homogeneous of degree $1 - \gamma$. That is, for any positive scalar $\lambda > 0$ and for any process $\{c\}$, we have

$$U(\lambda\{c\}) = \lambda^{1-\gamma}U(\{c\}).$$

Under these assumptions, we obtain that the marginal cost is higher than the total cost.

PROPOSITION 3. Assume that U is increasing, concave, and homothetic. Also assume that $\{C\}$ is preferred to $\{c\}$, that is, $U(\{C\}) > U(\{c\})$. Then $\Omega(\alpha)$ is concave, and thus

$$\omega_0 \equiv \Omega'(0) \geq \Omega(1).$$

Examples from the literature that satisfy this homogeneity property are the preferences used in Mehra and Prescott (1985), Epstein and Zin (1991), Abel (1999), and Tallarini (2000).

B.   *Example: Cost of All Uncertainty with Expected Utility*

Now we present some implications for the total and marginal cost Ω and Ω' with time-separable expected utility. We also assume that the trend $\{C\}$ is given by the expected value of consumption; that is, we evaluate the elimination of all uncertainty. We assume that consumption fluctuations are small. We show that for an approximation up to the order of the variance of consumption, the total cost of uncertainty equals half of the marginal cost; that is, $\Omega(1) = \frac{1}{2}\Omega'(0)$. In this case, the marginal cost is given by a weighted average of the product of risk aversion and the variance of consumption for different periods. We also consider a higher-order approximation to examine the role of skewness in consumption fluctuations. We show that if the period utility function u displays prudence, that is, $u''' > 0$, and if consumption fluctuations have

negative skewness, then we obtain a stronger inequality, that is, $\Omega(1) < \frac{1}{2}\Omega'(0)$.[12]

Consider the one-period case, where consumption is given by $c = \bar{c}(1 + \sigma\epsilon)$ for a zero-mean random variable ϵ. The parameter σ indexes the amount of risk. The trend is given by the expected value, that is, $C = \bar{c} \equiv E[c]$. Notice that the variance of c is proportional to σ^2—that is, $\text{Var}\,(c/\bar{c}) = \sigma^2 E\epsilon^2$—and that its third moment is proportional to σ^3. We include σ as an argument of the total and the marginal costs, which are given by

$$E[u(c[1 + \Omega(1, \sigma)])] \equiv E[u(\bar{c}(1 + \sigma\epsilon)[1 + \Omega(1, \sigma)])] = u(\bar{c}) \quad (14)$$

and

$$\Omega'(0, \sigma^2) = \frac{E[u'(c)(\bar{c} - c)]}{E[u'(c)c]} \equiv \frac{-E[u'(\bar{c} + \bar{c}\sigma\epsilon)(\bar{c}\sigma\epsilon)]}{E[u'(\bar{c} + \bar{c}\sigma\epsilon)(\bar{c} + \bar{c}\sigma\epsilon)]}. \quad (15)$$

PROPOSITION 4. If $E[u''''(\bar{c}(1 + \epsilon))\epsilon^4]$ is finite, then

$$\Omega'(0, \sigma) = 2\Omega(1, \sigma) - \frac{\sigma^3\bar{c}^3 u'''(\bar{c})}{6 \quad u'(\bar{c})}E\epsilon^3 + o(\sigma^3),$$

where $h(\sigma) = f(\sigma) + o(\sigma^p)$ means that $\lim_{\sigma\to0}[h(\sigma) - f(\sigma)]/\sigma^p = 0$.

The proof is standard, and together with additional examples and the multiperiod case, it can be found in our working paper (Alvarez and Jermann 2000).

VI.   Conclusion

The approach developed in this paper allows us to estimate the cost of consumption fluctuations directly from asset prices. Instead of specifying and calibrating a utility function, we use the idea of no arbitrage to compare the value of a claim to lifetime consumption and a claim to stabilized lifetime consumption. Our two main quantitative findings are that the elimination of all consumption uncertainty would be very valuable whereas the elimination of consumption fluctuations at business cycle frequencies is not.

The main reason we find such a large gain from the elimination of all uncertainty is that consumption and the pricing kernel (i.e., the marginal utility of wealth) have large permanent components. The main reason the cost of business cycles is so much smaller is that we define business cycles to comprise only transitory fluctuations, which are small

[12] Rietz (1988) assumes that there is a small probability of a large drop in consumption, motivated by the Great Depression, and he shows that this leads to a substantial increase in the equity premium.

relative to permanent fluctuations. In Alvarez and Jermann (2002), we directly estimate the importance of the permanent component in the pricing kernel, and we indeed find it to be large.

