

Uncertainty and Consumer Durables Adjustment

GIUSEPPE BERTOLA
Università di Torino and CEPR

LUIGI GUISO
University of Sassari and CEPR

and

LUIGI PISTAFERRI
Stanford University, Hoover Institution and CEPR

First version received March 2002; final version accepted November 2004 (Eds.)

We characterize infrequent durables stock adjustment by consumers who also derive utility from non-durable consumption flows in the presence of idiosyncratic income uncertainty. The data we analyse include subjective future income uncertainty measures, which we use as instruments in the estimation of relevant parameters of heterogeneous consumers' dynamic adjustment problems. The data feature two conceptually distinct sources of variation: cross-sectional heterogeneity of the sampled households' dynamic problems, and history-dependent heterogeneity in their situation during the observation period. We note that the latter should affect 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 insight.

1. INTRODUCTION

Over the last 10 years a growing literature has focused on models where first-order adjustment costs justify intermittent microeconomic adjustment. Small adjustment costs can imply wide ranges of inaction in a very uncertain environment, since the option of remaining inactive is valuable when deviations from frictionless optimum levels are likely to be erased soon by volatile exogenous developments rather than by costly action. This theoretical perspective is applicable to many adjustment decisions (including labour demand, investment, inventories, and cash balances) as well as to the durable goods consumption patterns on which this paper focuses. We develop relevant theoretical implications and bring them to bear on a data-set featuring extensive information on durable purchases and subjective measures of future income uncertainty.

Our theoretical model makes simple and testable functional form assumptions to obtain a set of characterization results regarding the role of drift, adjustment cost, and taste parameters in determining the frequency and size of durables adjustment, highlighting in particular a conceptually important distinction between conditional (on the current state) and unconditional adjustment behaviour. Higher uncertainty widens the range of inaction, but also increases the unconditional likelihood of wide deviations from the preferred durables stock. Conditioning on the current deviation, a more uncertain future evolution of the problem's forcing variables implies that adjustment is less likely in the immediate future, and larger if it does occur.

The deviations from the statically optimal durable/non-durable consumption bundle that are infrequently corrected by costly adjustment are driven in the model by fluctuations of consumers' intertemporally optimal consumption levels. Our data include subjective indicators of future

income uncertainty, which need not be structurally related to the relevant volatility of total consumption. In theory, prudent consumers facing more volatile income processes should obtain more self-insurance by accumulating larger buffer stocks of wealth. They might also hold state-contingent financial assets and alter their behaviour in ways that would be extremely difficult to model explicitly. A semi-structural approach allows us to sidestep these problems. To assess empirically the relationship between consumption volatility and labour income uncertainty in our sample, we use subjective income uncertainty measures as instruments for consumption volatility in estimation of optimality conditions on the intertemporal non-durable consumption margin. The results conform to theoretical predictions regarding the role of uninsurable income uncertainty in shaping consumption dynamics, and yield theoretically sound measures of heterogeneous consumption uncertainty and drift across our sample households.

We can then assess the effects of uncertainty (and of observed and unobserved taste heterogeneity and other factors) on the frequency and size of adjustment. The data we analyse include relatively detailed information on durable goods stocks and purchases, as well as some information on heterogeneous adjustment costs. Econometric identification of our controlled regression estimates is based on a key distinction between two conceptually different sources of variation in our data. While heterogeneity of consumers' dynamic environment and tastes bears on both the frequency and size of optimal adjustment policies, a consumer's decision to adjust during the observation period also depends importantly on the dynamic history that brought him or her close to the boundaries of the inaction range. The former (cross-sectional) variation can be controlled on the basis of household-specific consumption drift and volatility, and other characteristics such as demographic variables and adjustment costs. After conditioning on such variables the latter (dynamic) variation is summarized in cross-sectional data by the value of the stock of durables in relation to non-durable consumption prior to adjustment. This variable conveys information about the consumer's position within the inaction band, hence about the likelihood of adjustment. Conditionally on adjustment taking place, however, it is not expected to influence the size of durables stock adjustment, which should be based on forward-looking considerations and depend on uncertainty, adjustment costs, and other characteristics of individuals' environments rather than on their past history. Thus, an infrequent-adjustment perspective provides a theoretically sound exclusion restriction to standard selection-controlled estimation techniques.

Section 2 reviews existing work and introduces our data. Section 3 sets up a theoretical framework, first characterizing frictionless intertemporal choice of optimal durables stock/non-durables flow ratios, then showing that in the presence of adjustment costs the solution features action and return points in terms of (log) deviations of the durable/non-durable ratio from its no-adjustment-costs level. Section 4 reviews testable theoretical predictions. Section 5 discusses how we measure the uncertainty and drift of the optimal durable/non-durable consumption ratio in the absence of adjustment. Section 6 specifies semi-structural models for adjustment probabilities and selection-controlled adjustment sizes, tests theoretical predictions on estimated parameters, and computes aggregate measures of optimal adjustment's responsiveness to parameter changes. Section 7 concludes.

2. THE ISSUES AND AVAILABLE DATA

Despite the realism of infrequent adjustment models and their potential importance for explaining aggregate phenomena, relatively few studies test and estimate such models on microeconomic data, and fewer still focus on the empirical role of uncertainty. Lam (1991) uses Panel Study of Income Dynamics (PSID) data to estimate the parameters of a threshold adjustment model in an extended permanent income hypothesis model, and finds evidence of 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 Consumer Expenditure Survey (CEX). His specification focuses on characterization of trigger and return points rather than on the role of structural parameters. Martin (2003) considers the behaviour of non-durable consumption during periods when durable goods stocks are not adjusted, focusing on housing purchases (where such periods are very long) and on the role of possible non-separability between durable and non-durable consumption goods.

We are not aware of previous empirical studies of the joint effects of microeconomic uncertainty and other parameters on both the extensive and intensive margin of the durable adjustment decision. Eberly (1994) and Foote, Hurst and Leahy (2000) are perhaps the closest antecedent to our work. Eberly (1994) studies car purchases using panel data from the Survey of Consumer Finances and finds that higher uncertainty increases the width of the inaction bands, but does not characterize its effects on the probability of adjustment. Her theoretical framework is based on the Grossman and Laroque (1990) model and, like that of Hassler (2001), does not allow for idiosyncratic labour income risk. Foote *et al.* (2000) instead neglect the size of durable stock adjustments, and focus on their frequency. They find that the frequency of adjustment in the CEX is negatively related to the imputed variance of household income obtained from regressions estimated with PSID data, a proxy that may be contaminated by measurement error in income and by prediction errors due to the small information set available to the econometrician.

Since realistic adjustment costs and financial market imperfections prevent closed-form solutions, structural estimation can hardly be fruitful. Our alternative strategy aims at testing key qualitative implications of a structural theoretical model. The data we analyse makes it possible to control for non-durable consumption flows along with the stock of durables (rather than for the PSID car-age variable which, in Foote *et al.*'s regressions, may imperfectly proxy for the stock of vehicles), and also allow us to account explicitly for non-diversifiable risk. To the extent that durable and non-durable consumption are differently illiquid, borrowing constraints would give a more complex role to income fluctuations. Eberly (1994) discusses ways of detecting liquidity constraints in her sample, but only from the theoretical perspective of Grossman and Laroque (1990), where an integrated financial market prevents idiosyncratic risk from inducing precautionary behaviour. Semi-structural Euler equation approaches are also difficult to implement in the presence of uninsurable risk which, as noted by Carroll (2001), generally implies that expected income growth and consumption growth are related. This makes it difficult to find valid instruments under the null of no liquidity constraints, as well as to detect the presence of liquidity constraints which, like uninsurable risk, affect predictable consumption growth and may alter the consumer's within-period budget allocation if borrowing is easier towards durables purchases (Meghir and Weber, 1996). However, we exploit the theoretically sensible notion that, in the absence of credit market imperfections, the conditional volatility of consumption growth is a sufficient statistic for all risks faced by the household, and our empirical approach and data make it possible to test (and fail to reject in our sample) the hypothesis of no liquidity constraints.

2.1. *The data*

We analyse household data drawn from the 1995 Survey of Households Income and Wealth (SHIW), which collects income, durable and non-durable consumption expenditures, financial wealth, real estate wealth, and demographic variables for a representative sample of about 8000 Italian households at intervals of (normally) two years. The four most recent waves of the SHIW have been conducted in 1993, 1995, 1998, and 2000. Appendix A describes the survey contents,

its sample design, interviewing procedure and response rates, and Appendix B reports definitions of the variables used in the empirical analysis.

The data offer information on three categories of durable goods: 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 at the end of the year, the value of any sales and purchases during the year (the value of sales is not available for furniture, as in Italy there is virtually no second-hand market for household items), and the expenditure on non-durable goods during the year. Subtracting purchases and adding sales to the end-of-year stock, and abstracting from depreciation during the period, yields an observable counterpart to the consumer’s durable stock at the time of adjustment; for households that do not adjust, we use the stock at the end of the period. It is not immediately obvious whether and how idiosyncratic variation in depreciation rates would bias the empirical findings, but the discrete-time observations available to us are a better approximation to the relevant continuous-time variable than the beginning-of-year stock (used in other applied work), which also neglects depreciation during the year. Availability of information on various categories of goods makes it possible to study infrequent adjustment of durables other than cars, on which previous microeconomic studies have focused. As we shall see, some results are fairly robust across different categories of durable goods, but theory fits cars better than furniture or jewelry.

The data also offer very useful information on each household’s dynamic environment. In a special section of the SHIW survey, households are asked a set of questions designed to elicit the perceived probability of being employed over the 12 months following the interview and the variation in earnings if employed (see Appendix B, and Guiso, Jappelli and Pistaferri (2002), for details). The relevant portion of the survey is administered randomly (according to year of birth) to only half of the 1995 SHIW households who neither are retired, nor plan to retire in the following year. After excluding observations that miss subjective expectation, income, or durables information (and three observations with outlier—larger than 10—conditional subjective variance of income growth), the sample includes 1873 households. Below, we use subjective future income information to construct measures of the first two moments of the distribution of future income. These can be used to obtain a measure of non-diversifiable risk (see Section 5 below) if consumption growth data are available, as they are for observations in the rotating-panel component of the SHIW (approximately half of the sample units are re-interviewed in subsequent rounds). For 652 (little more than one third) of the households in our 1995 sample, data on 1998 non-durable consumption are available.¹

Table 1 reports summary statistics for the sample used in estimation and for the whole 1995 sample (excluding the retirees). Values for the two samples appear comparable and confirm the randomness of the sample selection.² In the sample we will be using in estimation, the average value of the stock of cars is little more than 6000 euro, compared to about 10,000 euro for furniture and 3000 euro for jewelry. The corresponding durables stock/non-durable flow ratios are 36%, 59%, and 18%, respectively. The fractions of households who report sales or purchases is 18% for cars, 30% for furniture, and 10% for jewelry. Average net family disposable income is about 25,000 euro.

