THE JOURNAL OF FINANCE • VOL. LIV, NO. 2 • APRIL 1999

Can Costs of Consumption Adjustment Explain Asset Pricing Puzzles?

DAVID A. MARSHALL and NAYAN G. PAREKH*

ABSTRACT

We investigate Grossman and Laroque's (1990) conjecture that costs of adjusting consumption can account, in part, for the empirical failure of the consumption-based capital asset pricing model (CCAPM). We incorporate small fixed costs of consumption adjustment into a CCAPM with heterogeneous agents. We find that undetectably small consumption adjustment costs can account for much of the discrepancy between the observed variance of nondurable aggregate consumption growth and the predictions of the CCAPM, and can partially reconcile nondurable consumption data with the observed equity premium. We conclude that the CCAPM's implications are nonrobust to extremely small adjustment costs.

THE CONSUMPTION-BASED CAPITAL ASSET PRICING MODELS (CCAPMs)[1] of Lucas (1978), Breeden (1979), and Grossman and Shiller (1982) have difficulty matching the high volatility of equity returns and the high mean equity premium found in U.S. data. First, the variance of the CCAPM asset pricing kernel is too low to generate plausible equity-return volatility. More precisely, Hansen and Jagannathan (1991) and Cochrane and Hansen (1992) show that, for any conjectured pricing kernel mean, the variance of the kernel is too low to satisfy the Hansen–Jagannathan bounds without implausibly high risk aversion or substantial habit formation.[2] Second, the covariance between the CCAPM asset pricing kernel and excess equity returns is too low to generate a plausible equity premium. More precisely, let us define EP' as follows:

$$EP' \equiv \frac{-\text{cov}(m_t, r_t^e)}{E(m_t)}, \tag{1}$$

* Marshall is with the Federal Reserve Bank of Chicago. Parekh is with Bear, Stearns & Co., Inc. An earlier version of this paper was circulated under the title, "The Effect of Costly Consumption Adjustment on Asset Price Volatility." Discussions with Larry Christiano, Kent Daniel, Bob McDonald, and Lew Segal, and the comments of an anonymous referee are gratefully acknowledged. Research assistance was provided by Kristin Stanton and Denise Duffy. The views in this paper are those of the authors, and do not necessarily reflect positions of the Federal Reserve Bank of Chicago, the Federal Reserve System, or Bear, Stearns & Co., Inc. All errors remain the responsibility of the authors.

[1] By "CCAPM" we refer to models where the asset pricing kernel (or stochastic discount factor) is a function of aggregate consumption.

[2] For example, in Grossman, Melino, and Shiller's (1987) estimation of the CCAPM, the two best-fitting specifications imply coefficients of relative risk aversion equal to 154 and 185, respectively.

where r_t^e denotes the excess equity return and m_t denotes the CCAPM pricing kernel. The variable EP' which we call the "theoretical equity premium," is the mean equity premium that would obtain if the CCAPM were true. As noted by Mehra and Prescott (1985), the covariance between m_t and r_t^e is very low (again, unless high risk aversion or habit formation is assumed). As a result, estimated values of EP' tend to be far lower than the equity premium actually observed in financial markets.

In this paper, we ask whether the observed behavior of asset returns can be reconciled with the statistical properties of aggregate consumption data *without* imposing extreme curvature assumptions on preferences. One approach (explored by Luttmer (1996) and He and Modest (1995)) is to impose adjustment costs on portfolio transactions. Our approach is somewhat in the same spirit. Following a conjecture of Grossman and Laroque (1990), we ask whether the empirical failure of the CCAPM may be due to nonconvex adjustment costs in consumption transactions.

To see why consumption adjustment costs can affect the implications of the CCAPM, note that the CCAPM has the following three properties:

1. The asset pricing kernel equals investors' intertemporal marginal rate of substitution (IMRS) in wealth.
2. The IMRS in wealth equals the IMRS in consumption.
3. The IMRS in consumption is a function of the aggregate consumption process.

The first of these properties is shared by all modern asset-pricing theories. The second is the *envelope property*, which holds only if wealth can be converted into consumption one-for-one. Consumption adjustment costs drive a wedge between the IMRS in wealth and the IMRS in consumption, invalidating the envelope property. The third property is the *aggregation property*. With fixed adjustment costs, aggregation fails: aggregate consumption does not resemble the optimal consumption path of any individual agent. Individual consumption is characterized by infrequent discrete jumps; aggregate consumption is smooth.

These considerations suggest that the CCAPM may be sensitive to fixed costs of consumption adjustment. However, it is not known whether *plausible* adjustment costs can reconcile the high variance of asset returns with the low variance of aggregate consumption growth or resolve the equity premium puzzle. We address this question directly by introducing fixed consumption adjustment costs into a simple consumption-based model with heterogeneous agents. We then compute the means and variances of asset returns and aggregate consumption growth implied by the model, as well as the theoretical equity premium defined in equation (1), and compare these statistics with estimates from postwar U.S. data.

What constitutes a "plausible" adjustment cost? The answer depends critically on what we mean by "consumption." Grossman and Laroque (1990) and Reider (1993) interpret consumption as the service flow from a stock of durable goods or housing. The cost of consumption adjustment is the cost of changing this stock. Under this interpretation, high adjustment costs are plausible. Consider, for example, the closing costs in a real estate transaction or the wholesale/retail spread paid when trading in a used car.

Alternatively, one could follow most of the empirical literature on consumption-based asset pricing, and interpret consumption as expenditures on nondurable goods. This interpretation poses an obvious difficulty for our purpose because there are no major costs to changing the flow rate of nondurable consumption. However, we argue that costs of adjusting nondurables are plausible provided the costs are undetectably small (i.e., costs so small that they would be swamped by simple measurement error in collecting aggregate consumption data). These costs might include the time spent implementing a new consumption/savings plan, search costs in finding vendors for the new, better (or worse) quality goods to be purchased, or even the psychological effort of consumption reoptimization. (This last category of adjustment costs is similar to what Cochrane (1989) calls "near-rational behavior.") Marshall (1994) provides evidence that undetectably small adjustment costs are sufficient to dramatically change the relative variances of consumption growth and asset returns. Intuitively, there is a relatively small utility gain from period-by-period adjustment of consumption in response to movements in asset returns. Even a small fixed cost of consumption adjustment would be sufficient to induce agents to forgo this small utility gain, and to change consumption only infrequently. When we evaluate our adjustment cost model, we will consider both the "durables" and the "nondurables" interpretation for consumption.

In this paper, we follow Constantinides (1990) and take the returns process as given. We calibrate this process to observed asset return data and compute the optimal consumption decisions given this returns process. We then ask whether the implied aggregate consumption process is sufficiently smooth to match observed consumption data, and whether the theoretical equity premium implied by this consumption process is sufficiently low to replicate the theoretical equity premium computed from the data.

Our results confirm that a very small fixed cost of consumption adjustment sharply reduces both the variance of aggregate consumption growth and the theoretical equity premium. In our benchmark case, adding an adjustment cost equal to 0.002 of current consumption reduces the standard deviation of aggregate monthly consumption growth by almost 35 percent, and reduces the theoretical equity premium by about 80 percent. This can account for almost 60 percent of the discrepancy between the standard deviation of observed nondurable consumption growth and of consumption growth as implied by the CCAPM. It explains more than 80 percent of the discrepancy between the actual equity premium observed in asset markets and the theoretical equity premium computed from aggregate consumption data. However, we find that even very large adjustment costs cannot reduce the variance of consumption growth sufficiently to match data measuring the service flow from the stock of durables or housing.

Our paper is related to recent studies of equilibrium asset pricing with heterogeneous agents; see, for example, Marcet and Singleton (1991), Heaton and Lucas (1996), Telmer (1993), Constantinides and Duffie (1996), and Krusell and Smith (1997). In these papers optimal risk sharing fails because idiosyncratic income shocks are uninsurable. In our model, the fixed cost of consumption adjustment precludes optimal risk sharing. In Section III.B, below, we interpret this failure as a particular form of market incompleteness.

Our paper extends the work of Lynch (1996), who imposes nonsynchronous consumption adjustment as an exogenous feature of the economy. His model has similar qualitative implications to our paper. In particular, as the time between consumption adjustments increases, the variance of aggregate consumption decreases, and the covariance between aggregate consumption and equity returns (which is at the heart of the "equity premium" puzzle) becomes more realistic. Lynch does not formalize the economic frictions that induce this nonsynchronous consumption. Rather, he assumes a fixed, nonstochastic time interval between individual consumption adjustments. In our model, nonsynchronous consumption emerges as the optimal strategy of rational agents facing frictions that are explicitly modeled. We find that the interval between consumption adjustments is endogenous and stochastic, varying both cross-sectionally and through time. Furthermore, our approach allows us to address the key question of interest: whether realistic frictions can generate sufficient nonsynchronism to match the data on asset returns and aggregate consumption.

The remainder of the paper is organized as follows. Section I presents a simple two-period example that illustrates how small adjustment costs can induce a large region of consumption inactivity. Section II introduces the model, which incorporates the consumption/portfolio problem of Grossman and Laroque (1990) into a heterogeneous-agent economy. In Section III, we characterize the individual consumption decision in our model and discuss its implications for aggregate consumption. In Section IV we compare the implications of the model with postwar U.S. data. Our conclusions are summarized in Section V.

I. The Impact of Small Adjustment Costs: A Simple Example

To see why very small consumption adjustment costs can induce large changes in individual decisions, consider the following two-period deterministic model. An agent with one unit of wealth maximizes $U(c_1,c_2) = \log(c_1) + \log(c_2)$, where c_i = consumption in period i. If the gross risk-free rate is unity, the budget constraint is $c_1 + c_2 \leq 1$. The optimal choice is to set $(c_1,c_2) = (0.5,0.5)$. Let function $D(c) \geq 0$ denote the cost, in units of wealth, of setting $(c_1,c_2) = (c,1-c)$, instead of choosing the optimum. Formally, $D(c)$ satisfies $U(0.5[1 - D(c)],0.5[1 - D(c)]) \equiv U(c,1 - c)$. If an agent is constrained to set $(c_1,c_2) = (c,1-c)$, $D(c)$ represents the maximum amount the agent will pay to lift this constraint and be permitted to optimize freely.