Appendix A

Proofs

Proof of Proposition 1

Start by collecting all the terms in $\{C\}_{t=1}^{\infty}$ that involve a c_t for some arbitrary $t \geq 1$. To do this, consider the dividend paid by the consumption trend asset at times t, $t+1$, ..., $t+K$: $C_t: a_0 c_t + \cdots$, $C_{t+1} = \cdots + a_1(1+g)c_t + \cdots$, and $C_{t+K} = \cdots + a_K(1+g)^K c_t$. Owing to the constant interest rates, we can assign a value to each of the terms that include a c_t through simple replication, so that

$$V_0[C_t] = a_0 V_0[c_t] + \cdots,$$

$$V_0[C_{t+1}] = \frac{a_1(1+g)V[c_t]}{1+y} + \cdots,$$

$$\vdots$$

$$V_0[C_{t+K}] = \cdots + \frac{a_K(1+g)^K V_0[c_t]}{(1+y)^K},$$

where $V_0[c_t]$ is the price at time 0 of a claim to c_t at time t. Clearly, $V_0[\{c_t\}_{t=1}^{\infty}] = \sum_{t=1}^{\infty} V_0[c_t]$. Thus, collecting the terms that have common factor $V_0[c_t]$, we get

$$V_0[c_t]\left[a_0 + a_1\frac{1+g}{1+y} + a_2\left(\frac{1+g}{1+y}\right)^2 + \cdots + a_K\left(\frac{1+g}{1+y}\right)^K\right].$$

There is an expression like this one for each $t \geq 1$. The remaining payoffs at time $t = 1, 2, \ldots, K$ that correspond to consumption values $c_0, c_{-1}, \ldots, c_{-K}$ are grouped in a similar fashion. Rearranging terms and using the assumption that $1 + g = c_0/c_{-1} = c_{2-K}/c_{1-K}$, we get

$$V_0\left[\left\{\frac{C_t}{c_0}\right\}_{t=1}^{\infty}\right] = \left[a_1\frac{1+g}{1+y} + a_2\left(\frac{1+g}{1+y}\right)^2 + \cdots + a_K\left(\frac{1+g}{1+y}\right)^K\right]$$

$$+ \left[a_2\frac{1+g}{1+y} + a_3\left(\frac{1+g}{1+y}\right)^2 + \cdots + a_K\left(\frac{1+g}{1+y}\right)^{K-1}\right]$$

$$+ \cdots + a_K\frac{1+g}{1+y}$$

$$+ V_0\left[\left\{\frac{c_t}{c_0}\right\}_{t=1}^{\infty}\right]\left[a_0 + a_1\frac{1+g}{1+y} + a_2\left(\frac{1+g}{1+y}\right)^2 + \cdots + a_K\left(\frac{1+g}{1+y}\right)^K\right].$$

Equation (4) is derived through the following steps. Using the definition of

ω_0, rearranging terms, and using the definition for r_0, $(r_0 - g)/(1 + g) = c_0/V_0[\{c_t\}_{t=1}^{\infty}]$, gives

$$1 + \omega_0 = a_0 + a_1 \left[\frac{r_0 - g}{1 + g}\left(\frac{1 + g}{1 + y}\right) + \frac{1 + g}{1 + y} \right]$$

$$+ a_2 \left\{ \frac{r_0 - g}{1 + g}\left[\frac{1 + g}{1 + y} + \left(\frac{1 + g}{1 + y}\right)^2\right] + \left(\frac{1 + g}{1 + y}\right)^2 \right\}$$

$$+ \cdots$$

$$+ a_K \left\{ \frac{r_0 - g}{1 + g}\left[\frac{1 + g}{1 + y} + \cdots + \left(\frac{1 + g}{1 + y}\right)^K\right] + \left(\frac{1 + g}{1 + y}\right)^K \right\}.$$

Defining

$$w_{0,t} = \left(\frac{r_0 - g}{1 + g}\right)\left(\frac{1 + g}{1 + r_0}\right)^t$$

and replacing it in the last expression gives (4) after some arrangement. Q.E.D.