1. The 1998 SHIW is the only other survey that contains information on subjective uncertainty comparable to that collected in 1995 as well as durables information, so that at least in principle we could use the 1995–1998 panel for fixed-effect estimation of the durable choice equations. However, a subsidy programme for early scrapping of cars and motorcycles—similar to the one set up in France (see Adda and Cooper (2000))—was in place in 1997–1998. Modelling the effect of scrapping incentives is beyond the scope of this paper. Using the 1993–1995 panel is also not an option because subjective uncertainty questions were not asked in the 1993 survey. Hence we focus on the 1995 data.

2. Demographic and socio-economic characteristics of the 652 panel households are similar to those of the other households in our sample. Tables are available on request.

TABLE 1
Summary statistics for the selected and whole samples

		Our sample	1995 sample
<i>Value of stock</i>	Vehicles	6335 (6477)	5944 (6437)
	Furniture	9980 (9587)	9451 (9580)
	Jewelry	3229 (7055)	3219 (10,281)
<i>X/C</i>	Vehicles	0.3577 (0.3888)	0.3409 (0.3652)
	Furniture	0.5929 (0.5581)	0.5770 (0.5770)
	Jewelry	0.1799 (0.4441)	0.1701 (0.4684)
<i>Frequency of adjustment</i>	Vehicles	0.1788 (0.3833)	0.1765 (0.3813)
	Furniture	0.3038 (0.4600)	0.2940 (0.4557)
	Jewelry	0.0998 (0.2998)	0.0995 (0.2994)
<i>Value of purchase</i>	Vehicles	7536 (6882)	7206 (6575)
	Furniture	2168 (3717)	2514 (5031)
	Jewelry	1259 (2139)	1081 (1795)
<i>Family income</i>		24,949 (17,409)	24,163 (18,087)
<i>Age</i>		42.67 (9.09)	43.94 (9.79)
<i>Years of schooling</i>		10.00 (4.09)	9.45 (4.20)
<i>Family size</i>		3.46 (1.23)	3.35 (1.23)
<i>South</i>		0.3466 (0.4760)	0.3477 (0.4763)
<i>Public transport</i>		4.66 (1.00)	4.68 (1.02)
<i>Local council offices</i>		4.48 (1.00)	4.50 (1.02)
<i>Traffic congestion</i>		5.92 (0.96)	5.93 (0.97)
<i>Accidents per 1000 cars</i>		11.04 (2.55)	11.04 (2.58)
<i>Income uncertainty</i>		0.0465 (0.2820)	--
<i>Consumption uncertainty</i>		0.0193 (0.1174)	--
<i>Drift</i>		-0.0095 (0.2153)	
Number of observations		1873	4775

The table shows summary statistics (sample means and standard deviations, the latter in parentheses) for the selected sample of households reporting the information on the subjective probability distribution of future earnings and for the entire 1995 sample. Both samples exclude the retirees for comparison purposes. Values are in euro and are 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. See Appendix B for a definition of the quality of life indicators.

Table 2 reports summary statistics for the sub-sample of households that adjust the stock of each durable good. The large size and relatively low frequency of vehicle stock adjustments is consistent with the fact that they are costly to adjust, not only because asymmetric information (“lemon”) problems plague the second-hand market, but also because sales and purchases of vehicles entail administrative costs. Furniture (and especially jewelry) also faces “lemon” and

TABLE 2
Summary statistics for the sample of those who adjust

		Adjust vehicles	Adjust furniture	Adjust jewelry
<i>Value of stock</i>	Vehicles	10,811 (8036)	7224 (6905)	9011 (8028)
	Furniture	11,079 (9492)	11,120 (9540)	13,904 (12,148)
	Jewelry	3586 (6055)	3967 (8058)	6593 (11,333)
<i>X/C</i>	Vehicles	0.5257 (0.3589)	0.3733 (0.3308)	0.4388 (0.3664)
	Furniture	0.5427 (0.4229)	0.5960 (0.5137)	0.7013 (0.6074)
	Jewelry	0.1679 (0.2475)	0.1938 (0.3005)	0.3045 (0.4283)
<i>Value of purchase</i>	Vehicles	7536 (6882)	7447 (7649)	9131 (8646)
	Furniture	1709 (2547)	2168 (3717)	2888 (4155)
	Jewelry	847 (835)	1055 (1279)	1259 (2139)
<i>Family income</i>		30,818 (22,123)	28,544 (19,450)	33,499 (43,884)
<i>Age</i>		42.49 (9.15)	42.06 (9.24)	40.32 (9.60)
<i>Years of schooling</i>		10.40 (3.77)	10.84 (4.20)	11.83 (4.17)
<i>Family size</i>		3.54 (1.24)	3.35 (1.29)	3.19 (1.14)
<i>South</i>		0.2539 (0.4359)	0.2922 (0.4552)	0.2499 (0.4341)
<i>Public transport</i>		4.48 (0.91)	4.28 (1.04)	4.61 (0.98)
<i>Local council offices</i>		4.74 (0.88)	4.55 (0.97)	4.87 (0.83)
<i>Traffic congestion</i>		6.77 (0.84)	6.87 (0.96)	6.67 (0.90)
<i>Accidents per 1000 cars</i>		10.57 (2.12)	11.08 (2.52)	10.40 (2.08)
<i>Income uncertainty</i>		0.0438 (0.4737)	0.0483 (0.3921)	0.0553 (0.6227)
<i>Consumption uncertainty</i>		0.0150 (0.0107)	0.0152 (0.0194)	0.0153 (0.0115)
<i>Drift</i>		-0.0134 (0.0329)	-0.0104 (0.0417)	-0.0127 (0.0342)
Number of observations (sample fraction)		322 (0.1641)	564 (0.3005)	193 (0.1018)

The table shows summary statistics (sample means and standard deviations, the latter in parentheses) for the selected sample of households reporting the information on the subjective probability distribution of future earnings and adjusting the stock of durables. Values are in euro and are 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. See Appendix B for a definition of the quality of life indicators.

intermediation costs, but entails little or no administrative costs upon purchase. The data do not offer a good measure of overall adjustment costs, but we can rely on a proxy which applies mainly to cars: the efficiency of local public administration. In order to get a licence plate and temporary registration, the buyer of a new car in Italy 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 of the purchase cars must be registered with the *Pubblico Registro Automobilistico* (PRA). This entails filling in an application form (with notarized signatures of both buyer and seller) as well as paying a fee. The costs of dealing with the PRA and notary publics may have a monetary equivalent for busy and well-to-do households who hire the services of specialized agencies, but for most households car purchase

entails search, time, and psychic utility costs. They certainly do vary across different parts of Italy, and the relevant variation is likely to be captured by the responses to the SHIW survey question on perceived efficiency of public council offices (see Appendix B). 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 (a mean difference test has a $|t|$ -statistic of 6). We will proxy adjustment costs with indicators of such dissatisfaction at the level of provinces.³ We will also make use of province-level indicators of dissatisfaction with public transport systems, traffic congestion, and frequency of car accidents, which may affect depreciation rates (detailed definitions are listed in Appendix B; in Table 1, we see that about 1% of cars are involved in an accident during the year). These variables may of course capture province-specific effects other than transaction costs and depreciation, especially as regards adjustment of durables other than cars. As we shall see, controlled regression specifications suggest that the adjustment-cost proxy appears to capture relevant features of the cars' adjustment process, while indeed fitting other goods' adjustment process rather poorly.

Tables 1 and 2 also report summary statistics for our measures of income and consumption risk and for the drift in the durable/non-durable consumption ratio. We defer discussion of these statistics to Section 5.

3. A THEORETICAL FRAMEWORK

If both durable and non-durable consumption yield utility, and both income and asset returns are random, a consumer's dynamic optimization programme is analytically intractable and even numerical analysis must rely on drastic simplifications. The classic (Grossman and Laroque, 1990) study of optimal durable consumption abstracts from non-durable goods and labour income, to obtain analytic and numerical results for the case where asset returns are described by Brownian increments and the utility function has constant elasticity. Other researchers have chosen different simplifications and approximations (see Attanasio (2000) and Padula (2000), for references and discussions).

We do not aim at offering a full characterization of all realistic features of optimal durable and non-durable consumption policies. Our theoretical approach is focused on qualitative insights conducive to empirical estimation and testing, and is based on variational arguments which offer characterization results of some generality and, for our purposes, are more useful than existence results and numerical construction techniques. Even if a tractable specification or extensive numerical results were available, empirical analysis would not be able to detect and test detailed features of realistic models. Since the data we analyse do contain information on labour income uncertainty, we follow Bertola and Caballero (1990) and Bar-Ilan and Blinder (1992) in characterizing durable consumption choices in the presence of labour income risk. And since detailed information is available on consumer characteristics and on more than one durable good stock, we pay particular attention to characterization results for stocks and adjustment patterns across heterogeneous observations.

Let time- t period utility be a function of non-durable consumption $C(t)$ and of the available durable goods stock $X(t)$; more than one durable good may yield utility, and we will briefly discuss below the implications of treating X as a vector of imperfectly substitutable goods. In continuous time, the consumer aims at maximizing

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

3. The country is divided into 95 provinces and 20 regions (a collection of provinces). A province size is roughly comparable to that of a U.S. county.

where $u(\cdot)$ is an increasing and strictly concave function and ζ the consumer's discount rate. The utility function may depend on the consumer's demographic characteristics and on other, observable or unobservable, determinants of preferences. For simplicity, however, the notation does not explicitly account for taste heterogeneity in this section.

The problem's budget constraint features a stochastic flow of labour income, $Y(t)$, and financial assets or liabilities, $W(t)$, at time t . Financial returns are independent of consumption and income dynamics, hence no insurance is available. In the light of the essentially cross-sectional character of the available data-set, we take relative prices and rates of return to be constant over time when modelling a given consumer's optimal consumption policy. The relative price p of durables and the depreciation rate δ of the durable good, however, may in principle differ across households and goods in the data. Denoting with r the financial rate of return, the accounting identity $C(t)dt + dW(t) + (dX(t) + \delta X(t)dt) p = Y(t)dt + rW(t)dt$ can be rearranged to read

$$C(t) + (r + \delta) pX(t) = r\tilde{W}(t) + Y(t) - \frac{d\tilde{W}(t)}{dt}, \forall t, \quad (2)$$

where $\tilde{W}(t) \equiv W(t) + pX(t)$ includes both financial assets and the durable goods stock, and the latter's flow of services is notionally rented at the "user cost" flow price $(r + \delta)p$ per unit time. The financial and durable stock component of total wealth $\tilde{W}(t)$ can both be discontinuous when adjustment costs imply that purchases are discrete, but $\tilde{W}(t)$ is continuous: its increments are given by the difference between the flow of purchasing power accruing from wealth returns and labour income and that imputed to the period's durable and non-durable services consumption.