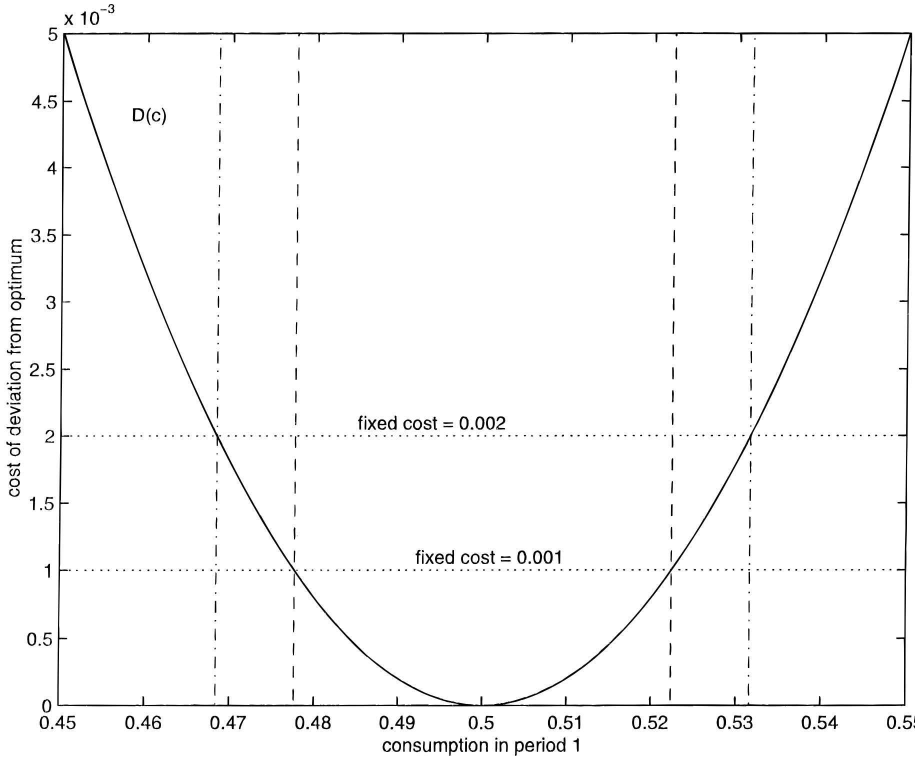

Figure 1. Simple example of an inactivity region induced by adjustment costs. An agent with unit wealth facing a unit gross interest rate maximizes $U(c_1,c_2) \equiv \log(c_1) + \log(c_2)$, subject to $c_1 + c_2 = 1$. The optimal choice is to set $(c_1,c_2) = (0.5,0.5)$. $D(c)$ (solid line) measures the disutility (in units of wealth) associated with setting $(c_1,c_2) = (c,1 - c)$. Formally, $D(c)$ satisfies $\log(c) + \log(1 - c) \equiv 2\log(0.5[1 - D(c)])$. In the figure, $D(c)$ is plotted against c. The region [0.478,0.522] bounded by the vertical dashed lines gives the region of inactivity when the fixed cost of adjusting c_1 equals 0.001. The region [0.468,0.532] bounded by the vertical dot/dash lines gives the region of inactivity when the fixed cost of adjusting c_1 equals 0.002.

The function $D(c)$ is plotted in Figure 1. The convex shape of D reflects the concavity of the utility function. Suppose an agent starts at some point $c \neq 0.5$, and there is a small fixed cost of adjusting c_1 away from c. How big does the fixed cost have to be to preclude adjustment? Clearly, if the fixed cost exceeds $D(c)$, the agent refrains from adjusting, since the cost of adjustment exceeds the utility gain (in consumption-equivalent terms) from moving to the (frictionless) optimum. Figure 1 illustrates this for a fixed cost equal to 0.001 (0.1 percent of wealth). According to the figure, this fixed cost is sufficient to preclude adjustment as long as c is in the interval [0.478,0.522] (indicated by the vertical dashed lines). We refer to this interval as the region of inactivity. Notice that this region is 45 times larger than the size of the adjustment cost itself.

This illustrates a general principle: Very small fixed costs of changing a decision variable can induce inactivity in a large region of the state space. The reason is that a differentiable objective function is nearly flat in the neighborhood of its maximum, so the value lost by deviating from the maximizing argument is quite small. Turning this logic around, if there were a very small cost of achieving the maximum, it may be optimal for agents to choose a value of the decision variable quite far from the value that maximizes the objective function. If the objective function is concave (as in this example), the objective function deteriorates at an accelerating rate as one moves further from the optimum. As a result, the region of inactivity grows at a decreasing rate as the fixed cost increases. This can be seen by the vertical dot/dash lines in Figure 1, which show that when the fixed cost is doubled to 0.002, the region of inactivity expands only 41 percent to [0.468,0.532].

II. A Dynamic Model with Fixed Costs of Consumption Adjustment

The static model of Section I illustrates how small costs of attaining the theoretical optimum can substantially distort agents' decisions. The model we now present extends this insight to a dynamic stochastic economy. As in the model of Section I, a fixed cost of consumption adjustment induces a region of inactivity in which consumption remains unchanged. As a result, individuals adjust consumption infrequently, and both the stochastic properties of aggregate consumption growth and the comovements between aggregate consumption growth and the return on risky assets are thereby altered.

A. Basic Structure

Time is continuous. There are N infinitely lived agents, who differ only in their initial consumption/wealth ratio. As in Merton (1971), each agent has access to a risk-free asset and a risky asset. The risk-free asset pays an instantaneous rate of return, r, which is constant through time. Let b_t denote the value of one unit of the risky asset at date t. It is assumed that b_t follows a geometric Brownian motion process:

$$db_t = b_t[(\mu + r)\,dt + \sigma\,dw_t], \tag{2}$$

where μ and σ are mean excess return and the standard deviation of the risky return respectively, and w_t is a standard Wiener process. Equation (2) implies that the return to holding the risky asset over any fixed time interval is Gaussian white noise. Let X_t^j denote the value of the jth agent's holdings of the risky asset, and let B_t^j denote the value of the jth agent's holdings of the risk-free asset. Agent j's total wealth, Q_t^j, is given by

$$Q_t^j \equiv B_t^j + X_t^j. \tag{3}$$

If consumption is not adjusted at date t, the law of motion for Q_t^j follows a standard wealth accumulation equation:

$$dQ_t^j = Q_t^j r\,dt + X_t^j(\mu\,dt + \sigma\,dw_t) - c_t^j\,dt, \tag{4}$$

where c_t^j denotes the consumption of agent j at date t. Changing the rate of consumption involves a fixed cost, which is proportional to the previous rate of consumption. Specifically, if the consumption rate immediately before a consumption change at date t for agent j is denoted $c_{t^-}^j$, the cost at date t of changing consumption is $\lambda c_{t^-}^j$, where $0 \le \lambda \le 1$. If consumption is adjusted at date τ,

$$Q_\tau = Q_{\tau^-} - \lambda c_{\tau^-}. \tag{5}$$

Finally, agents face a no-bankruptcy condition:

$$Q_t^j - \lambda c_t^j \ge 0. \tag{6}$$

Agents have time-separable preferences with constant relative risk aversion. That is, each agent solves the following optimal control problem:

$$\max_{X_t, c_t} E \int_0^\infty e^{-\delta t} \frac{c_t^a}{a}\,dt, \qquad \delta > 0, a < 1, a \ne 0, \tag{7}$$

subject to constraints equations (3), (4), (5), and (6), and initial conditions $\{Q_{0^-}^j, c_{0^-}^j\}$. Parameters δ and a are assumed to be the same for each agent.

B. Optimality Conditions

When $\lambda = 0$ (no adjustment costs) the optimization problem corresponds to Merton (1971). Optimal consumption and asset holdings are proportional to wealth:

$$X_t^j = \left[\frac{1}{1-a}\right]\frac{\mu}{\sigma^2} Q_t^j, \tag{8}$$

$$c_t^j = \left[\frac{\beta}{1-a}\right] Q_t^j, \tag{9}$$

where

$$\beta \equiv \delta - ar - \frac{1}{2}\left[\frac{a}{1-a}\right]\frac{\mu^2}{\sigma^2}, \tag{10}$$

and the value function (defined as the maximized value of equation (7), given initial wealth Q) is

$$V_{\lambda=0}(Q) = \frac{1}{a}\left[\frac{\beta}{1-a}\right]^{a-1} Q^a. \tag{11}$$

When $\lambda > 0$, the problem corresponds to that studied by Grossman and Laroque (1990). The derivation of the optimality conditions is provided in the Appendix. Here we summarize the optimality conditions. Let $V(Q,c)$ denote the maximized value of the objective in equation (7), starting from initial conditions (Q,c). V is homogeneous of degree a, so the dimensionality of the state space can be reduced. Define:

$$y \equiv \frac{Q}{c} - \lambda,$$

$$h(y) \equiv c^{-a}V(Q,c) \equiv V(\lambda + y,1),$$

$$M \equiv \sup_y (y + \lambda)^{-a} h(y), \tag{12}$$

(where the superscripts and subscripts have been suppressed for simplicity). Following Grossman and Laroque (1990), it can be proved that the transformed value function $h(y)$ satisfies the differential equation

$$-\frac{1}{2}\frac{h'(y)^2}{h''(y)}\frac{\mu^2}{\sigma^2} + h'(y)[r(y+\lambda) - 1] - \delta h(y) + \frac{1}{a} = 0. \tag{13}$$

The optimal consumption strategy is determined by three points $y_1 \le y^* \le y_2$, satisfying

$$h(y_i) = y_i^a M; \qquad h'(y_i) = a y_i^{a-1} M, \qquad i = 1,2, \tag{14}$$

and $y^* \equiv \text{argmax}_y (y + \lambda) h(y^a)$. The interval (y_1,y_2) is the region of inactivity for consumers (analogous to the region of inactivity in the static model of Section I). When $y \in (y_1,y_2)$, it is optimal not to adjust consumption. Whenever $y = y_1$ or $y = y_2$ (i.e., the wealth/consumption ratio is sufficiently low or sufficiently high), it is optimal to adjust consumption so that $y = y^*$, in which case

$$c_t = \frac{Q_t}{y^* + \lambda}. \tag{15}$$

We refer to this pattern as a "(y_1,y_2,y^*) policy." Intuitively, if $y = y_2$ the wealth-to-consumption ratio is very high, so the agent's consumption level is sufficiently low relative to her wealth that it is optimal to pay the fixed cost and to increase consumption to the level given by equation (15). (A similar intuition holds if $y = y_1$.) Finally, the optimal risky asset holdings are given by

$$X_t = c_t \frac{-h'(y_t)}{h''(y_t)} \frac{\mu}{\sigma^2}. \tag{16}$$

We have formulated the model in terms of consumption flow c_t, which can be interpreted either as purchases of nondurables or as a service flow from a stock of durables. Grossman and Laroque (1990) formulate their model over the stock of durable consumption goods, which may have nonzero depreciation. The two formulations are isomorphic. In particular, the model structure and equilibrium conditions derived by Grossman and Laroque (1990) can be recovered from equations (2) through (7) and (13) by a simple redefinition of variables. Details are provided in the Appendix.

III. Individual and Aggregate Choices

A. Simulation Methodology

This section describes the simulation methodology we use to explore the implications of the model. For a given set of parameter values, we compute the equilibrium by numerically solving equation (13) to obtain the function h. (Details are provided in the Appendix.) Once the function h has been computed, we use this function in equations (12) and (14) to compute y_1, y_2, and y^*. We then simulate the returns processes, and we use the optimal consumption and portfolio rules to trace the evolution of wealth and consumption for a population of 100 individuals. All individuals have the same initial wealth, but differ in their initial consumption. In particular, the starting cross-sectional distribution of individual consumptions is chosen to approximate the unconditional distribution of the wealth/consumption ratio. We approximate continuous time by evaluating the consumption and portfolio rules at discrete time intervals Δt. Equation (2) implies that the cumulative return to the risky asset between t and $t + \Delta t$ equals $\exp[(\mu + r - \sigma^2/2)\Delta t + \sigma(\Delta t)^{0.5}\epsilon]$, where ϵ is a standard normal variate. The only approximation in moving from continuous to discrete time is that agents are allowed to adjust their consumption and portfolios only at discrete intervals. In our simulations, we set $\Delta t = (1/120)$th of a year. (Experiments setting Δt equal to one day yield virtually identical results.) For each parameter set, the model is simulated for 300,000 periods (2500 years). We keep track of consumption, investment (defined as the change in wealth), and portfolio composition both for a single agent and for the aggregate (computed as the cross-sectional sum of the 100 individual quantities).