Proof of Proposition 2

Assuming that $a_0 = a_1 = \cdots = a_{K-1} = 0$ and $a_K = 1$, we can write the last equation in the proof of proposition 1 as

$$1 + \omega_0 = \sum_{t=1}^{K} w_{0,t}\left(\frac{1 + r_0}{1 + y}\right)^t + \left(\frac{1 + r_0}{1 + y}\right)^K \sum_{t=K+1}^{\infty} w_{0,t}.$$

Take the limit as $K \to \infty$:

$$1 + \lim_{K\to\infty} \omega_0 = \lim_{K\to\infty} \sum_{t=1}^{K} w_{0,t}\left(\frac{1 + r_0}{1 + y}\right)^t = \lim_{K\to\infty} \sum_{t=1}^{K} \left(\frac{r_0 - g}{1 + g}\right)\left(\frac{1 + g}{1 + r_0}\right)^t\left(\frac{1 + r_0}{1 + y}\right)^t$$

$$= \left(\frac{r_0 - g}{1 + g}\right)\lim_{K\to\infty} \sum_{t=1}^{K} \left(\frac{1 + g}{1 + y}\right)^t = \left(\frac{r_0 - g}{1 + g}\right)\left(\frac{1 + g}{1 + y}\right)\frac{1}{1 - [(1 + g)/(1 + y)]}$$

$$= \frac{r_0 - g}{y - g},$$

where we have used that

$$\lim_{K\to\infty} \left(\frac{1 + r_0}{1 + y}\right)^K \sum_{t=K+1}^{\infty} w_{0,t} = \lim_{K\to\infty} \left(\frac{1 + r_0}{1 + y}\right)^K\left(\frac{1 + g}{1 + r_0}\right)^K = \lim_{K\to\infty} \left(\frac{1 + g}{1 + y}\right)^K = 0.$$

Q.E.D.

Proof of Proposition 3

If U is increasing and concave in $\{c\}$, there must exist a utility function v that is homogeneous of degree one, positive, and quasi-concave and satisfies

$$U(\{c\}) = \frac{[v(\{c\})]^{1-\gamma}}{1-\gamma}.$$

To start, we show that $\Omega(\alpha)$ is concave in α. By homogeneity of U,

$$[1+\Omega(\alpha)]^{1-\gamma}\frac{[v(\{c\})]^{1-\gamma}}{1-\gamma} = \frac{[v((1-\alpha)\{c\}+\alpha\{C\})]^{1-\gamma}}{1-\gamma}.$$

Thus, after multiplying by $1-\gamma$, taking the $1/(1-\gamma)$ power, and dividing by $v(\{c\})$ on both sides, we obtain that

$$1+\Omega(\alpha) = \frac{v((1-\alpha)\{c\}+\alpha\{C\})}{v(\{c\})}.$$

Since $v(\cdot)$ is positive, quasi-concave, and homogeneous of degree one, it is concave. With $(1-\alpha)\{c\}+\alpha\{C\}$ linear in α, $v(\cdot)$ is also concave in α; thus $\Omega(\alpha)$ is concave. Now we use the concavity to obtain the desired relationships:

$$\Omega(1) = \Omega(0) + \int_0^1 \Omega'(\alpha)d\alpha \leq \Omega'(0),$$

where the inequality uses $\Omega(0) = 0$, the concavity of Ω, and $\alpha \leq 1$. Q.E.D.

Appendix B

Approximation for the Marginal Cost of Business Cycles

Starting with equation (4) and assuming $a \geq 0$, we obtain the following inequality:

$$
\begin{aligned}
\omega_0 &= \sum_{t=1}^{\infty} w_{0,t} \sum_{k=0}^{K} a_k \left(\frac{1+r_0}{1+y}\right)^{\min\{t,k\}} - 1 \\
&= \sum_{t=1}^{K} w_{0,t} \sum_{k=0}^{K} a_k \left(\frac{1+r_0}{1+y}\right)^{\min\{t,k\}} + \sum_{t=K+1}^{\infty} w_{0,t} \sum_{k=0}^{K} a_k \left(\frac{1+r_0}{1+y}\right)^{k} - 1 \\
&\leq \left(\sum_{t=1}^{K} w_{0,t}\right) \sum_{k=0}^{K} a_k \left(\frac{1+r_0}{1+y}\right)^{k} + \left(1 - \sum_{t=1}^{K} w_{0,t}\right) \sum_{k=a}^{K} a_k \left(\frac{1+r_0}{1+y}\right)^{k} - 1 \\
&= \sum_{k=0}^{K} a_k \left(\frac{1+r_0}{1+y}\right)^{k} - 1, \tag{B1}
\end{aligned}
$$

with equality if $a_0 = 1$ and $a_1 = \cdots = a_K = 0$. Thus, to the extent that not too much weight is given to the a's corresponding to long lags, the inequality is close to an equality. Moreover, using a first-order approximation around $r_0 = y = 0$, we get

$$\sum_{k=0}^{K} a_k \left(\frac{1+r_0}{1+y}\right)^{k} - 1 \cong \sum_{k=0}^{K} a_k[1 + k(r_0 - y)] - 1 = (r_0 - y) \cdot \sum_{k=0}^{K} a_k k.$$