3.1. No adjustment costs

We begin by characterizing optimal consumption patterns in the *absence* of adjustment costs.

Denoting with $v \equiv (r + \delta)p$ the user cost of durables, and with $M(t) \equiv C(t) + vX(t)$ the total purchasing power allocated to non-durable consumption and durables services at time t , consider the indirect utility function

$$\max_{X(t)} [u(M(t) - vX(t), X(t))] \equiv U(M(t)). \quad (3)$$

If the consumer's tastes are homothetic, they can be written in the form

$$u(C(t), X(t)) = f(h(C(t), X(t))) \quad (4)$$

for $f(\cdot)$ a strictly increasing and concave function and $h(\cdot, \cdot)$ a first-degree homogeneous function. The marginal rate of substitution only depends on the ratio of the quantities consumed, and along an optimal path with constant r , δ , and p the first-order conditions of (3) imply

$$\begin{aligned} vX(t) &= \chi(v)M(t) \\ C(t) &= (1 - \chi(v))M(t) \end{aligned} \quad (5)$$

where $\chi(\cdot)$ is a function of the user cost v whose form depends on that of the $h(\cdot, \cdot)$ function.⁴

4. For example, if $h(C, X) = ((1 - \beta)C^\mu + \beta X^\mu)^{1/\mu}$ then

$$\frac{X}{C} = \left(\frac{\beta}{(1 - \beta)v} \right)^{\frac{1}{1-\mu}} \text{ and } \chi(v) = \frac{\left(\frac{\beta}{1-\beta} \right)^{\frac{1}{1-\mu}}}{v^{\frac{\mu}{1-\mu}} + \left(\frac{\beta}{1-\beta} \right)^{\frac{1}{1-\mu}}}.$$

The special case of Cobb–Douglas $h(\cdot, \cdot)$ is obtained by letting $\mu \rightarrow 0$.

Hence, the utility flow accruing in period t , $U(M(t)) = f\left(M(t)h\left(1 - \chi(v), \frac{\chi(v)}{v}\right)\right)$, is a function $f(\cdot)$ of the purchasing power $M(t)$ allocated to finance that period's consumption, scaled by a constant that depends only on intratemporal preferences and v .

If the consumer can borrow and lend at rate r , intertemporal allocation of purchasing power obeys the standard Euler condition

$$U'(M(t)) = e^{(r-\delta)\tau} E_t [U'(M(t+\tau))] \quad \forall t, \tau \geq 0, \quad (6)$$

where $E_t[\cdot]$ denotes expectations taken at time t using the conditional probability measure induced on future variables by observation of the income process and by the consumer's optimal policy. In the absence of adjustment costs or borrowing constraints, each of the non-durable consumption and durables service flow marginal utilities also change unpredictably with respect to that filtration, because the allocation of $M(t)$ that solves (3) implies that the marginal utilities derived from the two goods are proportional to $u'(M(t))$.⁵

An optimal solution to the intertemporal $M(t)$ allocation problem exists under mild regularity conditions, and is immediately applicable to the simple homothetic preferences, constant-prices multiple-good setting considered here. In the presence of uninsurable risk, however, analytic expressions are only available for restrictive utility specifications, and numerical solution is practical only for simple specifications of the consumer's environment. Explicit solution is further complicated by first-order transaction costs which, if the utility function is differentiable, imply that small deviations from the frictionless durable/non-durable consumption ratio implied by (5) should not be corrected. For our empirical purposes, it will suffice to combine the familiar Euler characterization of intertemporal allocation with conceptually similar relationships between endogenous variables that link durable and non-durable consumption flows, rather than current and future non-durable consumption flows.

3.2. Separation of intertemporal and adjustment dynamics

For simplicity and in the light of data availability, we express adjustment costs in per-adjustment form, and in terms of utility. This specification is appropriate if adjustment entails lump-sum costs independent of the size of adjustment—as is the case when purchasing a car requires visiting dealers and registration bureaus. It does not account for the difference between the unit price applicable to consumers' durable sales and purchases: in reality, second-hand markets do generally discount used goods over and beyond economic depreciation, but our data offer no information as to whether and how such transaction costs may vary in the sample. Fixed per-adjustment costs imply that the unit price of small net purchases of durables is large, and the policy of frequent purchases and sales that would allow the consumer's durables stock to remain aligned with the current period's budget non-durable consumption as in (5) cannot be optimal. The optimal policy satisfies optimality conditions that restrict the total (rather than marginal) expected welfare effect of feasible infrequent adjustment decisions, and depend in general on all aspects of the consumer's dynamic environment.

We can however specify (and test empirically) functional form restrictions that allow different dynamic aspects of the consumer's optimization problem to be analysed separately. If the utility function is not only intratemporally homothetic, but also displays unitary intertemporal elasticity of substitution, then the optimality condition for intertemporal consumption patterns

5. Variability of the user cost of durables $[\delta(t) + r(t)]p(t) - E[dp(t)]/dt$ would bear on all aspects of the consumer's problem, which would need to be based on conditional expectations of future price and rate of return dynamics. We abstract from such complications because our essentially cross-sectional data offer no information on relative prices and rates of return or depreciation.

does not depend on the stock of durables implied by infrequent adjustment. To see this, note that if the function $f(\cdot)$ in (4) is logarithmic then

$$u(C(t), X(t)) = \ln(M(t)) + \ln\left(h\left(1 - \chi Z(t), \frac{\chi}{v} Z(t)\right)\right), \quad (7)$$

where $Z(t) = \frac{X(t)}{X^*(t)}$ for $X^*(t) \equiv \frac{\chi}{v} M(t)$ the frictionless durable stock. Thus, the marginal utility afforded by a higher $M(t)$ does not depend on its within-period allocation to durables and non-durables, as indexed by the ratio $Z(t)$ of the actual durable stock to the optimal one.

Additive separability makes it straightforward to extend the theoretical framework to allow the consumer to derive utility from more than one stock of durables. Just as it would allow two-stage budgeting if adjustment costs were absent but relative prices were allowed to vary over time, it yields an equally tractable framework for our empirical application. During periods when adjustment costs make inaction optimal, the actual composition of the consumption bundle does not coincide with the preferred one, which is constant under homothetic preferences and constant prices. Inaction causes fluctuations in the marginal utilities of the two goods but, by the additively separable form of (7), does not affect the marginal utility of the expenditure flow $M(t)$ —which therefore satisfies a logarithmic-utility version of the Euler equation (6) regardless of whether adjustment costs are present.

Logarithmic preferences imply that infrequent adjustment leaves unaffected the Euler equation characterizing the optimal intertemporal allocation of purchasing power, and since adjustment costs are viewed in terms of utility they do not appear in the consumer's budget constraint. Hence, the $\{M(t)\}$ process is the same for any adjustment costs. Under logarithmic (and other CRRA) preferences, marginal utility increases at an increasing rate as consumption declines towards zero, and makes it optimal for the consumer to engage in precautionary-savings behaviour. For our purposes, we do not need to characterize theoretically the strength of this and other influences on the volatility of marginal utility. Rather than postulating *ad hoc* approximate relationships between the consumption and income processes—as in, *e.g.* Banks, Blundell and Brugiavini (2001)—we will assess empirically below the extent to which self-reported income volatility appears to affect non-durable consumption volatility, and the extent to which the latter induces steeper “precautionary” consumption profiles for our sample households.

Whatever are the effects of income uncertainty (filtered by precautionary savings) on the non-durable consumption flow and stock of durables that the consumer would find optimal in the absence of adjustment costs, both components of the desired consumption bundle are identically affected by those phenomena if preferences are homothetic and prices are constant. Our specification of the consumer's problem lends itself naturally to a study of cross-sectional data with information about different durable goods and non-durable consumption flows, and makes it possible to focus empirically on specific issues (such as the role of uncertainty in shaping durables adjustment) on the basis of relationships between endogenous variables whose levels cannot be pinned down exactly by fully worked-out theory, without having to rely on a complete, and necessarily simplified, representation of reality. Of course, unitary intertemporal elasticity may unduly restrict consumer preferences. In the absence of information generated by fluctuations of intertemporal rates of transformation along the time-series dimension, however, data cannot offer much information about utility curvature, and our Euler equation estimates below offer statistical support to the logarithmic specification.

3.3. Infrequent adjustment

We will focus our empirical approach on effects of uncertainty on the width and frequency of durables adjustment. Here, we proceed to characterize adjustment of durable/non-durable

consumption ratios when income is the only source of uncertainty, preferences are homothetic with unitary intertemporal elasticity, and adjustment costs imply fluctuations of durable/non-durable ratios around the (constant) ratio implied by those assumptions. The consumer should adjust the composition of the consumption bundle when the cost of doing so is at least compensated (and, along an optimal path, exactly compensated) by the expected discounted welfare benefits produced by smaller future deviations of the utility flow accruing to the consumer from that which would be implied by unconstrained allocation to durables and non-durables of the intertemporally optimal $M(t)$ expenditure flow. At times when it is optimal to adjust the stock of durables, financial funds are used to purchase durables (or durable sales increase financial funds) and the composition of the consumer's total wealth $\tilde{W}(t) = W(t) + pX(t)$ changes discretely. Throughout periods when the durable stock is not adjusted, conversely, non-durable consumption is $C(t) = M(t) - vX(t)$, and homothetic preferences imply that flow utility losses from inaction are only a function of the $Z(t)$ ratio of the current durables stock to the statically optimal one.

If the expected present discounted value of those losses can be written as a function of that same state variable, the optimal adjustment policy must take the trigger/return form familiar from Bertola and Caballero (1990), Eberly (1994), and other related work: namely, two durable/expenditure ratios, L and U , should trigger adjustment to a return point s . For any linearly homogeneous function, a Taylor expansion in logarithms around the optimal ratio (where the first derivative vanishes by definition) reads

$$\ln(h(M - vX, X)) - \ln\left(h\left((1 - \chi(v))M, \frac{\chi(v)}{v}M\right)\right) \approx \frac{1}{2}b(v)(\ln(X) - \ln(X^*))^2,$$

where

$$b(v) \equiv \frac{d^2}{dx^2} \ln(h(M - ve^x, e^x)) \Big|_{x=\ln(X^*)} < 0 \tag{8}$$

may depend on the user cost, and is intuitively larger in absolute value when durables and non-durables are poor substitutes for each other.⁶ If the elasticity of substitution between durables and non-durables is unitary, the Cobb–Douglas specification $\ln(h(C, X)) = (1 - \beta)\ln(C) + \beta\ln(X)$ of the intratemporal utility function implies that $b(v) = \beta/(1 - \beta)$ depends only on the utility weight (and budget share) β of the durable stock: when durables do not yield much utility, then deviations from the small optimal frictionless stock matter little. We will focus below on this intuitive implication of the simple log-linear specification, which also straightforwardly allows multiple durable goods to make separable contributions to the consumer's utility.