We calibrate μ and σ^2 to match the mean and variance of the monthly real return to the CRSP value-weighted stock portfolio from 1959:2 through 1993:12. The risk-free rate r is calibrated to the mean of the monthly real return to one-month Treasury bills over this same period. In particular, the mean monthly continuously compounded return on the value-weighted stock portfolio is 0.502 percent, and the standard deviation of this return is 4.346 percent. The mean monthly continuously compounded return on one-month T-bills over this period is 0.0772 percent. The rate-of-return parameters μ, σ, and r are annualized rates, so the mean of the monthly log return to the risky asset is $(\mu + r - \sigma^2/2)/12$, the variance is $\sigma^2/12$, and the monthly risk-free log return is $r/12$. Setting these expressions equal to the moments in the observed monthly data, we obtain $\mu = 0.06227$, $\sigma = 0.1505$, and $r = 0.00926$.

For the remaining parameters $\{a, \lambda, \delta\}$ we focus initially on the following benchmark parameterization: $a = -4$ (which corresponds to a coefficient of relative risk aversion of 5 when $\lambda = 0$); $\delta = 0.04$ (a 4 percent annual subjective discount rate), and $\lambda \in [0,0.5]$. Later, in Section IV.D, we consider alternative values of a and δ. We also consider an alternative interpretation of the model in which consumption is a flow of services from a stock of durables.

B. Failure of Aggregation

In Figure 2 we display a typical individual consumption path taken from our simulation exercise (solid line), as well as the corresponding path for aggregate consumption (dot/dash line). As described in Section II, individual consumption changes by discrete jumps. In contrast, aggregate consumption evolves smoothly. It is clear that the aggregate consumption process does not resemble any individual's consumption path, so aggregation (in the sense of Rubinstein (1974)) does not hold; that is, there does not exist a "representative agent" whose optimal consumption and savings choices equal the corresponding equilibrium aggregate magnitudes.

This failure of aggregation is due to a subtle market incompleteness. One can reinterpret a "consumption transaction" as a private production technology that converts "wealth" (a durable intermediate good) into "consumption" (a perishable final good). This private technology incorporates the adjustment cost. As shown in Arrow (1964), markets are complete in this two-good economy if a complete set of securities are traded that pay off one unit of wealth at a given date in a given state, and a spot market exists at each date where wealth is traded for consumption. In our model, markets are effectively complete in wealth trades. That is, introducing Arrow–Debreu securities paying off in wealth does not change the equilibrium allocations. However, we have implicitly ruled out these "spot markets" in consumption. That is, we do not allow one agent to "produce" consumption from wealth (paying the fixed cost if necessary), and then to sell this consumption to a second agent (thereby sparing her the fixed cost). In this sense, markets are incomplete: The marginal rates of substitution between consumption and wealth differ across agents, so there are gains from consumption trades that remain unexploited in equilibrium. In the usual case, where production technology is unaffected by the cross-sectional distribution of consumption, the

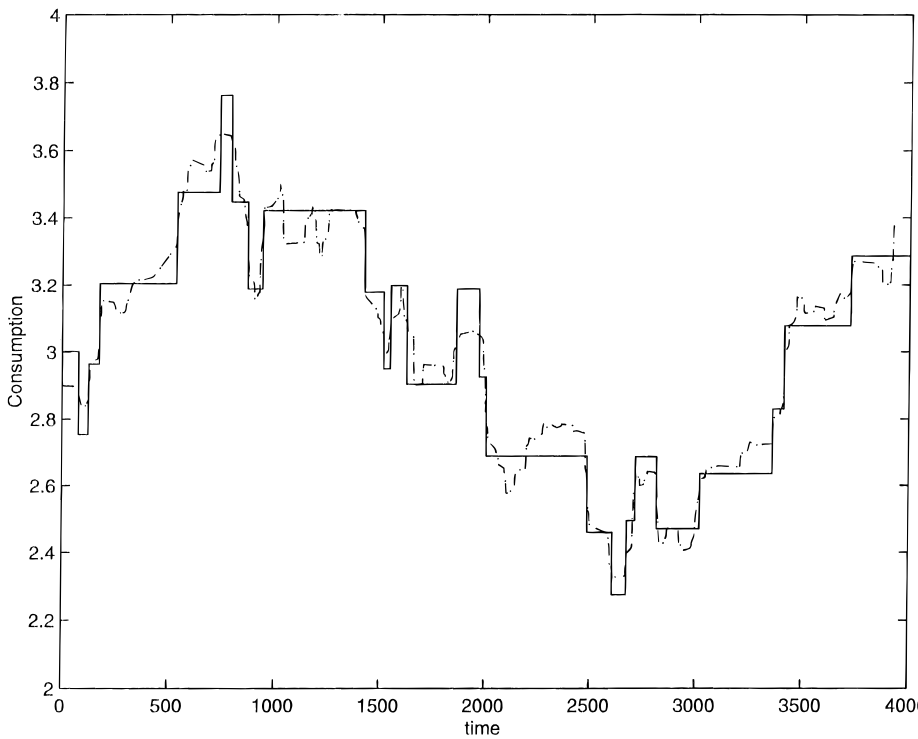

Figure 2. Example of individual and aggregate per capita consumption. This figure plots a typical sample path of individual consumption (solid line) and aggregate consumption (dot/dash line) from simulating the model. The parameters are set as follows: adjustment cost parameter $\lambda = 0.002$; preference parameters $a = -4$; $\delta = 0.04$; the mean excess return $\mu = 0.06227$; the standard deviation of the risky return $\sigma = 0.1505$; and the instantaneous risk-free rate of return $r = 0.00926$.

existence of spot markets in consumption is sufficient for aggregation to obtain. (The proof is analogous to that used in Rubinstein (1974), and relies on the assumption of hyperbolic absolute risk aversion utility.) In the absence of these spot markets, aggregation fails.

As it happens, the technology for transforming wealth into consumption implied by our particular fixed cost specification (equation (5)) does depend on the cross-sectional distribution of consumption. In particular, the most cost-efficient way to implement a given change in aggregate consumption is to have the entire change absorbed by the agent with the smallest current consumption rate. Replacing a distribution of N agents with a single representative agent increases the cost of a given change in the aggregate consumption, and therefore changes the equilibrium. This feature precludes aggregation even with complete markets. (This problem does not arise if the fixed cost is constant, rather than being proportional to current consumption.)

C. The Effect of Adjustment Costs on Consumption and Portfolio Decisions

What happens to individual and aggregate choice variables as the adjustment-cost parameter λ increases? In Figure 3, Panel A, we display the boundary points y_1 and y_2 of the region of inactivity, along with y^*, the target value of y. Note that the region of inactivity $[y_1, y_2]$ expands rapidly as λ increases from zero, but the rate of increase falls off as λ gets larger. As discussed in Section I above, this diminishing marginal impact of adjustment costs is a simple implication of concave preferences. As the region of inactivity expands, the time between adjustment increases. Panel B of Figure 3 displays the mean number of years between adjustments as a function of λ. The time between adjustments grows quite large, even for rather small values of λ. For example, when $\lambda = 0.002$, the average time between consumption adjustments is approximately 11 months. As consumption adjustments become less well synchronized, the cross-sectional dispersion of consumption increases. This is shown in Panel C of Figure 3, which displays the mean value of the (time-varying) cross-sectional coefficient of variation (std. dev./mean) of individual consumptions as a function of λ. This measure of cross-sectional variability is nontrivial even for the smallest values of the adjustment cost.

In Figure 4 we plot the standard deviation of consumption growth from t to $t + \Delta t$, where $\Delta t = (1/120)$th of a year. We display this statistic for both consumption (dashed line) and aggregate consumption (solid line) as a function of λ. While the standard deviation is not the best measure of the volatility of a highly non-Gaussian series like individual consumption growth, it is of interest that the impact of adjustment costs on the standard deviation of individual consumption is rather small. To see why, let $p_{\Delta t}^{up}$ and $p_{\Delta t}^{down}$ denote the unconditional probabilities of an upward and downward consumption adjustment, respectively. For small Δt,[3]

$$\mathrm{var}(c_{t+\Delta t}/c_t) \approx p_{\Delta t}^{up}\left[\frac{y_2}{y^* + \lambda}\right]^2 + p_{\Delta t}^{down}\left[\frac{y_1}{y^* + \lambda}\right]^2 - [E(c_{t+\Delta t}/c_t)]^2. \qquad (17)$$

As λ increases, both $p_{\Delta t}^{up}$ and $p_{\Delta t}^{down}$ decrease. However, the reduced probabilities of adjustment are almost exactly offset by the increased magnitude of the adjustments. (The mean growth rate of individual consumption $E(c_{t+\Delta t}/c_t)$ is virtually unaffected by increasing λ. There is a slight wealth reduction due to the infrequent payments of the adjustment costs.)

Although the variability of individual consumption growth is largely unaffected by increasing the adjustment costs, there is a pronounced reduction in the standard deviation of aggregate consumption growth. According to Figure 4, this standard deviation drops rapidly as λ increases from zero,

[3] In equation (17), we use the fact that the gross consumption growth rate in the event of a downward adjustment equals $y_1/(y^* + \lambda)$ (analogously for an upward consumption adjustment), and that the probability of two or more adjustments during interval Δt vanishes at rate $o(\Delta t)$.

