Uncertainty and consumer durables adjustment*

Giuseppe Bertola[†] Luigi Guiso[‡] Luigi Pistaferri[§]

First draft: March 2001.

This draft: July 2001.

Preliminary – Comments welcome

Abstract

We study the effect of uncertainty on optimal adjustment of durable good stocks, and confront the relevant theoretical insights with a data set featuring extensive information on disaggregated durable goods and subjective measures of future income uncertainty. We discuss the intertemporal optimization problem of a consumer who derives utility both from non-durable consumption flows and from use of durable goods, in the presence of adjustment costs and idiosyncratic income uncertainty. The model delivers implications for the cross-sectional distribution of the durable/non durable ratio, the probability of adjustment, and the size of adjustment. We discuss observable counterparts of such features and assess the extent to which empirical evidence conforms with theoretical predictions. In so doing, we identify and exploit exclusion restrictions based on the distinction between cross-sectional heterogeneity of the sampled households' dynamic problems, and history-dependend variation in their situation at the beginning of the observation period. We argue that the latter should bear on the likelihood but not the size of stock adjustment decisions, and find broad support for theoretical predictions in formal selection-controlled regressions based on this simple insight.

Keywords: Adjustment costs, uncertainty, durable demand.

JEL Classification: D11, D12, D91.

*This paper is part of the "Finance and Consumption in the EU" research project, sponsored by Findomestic SpA and CELETEM at the European University Institute. Guiso acknowledges financial support from the EUI "Finance and Consumption in the EU" Chair, MURST, and the European Commission; Pistaferri acknowledges financial support from the "Taube Faculty Research Fund" at SIEPR. This draft benefits from helpful comments by John Pencavel, Ed Vytlacil, participants at the first Economics of Consumer Credit workshop (EUI, March 2001), and our discussant, Fumio Hayashi. We are also very grateful to Mirko Wiederholt and Winfried Koeniger for competent research assistance and helpful suggestions.

[†]European University Institute and Università di Torino.

[‡]Ente Luigi Einaudi for Monetary, Banking and Financial Studies, University of Sassari, and CEPR.

[§]Stanford University and CEPR.

1 Introduction

Over the last ten years, a growing literature has focused on realistic models of intermittent adjustment at the microeconomic level. Examples include the adjustment of durable goods (the focus of this paper), labor demand, investment, inventories, and cash balances. Intermittent large adjustments can be explained in an optimizing framework by the observation that microeconomic adjustment cost functions are often *kinked* at the no-adjustment point (Bertola and Caballero, 1990). Inaction is costless *per se*, but does affect the objective function through the implications of an unchanged stock of endogenous durables for flow benefits. Conversely, at the microeconomic level action typically entails first-order adjustment costs, which interact importantly with the extent of uncertainty about future dynamics of the problem's forcing variables. Even small adjustment costs can justify very large deviations of durables stocks from statically optimal levels when such deviations are volatile, and likely to be erased soon by exogenous developments rather than by costly action. This intuitive point may also be expressed in terms of option values: action becomes less attractive in a more uncertain environment, where the option of remaining inactive is more valuable.

In this paper, we examine the relevant theory's implications for the relationship between optimal dynamic adjustment policies and a variety of microeconomic features of an individual consumer's problem, and we then bring them to bear on a data set featuring extensive information on durable purchases. The data set also contains subjective measures of future income uncertainty, which are of particular interest in our theoretical and empirical work. Theory predicts that higher uncertainty should increase the likelihood of wide deviations from the preferred durables stock, but also widens the range of inaction. Hence, conditionally on the current state, higher uncertainty about the future evolution of the problem's forcing variables implies that immediate adjustment is less likely, and also that adjustment is larger if it does occur. These are *ceteris paribus* statements, of course, and theory offers analogous comparative-dynamics results for the drift, the extent of adjustment costs, and taste parameters of optimal adjustment problems solved by the consumers sampled in our data. In keeping with the theoretical predictions, we consider the implications of a more pronounced drift in expected durable consumption (measured on the basis of optimality conditions on that margin) and of adjustment costs (proxied by measures of bureaucratic inefficiency) for optimal durables stock adjustment.

In our empirical work, controlled regression techniques make it possible to disentangle the conceptually different effects of uncertainty and other factors on the frequency and size of adjustment. Econometric identification is achieved by distinguishing two sources of variation in the intensity and size of durables adjustment decisions by the cross-section of consumers in our data set. Cross-sectional heterogeneity of consumers' and goods' characteristics bears on both the frequency and size of optimal infrequent adjustment policies. Whether a consumer adjusts the stock of a durable good during the observation period, however, also depends importantly on the dynamic history that (for given parameters of the dynamic problem and optimal adjustment policy) brought him or her close to the boundaries of the inaction range. The latter dynamic aspects are summarized in the data, after conditioning on observable characteristics that are relevant for the optimal adjustment policy, by the beginning-of-period value of the stock of durables in relation to nondurable consumption. This variable conveys information as to the consumer's position within the inaction band at the beginning of the observation period and to the likelihood of adjustment. Conditionally on adjustment taking place, however, the initial position in the inaction band does *not* influence the size of durables stock adjustment, which reflects forward-looking considerations and is therefore influenced by uncertainty, adjustment costs, and other observable and unobservable cross-sectional characteristics rather than by past history. Thus, an infrequent-adjustment perspective offers a theoretically sound exclusion restriction to standard selection-controlled estimation techniques.

We are not aware of previous empirical studies of the joint effects of uncertainty, adjustment costs, and drift on both the extensive and intensive margin of the durable adjustment decision. Despite the realism of infrequent adjustment models and their potential importance for explaining aggregate phenomena, in fact, relatively few studies test and estimate such models on microeconomic data, and fewer still focus on the empirical role of uncertainty. Lam (1991) uses PSID data to estimate the parameters of a threshold adjustment model in an extended permanent income hypothesis model, and finds evidence for liquidity constraints and resale market imperfections. More recently, Attanasio (2000) estimates a semi-structural model of car purchases on a sample of U.S. households drawn from the CEX, focusing on a characterization of trigger and return points rather than on the more structural features of the model. Eberly (1994) is one of the first attempts to test the effect of uncertainty on durable purchases decisions. Using SCF data, she finds that uncertainty affects positively the width of the inaction bands. However, her measure of uncertainty excludes idiosyncratic labor income risk, which is likely to be the most relevant one at the micro level. Foote, Hurst, and Leahy (2000) is perhaps the closest antecedent to our work, and offers a study of the impact of uncertainty on durable adjustment (from CEX data) using the imputed variance of household income obtained from regressions estimated with PSID data as a measure of uncertainty. They find that, controlling for the age of the car, uncertainty has (at least in some regressions) a negative effect on the probability of adjusting and interpret this result as evidence against standard models of optimal inaction. Their proxy for uncertainty, however, may be contaminated by measurement error in income, and their econometric procedure assumes that households rely on the same information set as the econometrician. Their specification hinges on fairly arbitrary exclusion restrictions on the variables that predict the variance of income but have no effect on the durable adjustment decision. Our work has also the advantage of using a measure of uncertainty that is likely to be less controversial than those used in the literature, and emphasis on the distinction between the effects of uncertainty on the decision to adjust and on the size of adjustment supports a neater empirical implementation of the model. In the context of the literature, our approach may also have important implications for understanding the role of uncertainty in the dynamics of aggregate *per capita* durable purchases. Since at any given point in time the latter is the product of the fraction of households adjusting (given their stock of durables at adjustment time) and the value of the adjustment for those adjusting, the effect of changes in uncertainty is composed at the aggregate level of conflicting effects on the two margins, and may well be of small empirical relevance.

The rest of the paper is organized as follows. Section 2 sets up the relevant theoretical framework. We first characterize frictionless intertemporal consumption choices, which would imply that the distribution across consumers of the durables stock/nondurables flow ratio depends on preferences and relative prices. We then introduce adjustment costs and show that an approximate solution can be expressed in terms of action and return points of the optimal policy terms of (log) deviations of the durable/nondurable ratio from its no-adjustment-costs level. In Section 3 we introduce our data, discuss available observable counterparts for theoretical uncertainty, adjustment-cost, and drift parameters, and review theoretical predictions and empirical evidence on the relationship of such parameters to the shape of cross-sectional distributions and to the frequency and size of stock adjustment. In Section 4 we estimate formal models for the probability of adjustment and for the size of adjustment conditional upon adjusting, using a Heckman selectivity model. We also evaluate unconditional effects of changes in uncertainty, drift, adjustment costs, and taste for durables. Section 5 concludes.

2 Theoretical framework

In this section we outline a general formal model and discuss how it may be reduced to an approximate, but still quite detailed, and relatively tractable form. Let period utility be a function of nondurable consumption flows, $C(t)$, and available stocks of durable goods, $X(t)$ at time t. In continuous time, the consumer solves

$$\max_{\{C(\tau),X(\tau)\}} E_t\left[\int_t^{\infty} e^{-\rho(\tau-t)}u(C(\tau),X(\tau))\,d\tau\right] \tag{1}$$

$$\begin{aligned}
\text{s.t. } A(\tau_+)-A(\tau) &= [r(\tau)A(\tau)+y(\tau)-C(\tau)]\,d\tau-\Delta X(\tau)\,p(X(\tau_+),X(\tau),\tau)\\
X(\tau_+)-X(\tau) &= -X(\tau)\,\delta d\tau+\Delta X(\tau)
\end{aligned}$$

where $u(\cdot)$ is an increasing and concave function, ρ the consumer's discount rate, $A(\tau)$ the level of assets and $y(\tau)$ the flow of labor income at time τ, and $r(\tau)$ is the rate of return on assets. The notation allows for unrestrained borrowing and lending on a non-contingent basis. The price of durables in terms of nondurables, $p(\cdot)$, may vary exogenously over time but, as discussed below, also depends on the size and sign of durables stock changes, $\Delta X(\tau)$, which result from sales or purchases (adjustment) rather than from depreciation in use, which occurs continuously at rate δ. The slightly nonstandard notation $A(\tau_+)$, $X(\tau_+)$ allows for possible discontinuity or non-differentiability of the durables stock path at times when discrete adjustment is undertaken. If $\Delta X(\tau)=0$, then the budget constraint is understood to be

$$\lim_{\tau_+\to\tau}\frac{(A(\tau_+)-A(\tau))}{\tau_+-\tau}=r(\tau)A(\tau)+y(\tau)-C(\tau),$$

and accumulation of durables stocks is similarly understood to obey the standard dynamics $dX(\tau)/d\tau=-\delta X(\tau)$ when $\Delta X(\tau)=0$. At times when the consumer adjusts durables stocks by purchases or sales, however, he or she does so by discrete or non-differentiable (singular) increments, and the dynamics of assets and durables stocks behave similarly.

The optimization program (1) realistically allows both durable and nondurable consumption to yield utility, and both income and asset returns to be random. As a result, the problem is analytically intractable, and numerical analysis must rely on drastic simplifications. The classic Grossman and Laroque (1990) study of optimal durable consumption, for example, abstracts from nondurable goods and labor income to obtain analytic and numerical results for the case where asset returns are described by Brownian increments, and utility has constant elasticity. Other researchers have chosen different simplifications and approximations (see Attanasio, 2000, and Padula, 2001, for references and discussions). Our empirical analysis, like those of Attanasio (2000), Eberly (1994), and others, would not be able to detect and test many of the properties of realistic models, even if a tractable specification were available. The data we analyze do contain detailed information on nondurable goods and stocks, and on labor-income expectations and uncertainty. Hence, we follow Bar-Ilan and Blinder (1992) and Bertola and Caballero (1990) in studying approximate durable consumption choices in the presence of labor income risk. And since the information in the data is mostly cross-sectional in nature, we pay particular attention to the implications of such approximations for the steady-state distribution of stocks and adjustments across a possibly heterogeneous population.

In the absence of adjustment costs, the optimality condition

$$\frac{\partial u(C(t),X(t))}{\partial C(t)}v(t)=\frac{\partial u(C(t),X(t))}{\partial X(t)} \tag{2}$$

would hold at all times, where

$$v(t)\equiv(r(t)+\delta)\,p(t)-\frac{1}{dt}E[dp(t)]$$

is the user cost of the durable good. Intuitively, if the unit price $p(\cdot)$ of durables purchased or sold does not depend on size and direction of stock changes, then sales and purchases of the durables stock absorb and release funds that could also finance nondurable consumption flows. For any strictly concave utility function condition (2) is necessary and sufficient to ensure that the consumer could not obtain a larger utility flow from the purchasing power expended at each point in time.

To interpret the approximations proposed below, it will be useful to consider specializations of the general problem. When $u(\cdot,\cdot)$ is homothetic, (2) implies that the optimal ratio of durable to nondurable consumption is a function $z(\cdot)$ of the user cost,

$$\frac{X(t)}{C(t)}=z(v(t)),\quad z'(\cdot)<0.$$

If the period utility flow is well approximated by a constant elasticity function $u(\cdot)=\alpha\log C+\beta\log X$ of its arguments, then

$$\frac{X(t)}{C(t)}=\frac{\beta}{\alpha}\frac{1}{v(t)}. \tag{3}$$

To characterize the behavior of nondurable (and durable) consumption over time, one may use the optimality condition implied by perfect borrowing and lending opportunities. Writing the rate of return on assets as a constant for simplicity, the Euler condition reads

$$\frac{\partial u(C(t), X(t))}{\partial C(t)} = e^{(r-\rho)\tau} E_t \left\{ \frac{\partial u(C(t+\tau), X(t+\tau))}{\partial C(t+\tau)} \right\}, \tag{4}$$

and requires that discounted (by $e^{(r-\rho)\tau}$) marginal utility changes be unpredictable. If marginal utility has (approximately) constant elasticity γ and the logarithm of future consumption is (approximately) normally distributed with conditional variance σ^2, for example, the expected growth rate of consumption can be written in the form (see e.g. Carroll, 1997)

$$E_t \left\{ \Delta \log(C(t)) \right\} \approx \frac{r-\rho}{\gamma} + \frac{\gamma}{2}\sigma^2. \tag{5}$$

A similar characterization is appropriate for the stock of durables when, as is the case under homotheticity, it is proportional to the nondurable consumption level.