To characterize infrequent adjustment, consider the state variable

$$z(t) \equiv \ln(Z(t)) = \ln(X(t)) - \ln(X^*(t)) \equiv x(t) - \kappa - m(t), \tag{9}$$

where $x(t) \equiv \ln X(t)$, $m(t) \equiv \ln M(t) = \ln(C(t) + vX(t))$, and $\kappa \equiv \ln \frac{\chi}{v}$.

If at times when the consumer refrains from purchases or sales of durables this variable's dynamics can be approximated by an arithmetic Brownian motion process,

$$dz(t) \equiv \vartheta dt + \sigma d\omega(t), \tag{10}$$

an explicit solution is available (see, e.g. Bertola and Caballero (1990) and Eberly (1994)) for the optimization problem

6. If durables and non-durables are strict complements, deviations from the optimal ratio have a strong negative impact on utility (there is little joy in owning a luxury car, but using no gasoline). Conversely, intratemporal utility is not strongly affected by non-adjustment if it is nearly linear, and non-durable consumption can perform much the same role as the durable stock in generating utility.

$$V(z_t) \equiv -\min_{\{z_\tau\}} E_t \int_t^\infty e^{-\rho(\tau-t)} \left(\left(\frac{b(v)}{2} z(\tau)^2 \right) d\tau + [\text{adjustment costs}] \right), \quad (11)$$

where $b(v)$ is the slope of marginal utility losses due to deviations from the within-period optimum, as defined and approximated above, and adjustment events all entail the same lump-sum cost in terms of utility: this, denoted A in the figures below, can literally represent the consumer's annoyance when shopping for durable goods; but since logarithmic preferences imply that the utility price of wealth is inversely proportional to the intertemporally optimal expenditure flow, it may also approximate adjustment costs in terms of goods that are proportional to the consumer's durables stock, or total wealth. Appendix C reports the optimality conditions satisfied by the L, s, U adjustment points for the state variable z , the piecewise exponential formulae for the long-run distribution of the controlled variable z within the $[L, U]$ optimal inaction interval, and the expressions for the process's adjustment intensities. Such long-run features are relevant to our empirical work on a set of cross-sectional observations which, after controlling for observable heterogeneity, may be interpreted as independent draws from the histories of decision makers facing similar problems.

The log-linear process (10), with ϑ and σ constant parameters, is approximately realistic for the $z(t)$ process defined in (9) if the logarithm of the durables stock depreciates linearly and a random walk with drift is a good approximation to the consumer's intertemporal expenditure pattern.⁷ Of course, in the presence of uninsurable risk optimal expenditure growth need not be i.i.d., and non-separability of durable and non-durable consumption may also shape the relevant processes during inaction periods (Martin, 2003). Qualitatively, however, the volatility parameter σ would still reflect uninsurable shocks, and the drift parameter ϑ would still reflect depreciation of the durables stock as well as predictable consumption growth. Our empirical analysis of different durables stocks and heterogeneous consumers will aim at exploiting heterogeneity of such phenomena across households.

4. THEORETICAL IMPLICATIONS

Like any model, the one outlined above should be viewed as a simple approximate representation, to be confronted with the available data, of more general preferences and dynamics. The preference specification underlying our empirical analysis relies on a log-linear utility function and on log-linear dynamics within an inaction period, with unpredictable increments depending on the volatility of the optimal intertemporal spending pattern, and a predictable drift component depending on expected marginal utility increments (which in turn may depend on volatility, via precautionary behaviour), as well as on the rate of depreciation of durable goods.

Before proceeding to formal empirical analysis, in this section we characterize the theoretical implications of various features of a consumer's problem for the probability and the size of adjustment and for the shape of the durable/non-durable consumption ratio's long-run distribution.

7. Formally, approximation of the Euler equation satisfied by $M(t)$ leads to

$$dm(t) = \vartheta_m dt + \sigma_m d\omega(t),$$

where σ_m^2 is the variance of $m(t)$ increments and ϑ_m its drift, or expected expenditure growth (which also depends on σ_m^2 via precautionary saving). Since $dx(t) = -\delta dt$ in the absence of adjustment, one can use (9) to obtain (10),

$$dz(t) = -(\vartheta_m + \delta) dt - \sigma_m d\omega(t) = \vartheta dt + \sigma d\omega(t).$$

The drift $\vartheta = -(\vartheta_m + \delta)$ is determined by the rates of depreciation, δ , and expected expenditure growth, ϑ_m . The volatility of z coincides with the volatility of expenditure growth if the user cost of durables is constant, as we assume. In more general models, changes in relative prices and stochastic depreciation would also be relevant to the dynamics of z . Given the cross-sectional nature of our data, we abstract from these complications.

4.1. Tastes

It is important to recognize at the outset that while the optimizing behaviour characterized above is that of a hypothetical infinitely lived consumer, the data we will use later to test the model's predictions are drawn from a cross-section of demographically heterogeneous consumers with possibly different tastes for durables (Attanasio (2000) finds strong evidence of demographic effects in his data). Our theoretical approach, explicitly based on a preference specification, makes it possible to give a structural interpretation to demographic variation in empirical work.

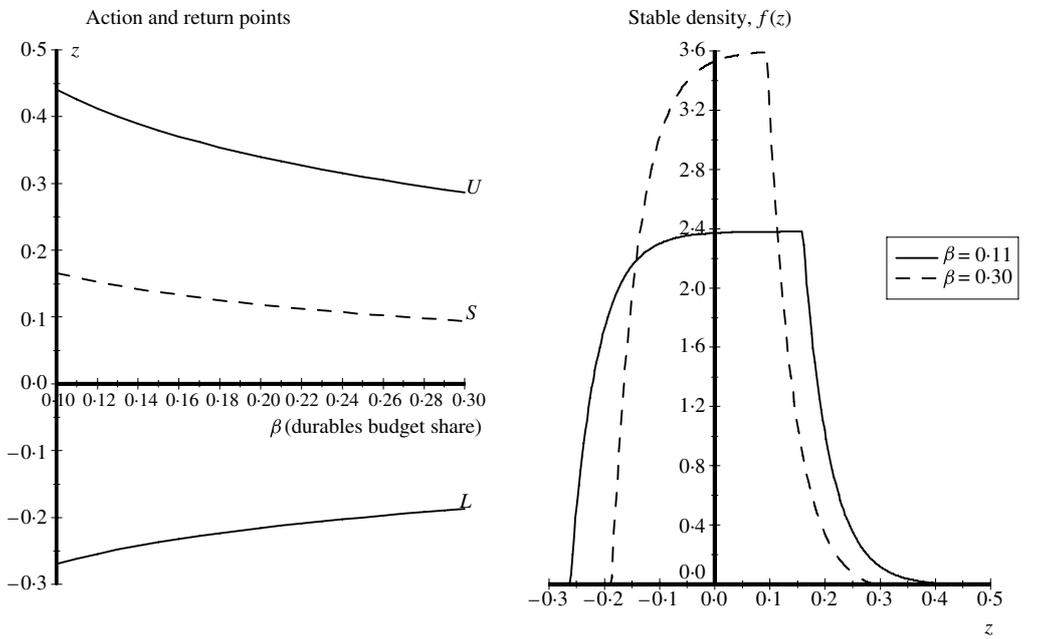
Figure 1 illustrates how several aspects of the dynamic problem's solution depend on the budget share of durables, denoted β , when the intratemporal elasticity of substitution is unitary. The values chosen for the drift ϑ , the variance σ^2 , and the slope of utility losses from misalignment of non-durable consumption flows and durables stocks are meant to be roughly realistic if a unit of time is a year. Adjustment cost parameters are in terms of utility units, or fractions of permanent income under the logarithmic utility approximation. Thus, the $A = 0.01$ lump sum cost in the figure would mean that direct and opportunity costs of shopping for the durable good amount to 1% of a year's utility flow. Tastes have an obvious impact on the consumer's optimal policy: larger values of β imply larger trigger and return points, and long-run average levels, for the durables/non-durables ratio. Less obviously, and more importantly for our purposes, this parameter also affects the frequency and size of adjustments. In the top-left diagram of Figure 1, a larger durables budget share β (on the horizontal axis) is associated with an increasingly narrower distance between the trigger and return values of the state variable z (on the vertical axis).⁸ By the definition of b in equation (8) above, a larger β increases the cost of departing from the statically optimal consumption bundle. Hence, stronger taste for durables implies a narrower inaction range and smaller optimal adjustment sizes in the bottom-left diagram of Figure 1. The long-run density of the consumer's deviation from the statically optimal durable/non-durable log ratio, displayed in the top-right diagram, is skewed in the direction of the drift and has wider domain when β is smaller. And in the bottom-right diagram, the unconditional probability intensity of adjustment at the lower boundary of the inaction range is an increasing function of β , since, for given volatility, the boundaries of a narrower inaction range are more likely to be reached frequently (here and in the following figures the drift is sufficiently negative, reflecting strong depreciation and/or fast average consumption growth, to imply that downward adjustment occurs with negligible probability).

The theoretical relevance of taste parameters motivates inclusion of demographic variables in our empirical specifications below. To assess the descriptive fit of our modelling perspective, we consider the empirical distribution of the durables/non-durables log ratio $\ln(X(t)/C(t))$, which is directly observable in our data and is closely related to the theoretically relevant variable defined in (9):

$$z(t) \equiv \ln\left(\frac{X(t)}{C(t) + vX(t)}\right) - \kappa \approx \ln\left(\frac{X(t)}{C(t)}\right) - \kappa,$$

where the approximation neglects the role of the constant v in shaping the level and dynamics of durable and non-durable consumption for each individual. In cross-section, the level of X/C depends, through κ , on the household's tastes. Using a procedure similar to that of Eberly (1994), we filter the data with a regression of the durables stock/non-durable consumption ratio on the same set of demographic characteristics we will include in the regressions below—education, age, number of children in three age bands, number of earners, dummies for city size and for area of residence—which may absorb determinants of that ratio other than dynamic variation of

8. Here, and in the following figures, non-zero drift and discount rates make it optimal for the return point to differ slightly from the static optimum, since this yields a better present discounted value of deviations in exchange for the adjustment cost.