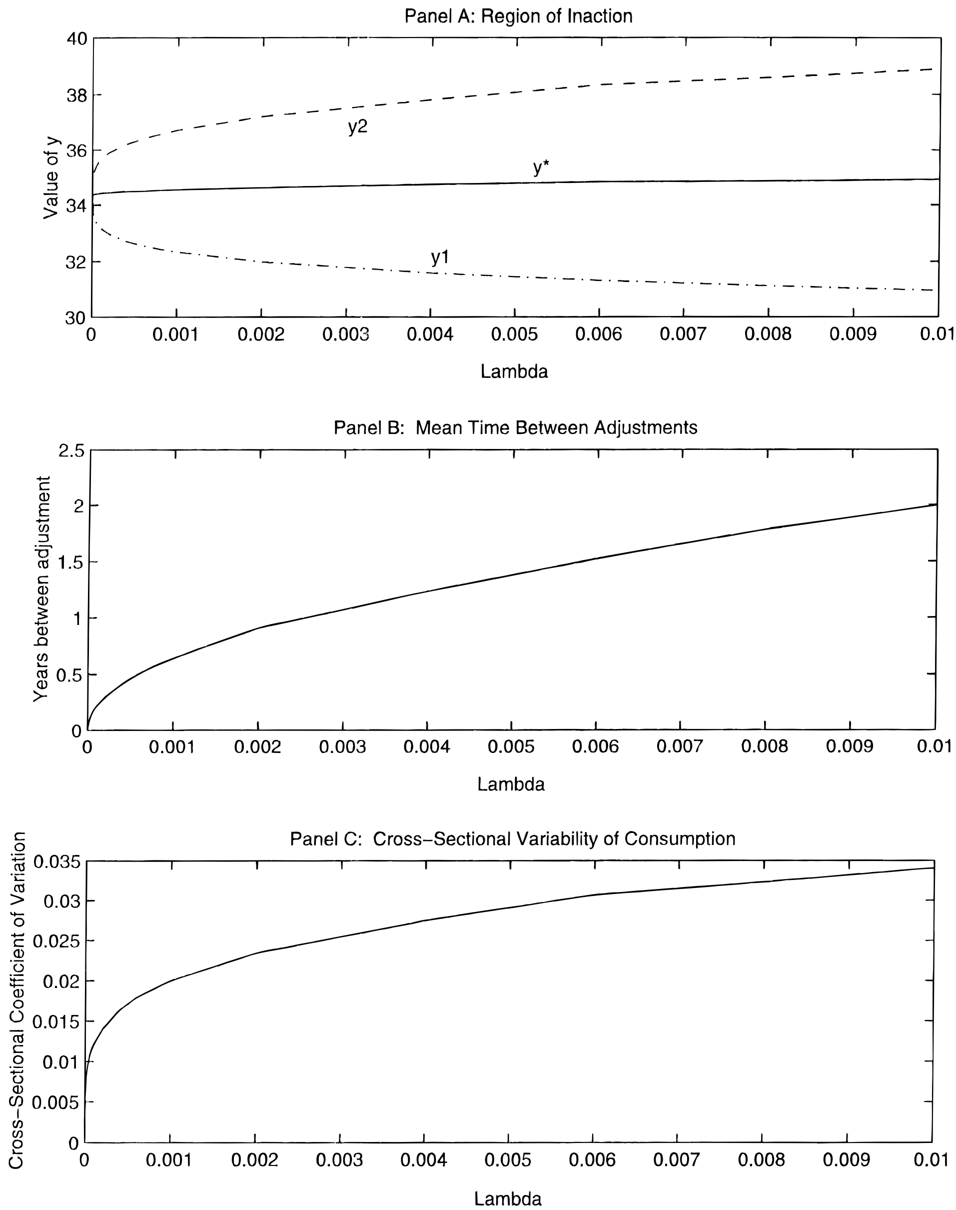

Figure 3. Consumption behavior as adjustment costs increase. For the benchmark parameterization, Panel A plots the lower bound y_1 (dot/dash line) and upper bound y_2 (dashed line) of the region of inactivity as the fixed cost parameter λ increases. The solid line in Panel A plots y^*, the target value of y following a consumption adjustment. Panel B plots the mean time in years between consumption adjustments. Panel C plots the mean value of the cross-sectional coefficient of variation (std. dev./mean) of individual consumption. In all three panels, the statistics are plotted as functions of $\lambda \in [0,0.01]$. The benchmark parameters are: preference parameters $a = -4$; $\delta = 0.04$; the mean excess return $\mu = 0.06227$; the standard deviation of the risky return $\sigma = 0.1505$; and the instantaneous risk-free rate of return $r = 0.00926$.

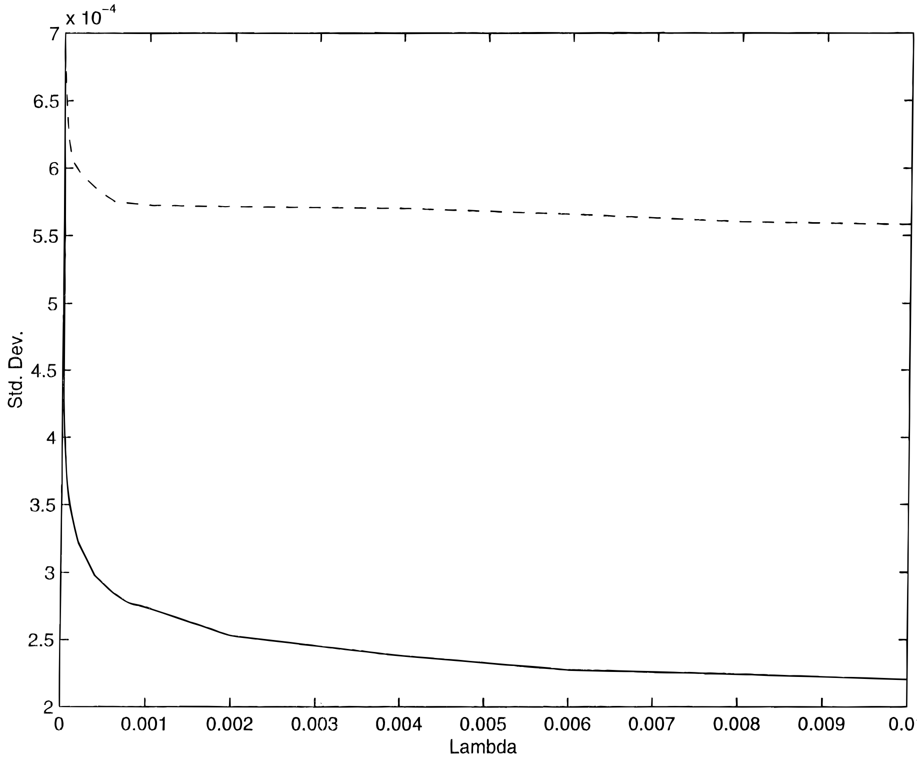

Figure 4. Standard deviations of individual and aggregate consumption versus fixed cost parameter λ. For the benchmark parameterization, this figure plots the standard deviation of the growth rate of individual consumption (dashed line) and cross-sectionally aggregated consumption (solid line) as functions of the fixed cost parameter λ. The growth rates are over a time interval equal to (1/120)th of a year. These standard deviations are plotted as functions of the fixed cost parameter $\lambda \in [0,0.01]$. The other parameters are set as follows: preference parameters $a = -4$; $\delta = 0.04$; the mean excess return $\mu = 0.06227$; the standard deviation of the risky return $\sigma = 0.1505$; and the instantaneous risk-free rate of return $r = 0.00926$.

with the rate of decrease falling off as λ grows. In other words, a small adjustment at the individual level induces a substantial reduction in consumption volatility at the aggregate level.

Figure 4 illustrates a result proved in a simpler setting by Caballero (1993).[4] He shows that when a continuum of agents follows a (y_1,y_2,y^*) strategy, the variability of the cross-sectionally aggregated process is smaller than the

[4] Caballero (1993) assumes that the *difference* between individual consumption and the frictionless target consumption (the optimal consumption level if λ were equal to zero) follows a (y_1,y_2,y^*) policy, and that this difference is a simple Brownian motion process between adjustments. Our model implies that the *ratio* between actual and target consumption follows the (y_1,y_2,y^*) policy. (This ratio is proportional to $y + \lambda$.) Furthermore, this ratio is not Brownian motion, but follows a nonlinear diffusion process that does not have an explicit representation.

variability of the frictionless process (i.e., the process that would be followed if the range of inactivity collapsed to zero). The intuition is as follows: The rate of change of aggregate consumption depends on the size of the individual adjustment when an agent hits a boundary of the inactivity region, and the *probability density* of agents in the neighborhood of these boundaries. The size of individual consumption adjustments is state-independent, so variability in aggregate consumption depends solely on the way the cross-sectional distribution of agents varies through time. Caballero shows that the variability of this cross-sectional distribution is smaller in magnitude and different in kind than the variability of consumption in the frictionless case.[5] Figure 4 suggests that a similar smoothing due to cross-sectional aggregation is present in our model. Since the marginal impact of the adjustment cost on the cross-sectional distribution is decreasing (see Figure 3, Panel C), it is not surprising that the marginal impact on aggregate consumption variability is also decreasing.

We conclude this section with a discussion of the individual's portfolio decision. In the absence of adjustment costs ($\lambda = 0$), the portfolio weight on the risky asset is constant (see equation (8)). In Figure 5 we display how this portfolio weight varies with y when $\lambda = 0.002$. Notice how agents behave in a less risk-averse fashion near the boundaries of the region of inactivity (when y is either very high or very low). In Figure 6 we plot the minimum, mean, and maximum values of this portfolio weight as a function of λ. As in Figure 3, the marginal effect on the range of admissible portfolio weights as λ is increased is biggest for the smallest λ's. There is also a slight reduction in the mean portfolio weight as λ increases.

IV. Comparison of the Model with the Data

A. Methodology

In this section, we ask how well the benchmark model replicates the patterns found in the data. Observed data on consumption and investment are time-averaged. To compare the behavior of the model with observed data, we time-average the aggregate series generated by our model simulations to create monthly or quarterly series. We then compute a number of statistics from these series, which we compare to the corresponding statistics computed using observed data. First, we compute the mean and variance of aggregate consumption growth. Second, we compute the theoretical equity premium EP', as characterized by equation (1). In our model, the theoretical equity premium over horizon Δt is computed as:

$$EP' = \frac{-\mathrm{cov}(e^{-\delta\Delta t}(c_{t+\Delta t}/c_t)^{a-1},(b_{t+\Delta t}/b_t))}{E(e^{-\delta\Delta t}(c_{t+\Delta t}/c_t)^{a-1})}. \tag{18}$$

[5] In particular, consumption in the frictionless case is a Brownian motion process, implying a nondifferentiable sample path. In contrast, cross-sectionally aggregated consumption in the (y_1,y_2,y^*) model displays a differentiable (and, in that sense, smoother) sample path.

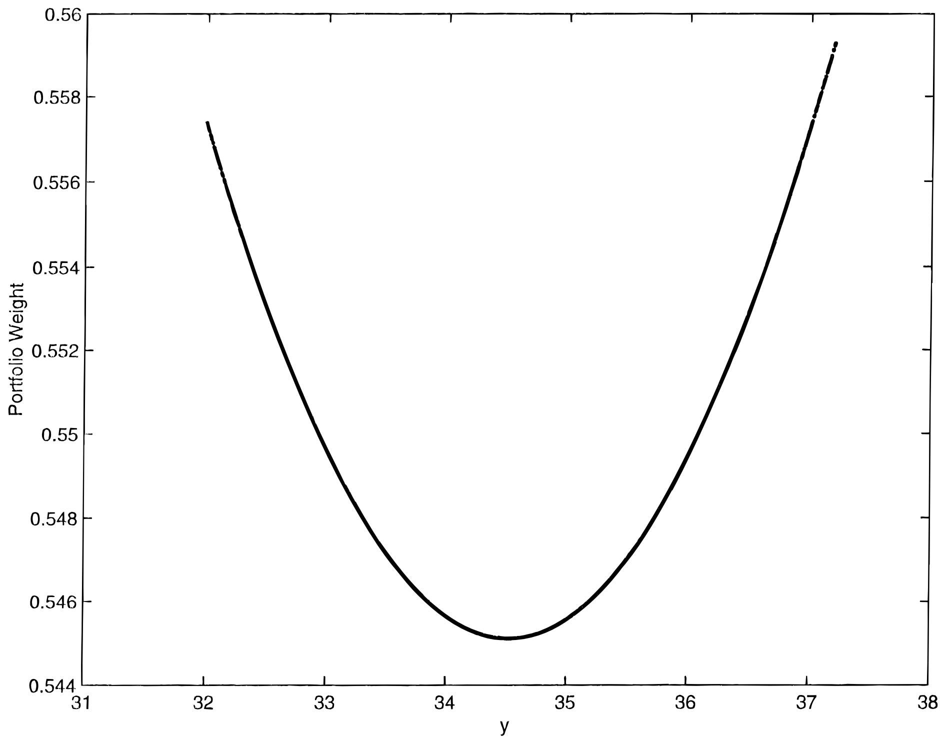

Figure 5. Portfolio weight on the risky asset, as a function of y. The portfolio weight on the risky asset is plotted as a function of y, defined as the wealth/consumption ratio minus λ. The parameters are set as follows: adjustment cost parameter $\lambda = 0.002$; preference parameters $a = -4$; $\delta = 0.04$; the mean excess return $\mu = 0.06227$; the standard deviation of the risky return $\sigma = 0.1505$; and the instantaneous risk-free rate of return $r = 0.00926$.