In the absence of adjustment costs, conditions (2) and (4) and the relevant version of the budget and accumulation constraints in (1) suffice in principle to determine optimal consumption choices, which are well defined under mild regularity conditions. An explicit solution is not available in general, but the specializations and approximations above offer useful qualitative insights. If utility flows are a constant-elasticity function of nondurable consumption flows and durables stocks, and the user cost of durables is constant, then the durable/nondurable consumption ratio should be constant as well. The exact consumption function is of course highly nonlinear in the presence of precautionary savings. If (5) is a satisfactory approximation, however, future expected consumption is proportional to current consumption, and an expected-value version of the consumer's budget constraint implies that each of the durable and nondurable consumption levels should be approximately proportional to the consumer's total wealth (the coefficient of proportionality depends on the rates of return, of utility discount, and of expected marginal utility growth). Hence, the volatility of consumption should be an increasing function of lifetime wealth volatility, which is in turn positively related to the volatility of permanent and transitory labor income shocks if the rate of return on assets and the user cost of durables are not random.

2.1 Adjustment costs

The above discussion of intertemporal and intratemporal choices provides a useful starting point for the characterization of optimal infrequent durable adjustment. In reality, the unit price applicable to durable purchases is different from that applicable to durable sales, even after taking depreciation into account, because second-hand markets generally price used goods at less than the seller's valuation from continued use. Moreover, adjustment can entail lump-sum costs, independent of the size of adjustment, for example because exchanging one's stock of durables requires a costly trip to a dealer's shop. This implies that the unit price of very small net purchases of durables is very large, and would tend to infinity if the consumer tried to track continuous changes over time of the optimal durables stock implied by (2); it also implies that sales of one's durable goods may only be optimal if their revenue (the unit price obtained on the second-hand goods market, times the quantity sold) is at least as large as the lump-sum transaction cost, for otherwise the consumer would forsake use of durable goods without any gain in purchasing power terms. Thus, the policy of potentially very frequent purchases and sales that would allow the consumer's durables stock to remain aligned with nondurable consumption as in (2) cannot be optimal: if the utility function is differentiable, small misalignment of the durable/nondurable ratio cannot have a first-order impact on utility flows, and should not be corrected if adjustment entails first-order transaction costs.

We proceed to formalize such intuitive insights, following closely the approach and methods of earlier theoretical studies. With an eye to our empirical application below, however, we discuss in some detail the nature of the approximation, and we express it in terms of deviations from the durable/nondurable consumption bundle given by (2) or (3) rather than in terms of wealth allocation.

Bertola and Caballero (1990) characterize the optimal timing and size of infrequent costly adjustment of durables (or other stock variables) towards a statically optimal level $X^*(t)$ when the deviations from that level follow a random walk with drift in continuous time, and entail quadratic flow losses. In our specification of the consumer's problem without adjustment costs, the consumer at each point in time allocates a flow E of purchasing power to expenditure on nondurable goods, C, and imputed durable user costs, υX. This static problem,

$$\max_{X} \left\{ \alpha \log(E - \upsilon X) + \beta \log(X) \right\},$$

selects

$$X^* = \frac{\beta}{\alpha + \beta} \frac{E}{\upsilon}$$

and achieves the optimal utility flow

$$\alpha \log(E - \upsilon \frac{\beta}{\alpha+\beta}\frac{E}{\upsilon}) + \beta \log(\frac{\beta}{\alpha+\beta}\frac{E}{\upsilon}) \equiv u^*(E).$$

Along the intertemporal dimension, the process followed by the "expenditure" flows $E(t)$ satisfies the budget constraint by definition and, since

$$C(t) = \frac{\alpha}{\alpha+\beta} E(t),$$

has the unpredictable-increment properties familiar from standard applications of Euler conditions in the form (4) or (5).

To obtain an approximate characterization of the problem in the presence of adjustment costs, we begin by noting that an arbitrary durables stocks process $X(t)$, for the same sequence of $E(t)$ flows, would generally yield smaller utility flows. Log-approximating utility flows around $X^*(t) = \frac{\beta}{\alpha+\beta}\frac{E(t)}{\upsilon(t)}$, the first-order term vanishes by optimality and neglecting third-order and higher terms yields

$$\alpha \log(E - \upsilon X) + \beta \log(X) - u^*(E) \approx -\frac{\beta}{\alpha}(\alpha+\beta)\left[\log(X) - \log\left(\frac{\beta}{\alpha+\beta}\frac{E}{\upsilon}\right)\right]^2.$$

By definition, the independent-increments sequence $E(t)$ of expenditures on nondurable goods and imputed user costs for durable goods satisfies the intertemporal budget constraint in the absence of transaction costs. It can also be viewed as an approximation to its budget-constrained counterpart when transaction costs are present. A more precise approximation should adjust the overall level of expenditures so as to finance transaction costs, whose frequency and size can only be ascertained in the context of the whole problem's solution.

Hence, the utility flow accruing to a consumer who optimally allocates purchasing power over time, continuously adjusts the nondurable consumption flow $C(t)$, and uses a given durables stock $X(t)$ can be approximated as follows:

$$-\frac{\beta}{\alpha}(\alpha+\beta)\left(\log(X(t)) - \log\left(\frac{\beta}{\alpha}\frac{C(t)}{\upsilon(t)}\right)\right)^2 = -\frac{b}{2}\left(x(t) - c(t) + \log \upsilon(t) - \kappa\right)^2, \qquad (6)$$

where

- $b = \frac{2\beta}{\alpha}(\alpha+\beta)$ is the slope of marginal utility losses due to deviations from the optimal durable/nondurable ratio along the optimal intertemporal expenditure pattern,
- $x(t) = \log X(t)$ is the log durables stock,

- $c(t) = \log C(t)$ is the log nondurable consumption flow,
- $v(t)$ again denotes the user cost of the durables stock (which may be constant, or vary along with the relative price of durables), and
- $\kappa = \log \frac{\beta}{\alpha}$.

In (6), the slope of marginal utility flow losses, b, is related to structural parameters of the consumer's problem, specifically to the ratio β/α of the durable and nondurable weights (or budget shares) in a log-linear utility flow expression. The steps outlined above yield similar expressions if the consumer derives utility from more than one stock of durables, with possibly different budget shares. Hence, this representation of the consumer's problem lends itself naturally to a study of cross-sectional data with information about different durable goods and nondurable consumption flows.

Of course, (6) is only an approximate representation of the consumer's general problem, and more approximations need to be introduced in order to bring results from Bertola and Caballero (1990) and others to bear on our empirical analysis. Standard solution methods are applicable if we define

$$z(t) \equiv x(t) - c(t) + \log v(t) - \kappa \tag{7}$$

and proceed to study the optimization problem

$$V(z_t) \equiv - \min_{\{z_\tau\}} E_t \int_t^\infty e^{-\rho(\tau - t)} \left(\left(\frac{b}{2} z(\tau)^2 \right) d\tau + [\text{adjustment costs}] \right) \tag{8}$$

where adjustment costs, expressed in terms of utility, have the "kinked" form studied in earlier work. Namely, let increasing x by $\Delta x > 0$ cost $C_l + c_l \Delta x$, for c_l the unit cost of upward adjustment of the durables stock, and C_l the lump-sum cost of purchasing any positive amount of durables. Similarly, let decreasing the state variable by $\Delta x < 0$ cost $C_u + c_u |\Delta x|$. All cost parameters are expressed in terms of utility, but an interpretation in terms of expenditure is available since, for a logarithmic specification of intertemporal utility, the utility price of wealth is inversely proportional to nondurable consumption and to the bundle of durable and nondurable goods (denoted by E above) around which the actual consumption pattern fluctuates over time. To ensure that adjustment costs do not become negligible or enormous as consumers become increasingly rich or poor over time, the model needs to assume that adjustment costs are proportional to the consumer's wealth, or consumption. In practice, constant adjustment costs in terms of utility can be viewed as approximate representation of a possibly realistic transaction cost technology whereby the lump-sum cost of adjustment is proportional to the consumer's durables stock (or wealth, or overall consumption, as may be realistic noting that wealthier consumers have higher opportunity costs of time), and proportional transaction costs are expressed in percentage terms of stock changes.[1]

An explicit solution is available if $z(t)$ follows an arithmetic Brownian motion process when the consumer refrains from purchases or sales of durables, i.e., if

$$[\triangle X(t) = 0] \quad \Rightarrow \; dz(t) \equiv \vartheta dt + \sigma dW(t)$$

where ϑ and σ are the drift and the standard deviation of the process. This representation of stochastic increments is approximately realistic, in light of the definition of $z(t)$ in (7), if the logarithm of the durables stock depreciates linearly, and the user cost $v(t)$ is either constant or influenced by geometric inflation of the durable's relative price. Along the intertemporally optimal path, the logarithm of nondurable consumption, $c(t)$, follows an unpredictable-increments process with drift. Hence, the parameter ϑ reflects both depreciation of the durables stock and the drift in nondurable consumption. The assumption of normally distributed percentage increments with constant variance per unit time, of course, can only be viewed as a useful approximation: in the exact solution of the general problem consumption increments may be heteroskedastic, and their predictable component need not be constant over time. Qualitatively, however, this representation has interesting empirical implications. The variance of consumption growth, while not necessarily constant over time, can in general be assumed to be monotonically increasing in the variance of uninsurable

[1]To express this specification of adjustment costs in terms of goods prices, as in the notation of problem (1), note that the formulation

$$p(X(\tau_+), X(\tau), \tau) = \begin{cases} \frac{C_l}{\triangle X} + c_l & \text{if } \triangle X > 0, \\ \frac{C_u}{\triangle X} - c_u & \text{if } \triangle X < 0 \end{cases}$$

would represent lump-sum (if $C_l > 0$ and/or $C_u > 0$) and proportional (if c_u and/or c_l are positive) adjustment costs in the budget constraint of (1). Lump-sum and proportional adjustment costs depend on logarithmic changes of the durables stock, as in (8), if the total goods-terms cost of a transaction of size ΔX is given by

$$p(X(\tau_+), X(\tau), \tau)\triangle X = \begin{cases} C_l X(\tau) + c_l \triangle X & \text{if } \triangle X > 0, \\ C_u X(\tau) + c_u \triangle X & \text{if } \triangle X < 0. \end{cases}$$

or, using $\triangle X / X \approx \Delta x$,

$$p(X(\tau_+), X(\tau), \tau) = \begin{cases} \frac{C_l}{\Delta x} + c_l & \text{if } \Delta x > 0, \\ \frac{C_u}{\Delta x} - c_u & \text{if } \Delta x < 0. \end{cases}$$

future income. The drift ϑ is negative for realistic parameterization of the problem: when analyzing data on different durables stocks and heterogeneous consumers, it will be helpful to keep in mind that ϑ should be larger in absolute value (more negative) for goods with fast depreciation and/or steeply declining relative price, and for consumers with more positive consumption drift.

The approximate optimization problem can be solved and characterized using standard continuous-time technical tools (see Appendix A). The optimal adjustment policy is described by action and return points (L, l, u, U). Adjustment only occurs when z is at *trigger* points L or U, with $L \leq U$: when z reaches L, control moves it instantaneously to the lower *return* point l, with $L \leq l < U$; and when z reaches U, control moves it back to point u, with $L < u \leq U$. The optimal L, l, u, U can be found as the solution of a relatively simple nonlinear system of equations. In our empirical exercise, we will analyze a set of cross-sectional observations, each of which may be interpreted as a draw from a history of infrequent adjustment similar to that characterized for a single decision maker. In the absence of time-series information on individual behavior, we will find it insightful to interpret the cross-sectional information available in terms of the long-run distribution of the controlled variable, z, within the $[L, U]$ optimal inaction interval. The stable distribution and adjustment intensities of the controlled Brownian motion process of interest are available in closed form (see Appendix A), but the formulae are complex and not particularly informative. The general qualitative properties of the solution, however, are quite clear and should be more general than the specific problem we have outlined. In the next section, we illustrate them graphically and confront them with simple empirical evidence.

3 Data and theoretical implications

A consensus has formed in the literature (see e.g. Attanasio, 2000) that attempts to estimate structural parameters of a realistic infrequent-adjustment model cannot be fruitful. The alternative strategy we adopt focuses on parameters of a semi-structural model instead, and aims at verifying more or less obvious theoretical implications on a microeconomic data set. We analyze household data drawn from the 1995 Survey of Households Income and Wealth (SHIW).[2] The SHIW collects data on income, consumption expenditures on durables and

[2]Measures of subjective uncertainty and detailed information on durables are also contained in the 1998 SHIW. However, in 1997-98 a subsidy program for early scrapping of cars and motorcycles (similar to the one set up in France, see Adda and Cooper, 2000) was in place. This would require modelling the effect of scrapping incentives which is beyond the scope of this paper. We thus focus only on 1995 data.

nondurable goods, financial wealth, real estate wealth, and several demographic variables for a representative sample of about 8,000 Italian households. Since 1989 the SHIW has also a rotating panel component whereby approximately half of the sample units are re-interviewed in the subsequent survey.[3] Information is provided for three durable-goods categories: means of transport (for brevity "cars" or vehicles in what follows); furniture, furnishings, household appliances and sundry articles ("furniture"); precious objects including jewelry, antiques, old and gold coins ("jewelry"). Households report the value of the stock and the value of any sales and purchases during the year. For furniture the value of sales is not available.[4] Stock values are reported as of the end of period while flows refer to the calendar year. In our empirical work, the consumer's state at the time when adjustment decisions are taken is measured as the difference between the end-of-year stock and the net amount purchased during the year (value of any purchases—value of any sales). For those who adjust, this measure reflects closely the value of the stock at adjustment time (neglecting the effect of depreciation), and is thus a better approximation to the continuous-time variable considered in the model than the beginning-of-year stock, typically used in applied work.

The availability of information on various categories of goods makes it possible to assess applicability of optimal inaction policies to durables other than cars, the only case studied in the empirical literature. As we shall see, some features of the model are fairly robust across different categories of durable goods, but others appear more relevant to households' choice of vehicles rather than that of furniture or jewelry. Jewelry purchase patterns are particularly hard to interpret, perhaps not surprisingly in light of the fact that the homotheticity assumption is less convincing, and accounting for taste heterogeneity more difficult, for this particular type of durable.

Table 1 reports summary statistics for the sample used in this paper and for the entire 1995 sample.[5] Values for the two samples appear comparable and confirm the randomness of our sample selection. In the sample we will be working on, the average value of the stock of cars is little more than 6,000 euro, compared to about 10,000 euro for furniture and 3,000 euro for jewelry. The corresponding durables stock-nondurable consumption ratios are 36, 59, and 18 percent, respectively. The share of households adjusting (either upward or downward) is 18 percent for cars, 30 percent for furniture and 10 percent for jewelry.