Appendix C

Recursive Pricing Approaches

We present here our recursive approaches to deriving price-dividend ratios v^* and $\underline{v}$. To obtain the price-dividend ratio v^*, we define the operator T^*: $R_+^n \to R_+^n$ given by

$$T^*(v)(s) = \boldsymbol{b}(s)^\top \cdot \boldsymbol{p}(s)$$

for each $s \in S$, where $\boldsymbol{b}(s)^\top \cdot \boldsymbol{x}(z')$ is the linear projection of $[1 + g(z')][1 + v(s')]$ into $\boldsymbol{x}(z')$; that is, it solves

$$[1 + g(z')][1 + v(s')] = \boldsymbol{b}(s)^\top \cdot \boldsymbol{x}(z') + u(z'),$$

$$0 = \sum_{s' \in S} \int x(s', \epsilon') u(s', \epsilon') dF(\epsilon'|s') \pi(s'|s)$$

for each $s \in S$, with u orthogonal to $\boldsymbol{x}$. The price-dividend ratio v^* of the spanned part of the consumption equity is given by the fixed point of T^*:

$$T^*(v^*)(s) = v^*(s).$$

More explicitly, substitute out $\boldsymbol{b}(s)$:

$$\boldsymbol{b}(s) = E_s[\boldsymbol{x}(z')\boldsymbol{x}(z')^\top]^{-1} E_s(\boldsymbol{x}(z')\{[1 + g(z')][1 + v(s')]\})$$

$$= \left[\sum_{s' \in S} \int \boldsymbol{x}(s', \epsilon') \boldsymbol{x}(s', \epsilon')^\top dF(\epsilon'|s') \pi(s'|s) \right]^{-1}$$

$$\times \sum_{s' \in S} \int (\boldsymbol{x}(s', \epsilon')\{[1 + g(z')][1 + v(s')]\}) dF(\epsilon'|s') \pi(s'|s).$$

We now describe a recursion whose fixed point is the price-dividend ratio $\underline{v}_t = \underline{V}_t/c_t$ in the Markov regime-switching setting described above. For this, we let the stochastic discount factor $m_{t+1} = m(\epsilon_{t+1}, s_{t+1})$ be a function of ϵ_{t+1} and s_{t+1} and the price-dividend ratio $v_t = v(s_t)$ be a function of s_t. We define the operator $T: R_+^n \to R_+^n$ as

$$T(v)(s) = \min_{m \in R_+^n} \sum_{s' \in S} \int m(\epsilon', s')[1 + g(\epsilon', s')][1 + v(s')] dF(\epsilon'|s') \cdot \pi(s'|s)$$

subject to

$$\boldsymbol{p}(s) = \sum_{s' \in S} \int [m(\epsilon', s')\boldsymbol{x}(\epsilon', s')] dF(\epsilon'|s') \cdot \pi(s'|s),$$

$$\sum_{s' \in S} \int m(\epsilon', s')^2 dF(\epsilon'|s') \cdot \pi(s'|s) \le \frac{h(s)^2 + 1}{(1 + y)^2},$$

where $h(s)$ is the bound on the conditional Sharpe ratio. The lower good-deal bound for the price-dividend ratio of the consumption equity is the fixed point of this operator, that is, $T(\underline{v})(s) = \underline{v}(s)$ for all $s \in S$.

Appendix D

Approximate Consumption Risk Premium

For equation (8), assuming lognormality for c'/c and R^M in the IID case, using $E[(c'/c)^\lambda[R^M - (1 + y)] = 0$ and the corresponding Euler equation for the return to the consumption claim, implies that

$$\log\frac{1 + r}{1 + y} = \left(\frac{\sigma^2_{\log c'/c}}{\sigma_{\log c'/c, \log R^M}}\right)\log\frac{1 + r^M}{1 + y}.$$

Using the approximation

$$\frac{\sigma^2_{\log c'/c}}{\sigma_{\log c'/c, \log R^M}} \cong \frac{\sigma^2_{c'/c}}{\sigma_{c'/c, R^M}},$$

and a first-order approximation around zero, r, y, and r^M gives the postulated expression. For equation (9), some algebra implies that

$$r - y = (1 + y)\left\{\frac{1}{1 - (\sigma_{c'/c, R^M}/\sigma^2_{R^M})(r^M - y)[1/(1 + g)]} - 1\right\},$$

and a first-order approximation around zero, g, y, and r^M gives the postulated expression.