If $\lambda = 0$ and consumption are point-in-time sampled, then equation (18) implies that $EP' = e^{(\mu+r)\Delta t} - e^{r\Delta t}$, the true equity premium over horizon Δt.[6] The theoretical equity premium differs from the true equity premium when time-averaged consumption is used and/or $\lambda > 0$. Finally, we also compute the mean and standard deviation of the growth rate of aggregate investment. While investment is not usually used as an explanatory variable in this class of asset pricing models, it is a natural variable to examine in a model where frictions distort consumption decisions. Income must be divided between consumption and investment. If consumption is distorted, one might think that investment will be distorted in the opposite direction. In our model, aggregate investment is simply the change in aggregate wealth.

[6] This can be seen algebraically by substituting equation (9) into equation (18), and using the laws-of-motion for wealth and equity value (equations (4) and (2)).

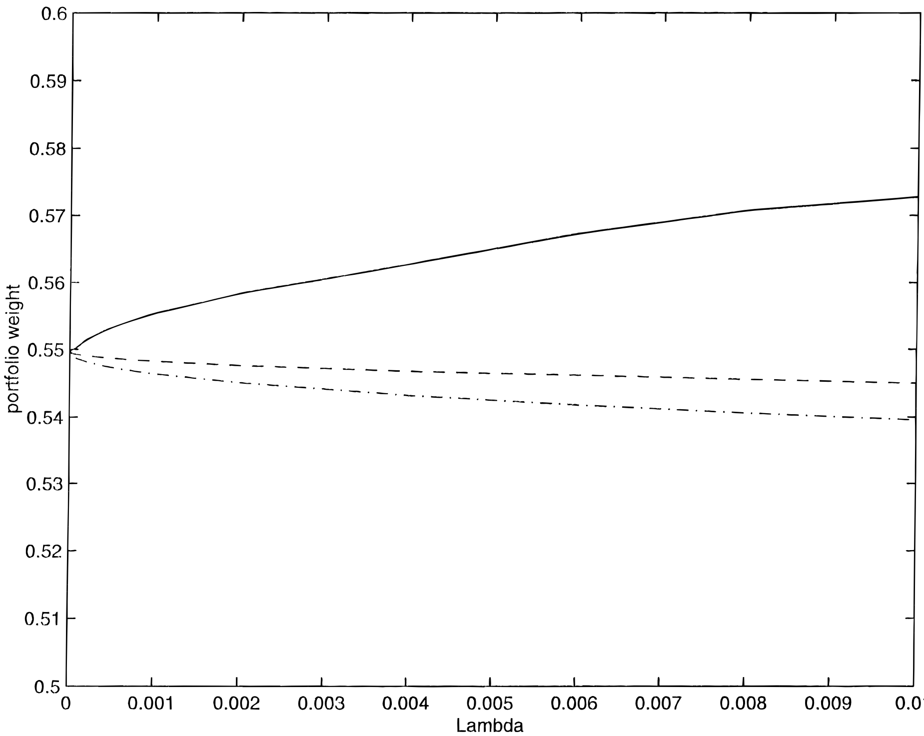

Figure 6. Minimum, mean, and maximum portfolio weight on the risky asset, as functions of fixed cost parameter λ. For the benchmark parameterization, this figure plots the minimum (dot/dash line), mean (dashed line), and maximum (solid line) values of the portfolio weight on the risky asset, as functions of fixed cost parameter λ, for $\lambda \in [0,0.01]$. The benchmark parameters are set as follows: preference parameters $a = -4$; $\delta = 0.04$; the mean excess return $\mu = 0.06227$; the standard deviation of the risky return $\sigma = 0.1505$; and the instantaneous risk-free rate of return $r = 0.00926$.

In comparing the model to the data, we wish to focus on reasonable values of the adjustment cost. When consumption is interpreted as the service flow from the stock of durables or housing, we set $\lambda = 0.5$ as the upper limit of the adjustment cost parameter. We regard this figure as conservative. For example, if the fixed cost of changing one's house is 6 percent of the value of the house being sold, and the annual service flow from a house is 12 percent of its value, then λ would equal 50 percent. Similarly, suppose the adjustment cost involved with buying a car is equal to the retail/wholesale spread that the consumer loses on the trade-in. If this spread equals 8 percent of the trade-in's value (a rather low figure), and if the service flow from a car equals 16 percent of the car's value, then λ would again equal 50 percent. When consumption is interpreted as the expenditure on nondurable goods, we consider only trivially small adjustment costs. For this case, we set $\lambda = 0.002$ as the upper limit of the adjustment cost parameter.

B. Quantitative Implications of the Benchmark Parameterization

Implications of this benchmark case for time-averaged aggregate consumption are displayed in Table I. The effect of adjustment costs on the mean rate of aggregate consumption growth is small. According to Panel A of Table I, increasing λ from 0 to 0.5 (the largest value we implement) only reduces the consumption-growth mean by 15 percent (from 0.1 percent/month to 0.085 percent/month). However, very small adjustment costs substantially reduce the variability of one-month consumption growth. For example, when we increase λ from 0 to 0.002 in this benchmark case, we observe a 35 percent reduction in the standard deviation of monthly aggregate consumption growth, from 1.95 percent per month to 1.27 percent per month. Furthermore, the most rapid reduction is for the very smallest values of the adjustment cost. This rapid fall-off in variability as λ increases is illustrated in Figure 7, which graphs the standard deviation as a function of $\lambda \in [0,0.01]$. As in Figure 4, very small adjustment costs have the biggest proportional impact on consumption-growth variability.

Similarly, we find a vast reduction in the theoretical equity premium EP' (defined in equation (18)) when extremely small fixed costs of consumption adjustment are imposed. These results are displayed in Table I, Panel B. The true monthly equity premium (computed as $\exp((\mu + r)/12) - \exp(r/12)$) is 52 basis points, or about 6.2 percent per year. When $\lambda = 0$, EP' replicates this value when it is computed using point-in-time-sampled consumption (the second column of Table I, Panel B). In our benchmark case, a value of $\lambda = 0.002$ reduces this theoretical equity premium to 2.4 percent per year, a reduction of 62 percent. However, this calculation does not replicate the way EP' is actually computed from observed data because aggregate consumption data are not point-in-time-sampled but, instead, represent a monthly average. This averaging reduces the covariance between equity returns and consumption growth. Even with $\lambda = 0$, temporal aggregation over the month reduces the theoretical equity premium to 3.5 percent per year. When a small adjustment cost of $\lambda = 0.002$ is introduced, the theoretical equity premium computed using time-averaged aggregate consumption falls to a mere 1.3 percent per year. To summarize, the combined effects of extremely small consumption adjustment costs and temporal aggregation induce a 79 percent reduction in the theoretical equity premium computed from consumption data.

As with the variance of aggregate consumption growth, the largest effect of increased λ on the theoretical equity premium is for very small adjustment costs. This is illustrated in Figure 8, which plots the theoretical equity premiums computed both with point-in-time sampled aggregate consumption (dashed line) and with time-averaged aggregate consumption (solid line). As can be seen, most of the reduction in the theoretical equity premium occurs when we move from $\lambda = 0$ to $\lambda = 0.002$.

Finally, Panel C of Table I shows how aggregate investment growth changes as adjustment costs increase. Investment data are available only at a quarterly frequency, so, to facilitate comparison with the data, we report the model's implications for moments of quarterly investment growth. Note that neither

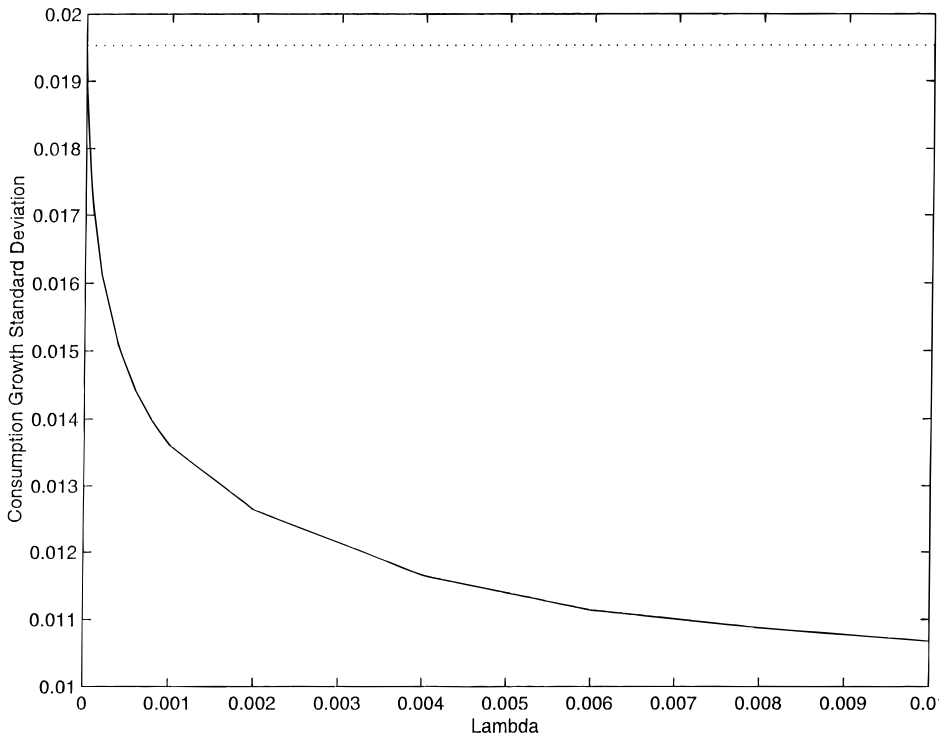

Figure 7. Standard deviation of aggregate monthly consumption growth as adjustment costs increase. The solid line plots the standard deviation of aggregate monthly consumption growth implied by the benchmark model against the value of the adjustment cost parameter λ, for $\lambda \in [0,0.01]$. (The dotted horizontal line gives the standard deviation of monthly consumption growth implied by the model when $\lambda = 0$.) Other model parameters are set as follows: preference parameters $a = -4$; $\delta = 0.04$; the mean excess return $\mu = 0.06227$; the standard deviation of the risky return $\sigma = 0.1505$; and the instantaneous risk-free rate of return $r = 0.00926$. The aggregate consumption process implied by the model was computed by simulating the optimal decisions of 100 individuals for 2,500 years, summing cross-sectionally, and temporally aggregating up to a monthly series.

the mean nor the standard deviation are affected substantially by adjustment costs. Recall that the variance of aggregate consumption is reduced by cross-sectional aggregation when individual consumptions follow a (y_1, y_2, y^*) pattern. Individual investment, however, does not follow this sort of pattern. Rather, it follows a Brownian motion path punctuated by infrequent downward jumps (when the adjustment cost is paid). Aggregating this type of individual process does not significantly reduce the variability, compared to the frictionless case.