[3]See Appendix B for a detailed description of the survey contents, its sample design, interviewing procedure and response rates.

[4]Questions on sales of furniture and household appliances are not asked because there is virtually no second hand market for such items in Italy.

[5]The 1995 sample excludes the retired for comparison purposes with the sample used in the estimation below.

Average net family disposable income equals about 25,000 euro; average age is 43 years, average schooling 10 years. The individuals in our sample have a family size of 3 and about one third of them live in the South. Table 2 reports summary statistics for the sub-sample of households that adjust the stock of either durable good. The value of purchases (conditional on adjusting) is much larger for cars than for any of the other two types of durables; just after adjustment, the average ratio of cars to nondurable consumption yearly flows becomes as large as that of furniture.

The data available to us also contain very useful, though certainly not perfect, information on the key parameters of the optimization problem introduced above: the drift and uncertainty of the consumer's income and consumption processes, adjustment costs, and depreciation rates. Table 1 reports summary statistics for some of the variables that convey such information. Definitions are listed in Appendix C; separate subsections below discuss measurement of uncertainty, drift, and adjustment cost parameters, and confront theoretical implications with simple empirical information on the distribution of durable/nondurable consumption ratios. These exercises impose a *ceteris paribus* condition that is likely to be violated. However, they help us connect theoretical insights with empirical evidence, offer information on some interesting unconditional features of the data, and pave the way for the more formal controlled regression framework proposed in Section 4.

3.1 Controlling for heterogeneity

Before proceeding to study the problem's dynamic aspects, however, it is important to recognize that our data set's cross-section of consumers is certainly heterogeneous along many dimensions that would bear on durable consumption patterns even in the absence of adjustment costs. While the predictions of the model refer to a hypothetical infinitely-lived consumer, the data are drawn from a cross-section of demographically heterogeneous consumers with possibly different tastes for durables. In the theoretical model, taste for durables is indexed by the relative budget share β/α. The four panels of Figure 1 plot several aspects of the dynamic problem's solution for a range of such taste parameters. Similar figures will illustrate other comparative results below. One advantage of the explicit approximation procedure outlined above is that other parameters can be roughly calibrated to represent adjustment problems for different durables, and are meant to be roughly realistic if a unit of time is a year (the drift ϑ, the variance σ^2, and the slope of utility losses from misalignment of nondurable consumption *flows* and durables stock*s* all need to be interpreted on a per-period basis). Adjustment cost parameters are in terms of utility units, or fractions of permanent income under the logarithmic utility approximation. Thus, the 0.005 fixed adjustment cost in the figure would indicate that the (opportunity) cost of shopping for the durable is 0.5% of a typical year's utility flow.

In the top-left panel of Figure 1, we see that a larger durables budget share (on the horizontal axis) is not surprisingly associated with increasingly large trigger and return durable/nondurable log ratios on the vertical axis. Also importantly, and less obviously, taste for durables affects not only the frictionless durables/nondurables ratio, but also the frequency and size of adjustments, because it determines the slope of the cost (denoted b above) of departing from the optimal, frictionless policy. Hence, the inaction range becomes narrower and the return points get closer to each other as β increases in the first panel of the Figure. The right-hand panels of the Figure display the predicted long-run density and cumulative distribution of the consumer's deviation from the statically optimal durable/nondurable log ratio, and the bottom-left panel plots the long-run probability intensity and size of adjustment at the upper and lower boundaries of the inaction range. Intuitively, the optimal policies for individuals with strong tastes for durables keeps them closer to the no-adjustment-cost durables/nondurable ratio. Hence, as the utility weight of durables increases the inaction band shrinks, and adjustment becomes more frequent and smaller in size.

This suggests that in order to highlight the implications of dynamic adjustment parameters it is empirically important to account for taste heterogeneity. Our data contain a wealth of information on several dimensions of heterogeneity: when comparing theoretical predictions with the empirical *shape* and *width* of the cross-sectional distribution of the X/C ratio of the stock of durables to the flow of nondurables goods, we will condition on a variety of controls meant to capture differences in tastes across consumers. Specifically, we first filter the data with a regression of X/C on a set of demographic characteristics (a polynomial in age, dummies for gender, education, area of residence, marital status, family size and household composition) and then compute the empirical density of the residuals, using Gaussian kernel nonparametric smoothers.[6]

Figure 2 plots the empirical density of the (residual) X/C for the three types of durable goods in our data set: vehicles (top-left panel), furniture (top-right), and jewelry (bottom). The stock of durables is defined as of the end of the period.[7] The empirical density of

[6]To highlight differences between groups, the densities are evaluated at 25 points over the range of X/C. We also exclude observations in the lower and upper percentile of the distribution.

[7]For households with no durables, the X/C ratio is zero. This is hard to rationalize by a theoretical framework where durables are infinitely divisible and preferences are homothetic. More sophisticated models

X/C for all types of durables is "whale-shaped," closely resembling the plots in the top right panel of theoretical figures. While such a shape could be spuriously generated by uncontrolled heterogeneity in raw data, finding it on filtered data suggests that, at least to a first approximation, the optimal inaction model provides a useful description of the dynamics of durable goods. We now turn to examining the effect of uncertainty, drift, and adjustment costs on the shape of the density of X/C.

3.2 Uncertainty

For our purposes, one important feature of the SHIW data is that, in a special section of the survey, households are asked a set of questions designed to elicit the perceived probability of being employed over the twelve months following the interview and the variation in earnings if employed (see Appendix B for details). We use this information to construct measures of the first two moments of the distribution of future earnings following the methodology developed in Guiso, Jappelli, and Pistaferri (2001), and thus obtain proxies for the idiosyncratic uncertainty faced by the household. This measure of uncertainty is superior to those an econometrician may construct from data realizations, since it is based on the individual's own information as to the evolution and riskiness of future earnings. This makes it possible to isolate risk from predictable (by the household) variability, and provides important overidentifying information.

Subjective income questions are addressed to only half the overall 1995 SHIW sample, after excluding the retired and those who plan to retire in the following year. The selection of the sub-sample reporting subjective expectations is random, based on whether the year of birth of the household head is even or odd. After excluding the households that are not in this sub-sample and those with missing values on the subjective expectations, the stock of durables or purchases and sales, our final sample consists of 1,877 households. In Table 1, the sample average of the coefficient of variation of expected future earnings, used here as a summary measure of uncertainty, is 4 percent. This is smaller than the standard deviation of the innovation to earnings growth obtained from longitudinal data on Italian workers as estimated by Guiso, Pistaferri and Schivardi (2001), who obtain a value of 11 percent. The main reason why subjective uncertainty may be smaller than "objective" estimates is that individuals are likely to condition on a larger information set than the econometrician and rely on variables that may not be observable to the latter. Comparing Table 1 and Table 2, we see that uncertainty in the sub-sample of those who adjust is generally lower than in the

could relax these assumption, at the cost of substantial complication.

whole sample (and so is the fraction of unemployed, of low educated, and of those living in the South).

The four panels of Figure 3 illustrate theoretical effects of different levels of uncertainty, as summarized by the variance σ^2 of the process $z(t)$ representing deviations from the optimal durable/nondurable consumption bundle. Along the vertical axis of the top left panel, we report the trigger and return control points of that process as a function of σ^2 on the horizontal axis (this parameter, like the other dynamic parameters considered below, has negligible effects on the average durables/nondurables ratio; hence, we report z on the vertical axis, rather than x as in Figure 1). Trigger and return points become increasingly spaced apart as uncertainty increases, reflecting the well known option value of inaction in problems with first-order adjustment costs. If a consumer knows that deviations from the optimal configuration of the durable/nondurable bundle are very volatile, he or she will find it optimal not to bear adjustment costs in order to correct large such deviations, and rather wait and see whether random events will correct them costlessly. Note that for a given state z on the vertical axis, the probability of falling outside the inaction range is higher at low than at high levels of uncertainty. Hence, a testable implication is that, conditioning on z, higher levels of uncertainty about future developments should be associated with lower probabilities of adjustment in the immediate future.

In the right-hand panels of Figure 3, larger values of σ are associated with wider dispersion of z deviations in steady state. The trigger and return points are not symmetric around the statically optimal bundle, because the drift implies that positive deviations tend to be corrected more quickly in the absence of adjustment effort. The drift also imparts an asymmetric "whale" shape to the density. Intuitively, optimal adjustment behavior induces largely symmetric probability-weighted deviations from the durables stock that would be optimal in the absence of adjustment costs; this implies that the derivations above, where the imputed expenditure flow was left unchanged when introducing adjustment costs, does provide a good approximation.

The above qualitative features are familiar from previous studies and our data, which include indicators of consumer-specific uncertainty, make it possible to verify them empirically. We can use our household-level measure of earnings uncertainty (which under fairly general conditions is monotonically related to the variability of X/C) to see whether the cross-sectional distribution of X/C becomes more spread-out as the variance of the state variable facing the consumer increases. Figure 4 reports the empirical density of X/C for households with low and high income uncertainty. To highlight differences between groups, we classify as high-income risk households with a standard deviation of future earnings above the 75^{th} percentile of the distribution. The empirical plots are similar to the theoretical ones: the density for high uncertainty households is below that for low uncertainty households everywhere but in the tails. While the difference between the two densities is not large, it appears to be qualitatively consistent with the theoretical predictions regardless of the type of durable considered.

The bottom-left panel of Figure 3 focuses on steady-state implications for the frequency (or intensity) and size of adjustment, and illustrates less well known and perhaps less intuitive features of the solution. As the degree of uncertainty σ increases, the range of inaction widens but the probability of action tends to increase. This phenomenon is not difficult to interpret intuitively. Larger uncertainty, as mentioned, makes inaction optimal in the face of larger z deviations, which are more likely to be corrected by random movements in any given period of time. Since marginal utility losses are increasing in z, however, it would not be optimal to expand the range of inaction so much as to imply that the frequency of adjustment remains unchanged. Such behavior would not affect the adjustment cost component of the objective function, but would imply larger flow losses. Higher variance reduces possible payoffs, but the optimal policy trades larger discounted average flow utility losses off more frequent adjustment. In the bottom-left panel of the figure, higher uncertainty is also associated with larger adjustment size (for given adjustment costs). Empirically, this implies that the *unconditional* probability of observing adjustment should be larger for high-variance samples than for low-variance samples and that, given adjustment, the size of adjustment should be larger in the former than in the latter.[8] It may be interesting to note that the uncertainty measure is indeed positively associated with adjustment in simple probit regressions with no other conditioning variables. In what follows, however, we will be focusing on controlled regression results and, in particular, on the distinction between the effects of uncertainty on the frequency and size of adjustment conditionally on

[8] There is some confusion in the literature about the two effects that changes in uncertainty have on the probability of adjustment. For instance, Foote, Hurst, and Leahy (2000) study the impact of uncertainty on durable (cars), and find that, controlling for the age of the car, uncertainty has (at least in some regressions) a negative effect on the probability of adjusting, and interpret this finding as evidence against increquent-adjustment models. However, if the age of the car in their regression proxies for the value of z, then this finding is consistent with those models. The model predicts that the probability of adjustment *declines* with uncertainty for a *given* durable stock-non durable consumption ratio (what can be termed a z-constant effect). On the other hand, the model predicts a positive correlation between uncertainty and the probability of adjustment when z is allowed to vary (the variable-z effect). In the bottom-left panel of Figure 1, the unconditional intensity of adjustment is upward sloping in the uncertainty indicator σ. In the top-left panel, the inaction band widens as a function of σ, so that more uncertainty makes it *less* likely that a *given* durables/nondurable ratio (a point on the vertical axis) triggers adjustment.

the beginning-of-period value of z.

3.3 Drift

As remarked in Section 2, the parameter ϑ (the drift for the X/C ratio) is larger in absolute value for goods with faster depreciation and faster decline in prices, and for consumers with a steeper nondurable consumption profile. Since information on the depreciation rate of durable goods is unavailable, we proxy the drift ϑ with predictable household nondurable consumption growth. The theory outlined in Section 2 predicts that the latter should be proportional to ϑ if preferences are homothetic. Obviously this proportionality assumption is only an approximation, as the frictionless level need not coincide with the level that would be chosen in the absence of transaction costs, and homotheticity may be violated. Moreover, the drift in durable consumption depends on unobservables (i.e., depreciation and relative price inflation rates), and these may vary in unpredictable ways with the economic environment and with household characteristics. As illustrated in equation (5), predictable consumption growth is affected both by preference parameters and by the extent to which borrowing and other forms of financial market access may allow households to smooth their consumption patterns over time.

To obtain an observable counterpart to consumption drift, we can exploit the rotating panel component of the SHIW survey. There are 851 households in our data set interviewed in both 1993 and 1995 (little less than half of our selected sample). For these households, consumption growth in the 1993-1995 period is directly observable. For the remaining households, however, such information is missing. Our strategy is, first, to estimate a consumption growth-Euler equation for the 1991-93 panel households. We assume that consumption growth is a function of *changes* in demographics (changes in the number of children, number of earners, family size, and homeownership status), and other characteristics (education, a quadratic in age, gender).[9] The predicted value of this regression is an estimate of expected nondurable consumption for the panel households. We impute the missing value of desired nondurable consumption growth to the remaining households in our sample using a best-subset regression strategy.[10] The average drift in consumption

[9]Importantly, the drift in nondurable consumption growth is estimated on the basis of *changes* in demographics, while regressions for the size of adjustment below exploit cross-sectional variation in the *levels* of demographics.

[10]The variables that are used to impute missing values include dummies for employed, self employed, region of residence, marital status, gender, education, a quadratic in age, number of children in three age bands, family size, and the number of earners.

growth is 0.56 percent in the subset of households with non-missing data, and 0.71 percent in the subset of households with missing data. The drift in nondurable consumption growth (with the opposite sign) is then used to proxy for the unobservable ϑ. As a further proxy for the drift in vehicles (but not for furniture or jewelry), we may use province-level measures of accident frequency (in Table 1, we see that about 1 percent of cars are involved in an accident during the year) and of dissatisfaction with the extent of traffic congestion,[11] as indicators of car depreciation. Comparing Table 1 and Table 2, it is interesting to note that the individuals adjusting the stock of vehicle appear to be living in provinces with less frequent car accidents and lower dissatisfaction with quality of public transports and traffic congestion.