$A = 0.010 \quad b = 2 \beta / (1 - \beta) \quad \rho = 0.05 \quad \sigma = 0.12 \quad \vartheta = -0.15$

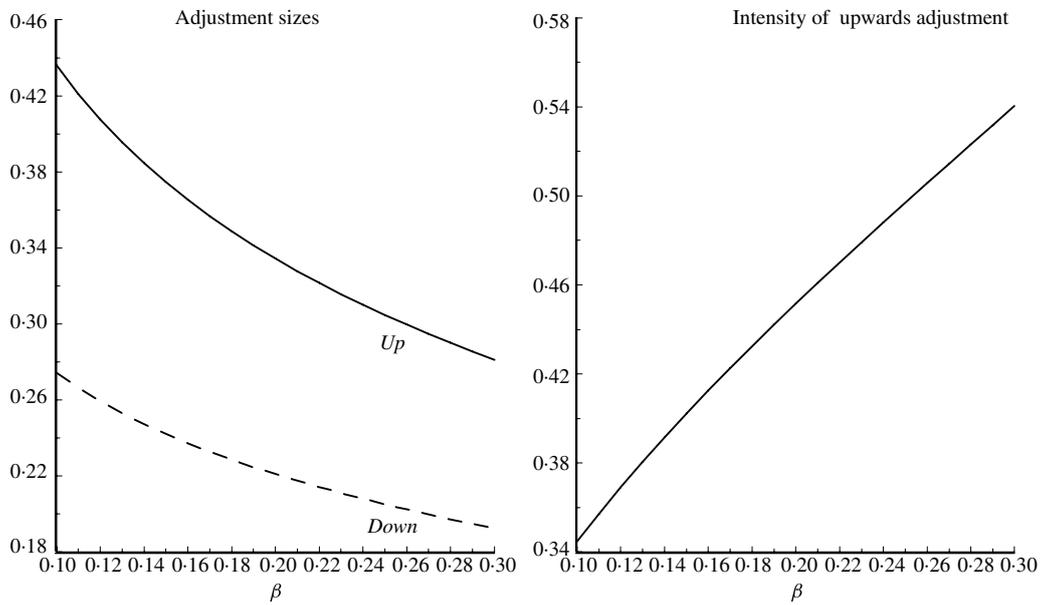


FIGURE 1
Implications of the durable budget share, β

the type featured by our representation of a typical consumer's problem: not only tastes, but also user costs may in reality vary across the households in our sample. However, in the absence of disaggregate information on prices and depreciation rates, we will subsume heterogeneity in user

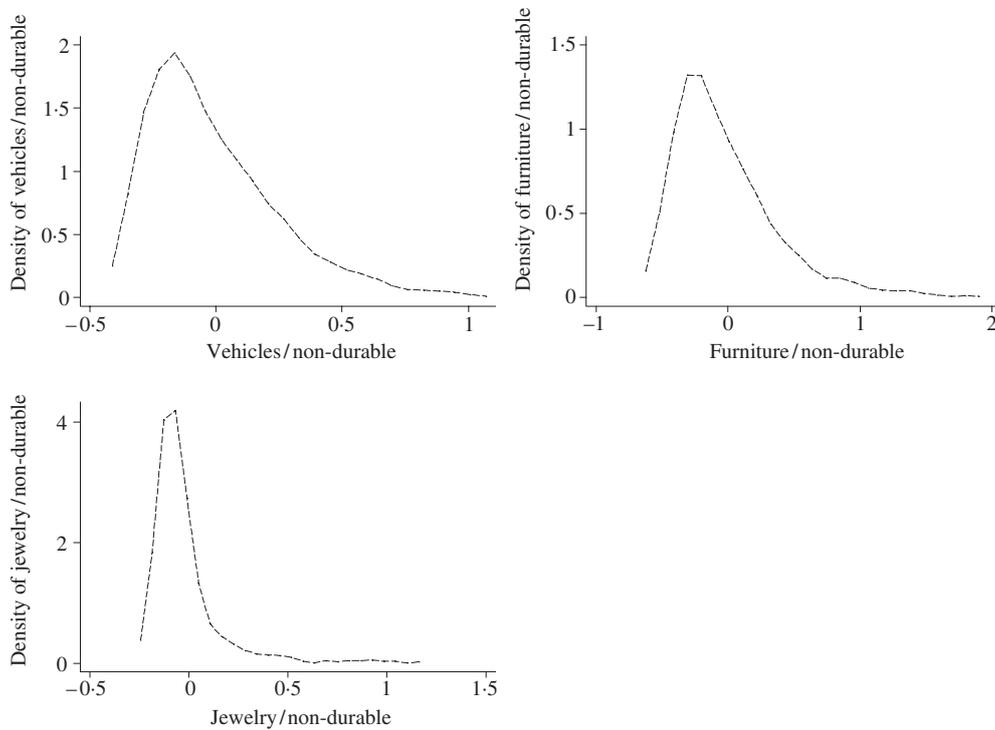


FIGURE 2

Empirical density of X/C for vehicles, furniture, and jewelry

costs in the component of the data variation controlled by observable characteristics. Then, we compute the empirical density of the residuals, using Gaussian kernel non-parametric smoothers evaluated at 25 points over the range of X/C .⁹ Figure 2 plots the empirical density of the log-ratio unexplained by demographic heterogeneity for vehicles (top-left diagram), furniture (top-right), and jewelry (bottom). In all cases, the empirical density of X/C is skewed, and qualitatively similar to the theoretical stable density plot in the previous figure. While such a shape could be spuriously generated by uncontrolled heterogeneity in raw data, its appearance in filtered data suggests that the optimal inaction model approximates the durable goods consumption problem well.

4.2. Uncertainty, drift, and adjustment costs

The four diagrams of Figure 3 illustrate theoretical effects of different levels of uncertainty, as summarized by σ (the standard deviation of innovations in the process, denoted $z(t)$ in Section 3, that represents deviations from the frictionless durable/non-durable consumption bundle in the absence of adjustment). The parameters are set at the same values as in the previous and following illustrations of theoretical insights.

9. The stock of durables is measured as of the end of the period, and observations in the lower and upper percentile of the distribution are excluded from the computation. The empirical X/C ratio is zero when a household owns no durables in a category. This is not literally consistent with our simple theoretical framework, where durables are infinitely divisible and preferences are homothetic. Attanasio's (2000) less structural model accounts for selection of observations into the zero-durables category.

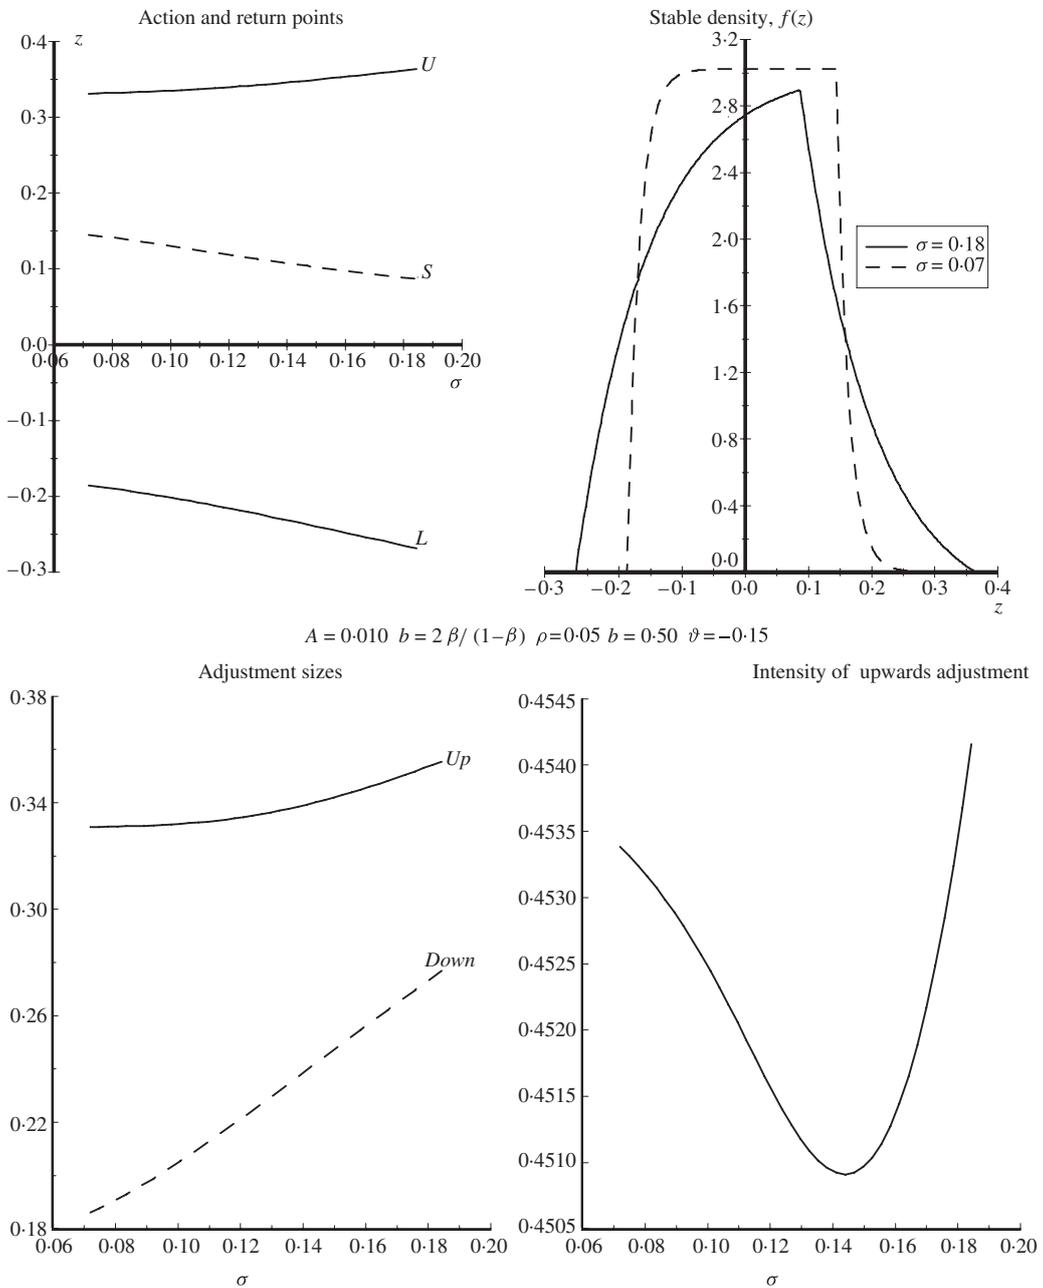


FIGURE 3
Implications of the extent of uncertainty, σ

Along the vertical axis of the top-left diagram, we see that the inaction range becomes wider as uncertainty increases, to imply that adjustment is larger when it does occur (see bottom-left diagram). Intuitively, if the consumer knows that deviations from the optimal configuration of the durable/non-durable bundle are very volatile it is optimal not to bear adjustment costs in order

to correct such large deviations, and rather wait and see whether random events correct them without cost. Note that for a given state z on the vertical axis, the inaction range is smaller for low than for high levels of uncertainty. Theory predicts that adjustment should occur whenever z falls outside the inaction range, never if it is inside it. Empirically, the stock of durables is measured (with error) at the time of adjustment, and non-durable consumption (also measured with error) refers to the average flow throughout a discrete-time interval when it is expected to grow at the trend rate, but is disturbed by news. In the formal regressions below, where the disturbance term captures these sources of noise, we will test the milder theoretical implication that, *conditioning* on the information summarized by z , higher levels of uncertainty about future developments should be associated with lower probabilities of adjustment.