C. How Big Are These Adjustment Costs?

The results in Section IV.B suggest that both the variability of consumption and the theoretical equity premium inferred from consumption are highly nonrobust to consumption adjustment costs. In fact, the costs needed to in-

Table I

Behavior of Benchmark Case as Adjustment Costs Increase

This table displays the statistical properties of the adjustment cost model with preference parameters $a = -4$, $\delta = 0.04$, and the adjustment cost parameter λ ranging from 0 through 0.5. In Panel A, columns 2 and 3, display the mean and standard deviation of the growth rate of per capita consumption, time-averaged to a monthly series; column 4 displays the percentage reduction in the standard deviation of aggregate consumption growth compared to $\lambda = 0$; and columns 5–7 are analogous to columns 2–4, except that consumption is time-averaged to a quarterly series. Panel B displays the theoretical equity premium,

$$EP' \equiv \frac{-\text{cov}((c_{t+\Delta T}/c_t)^{a-1},(b_{t+\Delta T}/b_t))}{E((c_{t+\Delta T}/C_t)^{a-1})},$$

(in units of percent per annum, where b_t is the unit value of the risky asset at time t) computed three ways: column 2 displays EP' computed using per capita consumption point-in-time sampled at a monthly frequency; column 3 displays EP' computed using per capita consumption time-averaged to a monthly frequency; and column 4 displays EP' computed using per capita consumption time-averaged to a quarterly frequency. Panel C gives the mean and standard deviation of quarterly investment growth, defined as

$$[Q_{t+2\Delta T} - Q_{t+\Delta t}]/[Q_{t+\Delta t} - Q_t]$$

for Δt = one quarter. Panel D, column 2, gives the average amount of wealth dissipated by the adjustment cost in equilibrium, as a fraction of economy-wide consumption; column 3, gives the maximum amount, on average, that an agent would pay (as a fraction of current wealth) to reduce λ to zero. This statistic is computed as

$$\frac{Q - c[h(y)a((1-a)/\beta)^{a-1}]^{1/a}}{Q}.$$

Parameters governing the asset return processes are: the mean excess return $\mu = 0.06227$; the standard deviation of the risky return $\sigma = 0.1505$; and the instantaneous risk-free rate of return $r = 0.00926$.

Panel A: Mean and Variance of Aggregate Consumption Growth

	Moments of Net Monthly Consumption Growth			Moments of Net Quarterly Consumption Growth		
λ	Mean	Standard Deviation	Reduction in Standard Deviation (%)	Mean	Standard Deviation	Reduction in Standard Deviation (%)
0	0.001	0.01952		0.003	0.03347	
0.0002	0.00094	0.01609	17.55	0.0029	0.03194	4.57
0.001	0.00091	0.01360	30.32	0.0029	0.02915	12.88
0.002	0.00089	0.01265	35.20	0.0028	0.02757	17.62
0.006	0.00088	0.01114	42.95	0.0028	0.02449	26.81
0.01	0.00087	0.01068	45.30	0.0027	0.02345	29.92
0.05	0.00086	0.00975	50.07	0.0027	0.02074	38.04
0.1	0.00086	0.00922	52.77	0.0026	0.01975	40.99
0.5	0.00085	0.00812	58.38	0.0026	0.01761	47.39

Table I—*Continued*

Panel B: Mean Theoretical Equity Premium (percentage per annum)

	Monthly Horizon		Quarterly Horizon
λ	Point-in-Time Sampled Consumption (%)	Time-Averaged Consumption (%)	Time-Averaged Consumption (%)
0	6.252	3.42	3.468
0.0002	3.672	1.884	2.672
0.001	2.712	1.464	2.136
0.002	2.388	1.344	1.924
0.006	1.932	1.152	1.604
0.01	1.776	1.08	1.488
0.05	1.392	0.924	1.204
0.1	1.272	0.876	1.096
0.5	1.056	0.804	0.936

Panel C: Mean and Standard Deviation of Aggregate Quarterly Investment Growth

λ	Mean	Standard Deviation
0.0000	0.0035	0.0416
0.0002	0.0035	0.0416
0.001	0.0035	0.0416
0.002	0.0035	0.0415
0.006	0.0035	0.0414
0.01	0.0035	0.0414
0.05	0.0035	0.0409
0.1	0.0035	0.0406
0.5	0.0034	0.0391

Panel D: Measures of Adjustment Cost Impact

λ	Average Aggregate Consumption Loss	Average Welfare Loss: Fractional Wealth Equivalent
0.0000	0.000000	0.0000
0.0002	0.000005	0.0016
0.001	0.000012	0.0034
0.002	0.000018	0.0049
0.006	0.000031	0.0084
0.01	0.000041	0.0109
0.05	0.000093	0.0247
0.1	0.000131	0.0352
0.5	0.000300	0.0809

duce major changes in these statistics are extremely small. We focus on values of λ around 0.002. Per capita consumption of nondurables and services in the United States in 1993 was approximately $14,000, so a value of λ equal to 0.002 implies a consumption adjustment cost of $28. Costs this small are virtually undetectable. Of course, agents modify their behavior to

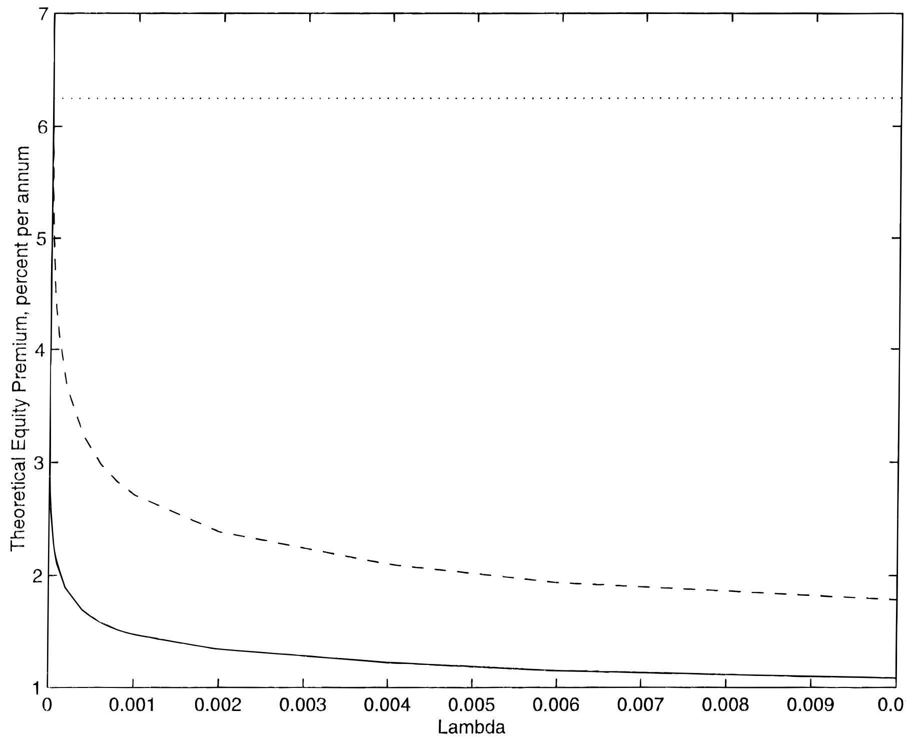

Figure 8. Theoretical equity premiums as adjustment costs increase. This figure plots the mean value of the monthly theoretical equity premium, EP' (in percentage per annum) implied by the model, against the value of the adjustment cost parameter λ. EP' is defined as

$$-\mathrm{cov}([c_{t+1}/c_t]^{a-1}, r^e_{t+1})/E([c_{t+1}/c_t]^{a-1}),$$

where r^e_{t+1} denotes the equity return, and c_t denotes the consumption series being considered (as described below). The dotted line gives the true equity premium $12 \cdot \exp((\mu + r)/12) - \exp(r/12)$. The dashed line plots EP' computed using point-in-time data on monthly consumption implied by the model. The solid line plots EP' computed using time-averaged data on monthly consumption implied by the model. Other model parameters are set as follows: preference parameters $a = -4$; $\delta = 0.04$; the mean excess return $\mu = 0.06227$; the standard deviation of the risky return $\sigma = 0.1505$; and the instantaneous risk-free rate of return $r = 0.00926$. The consumption process implied by the model was computed by simulating the optimal decisions of 100 individuals for 2,500 years and summing cross-sectionally. Point-in-time consumption data selects this consumption process once per month; time-averaged consumption data averages this consumption process over a one-month interval.

avoid paying this cost, so the equilibrium reduction in aggregate consumption due to the adjustment cost is even smaller. In particular, the second column of Table I, Panel D, reports that, in the benchmark case with $\lambda = 0.002$, the average fraction of economy-wide consumption paid out in adjustment costs is only 18 parts per million.

An alternative measure of the size of the adjustment cost is the welfare cost, defined as the maximum amount that an agent in an economy with $\lambda > 0$ would pay (as a fraction of current wealth) to reduce λ to zero from that time forth. Let this amount be denoted ζ. If $V(Q,c)$ denotes the value function with $\lambda > 0$, ζ must satisfy

$$V(Q,c) = \frac{1}{a}\left[\frac{\beta}{1-a}\right]^{a-1}(Q(1-\zeta))^{a}, \tag{19}$$

where we use equation (11) for the value function with $\lambda = 0$. Using equation (12), and making explicit the dependence of ζ on Q and c, we obtain

$$\zeta(Q,c) = \frac{Q - c[h(y)a(1-a)/\beta)^{a-1}]^{1/a}}{Q}. \tag{20}$$

Panel D of Table I reports the mean value of $\zeta(Q,c)$ for various values of λ. According to the third column of that panel, the maximum fraction of wealth that the average agent would pay to permanently reduce the adjustment cost from $\lambda = 0.002$ to $\lambda = 0$ in the benchmark case is 0.0049 (about 0.5 percent), less than a typical mutual fund's yearly management fee.

D. Results for Alternative Parameterizations

In this section, we briefly describe how deviations from the benchmark parameterization affect the model. (Details of the experiments described in this section are available upon request.) When risk aversion is reduced, both the mean and the variance of consumption growth increase because agents hold a larger fraction of wealth in the risky asset. For example, if we set $\lambda = 0.002$, but set $a = -1$ (which corresponds to a coefficient of relative risk aversion of 2 in the frictionless case), the mean of monthly aggregate consumption growth increases to 0.0026 (compared to 0.0009 in the benchmark case) and the standard deviation increases to 0.035 (as compared to the benchmark value of 0.012). In contrast, the results for the theoretical equity premium are nearly identical to the benchmark case. This should not be surprising because in the absence of adjustment costs or temporal aggregation, the results for the theoretical equity premium are unaffected by the value of a. (The theoretical equity premium simply equals the exogenously specified true equity premium.)

When the consumption flow is interpreted as a flow of services from consumer durables (as in Grossman and Laroque (1990)), it is of interest to consider the case where the durable stock depreciates. We experiment with a version of the benchmark parameterization in which consumption is a service flow from a stock of durables that depreciates at a rate of 10 percent per annum. The qualitative implications of the model are similar to those in Table I. Small adjustment costs still have a small impact on the mean of consumption growth, although the effect becomes somewhat larger when the largest adjustment costs are considered. The reason for this is that when durables depreciate, agents consume a larger fraction of output and invest less. However, the division of financial wealth between the risky and risk-free assets is virtually unaffected, so the variability of consumption growth and the theoretical equity premiums are nearly the same as in the benchmark parameterization.