Theoretical implications of a larger drift are illustrated in Figure 5. Quite intuitively, stronger drifts (i.e. more negative values of ϑ) enlarge the inaction range in the top-left panel, and in the top-right panel yield a steady-state distribution with wider support, which however becomes very strongly skewed and has increasingly negligible probability in the lower portion of the inaction range. Furthermore, in the bottom-left panel we see that as the drift becomes stronger the frequency of upward adjustment increases, while the size of adjustment decreases (this effect is rather mild for the parameters used in plotting the figure); similar effects of opposite sign apply to downward adjustments.

As shown in Figure 5 the main implication of a higher drift for the cross sectional distribution of X/C is that this becomes even more whale-shaped as the drift becomes larger in absolute value. One way to check this implication could be to exploit differences in drift across types of durable goods, classifying them according to their depreciation rates. However, the depreciation rate affects the frictionless optimum durables/nondurables basket through the user cost, as in equation (2), as well as the speed at which the actual basket traverses the consumer's state space. Furthermore, it is not obvious that focusing on depreciation as a source of drift heterogeneity would lead to the emergence of a clear ordering. In fact, our durable goods categories include different items with probably markedly different depreciation rates. For instance, what we call "furniture" also includes such high-depreciation items as household appliances and sundry articles as well as furniture proper, which depreciates slowly. Thus, it is not obvious whether these goods overall depreciate

[11] The 1993 SHIW asked each head to report (on a 0-10 scale) their satisfaction with the quality of various public and private services in the province or neighborhood of residence (there are 105 provinces in Italy, usually centered adound medium to large-sized cities). We construct the index reported in the text by taking a provincial average of households' responses. Standard errors of our estimates below are corrected for provincial clustering.

faster or slower than vehicles.

Accordingly, we explore the implications of different drifts using predictable consumption growth, as described above. Recall that faster expected consumption growth implies more negative drift in the actual state variable. Theory predicts that a more negative drift should make the density of X/C more spread-out and even more skewed to the left (more "whale-shaped"; see Figure 5, top right panel). To check this, we plot the empirical density of the X/C ratio (as deviation from the effect of observable household characteristics) for households with different predictable consumption growth rates. In Figure 6, a household is allocated to the "high drift" group regime if the drift of the durable/nondurable basket is *ceteris paribus* more negative, i.e. if consumption growth is above the 75^{th} percentile of the cross-sectional distribution. For vehicles, the empirical density resembles the theoretical one, and gives some confidence for use of the continuous measure in the regression analysis below. However, evidence for the other two types of durables is hard to detect.[12]

3.4 Adjustment costs

Different durable goods are likely to entail adjustment costs of widely different amounts and nature. In particular, cars and vehicles are likely to feature large adjustment costs because of the legal and administrative costs that the buyer incurs upon purchase. Adjusting the stock of such durable goods by selling the current car and buying a new one also implies large costs because *lemon* problems plague the second-hand market for vehicles while administrative costs are incurred upon purchase of the new car. In Table 2, the large size of vehicle stock adjustments suggests that the width of the inaction band is larger for cars than for the other durable goods categories, which is consistent with the plausible idea that buying a new car (or vehicle) entails various transaction costs which makes adjustment more costly for cars than for furniture and jewelry, which should thus have lower adjustment costs than cars. Furniture is likely to lie in between: like cars, household appliances face a *lemon* problem, but differently from cars and similarly to jewelry they carry little or no administrative costs in the primary market. These considerations lead us to conclude that an empirical measure of adjustment costs that is common to the three goods we study is hard to come by.

In the absence of a genuine and comprehensive measure of adjustment costs, we rely on a proxy which applies mainly to cars. Consider the process of buying a new car in

[12]One should notice, however, that the predictions of the model concerning the effect of the drift are complex. A larger drift widens the inaction band but also produces a more concentrated distribution within the inaction band.

Italy. In order to get a plate and a temporary registration, the buyer must provide the dealer with a certificate of residence and his/her social security number. Obtaining the former typically requires a visit to the local council office and a nominal fee. Within two months from the purchase, the car must be registered with PRA, a public registry for private vehicles. The buyer is required to fill in an application form, pay a fee, and have his/her and the dealer's signature authenticated by a public notary (which involves an additional trip and an additional, usually large, fee). While the monetary costs of this process are unlikely to vary across provinces, search, time and psychic costs of dealing with the local bureaucracy probably do vary across them (which may explain the proliferation of *We'll-do-it-for-you* agencies). To proxy for all these costs, we use an index of dissatisfaction with the efficiency of public council offices at provincial level. Comparing Tables 1 and 2, we see that dissatisfaction with the efficiency of the public administration is slightly lower among those who adjust the stock of cars. Transaction costs faced in the adjustment of the stock of furniture or jewelry are unlikely to be related to the efficiency of the local public administration; with this caveat in mind, the inclusion of this variable in the regressions for durables other than cars may capture province-specific effects.

Theoretical implications of larger adjustment costs (relative to the slope of marginal flow utility losses) are illustrated in Figure 7. All adjustment cost parameters are multiplied by the same constant (measured along the horizontal axis of the left-hand side panels), and range from very low levels to the configuration considered in the other figures. Increasingly large adjustment costs imply a wider inaction range, lower intensity of adjustment and (since the lump-sum cost is increasing along with costs that are proportional to the size of adjustment) larger adjustment steps.

All these effects are quite intuitive, and deserve to be confronted with the admittedly limited information available to us on variation in adjustment costs across the data set. In Figure 8, we plot the density of X/C (in deviations from the effect of observable characteristics), grouping households according to the level of adjustment costs that they face; for this we use our index of dissatisfaction with the efficiency of public offices at provincial level and allocate to the high-adjustment cost group all households living in a province with an index of dissatisfaction in excess of the 75^{th} percentile of the distribution. Since this index is likely to be a reasonable proxy of adjustment costs just for cars, we report the density function only for this type of durable. Figure 8 does not conform to the theoretical predictions illustrated in Figure 7. In particular, the distribution of the ratio of car-stock to nondurable consumption is more spread-out for households bearing smaller adjustment costs. There may be two reasons for this. First, our measure is a very crude indicator of transaction costs. Second, the same indicator may be correlated with characteristics that impart an opposite (and counteracting) effect on the shape of the cross-sectional density. We will account for such violations of the *ceteris paribus* requirement in the more formal regression analysis of the next section.

4 Regression analysis

The descriptive evidence and theoretical predictions reviewed above are in broad agreement with each other, but the evidence is not surprisingly mixed. In fact, each of the implications considered above holds on a *ceteris paribus* basis, i.e., all other characteristics that may affect the moments of the distribution of X/C should be held constant. This is obviously not the case in the comparisons that we have made in the previous section. For instance, when comparing the implications of uncertainty for the shape of the cross sectional distribution of X/C, we have not controlled for other differences between the high- and low-uncertainty subsamples that may be relevant for adjustment and that may even be correlated with uncertainty. In this section we examine the validity of the theoretical implications of the model outlined above. We focus on regressions for the probability of upgrading the existing stock of durable goods and the net size of adjustment.[13] Separate specifications are run for vehicles, furniture, and jewelry, allowing for possible interactions among the three types of durable-adjustment decisions.

In the class of infrequent adjustment models we consider, observing a larger value of X/C prior to adjustment makes subsequent upward adjustment less likely, and subsequent downward adjustment more likely. The model, in fact, predicts adjustment with certainty as soon as the deviation of X/C from the frictionless optimum exceeds the optimal inaction band. In our empirical strategy, we treat both the frictionless optimum and the inaction band as systematically different across individuals and unobservable as such. We focus on how observable variables bear on these when interpreting the results of regressions conditional on the pre-adjustment X/C : this is essentially the variable displayed by the smoothed distribution functions of the empirical illustrations above, which—after controlling for observable characteristics—we want to interpret as the history-dependent component of the determinants of optimal adjustment decisions.

[13] Very few households downgrade the stock of vehicles (a little more than 1 percent of the whole sample) or that of jewelry (0.2 percent). Accordingly, we thus focus on upward adjustment only.

For example, and most crucially, consider the implications of higher uncertainty. The model predicts that adjustment is less likely to be observed, for a given X/C, when a more uncertain outlook implies a wider (unobservable) band. The implications of the current X/C depend on the position of the frictionless optimum as well as on the width of the band, but the former (controlled by, e.g., demographics) is not affected by variance in the linear-quadratic approximation. Unconditionally, however, a higher expected variability should increase the frequency of adjustments. With appropriate measures of drift and adjustment costs, one can test the additional implications that a higher drift (a higher depreciation rate) increases the probability of adjustment, while higher adjustment costs make adjustment less frequent.

We assume that the decision of adjusting the current stock of durables occurs when a latent variable D_i^* (which, from our theory section above, may be interpreted as the distance between the action point and the durable-non durable ratio) crosses a threshold, $D_i^* > 0$, and model this variable as:

$$D_i^* = W_i' \gamma + u_i$$

where W_i is a vector of explanatory variables. The assumption that $u_i \sim N(0,1)$ yields the probit model:

$$\Pr(D_i^* > 0) = \Phi\left(W_i' \gamma\right) \tag{9}$$

where $\Phi(W_i' \gamma)$ is the standard normal cumulative density function evaluated at $W_i' \gamma$. The theory also delivers predictions in terms of the size of the adjustment conditional on adjusting (or equivalently band width), which we represent as:

$$E(A_i | Z_i, D_i^* > 0) = Z_i' \beta + E(\varepsilon_i | Z_i, D_i^* > 0) \tag{10}$$

where A_i is the size of adjustment for individual i, Z_i a vector of explanatory variables and ε_i a Gaussian error term ($\varepsilon_i \sim N\left(0, \sigma_\varepsilon^2\right)$) that captures measurement error and unobserved heterogeneity. The theory predicts that higher uncertainty and adjustment costs both *increase* the size of adjustment, while a higher drift reduces it. If $E(\varepsilon_i | Z_i, D_i^* > 0)$ is not zero, simple OLS regressions conducted on the sample of those who adjust will provide inconsistent estimates of the parameters of interest.

This is a well known problem in microeconometrics. In what follows, we present the results of Heckman-selectivity adjusted regressions for the size of adjustment. The use of a Heckman-selectivity model is called for by two orders of considerations. First, unobserved heterogeneity in the two margins is likely to exist along a variety of dimensions (tastes for durables, transaction costs, etc.) that obviously affect both the likelihood of observation in the neighborhood of trigger points and the width of the adjustment bands. Hence, the relevant self-selection mechanism implies that unobservable heterogeneity in the extensive margin is correlated with unobserved heterogeneity in the intensive margin (i.e., $\operatorname{cov}(u_i, \varepsilon_i) \neq 0$), thus precluding use of a Cragg (1971) model (see Lee and Maddala, 1985).

Second, the decision to adjust does not depend on the same variables affecting the decision about how much to adjust (if this is done at all). This precludes use of a simple Tobit model, and provides exclusion restrictions, i.e. it identifies variables that affect the likelihood of adjustment but not the size of adjustment: the value of the X/C state variable *prior to* adjustment affects the decision as to whether to adjust but, conditional on adjustment, the size of the adjustment does not depend on that variable, and is only affected by the other relevant parameters, namely the level of uncertainty, drift, adjustment costs, and heterogeneous taste for durables.

4.1 The probability of adjustment

Table 3 reports the results of probit regressions for the upgrading of the stock of vehicles, jewelry, and furniture, respectively. The table reports marginal effects throughout. This model estimates the probability that a household adjusts the stock of the specific durable conditioning on the characteristics that we measure in our data set. Since very few households downgrade their stock of the three durable goods, we focus on upgrades of the stock (a positive net purchase of the good). For example, in the case of vehicles adjustment is observed if the household has made a purchase in the year and its value exceeds that of a possible sale. A similar definition holds for jewelry and furniture, taking into account that for the latter no information on the revenue from sales in the secondary market is available.

Our specification includes the value of the durables stock-nondurable consumption ratio measured at the time of adjustment separately for vehicles, furniture and jewelry.[14] This accounts for non-separability between vehicles and other durable goods in a conditional

[14] This is true only for those who adjust. For those who do not adjust the X/C ratio is that at the end of the period. *A priori* it is not obvious which measure is better for this group, whether the end-of-period or the beginning-of-period stock. The two will differ because of depreciation.

demand framework. Our measure for uncertainty is the standard deviation of the individual distribution of future earnings. Since the standard deviation of future earnings is not dimension free, we include as an additional regressor the value of expected earnings; an alternative would be to use the coefficient of variation as an indicator of uncertainty (an alternative we have implemented finding qualitatively similar results, available on request), but this would perhaps be unduly restrictive, as it implicitly restricts expected future income to have a negative impact on the probability of adjustment.[15] Our main measure for the drift is expected household nondurable consumption growth.

As pointed out in Section 2 individuals with a particular taste for durable goods will, *ceteris paribus*, tend to depart less from the optimal, frictionless X/C ratio. As a consequence they will adjust more frequently and by smaller amounts. Since preference parameters are not observed, we proxy for taste heterogeneity inserting in all our regressions a vector of demographic variables (education, age, family size, the number of children in three age bands, the number of earners, and dummy indicators for region of residence and city size). Finally, we insert a set of variables intended to capture depreciation (the frequency of car accidents and the dissatisfaction with the extent of traffic congestion in the province where the household lives), adjustment costs (the dissatisfaction with the efficiency of the local public administration), and the opportunities to obtain an equivalent flow of services if no adjustment is undertaken (an index of dissatisfaction with the quality of public transports in the household province). The quality of public transports may have an impact on the frequency of adjustment since for individuals living in provinces with highly inefficient public transports the benefits from more frequent adjustment are larger than for households that can rely, if no adjustment is undertaken, on high quality public transport. Since for these variables we exploit province-level variability, we correct standard errors for provincial clustering.

Consider first the probability of upgrading cars, reported in the first column of Table 3. Consistent with the predictions of the theoretical model, the probability of upgrading decreases with the initial value of the ratio of the stock of vehicles to nondurable consumption, and the coefficient is highly statistically significant. It is also economically important: a 10 percent decline in the durable-nondurable ratio would increase the probability of adjusting by about 7 percent on average. On the other hand, the stocks of the other two durable

[15] Of course, our approximate model and data provide less than fully structural support for this empirical prediction. Qualitatively, however, the extent of the relevant uncertainty (in the proportional growth rate of nondurable consumption growth) should be monotonically related to the income-uncertainty variable after controlling for income levels. We have also experimented with normalized uncertainty measures.

goods (scaled by nondurable consumption) are statistically insignificant. One interpretation of this result is that it provides evidence in favor of separability in preferences and/or adjustment costs of vehicle purchase from either furniture or jewelry purchase.