Appendix E

Filter Coefficients

$$a = [0.6250\ 0.2251\ 0.1592\ 0.0750\ -0.0000\ -0.0450\ -0.0531$$
$$-0.0322\ 0.0000\ 0.0250\ 0.0319\ 0.0205\ -0.0000\ -0.0173$$
$$-0.0228\ -0.0150\ 0.0000\ 0.0133\ 0.0177\ 0.0119\ -0.0191].$$

Appendix F

Durable Consumption Shares

Rearranging the expression in the text gives

$$s_0 = \left(\frac{\frac{\partial U/\partial c_0^d}{\partial U/\partial c_0^{ns}}c_0^d}{c_0^{ns}}\right)\frac{V_0^d(\{c^d\})}{c_0^d}\bigg/\left[\frac{V_0^{ns}(\{c^{ns}\})}{c_0^{ns}} + \left(\frac{\frac{\partial U/\partial c_0^d}{\partial U/\partial c_0^{ns}}c_0^d}{c_0^{ns}}\right)\frac{V_0^d(\{c^d\})}{c_0^d}\right]$$

for the share of value of durable consumption equity to aggregate consumption equity. The following steps explain how we find an empirical counterpart to s_0. Tables 1 and 2 provide estimates for $V_0^{ns}(\{c^{ns}\})/c_0^{ns}$. In the text we describe how to estimate $V_0^d(\{c^d\})/c_0^d$, which is implemented in table 4. Thus the remaining task is to estimate $[(\partial U/\partial c_0^d)/(\partial U/\partial c_0^{ns})](c_0^d/c_0^{ns})$, which is the ratio of the value of con-

sumption of durables to the value of consumption of nondurables and services. To do this, we assume that the stock of durables evolves as

$$c_t^d = c_{t-1}^d (1 - \delta) + e_t, \tag{F1}$$

where δ is the depreciation rate. Rearranging the equation, we get

$$c_t^d = e_t \left[\frac{1}{1 - (1 - \delta)/(1 + g_t^d)} \right],$$

with $1 + g_t^d \equiv c_t^d / c_{t-1}^d$.

In this setting, the per period user cost of the stock of durables, that is, the cost of having one more unit of durables for one period, measured in units of the stock of durable goods, is $(y^d + \delta)/(1 + y^d)$, where y^d is the interest rate of durable goods. A consumer's first-order condition for the choice of durables versus nondurables is

$$\frac{\partial U/\partial c_t^d(z^t)}{\partial U/\partial c_t^{ns}(z^t)} = P_t^{de}(z^t) \frac{y^d + \delta}{1 + y^d}, \tag{F2}$$

where $P_t^{de}(z^t)$ is the price of durable expenditure goods relative to nondurable goods. The quantity $P_t^{de}(z^t)(\delta + r)/(1 + y^d)$ is the relative price of one durable in period t in terms of period t nondurable goods. Multiplying (F2) by $c_t^d(z^t)/c_t^{ns}(z^t)$ and substituting $c_t^d(z^t)$ in terms of expenditures $e_t(z^t)$, depreciation rate δ, and growth rate of durable consumption g_t^d, we obtain

$$\frac{\partial U/\partial c_t^d(z^t)}{\partial U/\partial c_t^{ns}(z^t)} \frac{c_t^d(z^t)}{c_t^{ns}(z^t)} = \frac{P_t^{de}(z^t)e_t(z^t)}{c_t^{ns}(z^t)} \frac{(y^d + \delta)/(1 + y^d)}{1 - (1 - \delta)/(1 + g_t^d)}.$$