Regarding *unconditional* implications, a broader inaction range obviously implies wider steady-state dispersion of z deviations when σ is larger, as illustrated in the top-right diagram of Figure 3. The negative drift imparts an asymmetric shape to the density, and the trigger and return points are offset in the opposite direction around the frictionless optimum: the consumer's infrequent-adjustment policy lets the durables/non-durables ratio fluctuate around that reference point, leaving its average value broadly unchanged for different levels of uncertainty (recall that, in contrast, different taste-for-durable parameters had sharp implications for the average bundle consumed). The bottom-right diagram of the figure plots the long-run average frequency (or intensity) of adjustment. Since drift and variance interact in determining the long-run density of the process in the neighbourhood of the adjustment trigger (see the formal expressions in Appendix C), the relationship between the probability of action and σ is not monotonic when the drift is not zero, as in the figure. In the absence of drift, adjustment would be unambiguously more frequent for larger values of σ : higher uncertainty makes inaction optimal in the face of larger z deviations, but also increases the likelihood of wider swings in that process's realizations, hence the probability that action be triggered by any given barrier; since marginal utility losses are increasing in z , the consumer should trade smaller average flow losses against larger adjustment costs, and should not expand the range of inaction so much as to imply an unchanged frequency of adjustment in the presence of more uncertainty. The width of the band also has implications for the average intensity of adjustments triggered by the process's drift, however, and for the parameters used in plotting the figures the relationship between the unconditional intensity of adjustment and σ is markedly U-shaped.

The conceptual distinction between conditional and unconditional predictions is important for understanding the relationship between our results and those of studies based on more stylized theoretical perspectives. Unlike Foote *et al.* (2000), we will not test the theory on the basis of relationships between uncertainty measures and adjustment frequencies. As discussed above, the relationship between these theoretical variables is not monotonic for given drift and this, together with the fact that drifts are heterogeneous across households and imperfectly observable, makes it difficult if not impossible to formulate and test unconditional predictions. Accordingly, our testing strategy below will be based on more precise and informative conditional predictions regarding the effects of uncertainty on the probability and size of adjustment.

Theoretical implications of the drift parameter are illustrated in Figure 4, where $\vartheta < 0$ as is realistic in the presence of depreciation and positive expected non-durable expenditure growth. A more negative value of ϑ , due to, *e.g.* faster depreciation or a larger difference between the consumer's rates of return and discount, represents a larger (in absolute value) drift of the state variable z . Consider the empirically relevant case of upward adjustment. As the consumer trades larger utility losses off the larger expected adjustment costs entailed by faster travel between the boundaries of a given inaction band, a larger drift (a more negative value of ϑ) enlarges the inaction range in the top-left diagram, increases the size of upward adjustments in the bottom-left diagram, yields an increasingly skewed steady-state distribution with wider support in the

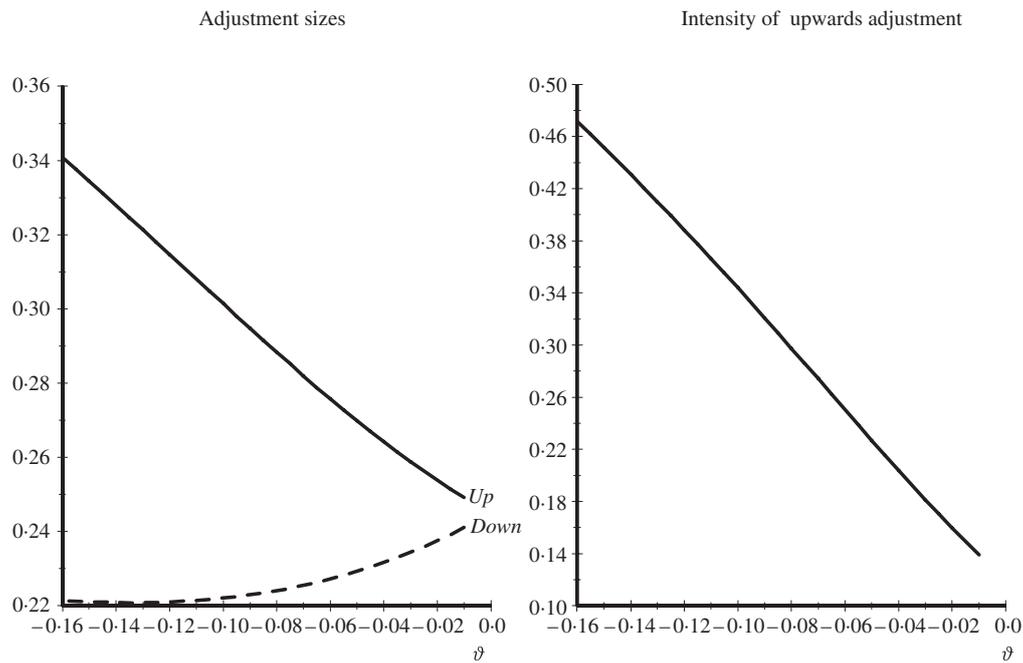
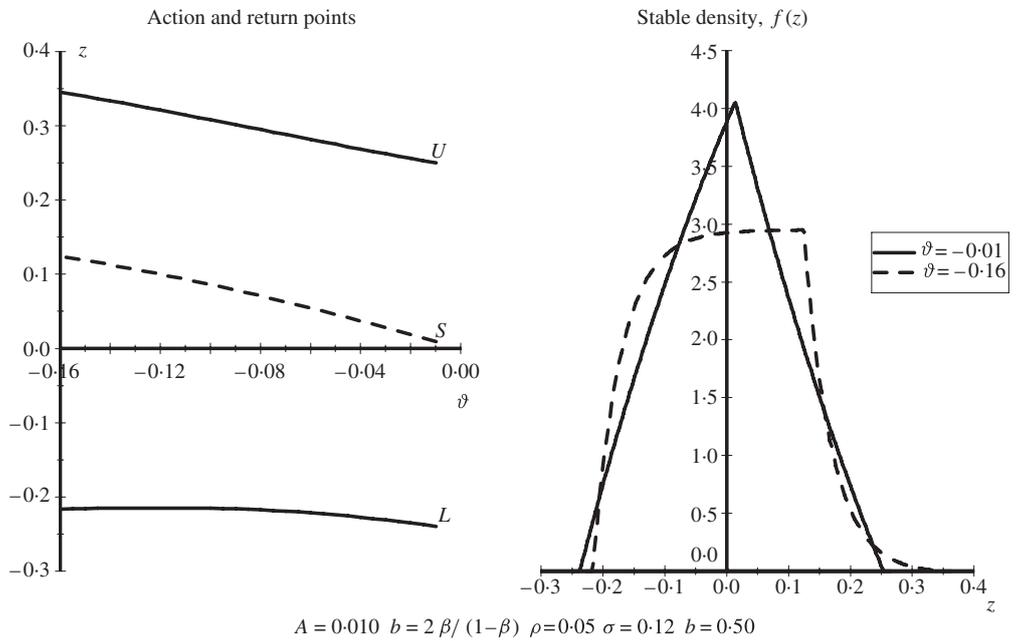


FIGURE 4
Implications of the drift, ϑ

top-right diagram, and increases the frequency of upward adjustment in the bottom-right diagram. Conditioning on the information summarized by z , theory predicts that an increase in the drift (a more negative value of ϑ) should increase the probability of upward adjustment.

Finally, larger adjustment costs (for given slope of marginal flow utility losses) imply a wider inaction range, lower intensity of adjustment, and larger adjustment steps. These intuitive effects do not need to be illustrated graphically, but deserve to be confronted with the admittedly limited information available to us as to variation of adjustment costs across households in our data-set.

5. MEASURING UNCERTAINTY AND THE DRIFT

An important aspect of our empirical analysis is the construction of measures of uncertainty and drift that match as closely as possible the theory's insights. The relevant drift and uncertainty are those of the difference z between the actual and optimal logarithms of the durable/non-durable consumption ratio during inaction periods, when only depreciation and relative price changes may drive actual and desired durable stocks. As noted above, the drift (volatility) of z should be larger in absolute value for goods with faster (more volatile) depreciation and for consumers with steeper (more volatile) non-durable consumption profile.

In principle, it would be interesting to exploit differences in drift across types of durable goods with heterogeneous depreciation rates. The depreciation rate, however, affects the frictionless optimum durables/non-durables basket through the user cost, as in equation (5), as well as the speed at which the actual basket traverses the consumer's state space. Too little relevant information is available in our data (see Section 2 above) for empirical work to yield robust structural insights. Our durable goods categories include different items with probably markedly different depreciation rates, but are not grouped so as to yield a clear ordering in that respect: what we call "furniture" includes fast-depreciating items, such as household appliances and sundry articles, as well as beds, chairs, and tables which depreciate slowly. Some observable variability is available as regards depreciation for vehicles: province-level measures of accident frequency are available, which however have implications for the volatility of actual durable stocks as well as for their expected rate of change.

Lacking information about depreciation and relative price dynamics, we view such determinants of the user cost of durables as essentially constant (or absorbed by observable characteristics) across households. In order to test theoretical prediction regarding the role of drift and uncertainty in shaping durable adjustment, we focus on the information available to us regarding the heterogeneous steepness and volatility of the sample households' consumption profiles. Under our homotheticity and separability assumptions, the drift of the durable/non-durable ratio in the absence of adjustment is larger in absolute value for households whose consumption profile is steeper, reflecting differences in return and discount rates as well as (if the third derivative of the utility function is not zero) different intensity of uninsurable consumption fluctuations. The relevant volatility (of the $\{M(t)\}$ sequence, in our application) is not directly related to that of uninsurable income. Unless preferences display constant absolute risk aversion (and the consumption function is linear in the expected present value of income), as pointed out by Carroll (2001), consumers faced with more volatile income processes tend to obtain more self-insurance by accumulating larger buffer stocks of wealth, and may also otherwise alter their behaviour in ways that offset the impact of labour income uncertainty on their consumption. However, we can estimate individual-specific proxies of predictable consumption growth and of consumption risk on the observations (roughly one third of the sample) that feature non-durable consumption information for both 1995 and 1998, and impute the results prediction to all households in our whole 1995 sample.