As predicted by the theory —and controlling for expected earnings (whose coefficient is highly statistically significant but economically quite small)— a higher level of uncertainty reduces the probability of adjusting (a p-value of 1.7 percent). Removing uncertainty will increase the probability of adjustment by 1.4 percentage points, or about 10 percent of the unconditional probability of adjustment. A 10 percent increase in the standard deviation of future earnings decreases the probability of adjustment by about 1 percent on average.

Our main control for the drift (expected household nondurable consumption growth) is statistically significant and displays a sign that is in agreement with the model's prediction: households with faster growing nondurable consumption adjust more frequently than those with a flatter profile; if the drift in consumption is increased by 10 percent, the probability of adjusting the stock of vehicles rises by little less than 0.4 percent. The alternative control for the drift (traffic congestion) is instead poorly measured. Despite the admittedly less than ideal character of these controls, this evidence suggests that expected changes have a strong effect on the probability of adjustment.[16]

The index of inefficiency of the public administration, which we view as a proxy for adjustment costs for vehicles, has a negative and statistically significant effect on the probability of adjustment, as predicted by optimal inaction models. Recall that in the descriptive evidence above we found that the relationship between durable adjustment and our measure of adjustment costs was not consistent with the predictions of simple infrequent-adjustment models. Conditioning on other relevant variables makes it possible to highlight the effect of interest, which is rather precisely estimated and appears to be quite large. A 10 percent increase in the efficiency of the local public administration increases the probability of adjustment by about 16 percent. Bringing all provinces to the level of efficiency of the most efficient province (a very extreme experiment) would increase the probability of adjustment by almost 10 percentage points.

As we argued, the frequency and size of adjustments is also affected by the utility cost that a household incurs for not adjusting the stock of the durable continuously in order to

[16]If our imputed drift measure is removed, the results are very similar: the point estimate of σ^e is -0.0210 with a standard error of 0.0092. If we focus only on the original panel households, again we obtain similar results (although less precisely measured due to a reduced sample size). In particular, the point estimate of σ^e is -0.0187 with a standard error of 0.0155. The point estimate of the drift is 0.5581 with a standard error of 0.1576. Thus it appears that our imputation procedure for the drift is not driving the results.

keep it as close as possible to the no-adjustment-cost level. This utility cost is increasing in the importance that the durable has in consumer preferences, here proxied by a set of demographic variables. Among the demographics, the likelihood of adjusting declines with age, and is lower for residents in the South. Adjustment is relatively more likely for individuals with high education, for multiple-earners, and for large family units. Although it is hard to interpret the effects of these demographic variables in terms of taste parameters, it is worth noticing that most of the effects are at least consistent with this interpretation: for instance it is conceivable that cars are at the margin more important for large households and less so for the elderly, leading the first to adjust (relatively) more and the second less frequently.

In addition, those living in provinces with bad public transport are more likely to upgrade and those living in provinces with a high frequency of car accidents less likely (although the estimate has a somewhat large standard error). The first effect is consistent with the idea that if a good substitute for private transport is available, the pressure to adjust when the stock of vehicles depletes is lessened leading to less frequent upgrading. One explanation for the second effect is that a high probability of car accident not only raises the drift but also increases uncertainty about the X/C ratio and the two have opposite effects on the probability of adjustment.

Thus, taken together, the estimated effects on the probability of adjustment of the main variables that theory predicts should affect the adjustment decision, i.e. the initial stock, the value of uncertainty, drift, and adjustment cost, lend considerable support to the model, at least as far as vehicles are concerned. The effects of drift is much less apparent in our regressions, or perhaps is absorbed within our demographic controls and unobserved heterogeneity.

In the second and third columns we report the results of probit regressions for the decision of upgrading the stock of furniture and that of jewelry, respectively. We leave the measure of inefficiency of the public administration as a proxy for adjustment costs, but note that this is likely to be appropriate for cars but not for furniture or jewelry.

Controlling for expected earnings (which, as in the case of vehicles, is statistically significant for both furniture and jewelry), uncertainty reduces the likelihood of upgrading, though standard errors are higher than in the vehicles case (p-values of 15 percent and 12 percent, respectively for furniture and jewelry). The drift term has no significant impact on the probability of adjustment. The index of inefficiency of the public administration—our measure of adjustment costs for cars—has no statistically significant effects for the probability of adjusting jewelry or furniture. This is consistent with the idea that adjustment costs are specific to the type of durable good under consideration. The own durables stock-nondurable consumption ratio is statistically significant in both equations, but while it is plausibly negative for furniture, it carries a positive sign for jewelry. This suggests that the theoretical framework we consider in this paper is not fully appropriate to characterizing the adjustment of precious objects. Another interpretation is that our controls for taste heterogeneity are unable to reflect idiosyncratic tastes for durables which are positively correlated with the initial stock biasing upward the coefficient of the latter. There is some evidence that preferences for furniture are non-separable with respect to jewelry; in particular, there is some weak evidence for complementarity. As for the demographics, education increases and age decreases the probability of upgrading. Family size does not affect the probability of buying furniture, but decreases the probability of buying jewelry. The number of earners increases the probability of both. The probability of replacing the stock of furniture is lower in the South.

4.2 The size of adjustment

The equation for the size of the adjustment contains an error term that captures measurement error as well as unobserved heterogeneity. The mean of this error term, conditional on adjusting the stock of the durable, is not necessarily zero. If this is the case, one is confronted with a standard selectivity problem. To account for this we estimate Heckman-selectivity regressions allowing for general correlation between unobserved heterogeneity in the extensive margin and unobserved heterogeneity in the intensive margin. As mentioned, theory predicts that the value of X/C at the beginning of the period (i.e., prior to adjustment) affects the likelihood of adjusting but not the size of the adjustment if it occurs. This is an important exclusion restriction, and the evidence presented below suggests that it is powerful enough for identification purposes.

The results of the second stage of the Heckman selectivity regressions are reported in Table 4. For furniture we lack the value of sales, so the variable we use is just the value of purchases. For cars, we find that income uncertainty increases the size of adjustment, as predicted by the theoretical model. The effect is strongly significant (a p-value of 1.1 percent) and economically sizable (see below).

The effect of expected earnings is positive but statistically poorly measured. Similar evidence holds for the drift, with the opposite sign. Of the demographic variables, we find that the only ones that are statistically significant are: city size dummies (with people living in a metropolitan area buying smaller cars), and the number of adolescents, which reduces the size of adjustment. The availability of efficient public transports increases the size of adjustment. There is evidence for self-selection (unobserved tastes for durables) which justifies our specification. The Wald test for independent equations has a p-value below 1 percent. The estimate of ρ (which measures the correlation between the error term of the adjustment size equation, ε_i, and the error term of the selection equation, u_i) is negative and statistically significant, and has an interesting economic interpretation: the unobserved heterogeneity that leads a household to adjust the stock of durables is negatively correlated with the unobserved heterogeneity that affects the size of adjustment. Our index for the inefficiency of the public administration does not have statistically significant effects on the size of automobile purchase. This can be rationalized considering that the adjustment costs captured by that measure are indeed lump-sum in nature, and have to be borne for each transaction regardless of its size.[17]

For jewelry, results are less in agreement with the theory, or poorly measured. For instance, our measure of income uncertainty has the wrong sign and displays a large standard error. Expected income increases the size of the adjustment. The drift has the predicted sign and is statistically well measured. Measures of adjustment costs (inefficiency of the public administration) are negative and precisely measured. Among the demographics, age and residence in the South increase the size of adjustment. Residence in small town decreases it. Finally, we find again that the estimate of ρ is negative and statistically significantly different from zero, consistent with a feedback story due to selection on the unobservables.

For furniture, results are more in line with theoretical predictions. More uncertainty increases the size of the adjustment (but the p-value is high, 13 percent). The drift has the right sign but it is insignificant, while expected income increases the value of the purchase. The elderly buy smaller items, and so do households with young children. Households living in the South adjust to a greater extent, and so do those who live in provinces with good public transports, probably a reflection of the better living standards enjoyed by these households. Once more, there is evidence for selection and negative correlation between unobservables in the adjustment equation and the selection equation. In summary, our regression evidence shows that the model seems to be appropriate for modelling the behavior

[17] If we focus only on the original panel households, we obtain similar results. In particular, the point estimate of σ^e is 1.57 and it is marginally significant (a standard error of 0.97). The point estimate of the drift is 8.72 with a large standard error of 12. The estimate of ρ is -0.58 (0.09). Also in this case, as in the probit case, point estimates of the main effects in the sub-sample of panel households are comparable to those obtained in the whole sample, just less precisely measured due to a smaller sample size. Removing the imputed drift measure from the regression has virtually no effect.

of car adjustment and, to some extent, furniture, but evidence for the jewelry is more mixed.

Interestingly, variables that have a positive effect on the probability of adjustment tend to have a negative effect on the size of the adjustment, and vice versa. This is true also for the unobservables, regardless of the durable good considered. This negative correlation is quite consistent with theoretical predictions of infrequent-adjustment models. The size of the adjustment depends on the location of the triggers and of the return point, and parameters that increase the width of the inaction band also tend to increase the size of adjustment. This is particularly clear in the case of lump-sum adjustment costs only, where all adjustments bring the consumer back to the same return point. When adjustment costs are proportional to transaction size, of course, the width of the inaction band and the size of transactions can to some extent vary independently from each other. Our findings of negative correlation between the determinants of the likelihood and size of adjustment suggest that lump-sum adjustment costs are the predominant source of "optimal inaction" in our data.

Evidence of negative correlation between the unobservables in the selection equation and the unobservables in the outcome equation is particularly intriguing, and consistent with the predictions of the approximate model outlined in Section 2. Consider individuals with similar target levels but different unobserved tastes for vehicles. Individuals with stronger tastes for car are likely to face a lower depreciation of their durable (e.g., because of better maintenance). This means that conditioning on a similar durable-nondurable ratio, individuals with stronger tastes for vehicles have narrower inaction bands (see Figure 1), which triggers adjustment more frequently. But this also implies that these individuals will adjust by a smaller amount vis-à-vis individuals with weaker tastes for vehicles (if return points are constant). Thus, unobserved heterogeneity in the intensive margin equation should indeed be negatively correlated with unobserved heterogeneity in the extensive margin equation.

4.3 Unconditional effects

In the introduction we mentioned that our approach may have important implications for understanding the role of uncertainty (and of other variables) on the dynamics of aggregate *per capita* durable purchases. This issue can be evaluated considering the expression for average household durable expenditure in the whole sample:

$$\begin{aligned} E(A_i) &= E(A_i | D_i^* > 0) \Pr(D_i^* > 0) + E(A_i | D_i^* \leq 0) \Pr(D_i^* \leq 0) \\ &= E(A_i | D_i^* > 0) \Pr(D_i^* > 0), \end{aligned}$$

where $E(A_i | D_i^* > 0)$ is average expenditure in the subsample of those who adjust, $\Pr(D_i^* > 0)$ the probability of adjusting, and the second equality follows from the fact that $A_i | D_i^* \leq 0 \equiv 0$ (individuals who do not adjust make no purchase). The marginal effect of a given variable x_{ik} (affecting both the size of adjustment and the probability of adjustment) on average household expenditure is then:

$$\frac{\partial E(A_i)}{\partial x_{ik}} = \frac{\partial E(A_i | D_i^* > 0)}{\partial x_{ik}} \Pr(D_i^* > 0) + \frac{\partial \Pr(D_i^* > 0)}{\partial x_{ik}} E(A_i | D_i^* > 0).$$

Intuitively, the marginal effect is given by the sum of two components. The first is the effect of x_{ik} on the size of the adjustment, weighted by the probability of adjusting. The second is the effect of x_{ik} on the probability of adjusting, weighted by the average amount purchased. Using (9) and (10), it follows that the marginal effect of x_{ik} on average expenditure (for those who adjust and for those who don't) is given by the expression:

$$\frac{\partial E(A_i)}{\partial x_{ik}} = \beta_k \Phi(W_i' \gamma) + \gamma_k \phi(W_i' \gamma) \left[Z_i' \beta - \rho \sigma_\varepsilon W_i' \gamma\right]$$

We calculate this marginal effect (and the breakdown in intensive and extensive margin effects) and the corresponding elasticity $\frac{\partial E(A_i)}{\partial x_{ik}} \frac{x_{ik}}{A_i}$ (evaluated at the mean of the variables),[18] and report the results of this exercise in Table 5. For simplicity, we limit our analysis to the case of vehicles.[19]

The overall marginal effect of uncertainty is small, due to intensive margin effects being almost exactly matched by extensive margin effects of opposite sign. Note also that the overall effect is positive. A 10 percent increase in the standard deviation of future earnings increases household car expenditure by 0.2 percent on average. The effect of a larger drift is

[18] The elasticity is estimated at the mean values using:

$$\left\{\widehat{\beta}_k \overline{\Phi} + \widehat{\gamma}_k \overline{\phi} \left[\overline{Z}' \widehat{\beta} - \widehat{\rho} \widehat{\sigma}_\varepsilon \overline{W}' \widehat{\gamma}\right]\right\} \frac{\overline{x}_k}{\overline{A}}$$

where $\overline{\Phi}$ is the average of the individual $\Phi(W_i' \widehat{\gamma})$ and $\overline{\phi}$ is the average of the individual $\phi(W_i' \widehat{\gamma})$.

[19] Results for the other two types of durable are available on request. Note that the calculated marginal effects and elasticities lack a rigorous statistical interpretation because of the absence of standard errors. However, they remain informative about the aggregate effects of interest.

similarly quite small, and has an overall negative impact. A 10 percent increase in the drift decreases average household expenditure by 0.3 percent. The inefficiency of the local public administration (which we interpret as a measure of the extent of adjustment costs faced in the adjustment process) has instead a quite sizable negative effect. A 10 percent decline in inefficiency (or a 10 percent increase in efficiency, corresponding to lower adjustment costs) brings about an increase in average household expenditure of approximately the same magnitude. More inefficient alternative public transports increase expenditure: a 10 percent decline in inefficiency of public transports decreases expenditure by almost 7 percent.