We also consider alternative values of the discount parameter δ. Of the statistics we are examining, δ only affects the mean growth rate of consumption. Reducing δ increases the mean growth rate of consumption, as consumers reduce consumption and increase savings each period. For example, in the case $\{a = -4, \lambda = 0\}$, when δ is reduced from 4 percent down to 0 percent[7] the mean growth rate of monthly aggregate consumption increases from 0.10 percent/month to 0.17 percent/month. When δ is increased to 0.06, the mean growth rate of monthly consumption falls to 0.07 percent/month. However, the effect of adjustment costs on consumption-growth variance or on the theoretical equity premium is virtually unaffected by changing δ.

E. Can Adjustment Costs Explain the Patterns Found in the Data?

We now ask how well the benchmark case replicates the moments of consumption growth and asset returns observed in the data. The risk-free rate and the mean and variance of the risky return are already calibrated to match observed asset returns. The key questions are whether plausible adjustment costs can reduce the variance of aggregate consumption growth and the implied theoretical equity premium to the levels estimated in postwar U.S. data. Of course, to make this sort of empirical assessment we must specify what sort of consumption data we wish to match. Most empirical studies of consumption-based pricing models (e.g., Hansen and Singleton (1982), Grossman et al. (1987)) use data on consumption of nondurables plus services. In contrast, Grossman and Laroque (1990) interpret consumption as a service flow from a stock of durables. Rather than taking a stand on the appropriate measure of consumption, we compare our model to five different consumption data series.[8] The first three are measures of nondurable consumption purchases: expenditures on nondurable consumption goods ("CEN"), expenditures on consumer services excluding the imputed rental value of owner-occupied housing ("CES"), and the sum of CEN and CES ("NDS"). The remaining two consumption series we explore are estimates of the service flows from the stock of consumer durables ("YCD") and from the stock of

[7] As noted by Grossman and Laroque (1990), the objective function (equation (7)) is well-defined with $\delta = 0$ as long as $a < 0$, since β (defined in equation (10)) is still positive.

[8] Ideally, one would like to include all types of consumption in a single model, with different adjustment costs for each consumption type. This task would vastly complicate the analysis, and is not attempted in this paper. Our procedure can be justified by assuming that the different types of consumption flow are additively separable in the utility function.

owner occupied housing ("CESH") (the latter measured as the imputed rental value of owner-occupied housing).[9] The Bureau of Economic Analysis's data for consumer expenditures on services lump together our two series CES and CESH. We separate these two data series because the adjustment costs associated with these two types of consumption flow are likely to be quite different. Clearly, the cost of adjusting the service flow from one's residence are quite large, whereas the cost of adjusting expenditures directly paid to service providers such as gardeners is likely to be nearly zero.

In Table II, Panels A and B, we display the mean and variance of consumption growth and the theoretical equity premium[10] computed using our five consumption measures. Robust standard errors are computed as in Hansen (1982). Consider first Grossman and Laroque's (1990) original interpretation of consumption as a service flow from a stock of durables or housing. Reider (1993) suggests that the empirical performance of the CCAPM may improve if consumption is interpreted in this way. As discussed in Section IV.A above, rather high adjustment costs would be plausible under this interpretation, so we compare the statistics in Table II for service flows (YCD and CESH) with the implications of the model for values of λ up to 0.5. Because our data on YCD are available only on a quarterly basis, we compare these data with the model's consumption process aggregated up to a quarterly frequency.

When compared to these service-flow measures of consumption, the performance of the model is quite poor. First, Table II reports that the mean of the quarterly growth rate of the service flow imputed from the stock of consumer durables is 0.0089 and the standard deviation of this growth rate is 0.005. The corresponding moments for the growth rate of the service flow from the stock of housing are 0.0062 and 0.0043. The mean consumption growth rates in the benchmark model (reported in Table II) are considerably lower than the means estimated from these data.[11] More important, for our purposes, the variability of the growth rate of either YCD or CESH is far too

[9] All of these consumption series except for YCD are monthly data compiled by the Bureau of Economic Analysis. The series YCD is available only quarterly, and is obtained from the Board of Governors of the Federal Reserve System. Brayton and Mauskopf (1985) describe how YCD is constructed. For the monthly data, the sample period used is February 1959 through December 1993. For quarterly data, the sample is 1959:Q2 through 1993:Q4. All these data are measured in constant dollars, seasonally adjusted, and are expressed as per capita magnitudes.

[10] The theoretical equity premiums are computed from a given consumption series $\{c_t\}$ as

$$-\text{cov}([c_{t+1}/c_t]^{\alpha-1}, r^e_{t+1})/E([c_{t+1}/c_t]^{\alpha-1}),$$

where the excess equity return r^e_t is the monthly return to the CRSP value-weighted portfolio minus the return to the one-month Treasury bill. Consistent with our benchmark parameterization, we set $\alpha = -4$.

[11] A substantial reduction in δ would improve the fit of the model somewhat along this dimension. In particular, reducing δ to zero raises the mean quarterly growth rate of consumption in the model to 0.0051.

Table II

Properties of Observed Data

Panels A and B compute statistics for the following five consumption series: CEN = consumer expenditures on nondurables; CES = consumer expenditures on services excluding the service flow from owner-occupied housing; NDS = CEN + CES; YCD = service flow from the stock of consumer durables; CESH = the service flow from owner-occupied housing. Panel A displays the means and variances of the growth rates of the five consumption measures; Panel B displays the theoretical equity premium $EP' = -\text{cov}([c_{t+\tau}/c_t]^{a-1}, r^e_{t,t+\tau})/E([c_{t+\tau}/c_t]^{a-1})$, where r^e_t = monthly return to the CRSP value-weighted portfolio minus the return to the one-month Treasury bill. and parameter $a = -4$. Panel C computes statistics for GPDI = real gross private domestic investment and GS = real gross savings. All data are from the Bureau of Economic Analysis except YCD, which is from the Board of Governors. All data series are monthly from 1959:2 through 1993:12, except YCD, GPDI, and GS, which are quarterly data over the same period. Hansen (1982) standard errors are in parentheses, computed using the Newey–West (1987) covariance estimate with six lags.

Panel A: Moments of Net Consumption Growth Rate

	Monthly		Quarterly	
Consumption Measure	Mean	Std. Dev.	Mean	Std. Dev.
CEN	0.00072	0.0077	0.00282	0.0071
	(0.00026)	(0.0004)	(0.00082)	(0.0006)
CES	0.00194	0.0048	0.00682	0.0059
	(0.00020)	(0.0003)	(0.00067)	(0.0005)
NDS	0.00134	0.0049	0.00483	0.0056
	(0.00020)	(0.0002)	(0.00070)	(0.0004)
YCD			0.00886	0.0050
			(0.00102)	(0.0004)
CESH	0.00170	0.0018	0.00618	0.0043
	(0.00014)	(0.0001)	(0.00068)	(0.0002)

Panel B: Theoretical Equity Premiums Computed from Consumption Data ($a = -4$)

Consumption Measure	Monthly	Quarterly
CEN	0.0026	0.0015
	(0.0009)	(0.0011)
CES	0.0013	0.0017
	(0.0006)	(0.0010)
NDS	0.0019	0.0016
	(0.0006)	(0.0009)
YCD		−0.0010
		(0.0006)
CESH	−0.00033	−0.0008
	(0.00027)	(0.0005)

Panel C: Moments of Quarterly Net Growth in Investment and Savings

	Mean	Standard Deviation
GPDI	0.01032	0.0512
	(0.00464)	(0.0047)
GS	0.00584	0.0399
	(0.00359)	(0.0033)

low to be matched by this model with any plausible adjustment cost. At the highest level of adjustment costs that we consider ($\lambda = 0.5$), the standard deviation of quarterly consumption growth predicted by the model is 0.0176, which is more than three times larger (and several standard errors higher) than the point estimates from our service flow data.

The model's performance becomes even more problematic when we look at theoretical equity premiums implied by this service-flow data. For both of these data series, the theoretical equity premiums are negative; that is, the growth rates of these service flows are negatively correlated with equity returns. Furthermore, these negative correlations are not an artifact of the short horizon. When we look at horizons as long as four years, the negative correlations remain. It is uncertain whether a consumption-portfolio choice model can explain these negative correlations, but simply including frictions in consumption transactions is clearly inadequate. The anomalous behavior of the series on services from owner-occupied housing is particularly disturbing, because this series is a component of the standard nondurables-plus-services measure of consumption that has been used to test consumption-based pricing models since Hansen and Singleton (1982). It would be of interest to see whether the poor performance of this class of models is attributable, in part, to the inclusion of this anomalous consumption series.

The benchmark model appears more successful when compared to the statistics for the nondurable consumption measures, displayed in the first three rows of Table II. As noted in Section IV.A, plausible adjustment costs for nondurable consumption must be extremely small, so we only consider values of λ less than or equal to 0.002. Consider first the statistics computed using CEN, our data on consumer expenditures on nondurable goods. The mean of monthly aggregate consumption growth in the benchmark case for all values of λ is within two standard errors of the value of 0.0007 estimated using CEN data. (It is barely outside the two standard deviation band for the NDS series.) The standard deviation of consumption growth is implausibly high in the absence of adjustment costs. In particular, the benchmark model with $\lambda = 0$ implies a monthly consumption-growth standard deviation of 0.0195, more than 2.5 times the value of 0.0077 computed from nondurable consumption growth data. Increasing λ to 0.002 reduces in the model's consumption-growth standard deviation to 0.0127, accounting for about 58 percent of the gap between the standard deviation in the frictionless model and that computed from CEN data. However, this improvement still leaves the model more than two standard errors above the estimated value. The discrepancy between the model and the data is somewhat larger when the model is compared to ether the CES or the NDS series.

We obtain similar results for the theoretical equity premium. The equity premium we compute from financial data is 6.25 percent per year, more than an order of magnitude higher than the theoretical equity premium of 0.26 percent computed from data on nondurable consumption. When the adjustment cost parameter λ is increased to 0.002, and temporal aggregation is accounted for, the model's prediction is reduced to 1.32 percent per year. That is, approximately 82 percent of the gap between the observed equity premium and the theoretical equity premium computed from CEN data can be explained by temporal aggregation and extremely small consumption adjustment costs. Nonetheless, this statistic in the model remains more than two standard deviations above the value estimated from the data. As with the consumption-growth variance, the model fares somewhat worse when compared to the CES or NDS series.

Finally, for completeness, we consider how well the model matches the moments of investment in the data. In our model, investment is unconsumed output (net of adjustment costs), so it corresponds to data on gross investment. Panel C of Table II displays means and standard deviations for quarterly growth rates of real gross private domestic investment and for real gross savings. Our benchmark model comes rather close to matching both the mean and the standard deviation of savings growth. The benchmark model understates the mean and standard deviation of aggregate investment growth, but the standard errors on these statistics are quite large. In short, the implications of our model for aggregate savings and investment do not appear strongly counterfactual.

V. Conclusions

The results of this paper support our conjecture that fixed costs of consumption adjustment have quantitatively important implications for the comovements of aggregate consumption and asset returns. In particular, undetectably small costs of consumption adjustment can account for a nontrivial portion of the discrepancy between the CCAPM and the behavior of asset markets, but only if consumption is interpreted as purchases of nondurables. However, adjustment costs alone are not sufficient to fully reconcile the statistical behavior of asset returns with the CCAPM. This result is consistent with Daniel and Marshall (1996, 1997), who find that the consumption-based pricing model can perform fairly well at a two-year horizon (a time horizon at which the effects of small adjustment costs are attenuated), but only if time-nonseparabilities are introduced into preferences. It would be of interest to examine a version of this paper's model with time-nonseparable preferences. However, such an exercise would require introducing an additional state variable into the model (a variable that measures the effect of lagged consumption on current utility), which would vastly complicate the analysis. We do not attempt such an extension in this paper.