Our empirical strategy, like that of Dynan (1993), starts from an Euler equation in the form of (6) and considers a second-order Taylor expansion of $E_t U'(M(t + \Delta t))$ around $U'(M(t))$ along the optimal consumption path of a consumer who may borrow and lend,

but is subject to uninsurable consumption risk. Let N_i be a vector of household characteristics and μ a vector of coefficients indexing their relevance to consumption profiles (*e.g.* through household-specific differences in the $r - \zeta$ difference between rates of return and of discount or preference heterogeneity). Denoting with Δt the time interval between available observations, the continuous-time expression (6) implies a relationship between the first and second moments of consumption growth:

$$E_t [\Delta m_i] \simeq N_i' \mu + \phi E_t [(\Delta m_i)^2], \quad (12)$$

where $\Delta m_i \equiv [M(t + \Delta t) - M(t)]/M(t)$ and $\phi \equiv -\frac{1}{2} M(\cdot) U'''(\cdot) / U''(\cdot)$ (half the coefficient of relative prudence as defined by Kimball (1990)) is constant for any CRRA specification of utility and, in particular, equals unity for the logarithmic utility function assumed in our theoretical derivations.

In our application, a (logarithmic) Euler equation is indeed satisfied by the purchasing power allocated to the consumption bundle, $M(t) = C(t) + vX(t)$. Since this includes the non-durable consumption flows and the user cost of current durable stocks, it is not directly observable in the absence of complete information about durable stocks and user costs. We approximate Δm_i by the non-durable consumption growth observed for the panel households.¹⁰ This is not exact in continuous time, because $M(t)$ and $C(t) = M(t) - vX(t)$ are not proportional to each other during periods of inaction and the $M(t)/C(t)$ jumps discontinuously upon adjustment events. The $M(t)/C(t)$ ratio, however, fluctuates around a constant in the long run: since adjustment of durable stocks is likely to take place during the two-year observation period of non-durable consumption, the latter's drift and volatility should be quite similar to those of the theoretically relevant $M(t)$ construct.

5.1. Income uncertainty as an instrument

To see how the parameters of the Euler equation (12) are identified, define the expectation error $\xi_i = \Delta m_i - E_t[\Delta m_i]$ of the first conditional moment of consumption growth and write

$$\Delta m_i \simeq N_i' \mu + \phi E_t [(\Delta m_i)^2] + \xi_i. \quad (13)$$

Consumption risk, $E_t[(\Delta m_i)^2]$, is not directly observed. What we observe is the realization $(\Delta m_i)^2$. Defining the expectation error $\zeta_i = (\Delta m_i)^2 - E_t[(\Delta m_i)^2]$ of the second conditional moment of consumption growth, we can rewrite (13) as

$$\Delta m_i \simeq N_i' \mu + \phi (\Delta m_i)^2 + v_i, \quad (14)$$

where $v_i = \xi_i - \phi \zeta_i$. Since $(\Delta m_i)^2$ is clearly correlated with the composite expectation error v_i , estimating (14) by OLS yields biased and inconsistent estimates. To assess the relevance of precautionary behaviour in a short CEX panel, Dynan (1993) addresses the endogeneity problem by instrumenting consumption growth variability with education, the number of earners, initial assets, and occupational and industry dummies. All these instruments, however, have fairly low power in her data. Our strategy exploits a theoretically more appealing and empirically more powerful instrument, namely the subjective measure of income uncertainty available in our data.

Subjective income risk is indeed the ideal instrument in this setting (Manski (2003) argues forcefully in favour of using subjective expectation data in the estimation of structural models of individual behaviour). Its orthogonality to the expectation errors defined above is

10. Since subjective expectation data refer to the 1995–1996 period, we convert 1995–1998 consumption growth to the 1995–1996 period assuming annual growth is approximately constant over the 1995–1998 period.

soundly justified in that income uncertainty, like all information available at the beginning of the observation period, should not affect consumption growth after controlling for the latter's conditional volatility, which is a sufficient statistic for the relevant risk in the absence of liquidity constraints.¹¹ Its explanatory power for consumption's conditional volatility may be weak if individuals isolate consumption from income shocks via wealth accumulation and decumulation, or via formal insurance. Subjective uncertainty measures condition on more information than that contained in observable household characteristics level, however, and their predictive power is unlikely to be as weak as that of Dynan's (1993) instruments.

Since both the mean and the variance of consumption growth are heterogeneous in the cross-sectional sample, as econometricians we need to relate observed consumption growth to more than one observable variable relevant to the individuals' beginning-of-period assessment of their own future consumption's distribution. Our data feature a rich set of suitable demographic variables. The procedure described below makes it possible to estimate both the first and second moment of individual consumption growth distributions as long as collinearity is sufficiently low between the projections of consumption growth rates and of their squares on individual characteristics used to predict individual-specific drift and risk variables (subjective income risk and demographics). In our data, the predictive power of the instruments turns out to be high, and collinearity does not prevent our regressions from disentangling first- and second-moment heterogeneity.¹²

5.2. Consumption uncertainty

In practice, we obtain an estimate of $E_t[(\Delta m_i)^2]$ as the predicted value of a first-stage regression in the form

$$(\Delta m_i)^2 = N'_i \eta_0 + \eta_1 \text{var}_t(\Delta y_i) + \zeta_i, \quad (15)$$

where $\text{var}_t(\Delta y_i)$ is the conditional subjective variance of income growth. This is calculated as the household-specific subjective variance of income *levels* reported in 1995 for 1996 (see Appendix B), divided by the square of 1995 income. The subjective income uncertainty may in principle, and does in practice, convey information as to the consumer's *ex ante* outlook for the whole period over which non-durable consumption growth is measured.

Estimates of the η_0 and η_1 parameters from the panel sub-sample make it possible to compute the predicted value $\hat{\sigma}^2 = N'_i \hat{\eta}_0 + \hat{\eta}_1 \text{var}_t(\Delta y_i)$, an estimate of consumption uncertainty, for *all* households with subjective income uncertainty data in the initial cross-section. Since income risk may be non-linearly related to consumption risk, and their relationship may depend on wealth, permanent income, or other observable characteristics, we extend (15) to a more flexible function of income risk: a cubic polynomial in $\text{var}_t(\Delta y_i)$, the same polynomial interacted with education, and the interaction between $\text{var}_t(\Delta y_i)$ and the wealth-income ratio at the beginning of the period. It is worth stressing that availability of subjective information relieves us from being specific about the exact form of the relationship between consumption risk and income risk, which is far from trivial outside simple and perhaps unrealistic cases.

In the first column of Table 3 we report the results of the first-stage (equation (15)) estimates. We control for the same set of demographics that appear in the adjustment size equations below. We also include indicators for liquidity constraints (described in Appendix C), expected income

11. Credit limitations would invalidate the exclusion restrictions, introducing a direct link between income volatility and the timing of non-durable consumption and durables expenditures: as we shall see, however, the data do not reject the null hypothesis that credit is unrestricted in the sample we analyse.

12. The information available to us regarding cross-sectional variation of prudent households' dynamic environment suffices to estimate their utility function's curvature even in the absence of time-series information of the type considered by Attanasio and Low's (2004) numerical study.

growth, marital status, dummies for employment and self-employment, and the (financial) wealth–income ratio at the beginning of the period. Finally, we include income uncertainty (up to a third-order polynomial) and the interactions with household characteristics listed above. These are the instruments for the endogenous squared term on the R.H.S. of the Euler equation above (equation (13)): with a partial R^2 of 0.14 (and a p -value of the F -test for their joint significance of less than 0.01%) there is little doubt about the identifying power of these instruments. The relationship between income and consumption risk appears to be significantly influenced by wealth and education: in particular, the negative sign of the wealth–risk interaction corroborates the theoretical implication that the association between income uncertainty and consumption volatility should become weaker as wealth increases and self-insurance improves. There is a broadly positive association between income uncertainty and consumption risk:¹³ at the mean values of income uncertainty, education, and wealth, consumption growth volatility is predicted to be 0.01 (on an annual basis); increasing income uncertainty by 10% (for the same values of the interacted variables) increases it to 0.011, and a 50% increase brings it to 0.0146.

The last rows of Table 2 report summary statistics for income risk and consumption volatility measures.¹⁴ Consumption volatility estimates are on average half as large as income volatility measures, reflecting consumption smoothing. Their pair-wise correlation is high (0.44), and significant (p -value below 0.01%). As discussed above, the implications of uncertainty for the unconditional probability of adjustment are theoretically ambiguous: it may be interesting to note that consumption uncertainty is lower on average in the sub-sample of those who adjust than in the whole sample.

5.3. Consumption drift

Instrumental variable estimation of the Euler equation (14)

$$\Delta m_i \simeq N_i' \mu + \phi (\Delta m_i)^2 + v_i$$

uses data on the households observed in the 1995–1998 period, and the resulting estimates allow us to construct predicted non-durable consumption growth for all households in our sample.¹⁵

Regression results are shown in the second column of Table 3. The Sargan test does not reveal evidence against our specification. The coefficient of consumption volatility, $\phi = 1.598$ with a standard error of 0.272, is largely significantly different from zero and statistically close (p -value = 2.8%) to the unitary value predicted by our theory's logarithmic specification. Expected income growth, the liquidity constraints indicator, and the beginning-of-period wealth–income ratio are not significant, as are most demographics (the employment dummy is marginally significant).

Interestingly, subjective income growth expectations are statistically significant in regressions (not reported) which exclude consumption risk from the specification. Traditionally, such “excess sensitivity” has been viewed as an indication of binding liquidity constraints. Finding that expected income growth becomes insignificant when we control for consumption risk in Table 3 is consistent with Carroll's (2001) view that excess sensitivity is more likely to

13. This becomes more transparent if we drop the interactions and the higher-order terms of income uncertainty (while keeping the N_i variables in the regression). In this case the subjective variance of income growth has a coefficient estimate of 0.04 with an S.E. of 0.01; the estimate of ϕ is 1.69 (S.E. 0.67); the results reported in Tables 4 and 5 are qualitatively similar although predictably less precisely measured.

14. Our first-stage equation does not impose any restrictions on the parameters, so negative predicted values are possible for the latter. This happens in less than 5% of the cases. We assume absence of consumption risk in this case. Results obtained dropping these households are very similar and are available on request.