5 Concluding remarks

This paper has explored a set of novel theoretical predictions concerning the adjustment of durable goods according to optimal inaction and infrequent adjustment rules. First, we have derived the approximated behavior of the durables stock/nondurable flow ratio for consumers who adopt infrequent-adjustment optimal policies in the presence of first-order adjustment costs, and examined how the ergodic distribution of this ratio changes in response to changes in uncertainty, the level of the durable drift, adjustment costs, and tastes for durable relatively to nondurable goods. Second, we have examined the effect of the same variables on the intensity of adjustment (i.e., the probability of making a transaction), and on the size of the adjustment (given that a transaction is observed), two margins of the adjustment policy that depend on structural features of the consumers' problem in conceptually distinct ways.

We have studied empirically the relationship between the relevant indicators in our data and the probability and size of adjustment of three categories of durable goods, focusing in particular on the implications of different degrees of uncertainty across consumers. The availability of information on different types of durable goods allows us to go beyond the study of cars on which the literature has focused thus far, and to exploit differences in depreciation rates and access to second-hand market across goods to test some additional implications of our model.

Our detailed theoretical framework makes it possible to explore relatively subtle features of the data. In particular, we have emphasized a conceptual distinction between stable cross-sectional heterogeneity of the sampled households' dynamic problem, controlled by a battery of demographic and geographic control variables; and history-dependent heterogeneity of their situation at the beginning of the observation period, captured in the data by variation in their durables/nondurable ratio that is orthogonal to those control variables.

The different theoretical implications of these distinct sources of empirical variation support a structural exclusion restriction in selection-controlled regression analysis. The results show that theoretical predictions are broadly supported by the data. In particular, the behavior of households when adjusting the stock of cars conforms closely to the theoretical predictions of our simple framework of analysis, while evidence for other types of durables (especially jewelry) appears to call for further theoretical refinements.

A Solution technique

During periods when inaction is optimal and the $\{z_t\}$ process follows a Brownian motion process with drift ϑ and standard deviation σ, the expected present discounted value $V(z)$ of quadratic flow losses must satisfy the differential equation

$$\frac{1}{2}V''(z)\sigma^2 + V'(z)\vartheta - \rho V(z) - \frac{bz^2}{2} = 0, \tag{11}$$

with solution

$$V(z_t) = -\frac{b}{2}\left(\frac{z_t^2}{\rho} + \frac{\sigma^2 + 2z_t\vartheta}{\rho^2} + \frac{\vartheta^2}{\rho^3}\right) + K_1 e^{\alpha_1 z_t} + K_2 e^{\alpha_2 z_t} \tag{12}$$

where α_1, α_2 are solutions of the characteristic equation $\alpha\vartheta + \frac{1}{2}\alpha^2\sigma^2 - \rho = 0$ and K_1, K_2 are constants of integration.

The optimality conditions require that costs and benefits of any action be equal along the optimal path, and that costs of potential actions be weakly larger than their benefits when the optimizer is inactive, on the other. Formally, $V(.)$ and (L, l, u, U) must be such that

$$\begin{aligned}
V(l) - V(L) &= C_l + c_l(l - L), &(13)\\
V(u) - V(U) &= C_u + c_u|U - u|, &(14)\\
V(x) - V(y) &\leq C_l + c_l(x - y) \quad \forall x > y &(15)\\
V(x) - V(y) &\leq C_u + c_u|y - x| \quad \forall x < y, &(16)
\end{aligned}$$

which also imply the "smooth-pasting" conditions

$$\begin{aligned}
V'(l) &= V'(L) = c_l &(17)\\
V'(u) &= V'(U) = -c_u. &(18)
\end{aligned}$$

Intuitively, optimal action and return points must then be such that $V'(x)$ equals the marginal cost of action whenever action is in fact undertaken ("smooth pasting"), and the value function at the trigger and return points must differ by the total cost of adjusting between the two points ("value matching"). Inserting the functional form (12) in these conditions forms a system of equation to be solved for the constants of integration and the action and return points. The solution is not available in closed form in general (but only when adjustment costs are prohibitive in one direction, to imply that adjustment decisions are irreversible).

A.1 Stable distribution

In our empirical exercise, we will analyze a set of cross-sectional observations, each of which may be interpreted as a draw from a history of infrequent adjustment similar to that characterized above for a single decision maker. In the absence of time-series information on individual behavior, we will find it insightful to interpret the cross-sectional information available in terms of the long-run distribution of the controlled variable, z, within the $[L, U]$ optimal inaction interval.

The long-run distribution of the approximate control problem is readily available (see Bertola and Caballero 1990, and their references, for derivations). The Kolmogorov equation for the steady-state density reads

$$\frac{\sigma^2}{2} f''(z) = \vartheta f'(z), \tag{19}$$

and is solved by a piecewise linear function if $\vartheta = 0$, a piecewise exponential function otherwise:

$$f(z) = \begin{cases} Az + B & \text{if } \vartheta = 0, \\ Ae^{\frac{2\vartheta}{\sigma^2} z} + B & \text{otherwise.} \end{cases} \tag{20}$$

The constants of integration A and B are determined for each of the state-space segments by boundary conditions, which ensure that inflows and outflows of probability mass balance out at the action and return points as well as within the inaction ranges. The stable density must satisfy

$$f'_{(-)}(l) = f'_{(+)}(l) + f'_{(+)}(L) \tag{21}$$

if $L < l$, where $f'_{(+)}(L)$ is the right-hand derivative evaluated at the lower trigger boundary; at the return point, $f'_{(-)}(l) \neq f'_{(+)}(l)$ because probability flows into l from the trigger point as well as from its immediate neighborhood.[20] Similar boundary conditions are satisfied by the stable density at the upper trigger (U) and return (u) points. These boundary conditions and the adding-up constraint

$$\int_L^U f(z)\, dz = 1, \tag{22}$$

[20] If $L = l$ (reflecting barrier), which is the case when control incurs linear adjustment costs with no fixed component, then the boundary condition reads

$$f'_{(+)}(L) = \frac{2\vartheta}{\sigma^2} f(L).$$

form a rank-deficient system of linear equations in (at most six) A and B constants, for which a solution can always be derived.

On the basis of these derivations, it is also straightforward to compute the probability (intensity, per unit time) of adjustment at the lower and upper boundaries of the inaction range. The rate at which adjustment events occur is the same as the rate of probability outflow from the lower trigger point, L, towards the return point l. The same derivations that lead to (21)—outlined in the Appendix of Bertola and Caballero (1990), and discussed more formally in their references—establish that the relevant probability flow is given by[21]

$$\frac{\sigma^2}{2} f'_{(+)}(L) = \frac{\sigma^2}{2} \frac{d}{dz} \left(A_l e^{\frac{2\vartheta}{\sigma^2} z} + B_l \right)_{z=L} = A_l \vartheta e^{\frac{2\vartheta}{\sigma^2} L},$$

where A_l is the constant of integration determined by the stable distribution's boundary conditions. The probability (intensity) of adjustment events at the upper boundary of the inaction range has a similar form, and also corresponds to the product of the infinitesimal likelihood of finding the process in the immediate neighborhood of the trigger point, $f'_{(-)}(U)$, and of the intensity of Brownian movements that may push the process towards that point, $\sigma^2/2$.

B Data: the SHIW

The Bank of Italy Survey of Household Income and Wealth (SHIW) collects detailed data on demographics, households' consumption, income and balance sheet items. The survey was first run in the mid-60s but has been available on tape only since 1984. Over time, it has gone through a number of changes in sample size and design, sampling methodology and questionnaire. However, sampling methodology, sample size and the broad contents of the information collected have been unchanged since 1989. Each wave surveys a representative sample of the Italian resident population and covers about 8,000 households, - although at times specific parts of the questionnaire are asked to only a random sub-sample. Sampling occurs in two stages, first at municipality level and then at household level. Municipalities are divided into 51 strata defined by 17 regions and 3 classes of population size (more than 40,000, 20,000 to 40,000, less than 20,000). Households are randomly selected from registry

[21] If $L = l$ (reflecting barrier), then

$$\frac{\sigma^2}{2} f'_{(+)}(L) = \frac{\sigma^2}{2} \frac{d}{dz} (A_l z + B_l) = \frac{\sigma^2}{2} A_l.$$

office records. They are defined as groups of individuals related by blood, marriage or adoption and sharing the same dwelling. The head of the household is conventionally identified with the husband, if present. If instead the person who would usually be considered the head of the household works abroad or was absent at the time of the interview, the head of the household is taken to be the person responsible for managing the household's resources. The net response rate (ratio of responses to households contacted net of ineligible units) was 57 percent in the 1995 wave. Brandolini and Cannari (1994) present a detailed discussion of sample design, attrition, and other measurement issues and compare the SHIW variables with the corresponding aggregate quantities.

The survey questions that elicit expectations about future income are simple but powerful. The employed and job seekers are asked to report, on a scale from 0 to 100, their chances of keeping their job or finding one in the next twelve months. Each individual assigning a positive probability to being employed is then asked to report the minimum and the maximum income he expects to earn if employed and the probability of earning less than the midpoint of the distribution of future earnings conditional on working. These data can be combined to obtain an estimate of expected earnings and their variance or standard deviation, which we use as a gauge of consumer income risk. For more details on the actual construction of the two first moments of the distribution of future income, distributional assumptions, and descriptive evidence, we refer the interested reader to Guiso, Jappelli, and Pistaferri (2001).

C Definitions of the variables

In the empirical analysis all demographic variables refer to the household head.

Nondurable consumption: nondurable consumption is the sum of the expenditure on food, entertainment, education, clothing, medical expenses, housing repairs and additions, and imputed rents.

Durables flows: includes expenditures and revenues from sales on three categories separately. "Means of transport" (includes cars, motorbikes, caravans, motor boats, boats, bicycles); "Furniture, furnishing, household appliances and sundry articles" (includes furniture, furnishing, carpets, lamps, household appliances, washing machines, dishwashers, TVs, PCs, Hi-Fi, CDs etc.); "precious objects" (including jewelry, old and gold coins, works of arts, antiques and antiques furniture).

Durable stock: Value of end of period stock of durables for each of the three categories.

The beginning of period value is computed subtracting purchases and adding sales to the end of period stock.

Education of the household head : This variable is coded as: no education (0); completed elementary school (5 years); completed junior high school (8 years); completed high school (13 years); completed university (18 years); graduate education (more than 20 years).

Quality of life indicators: All household heads are asked to assign a score between 1 (worst) and 10 (best) to various quality of life indicators, including: the functioning of public transports, health services, kindergartens, primary and secondary schools, Universities, public council offices, the availability of rentals, job opportunities, shopping facilities, leisure and public park facilities, the extent of traffic congestion, air and water quality, crime control and street safety, street cleaniness and noise pollution.

City size: This variable is coded as: 0-20,000 inhabitants (small town); 20,000-40,000 inhabitants (medium town); 40,000-500,000 inhabitants (large town); and more than 500,000 inhabitants (metropolitan area).

Indicators of earnings uncertainty: We use the standard deviation of expected earnings at the individual level. This is computed directly from survey questions asking: (a) the probability of keeping one's job (if employed) or of finding one (if unemployed) in the twelve months following the interview; (b) the minimum and maximum earnings expected conditional on being employed; and (c) the probability that future earnings will be less than the mid-point of the subjective distribution of future earnings. After making some assumptions on the shape of the on-the-job probability distribution of earnings (triangular distribution) and on the value of the unemployment compensation to each individual in the sample, Guiso, Jappelli, and Pistaferri (2001) use this information to recover measures of expected earnings and their dispersion.

References

[1] Adda, Jerôme, and Russell Cooper (2000), "Balladurette and Juppette: A discrete analysis of scrapping subsidies," *Journal of Political Economy* **108**(4), 778-806.

[2] Attanasio, Orazio (2000) "Consumer durables and inertial behaviour: Estimation and aggregation of (S,s) rules for automobile purchases," *Review of Economic Studies* 67, 667-96.

[3] Bar-Ilan, Avner and Alan Blinder (1992), "Consumer durables: Evidence on the optimality of usually doing nothing," *Journal of Money, Credit, and Banking* **24**, 258-72.

[4] Beaulieu, Joseph (1993), "Optimal durable and nondurable consumption with trnsaction costs," Finance and Economics Discussion Papers, Federal Reserve Board.

[5] Bertola, Giuseppe and Ricardo Caballero (1990) "Kinked adjustment costs and aggregate dynamics," *NBER Macroeconomics Annual*, 237-88.

[6] Brandolini, Andrea, and . Cannari (1994)

[7] Caplin, Andrew and John Leahy (1994), "Business as usual, market crashes, and wisdom after the fact," *American Economic Review* 84, 548-65.

[8] Carroll, Christopher D. (1997), "Buffer-stock saving and the lifecycle/permanent income hypothesis," *Quarterly Journal of Economics* 112, 1–55.

[9] Cragg, John (1971), "Some statistical models for limited dependent variables with application to the demand for durable goods," *Econometrica* 39, 829-44.

[10] Eberly, Janice (1994), "Adjustment of consumers' durables stocks: Evidence from automobile purchases," *Journal of Political Economy* 102, 403-36.

[11] Foote, Christopher, Erik Hurst and John V. Leahy (2000), "Testing the (S,s) Model," *American Economic Review Papers and Proceedings* 90, 116-19.

[12] Grossman, Sanford and Guy Laroque (1990), "Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods," *Econometrica* 58, 25-51.

[13] Guiso, Luigi, Tullio Jappelli, and Luigi Pistaferri (2001), "An empirical analysis of earnings and employment risk," CEPR Discussion Paper *2043*.

Table 1
Summary statistics for the whole sample

The table shows summary statistics for the selected sample of households reporting the information on the subjective probability distribution of future earnings. All figures are in euro and weighted by sample weights. An adjustment is defined as an action: upgrade, downgrade or both. For furniture, only upgrade is available. The value of purchase is calculated only for buyers. Standard deviations are reported in parenthesis. The 1995 sample excludes the retired for comparison with the values of our sample. 'Public transports', 'Local council offices', and 'Traffic congestion' are 1-10 indexes of dissatisfaction with the efficiency of public transports, public offices, and the extent of traffic congestion at the province level, respectively (10 corresponds to the highest level of dissatisfaction).