We generate a series for g_t^d using durable goods expenditures from the National Income and Product Accounts (NIPA) starting from a level that gives us the same average growth rate over the sample as for expenditures. For the ratio of the expenditure of durables to the expenditure share of nondurables and services, $P_t^{de}(z^t)e_t(z^t)/c_t^{ns}(z^t)$, we generate a series from the NIPA counterpart covering the whole period. The average of this series is 0.15, corresponding to a durable expenditure share of $.13 = .15/(1 + .15)$. On the basis of depreciation rates published by the Bureau of Economic Analysis, we choose a constant annual depreciation rate of 17.5 percent. Note that the main components of the bureau's reported durable good expenditures based on the first quarter of 2001 are motor vehicles and parts (about 43 percent) and furniture and household equipment (about 37 percent). Combining the series for

$$\frac{\partial U/\partial c_t^d(z^t)}{\partial U/\partial c_t^{ns}(z^t)} \frac{c_t^d(z^t)}{c_t^{ns}(z^t)}$$

with the price-dividend ratios in tables 1 and 4 for the IID cases, we report the sample average for s_0. Note that the interest rate in durables, $y^d = 5.08$ percent, is estimated as the sample average of the nominal interest rate minus the inflation of durable good prices; and the growth rate of durables stocks, $g = 4.34$ percent, is taken to be the average growth rate of durables expenditures.

References

Abel, Andrew B. 1999. "Risk Premia and Term Premia in General Equilibrium." *J. Monetary Econ.* 43 (February): 3–33.

Alvarez, Fernando, and Urban J. Jermann. 2000. "Using Asset Prices to Measure the Cost of Business Cycles." Working Paper no. 7978 (October). Cambridge, Mass.: NBER.

———. 2002. "Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth." Working paper. Philadelphia: Univ. Pennsylvania, Wharton School. http://finance.wharton.upenn.edu/~jermann/research.html.

Atkeson, Andrew, and Christopher Phelan. 1994. "Reconsidering the Costs of Business Cycles with Incomplete Markets." In *NBER Macroeconomics Annual 1994*, edited by Stanley Fischer and Julio J. Rotemberg. Cambridge, Mass.: MIT Press.

Baxter, Marianne, and Robert G. King. 1998. "Band-Pass Filters for Financial and Current Analysis." Manuscript. Charlottesville: Univ. Virginia.

———. 1999. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *Rev. Econ. and Statis.* 81 (November): 575–93.

Burns, Arthur F., and Wesley C. Mitchell. 1946. *Measuring Business Cycles.* New York: NBER.

Campbell, John Y., and John H. Cochrane. 1995. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." Working Paper no. 4995 (January). Cambridge, Mass.: NBER.

Cochrane, John H., and Jesus Saa-Requejo. 2000. "Beyond Arbitrage: Good-Deal Asset Price Bounds in Incomplete Markets." *J.P.E.* 108 (February): 79–119.

Dolmas, Jim. 1998. "Risk Preferences and the Welfare Cost of Business Cycles." *Rev. Econ. Dynamics* 1 (July): 646–76.

Epstein, Larry G., and Stanley E. Zin. 1991. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis." *J.P.E.* 99 (April): 263–86.

French, Ken. 2002. Data on Industry Portfolios. http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

Hansen, Lars Peter, Thomas J. Sargent, and Thomas D. Tallarini, Jr. 1999. "Robust Permanent Income and Pricing." *Rev. Econ. Studies* 66 (October): 873–928.

Ibbotson Associates. 2002. *Stocks, Bonds, Bills, and Inflation Yearbook.* Chicago: R. G. Ibbotson Assoc.

Krusell, Per, and Anthony A. Smith, Jr. 1999. "On the Welfare Effects of Eliminating Business Cycles." *Rev. Econ. Dynamics* 2 (January): 245–72.

Lewis, Karen K. 1996. "Consumption, Stock Returns, and the Gains from International Risk-Sharing." Working Paper no. 5410 (January). Cambridge, Mass.: NBER.

Lucas, Robert E., Jr. 1987. *Models of Business Cycles.* New York: Blackwell.

———. 2003. "Macroeconomic Priorities." *A.E.R.* 93 (March): 1–14.

Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *J. Monetary Econ.* 15 (March): 145–61.

Obstfeld, Maurice. 1994. "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability." *European Econ. Rev.* 38 (August): 1471–86.

Otrok, Christopher. 2001. "Spectral Welfare Cost Functions." *Internat. Econ. Rev.* 42 (May): 345–67.

Rietz, Thomas A. 1988. "The Equity Risk Premium: A Solution." *J. Monetary Econ.* 22 (July): 117–31.

Tallarini, Thomas D., Jr. 2000. "Risk-Sensitive Real Business Cycles." *J. Monetary Econ.* 45 (June): 507–32.
van Wincoop, Eric. 1999. "How Big Are Potential Welfare Gains from International Risksharing?" *J. Internat. Econ.* 47 (February): 109–35.