Our results do not support Reider's (1993) conjecture that illiquidity in *durable* goods can explain the poor empirical performance of the CCAPM. We find that even high adjustment costs cannot reconcile the high volatility of asset returns with the low volatility of consumption services from durables or housing. More interestingly, we find that measures of the growth rate of consumption services from durables or housing covary negatively with excess equity returns. In the context of the CCAPM, this would imply a negative equity premium. These negative covariances persist even when the horizon is lengthened up to four years, suggesting that the problem is not simply a mistiming between the consumption-services data and the equity-return data. We conclude that the CCAPM cannot be saved by including the service flow from durables into the measure of consumption. Indeed, the poor empirical results of the CCAPMs estimated by Hansen and Singleton (1982) and Grossman et al. (1987) may be driven, in part, by the inclusion of the service flow from housing in the consumption data used in these papers.

Our results suggest that predictions of consumption-based models for the short-run movements of asset returns are nonrobust to extremely small perturbations. It follows that the consumption-based approach is not well-suited to modeling high frequency asset-price behavior. However, this class of models still may be useful for understanding the comovements of asset returns and macroeconomic quantity variables at cyclical or longer frequencies.

Appendix: Derivation of Optimality Conditions; Discussion of Simulation Procedure

This sketch of the derivation of the optimality conditions follows Grossman and Laroque (1990). For a more complete description, the interested reader can refer to that paper.

Let $V(Q,c)$ denote the maximized value of the objective in equation (7), starting from initial conditions (Q,c). If τ denotes the first date at which it is optimal to adjust consumption, and c^* denotes the consumption level immediately following that adjustment, then $V(Q,c)$ satisfies

$$V(Q,c) = \sup_{c^*,\tau,(X_t,t\in[0,\tau])} E\left[\int_0^\tau e^{-\delta t}\frac{c_t^a}{a}\,dt + e^{-\delta\tau}V(Q_{\tau^-} - \lambda c_{\tau^-}, c^*)\right], \tag{A1}$$

subject to equations (3), (4), (5), and (6). Use equation (12) to transform equation (A1) in terms of the single variable y:

$$h(y) = \sup_{\tau,x_t} E\left[\int_0^\tau \frac{e^{-\delta t}}{a}\,dt + e^{-\delta\tau}My_{\tau^-}^a\right], \tag{A2}$$

where $x_t \equiv X_t/c$, and the maximization is subject to the appropriately transformed constraints. In particular, the law of motion between consumption adjustments for the new state variable, y, can be obtained from equation (4):

$$dy = [r(y+\lambda) - 1]dt + x[\mu\,dt + \sigma\,dw], \tag{A3}$$

and the no bankruptcy condition (6) becomes $y_t \geq 0$.

Grossman and Laroque's (1990) key insight is that, for fixed M, control problem (A2) is an optimal stopping problem in which the payoff upon stopping in state y is simply My^a. Let $h(y;M)$ denote the function h which solves problem (A2) for a particular fixed M. This function can be characterized using results from the optimal control literature. Once $h(y;M)$ has been determined, h can be computed as the function $h(\cdot\,;M^*)$, where M^* is implicitly defined by:

$$M^* \equiv \sup_y (y + \lambda) h(y;M^*). \tag{A4}$$

The function $h(y;M)$ has the following properties: $h(y;M) \geq y^a M$ (because it is always feasible to stop immediately). If $h(y;M) > y^a M$, then it is not optimal to stop. In this case, consumption equals c for a time interval $[0,\Delta t]$, and the function $h(y;M)$ must satisfy the Bellman equation,

$$h(y_0) = \sup_{x_t} E\left[\int_0^{\Delta t} \frac{e^{-\delta s}}{a}\, ds + e^{-\delta \Delta t} h(y_{\Delta t})\right], \tag{A5}$$

subject to equation (A3) (where the dependence of $h(y;M)$ on M is suppressed for simplicity). Taking limits as $\Delta t \to 0$ and applying Ito's lemma, one obtains the Hamilton-Jacobi–Bellman equation,

$$\sup_x \left[\frac{h''(y)}{2} x^2 \sigma^2 + h'(y)[x\mu + r(y+\lambda) - 1] - \delta h(y) + \frac{1}{a}\right] = 0. \tag{A6}$$

The optimal portfolio choices can be found by maximizing equation (A6) with respect to x. Doing so, one obtains equation (16). By substituting equation (16) into equation (A6), one obtains equation (13). If $h(y) = y^a M$, then it is optimal to stop. At these stopping times, the smooth pasting condition $h(y) = ay^{a-1}$ must hold.

Finally, it can be shown, using a proof strategy analogous to that used in Grossman and Laroque (1990, Theorem 3.3), that the set $\{y \mid h(y) \geq y^a M\}$ is a closed interval $[y_1, y_2]$, so the optimal strategy is to keep consumption constant until $y = y_1$ or $y = y_2$, at which point consumption is adjusted to set $y = y^* \in \operatorname{argmax}_y (y + \lambda)^{-a} h(y)$. Equations (14) immediately follow.

Grossman and Laroque's (1990) model is formulated in terms of the stock of a durable consumption good, denoted K, which depreciates at rate α. Their model is equivalent to ours under the following change of variables (where a given variable in the Grossman–Laroque setup has a "GL" superscript): $r^{GL} \equiv r + \alpha$; $K \equiv c/r^{GL}$; $y^{GL} \equiv r^{GL} y$; $\lambda^{GL} \equiv r^{GL}\lambda$. By substituting these variable definitions into equations (4), (5), and (6), one obtains Grossman and Laroque's equations (2.3), (2.4), and (2.5). Finally, to show that equilibrium condition (13) is equivalent to Grossman and Laroque's equation (3.5), one uses the facts that $h(y) = h(y^{GL})$ (since the transformed value function h is homogeneous of degree zero in $y + \lambda$),

$$h'(y) \equiv h'(y^{GL}/r^{GL}) = r^{GL} \frac{\partial h(y)}{\partial y^{GL}}, \tag{A7}$$

and

$$h''(y) \equiv h''(y^{GL}/r^{GL}) = (r^{GL})^2 \frac{\partial^2 h(y)}{\partial (y^{GL})^2}. \tag{A8}$$

To solve the model we choose an initial guess for M and an initial guess for y_1, set $h(y_1) = y_1^a M$ and $h'(y_1) = a y^{a-1} M$, continue h by solving differential equation (13) using a fifth-order Runge–Kutta algorithm, until a value for $y > y_1$ is found where $h'(y) = a y^{a-1} M$. If, at this value of y, it is also the case that $h(y) = y^a M$, then this value of y corresponds to y_2, and the computed path for $h(\cdot)$ corresponds to $h(\cdot;M)$ for the given M. If not, we adjust y_1 and repeat the process until we have an $h(\cdot,M)$ process for which all boundary conditions hold. We then evaluate our initial guess for M by computing, as a criterion function, $S(M) \equiv [M - \max_y (y + \lambda)^{-a} h(y;M)]^2$. The algorithm searches for a value M^* for which $S(M^*) = 0$.

REFERENCES

Arrow, Kenneth, 1964, The role of securities in the optimal allocation of risk-bearing, *Review of Economic Studies* 31, 91–96.

Brayton, Flint, and Eileen Mauskopf, 1985, *The MPS Model of the United States Economy* (Division of Research and Statistics, Board of Governors of the Federal Reserve System, Washington, D.C.).

Breeden, Douglas T., 1979, An intertemporal asset pricing model with stochastic consumption and investment opportunities, *Journal of Financial Economics* 7, 265–296.

Caballero, Ricardo J., 1993, Durable goods: An explanation for their slow adjustment, *Journal of Political Economy* 101, 351–384.

Cochrane, John H., 1989, The sensitivity of tests of the intertemporal allocation of consumption to near-rational alternatives, *American Economic Review* 79, 319–337.

Cochrane, John H., and Lars P. Hansen, 1992, Asset pricing explorations for macroeconomics; in Olivier Blanchard and Stanley Fischer, eds.: *1992 NBER Macroeconomics Annual* (MIT Press, Cambridge, Mass.).

Constantinides, George M., 1990, Habit formation: A resolution of the equity premium puzzle, *Journal of Political Economy* 98, 519–543.

Constantinides, George M., and Darrell Duffie, 1996, Asset pricing with heterogeneous consumers, *Journal of Political Economy* 104, 219–240.

Daniel, Kent D., and David A. Marshall, 1996, Consumption-based modeling of long-horizon returns, Unpublished manuscript, Northwestern University.

Daniel, Kent D., and David A. Marshall, 1997, The equity premium puzzle and the risk-free rate puzzle at long horizons, *Macroeconomic Dynamics* 1, 452–484.

Grossman, Sanford J., and Guy Laroque, 1990, Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods, *Econometrica* 58, 25–52.

Grossman, Sanford J., Angelo Melino, and Robert J. Shiller, 1987, Estimating the continuous-time consumption-based asset-pricing model, *Journal of Business and Economic Statistics* 5, 315–327.

Grossman, Sanford J., and Robert J. Shiller, 1982, Consumption correlatedness and risk measurement in economies with non-traded assets and heterogeneous information, *Journal of Financial Economics* 10, 195–210.

Hansen, Lars P., 1982, Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029–1054.

Hansen, Lars P., and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 225–262.

Hansen, Lars P., and Kenneth Singleton, 1982, Generalized instrumental variables estimation of nonlinear rational expectations models, *Econometrica* 50, 1269–1286.

He, Hua, and David M. Modest, 1995, Market frictions and consumption-based asset pricing, *Journal of Political Economy* 103, 94–117.

Heaton, John, and Deborah J. Lucas, 1996, Evaluating the effects of incomplete markets on risk-sharing and asset pricing, *Journal of Political Economy* 104, 443–487.

Krusell, Per, and Anthony A. Smith, Jr., 1998, Income and wealth heterogeneity in the macroeconomy, *Journal of Political Economy* 106, 867–896.

Lucas, Robert E., Jr., 1978, Asset prices in an exchange economy, *Econometrica* 46, 1929–1946.

Luttmer, Erzo G. J., 1996, Asset pricing in economies with frictions, *Econometrica* 64, 1439–1468.

Lynch, Anthony, 1996, Decision frequency and synchronization across agents: Implications for aggregate consumption and equity return, *Journal of Finance* 51, 1479–1498.

Marcet, Albert, and Kenneth Singleton, 1991, Equilibrium asset prices and savings of heterogeneous agents in the presence of incomplete markets and portfolio constraints, Unpublished manuscript, Stanford University.

Marshall, David A., 1994, Asset return volatility with extremely small costs of consumption re-optimization, Working paper WP-94-22, Federal Reserve Bank of Chicago.

Mehra, Rajnish, and Edward Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.

Merton, Robert C., 1971, Optimal consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* 3, 373–413.

Newey, Whitney, and Kenneth West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Reider, Robert L., 1993, Can transaction costs on durables explain high estimates of risk aversion?, Unpublished manuscript, Wharton School, University of Pennsylvania.

Rubinstein, Mark, 1974, An aggregation theorem for securities markets, *Journal of Financial Economics* 1, 225–244.

Telmer, Chris I., 1993, Asset pricing puzzles and incomplete markets, *Journal of Finance* 48, 1803–1832.