		Our sample	1995 sample
Value of stock	Vehicles	6,324	5,936
		(6,476)	(6,436)
	Furniture	9,967	9,439
		(9,584)	(9,577)
	Jewelry	3,225	3,214
		(7,050)	(10,274)
X/c	Vehicles	0.3572	0.3405
		(0.3887)	(0.3651)
	Furniture	0.5924	0.5766
		(0.5578)	(0.5768)
	Jewelry	0.1798	0.1700
		(0.4438)	(0.4681)
Frequency of adjustment	Vehicles	0.1785	0.1762
		(0.3831)	(0.3810)
	Furniture	0.3033	0.2936
		(0.4598)	(0.4555)
	Jewelry	0.0996	0.0994
		(0.2996)	(0.2992)
Value of purchase	Vehicles	7,536	7,206
		(6,882)	(6,575)
	Furniture	2,168	2,514
		(3,717)	(5,031)
	Jewelry	1,259	1,081
		(2,139)	(1,795)
Family income		24,907	24,125
		(17,425)	(18,098)
Age		42.66	43.94
		(9.10)	(9.79)
Earnings uncertainty		0.0398	–.–
		(0.0425)	
Accidents per 1,000 cars		11.04	11.04
		(2.55)	(2.58)
Years of schooling		9.99	9.44
		(4.09)	(4.20)
Family size		3.46	3.35
		(1.23)	(1.24)
South		0.3472	0.3485
		(0.4762)	(0.4765)
Public transports		5.66	5.31
		(1.00)	(1.05)
Local council offices		5.48	5.51
		(1.00)	(1.00)
Traffic congestion		6.91	4.06
		(0.96)	(0.97)
Number of observations		1,877	4,782

Table 3
Probit for the upgrading of durables stocks

The variable 'Car accidents' is the frequency of car accidents per 1,000 cars at province level. 'Public transports', 'Local council offices', and 'Traffic congestion' are 1-10 indexes of dissatisfaction with the efficiency of public transports, efficiency of local council offices, and the extent of traffic congestion at the province level, respectively (10 corresponds to the highest level of dissatisfaction). Standard errors are adjusted for provincial clustering. Both σ^e and μ^e are expressed in 1,000 euro.

	Vehicles	Furniture	Jewelry
σ^e	−0.0218	−0.0198	−0.0106
	(0.0092)	(0.0136)	(0.0064)
μ^e	0.0039	0.0040	0.0022
	(0.0014)	(0.0017)	(0.0008)
Drift	0.7544	0.0381	0.2726
	(0.1490)	(0.2262)	(0.1537)
Education	0.0041	0.0111	0.0083
	(0.0021)	(0.0034)	(0.0018)
Age	−0.0043	−0.0067	−0.0016
	(0.0011)	(0.0016)	(0.0009)
Family size	0.0193	0.0028	−0.0229
	(0.0110)	(0.0166)	(0.0087)
Kids 0-5	−0.0262	−0.0346	0.0291
	(0.0232)	(0.0318)	(0.0142)
Kids 6-13	−0.0156	−0.0103	0.0217
	(0.0157)	(0.0231)	(0.0133)
Kids 14-17	0.0218	0.0028	−0.0070
	(0.0219)	(0.0327)	(0.0156)
Number of earners	0.0509	0.0326	0.0435
	(0.0139)	(0.0179)	(0.0098)
Small town	0.0290	−0.0478	0.0341
	(0.0296)	(0.0590)	(0.0391)
Medium town	0.0123	−0.0021	0.0260
	(0.0236)	(0.0515)	(0.0351)
Large town	0.0197	−0.0042	0.0353
	(0.0199)	(0.0461)	(0.0317)
Car accidents	−0.0097	−0.0025	−0.0001
	(0.0032)	(0.0075)	(0.0043)
Public transports	0.0456	0.0695	0.0088
	(0.0147)	(0.0318)	(0.0172)
Local council offices	−0.0461	−0.0375	−0.0101
	(0.0197)	(0.0407)	(0.0202)
Traffic congestion	−0.0003	0.0028	−0.0024
	(0.0118)	(0.0212)	(0.0132)
South	−0.0551	−0.1366	−0.0181
	(0.0299)	(0.0536)	(0.0276)
Center	−0.0299	−0.0366	−0.0076
	(0.0497)	(0.0357)	(0.0226)
X/c, vehicles	−0.2994	0.0150	0.0082
	(0.0497)	(0.0340)	(0.0182)
X/c, jewelry	0.0081	0.0705	0.0367
	(0.0206)	(0.0348)	(0.0186)
X/c, furniture	0.0194	−0.0986	−0.0036
	(0.0141)	(0.0310)	(0.0139)

Table 4
Heckman selectivity model

The variable 'Car accidents' is the frequency of car accidents per 1,000 cars at province level. 'Public transports', 'Local council offices', and 'Traffic congestion' are 1-10 indexes of dissatisfaction with the efficiency of public transports, efficiency of local council offices, and the extent of traffic congestion at the province level, respectively (10 corresponds to the highest level of dissatisfaction). Standard errors are adjusted for provincial clustering. The dependent variable, σ^e and μ^e are all expressed in 1,000 euro.

	Vehicles	Furniture	Jewelry
σ^e	1.2898 (0.5067)	0.4496 (0.2972)	−0.0333 (0.0965)
μ^e	0.0614 (0.0618)	0.1152 (0.0379)	0.0426 (0.0087)
Drift	−2.1648 (9.9772)	−2.7510 (2.2195)	−3.8753 (1.3718)
Education	−0.1108 (0.0973)	−0.0935 (0.0561)	−0.0179 (0.0191)
Age	0.0417 (0.0553)	−0.0954 (0.0175)	0.0255 (0.0111)
Family size	−0.1568 (0.4104)	−0.0185 (0.2321)	−0.0584 (0.1251)
Kids 0-5	−0.2652 (0.7492)	−0.8492 (0.3492)	0.3644 (0.2447)
Kids 6-13	0.0577 (0.6676)	−0.6170 (0.2634)	0.0829 (0.1691)
Kids 14-17	−1.9568 (0.8541)	−0.4855 (0.3105)	0.0639 (0.1976)
Number of earners	−0.3702 (0.6833)	−0.1350 (0.2205)	−0.0585 (0.1725)
Small town	2.0041 (1.0256)	0.2759 (0.5360)	−0.7244 (0.4255)
Medium town	2.5132 (1.0941)	0.1920 (0.4965)	−0.3148 (0.5260)
Large town	1.3418 (0.7339)	−0.0435 (0.4525)	−0.3589 (0.3038)
Car accidents	0.3072 (0.2821)	−0.0685 (0.0756)	0.0454 (0.1007)
Public transports	−1.3203 (0.6464)	−0.6832 (0.2257)	0.2888 (0.1601)
Public offices	0.6879 (0.9227)	0.0294 (0.3700)	−0.7094 (0.2477)
Traffic congestion	0.0002 (0.4203)	0.1434 (0.2407)	0.0480 (0.1394)
South	−0.6063 (1.4795)	1.7695 (0.6464)	1.2019 (0.8436)
Center	1.2055 (1.0062)	0.4416 (0.4181)	0.5677 (0.2375)
ρ	−0.5555 (0.0843)	−0.3221 (0.0817)	−0.4321 (0.1681)
Wald-test (p-value)	0.0000	0.0002	0.0252

Table 5
Unconditional marginal effects and elasticities

The variable ‘Car accidents’ is the frequency of car accidents per 1,000 cars at province level. ‘Public transports’, ‘Local council offices’, and ‘Traffic congestion’ are 1-10 indexes of dissatisfaction with the efficiency of public transports, efficiency of local council offices, and the extent of traffic congestion at the province level, respectively (10 corresponds to the highest level of dissatisfaction). Marginal effects ($\frac{\partial E(A_i)}{\partial x_{ik}}$) are in 1,000 euro.

	Marginal Effect	Intensive margin	Extensive margin	Elasticity
σ^e	0.0306	0.2053	−0.1747	0.016
μ^e	0.0410	0.0098	0.0312	0.527
Drift	5.7069	−0.3446	6.0516	−0.034
Education	0.0149	−0.0176	0.0325	0.138
Age	−0.0280	0.0066	−0.0346	−1.019
Family size	0.1299	−0.0250	0.1549	0.417
Kids 0-5	−0.2523	−0.0422	−0.2101	−0.068
Kids 6-13	−0.1162	0.0092	−0.1254	−0.046
Kids 14-17	−0.1368	−0.3115	0.1747	−0.031
Number of earners	0.3495	−0.0589	0.4085	0.583
Small town	0.5434	0.3190	0.2244	0.128
Medium town	0.4967	0.4000	0.0967	0.104
Large town	0.3700	0.2136	0.1564	0.140
Car accidents	−0.0288	0.0489	−0.0777	−0.292
Public transports	0.1557	−0.2102	0.3659	0.804
Local council offices	−0.2602	0.1095	−0.3696	−1.316
Traffic congestion	−0.0020	0.0000	−0.0020	−0.013
South	−0.5545	−0.0965	−0.4580	−0.195
Center	−0.0285	0.1919	−0.2204	−0.005

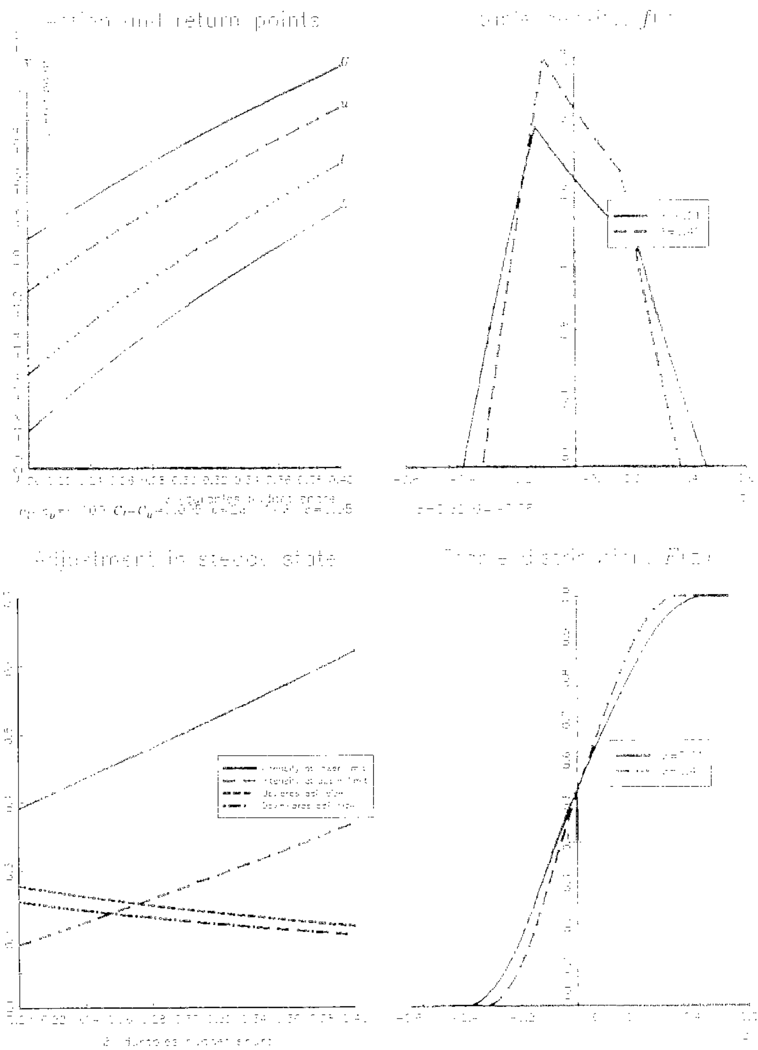

Figure 1: The implications of different durable budget shares for the durable/nondurable ratio trigger and return points and for the long-run stable distribution of its deviations from the static optimum.

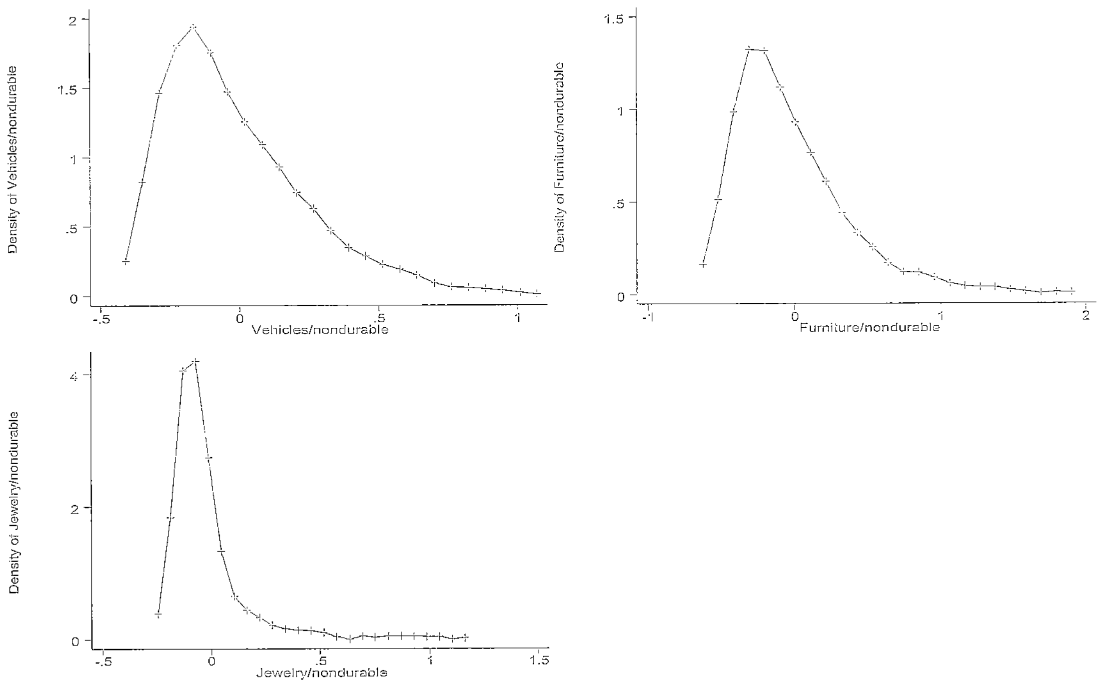

Figure 2: The empirical density of X/C for vehicles, furniture, and jewelry.

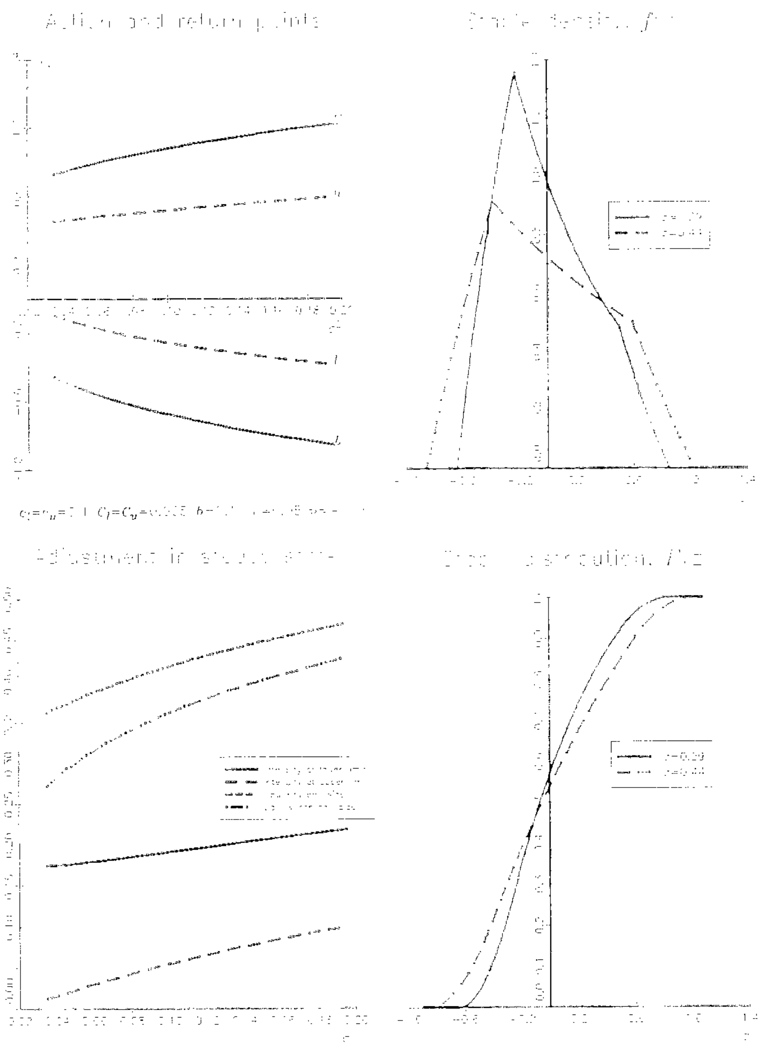

Figure 3: The implications of different degrees of uncertainty for the optimal adjustment policy and stable distribution.

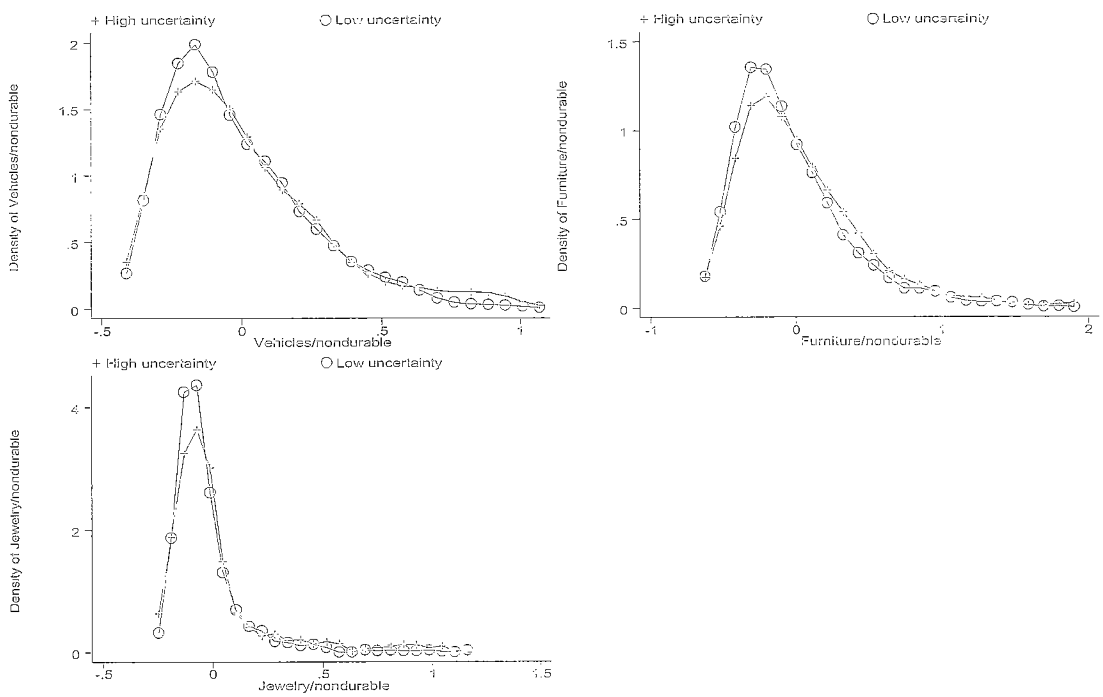

Figure 4: The empirical density of X/C for high- and low-risk households.

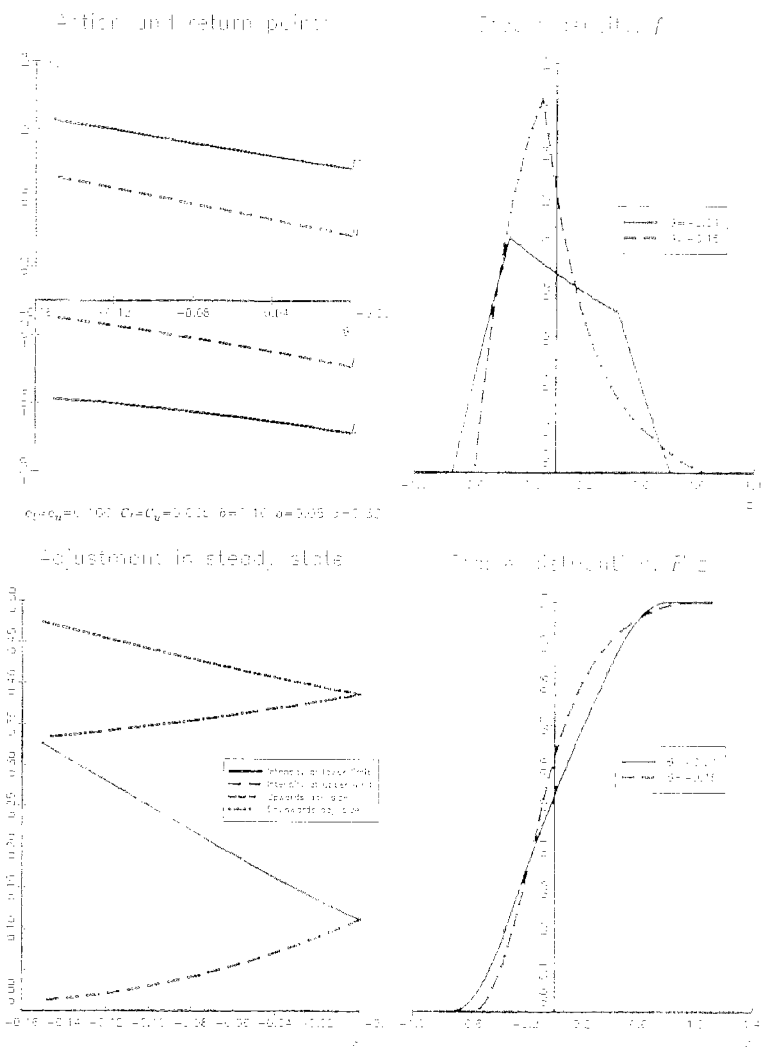

Figure 5: The implications of different drifts for the optimal adjustment policy and stable distribution.

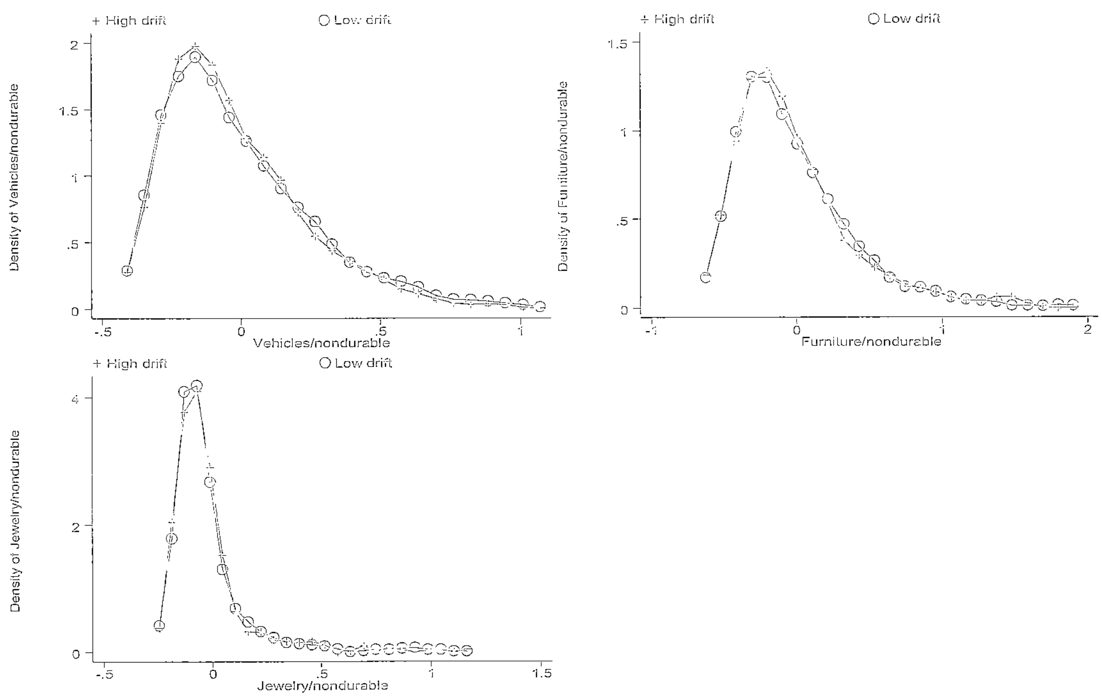

Figure 6: The empirical density of $\frac{X}{c}$ for high- and low-drift households.

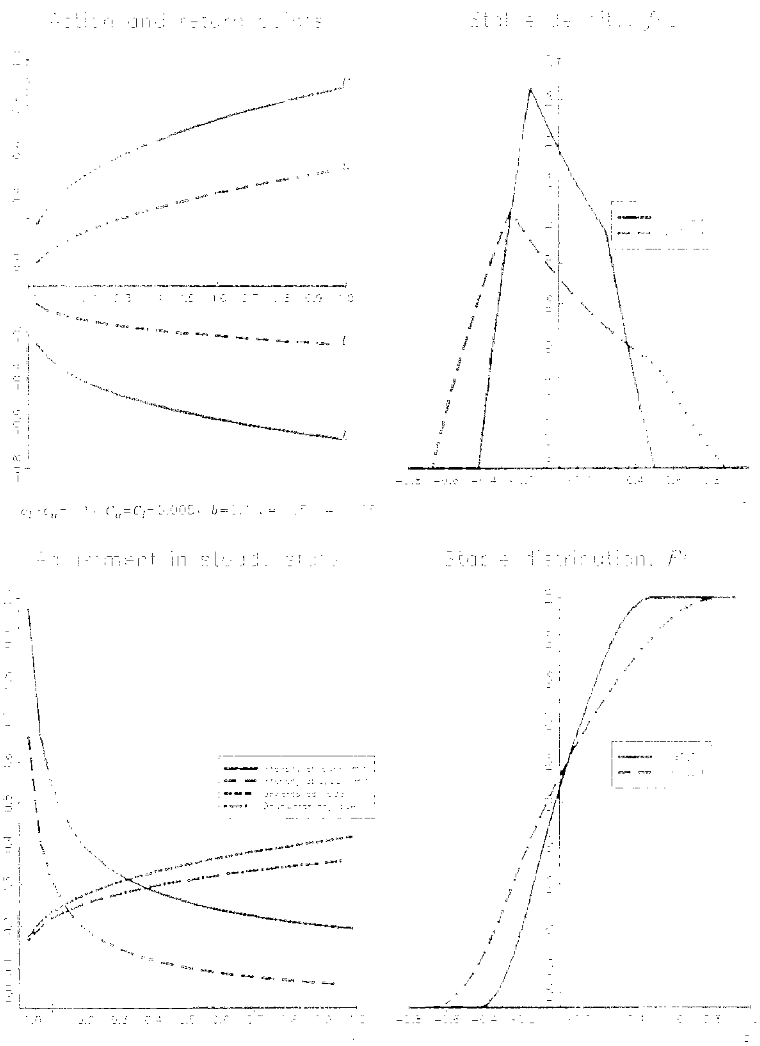

Figure 7: The implications of different adjustment costs for the optimal adjustment policy and stable distribution.

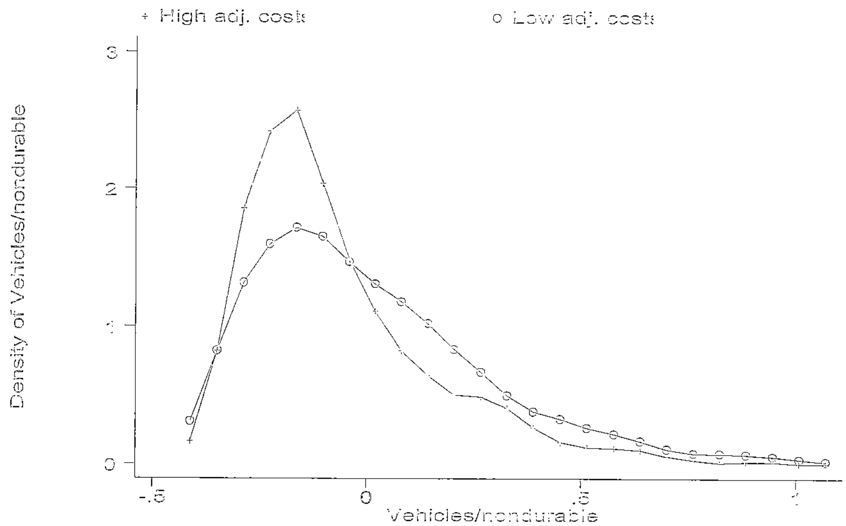

Figure 8: The empirical density of X/C (vehicles) for high- and low-adjustment costs households.