15. Jappelli and Pistaferri (2001) also use subjective income expectations to estimate the parameters of the Euler equation. Their findings are similar to those reported here.

TABLE 3
Euler equation estimates

	First-stage results	Second-stage results
Consumption uncertainty		1.5978 (0.000)
Expected income growth	0.0049 (0.492)	0.0241 (0.206)
Education	-0.0011 (0.017)	0.0010 (0.403)
Age	-0.0001 (0.696)	-0.0010 (0.179)
Family size	0.0009 (0.729)	0.0036 (0.622)
Number of children aged 0-5	-0.0021 (0.625)	-0.0141 (0.234)
Number of children aged 6-13	-0.0020 (0.531)	0.0061 (0.497)
Number of children aged 14-17	-0.0023 (0.584)	-0.0102 (0.377)
Number of earners	-0.0004 (0.897)	-0.0138 (0.089)
South	0.0050 (0.469)	-0.0327 (0.089)
Centre	-0.0035 (0.551)	-0.0191 (0.246)
Self-employed	-0.0086 (0.326)	-0.0328 (0.151)
Employed	-0.0072 (0.405)	-0.0395 (0.069)
Married	-0.0168 (0.005)	0.0322 (0.063)
Wealth-income ratio	0.0014 (0.426)	0.0004 (0.819)
Liquidity constraint	-0.0043 (0.742)	0.0117 (0.749)
Conditional variance income growth	-0.3040 (0.035)	
Conditional variance income growth ²	0.3347 (0.384)	
Conditional variance income growth ³	-0.1027 (0.655)	
Conditional variance income growth × Educ.	0.0545 (0.001)	
Conditional variance income growth ² × Educ.	-0.0597 (0.175)	
Conditional variance income growth ³ × Educ.	0.0204 (0.459)	
Conditional variance income growth × wealth-income ratio	-0.0615 (0.000)	
Adj. R ² excluded instruments	0.1424	
F-test of excluded instrument	14.75 (df: 7 622; p-value: 0.0000)	
Sargan test		2.83 (df: 6; p-value: 0.8294)

In the first-stage regression the L.H.S. variable is the square of consumption growth for the panel households in the sample; in the second-stage regression it is the rate of growth of household consumption. Both stages also control for the frequency of car accidents, indices of dissatisfaction with the efficiency of public transports, local council offices and traffic congestion at the province level, and city size dummies. Asymptotic *p*-values are reported in parentheses.

reflect omission of consumption risk from the Euler equation. Lack of evidence of liquidity constraints makes it unnecessary in our sample to account for the potentially very complex interactions between financial market imperfections and transaction costs.

Our measure of the unobservable drift of the durable/non-durable log ratio in the absence of adjustment is given by the predicted value ($N_i' \hat{\mu} + \hat{\phi} \hat{\sigma}^2$). Households with high values of this term have steeper non-durable consumption growth and therefore larger drift (a more negative value of the theoretical parameter ϑ). The empirical analysis below accounts for the fact that both

the consumption drift and the consumption uncertainty measures are generated regressors. The average estimated consumption drift is about -1% (annualized) in the whole sample (Table 1), and a little larger in absolute value in the sub-sample of households that adjust durables stocks (Table 2).

6. TESTING THE MODEL'S PREDICTIONS

The theoretical predictions reviewed above broadly agree with some simple descriptive evidence from our data-set. For instance, the (taste-adjusted) empirical cross-sectional distribution of X/C (not reported for brevity) tends to be more spread out for the subgroups of households with high uncertainty (consumption risk above the 75-th percentile) for each type of durables. This simple evidence, however, does not conform to the *ceteris paribus* assumption of theoretical comparative-dynamics predictions. To assess the impact of each parameter in isolation, all other characteristics that may affect the distribution of X/C should be held constant, as we proceed to do in a formal controlled regression framework.

Very few households downgrade the stock of vehicles (a little more than 1% of the whole sample) or that of jewelry (0.2%), and as mentioned no information is available on the value of furniture sales. Hence, we focus on regressions for the probability of upgrading the existing stock of durable goods, and for the net size of adjustment (the value of purchases minus that of sales, if any). Separate specifications are run for vehicles, furniture, and jewelry, allowing for possible interactions among the three types of durable-adjustment decisions.

6.1. The probability of adjustment

We let adjustment of the current stock of durables occur when a latent variable D_i^* ,

$$D_i^* = H_i'\theta + u_i,$$

is driven to be larger than zero. The assumption that $u_i \sim N(0, 1)$ yields the probit model

$$\Pr(D_i^* > 0) = \Phi(H_i'\theta), \quad (16)$$

where $\Phi(H_i'\theta)$ is the standard normal cumulative density function evaluated at $H_i'\theta$. In our theoretical framework, such a latent variable is readily interpreted as the distance between the action point and the durable/non-durable ratio. In continuous time, adjustment need not take place at the beginning or the end of observation periods, and should occur with certainty whenever the deviation of X/C from the frictionless optimum equals (or exceeds) the boundaries of the optimal inaction band. Our data provide a measure of the stock at the time of adjustment, rather than at the beginning of the period as in other data-sets, but discrete-time observations allow only imperfect assessment of the consumer's position within the band. The (normally distributed) disturbance term in equation (16) accounts for measurement errors due to non-durable consumption measurement throughout the period, within-period durable-stock depreciation, possible repeated durable purchases within the period, and for the fact that some regressors (consumption uncertainty and drift) are estimated.

In the class of infrequent adjustment models we consider, observing a large value of X/C prior to adjustment makes subsequent upward adjustment less likely. 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 interpret as the history-dependent

component of the determinants of optimal adjustment decisions. 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, for example, demographics) is not affected by variance in the linear-quadratic approximation. With appropriate measures of drift and adjustment costs, one can test the additional implications that stronger drift (a higher depreciation rate or steeper non-durable consumption growth) increases the probability of adjustment, while larger adjustment costs make adjustment less frequent. The drift and the measure of uncertainty we use are obtained using the Euler equation procedure described in Section 5. We estimate a model for the probability that a household upgrades each durable stock, conditioning on observable characteristics. Our theoretical characterization of the relationship between non-durable and durable consumption dynamics remains valid as long as utility is log-linear across goods, with different Cobb–Douglas budget shares, as well as over time. If utility is not separable across types of durable goods, however, then adjustment of one durable stock changes the desirability of others' adjustment: to allow for this possibility, we include all three durable/non-durable ratios in each of the equations. As pointed out in Section 4.1, individuals who are relatively more inclined to consume durable goods should tend to remain closer to the frictionless X/C ratio, and therefore adjust more frequently and by smaller amounts. We proxy for taste heterogeneity by 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 (frequency of car accidents and dissatisfaction with traffic congestion in the province where the household lives), adjustment costs (dissatisfaction with the efficiency of 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 transport in the household province). The quality of public transport may have an impact on the frequency of adjustment since for individuals living in areas with highly inefficient public transport the benefits from more frequent adjustment are larger than for households that can rely, if no adjustment is undertaken, on high quality public transport. These variables are expected to be most relevant for the vehicles regressions.

Table 4 reports marginal effect estimates from probit regressions for the upgrading of the stock of vehicles, jewelry, and furniture, respectively. For inference purposes, the computation of standard errors and test statistics must take into account that we use a multiple-stage estimation strategy and that in the later stages we use regressors generated from the earlier ones. We use a block bootstrap procedure to compute p -values for the test of significance of the coefficients, accounting for serial correlation and heteroscedasticity of arbitrary form, for the presence of panel and non-panel households in our sample, as well as for the fact that we use generated regressors. We should point out that this procedure is conservative, and thus the precision of our parameters is likely to be underestimated.

In the first column, the probability of upgrading decreases with the initial value of the ratio of the stock of vehicles to non-durable consumption, as predicted by the theoretical model. The coefficient is highly statistically significant and economically important: a one standard deviation decline in the durable/non-durable ratio would increase the probability of adjusting by about 11 percentage points (60% of the unconditional probability). The stocks of the other two durable goods (scaled by non-durable consumption) are statistically insignificant. Hence, the data do not reject separability in preferences and adjustment costs across the three durable goods considered.

TABLE 4
Probit for the upgrading of durables stocks

	Vehicles	Furniture	Jewelry
<i>X/C</i> , vehicles	-0.2808 (0.0000)	0.0087 (0.8782)	0.0064 (0.7864)
<i>X/C</i> , furniture	0.0181 (0.3393)	-0.0961 (0.0040)	-0.0023 (0.8902)
<i>X/C</i> , jewelry	0.0094 (0.5868)	0.0723 (0.1277)	0.0362 (0.1317)
Consumption uncertainty	-0.0325 (0.0080)	-0.0148 (0.2076)	-0.0150 (0.0758)
Drift	0.9194 (0.0559)	0.1188 (0.5429)	0.0446 (0.7585)
Adjustment costs	-0.0749 (0.0080)	-0.0511 (0.1277)	-0.0178 (0.3114)
Education	0.0030 (0.3593)	0.0124 (0.0000)	0.0073 (0.0000)
Age	-0.0031 (0.0319)	-0.0066 (0.0040)	-0.0014 (0.1397)
Family size	0.0117 (0.5589)	0.0040 (0.8224)	-0.0216 (0.0200)
Number of children aged 0–5	-0.0352 (0.1477)	-0.0430 (0.1238)	0.0152 (0.3553)
Number of children aged 6–13	-0.0225 (0.1158)	-0.0140 (0.4990)	0.0171 (0.2156)
Number of children aged 14–17	0.0305 (0.1357)	0.0014 (0.9900)	-0.0085 (0.6387)
Number of earners	0.0425 (0.0399)	0.0269 (0.2236)	0.0319 (0.0240)
Small town	0.0789 (0.0798)	-0.0523 (0.3353)	0.0312 (0.4152)
Medium town	0.0501 (0.1317)	-0.0020 (0.9541)	0.0273 (0.4551)
Large town	0.0570 (0.0719)	0.0004 (0.9102)	0.0382 (0.2555)
Car accidents	-0.0040 (0.8822)	-0.0007 (0.9820)	0.0012 (0.7305)
Public transport	0.0385 (0.0279)	0.0678 (0.0000)	-0.0023 (0.8144)
Traffic congestion	0.0070 (0.5110)	0.0033 (0.09182)	-0.0012 (0.8902)
South	-0.0238 (0.9581)	-0.1308 (0.0080)	-0.0124 (0.7385)
Centre	-0.0150 (0.7705)	-0.0373 (0.3034)	-0.0098 (0.6427)

Our measures of consumption uncertainty, drift, and adjustment costs are described in the text. The variable “Car accidents” is the frequency of car accidents per 1000 cars at province level. “Public transport”, “Local council offices”, and “Traffic congestion” are 1–10 indices of dissatisfaction with the efficiency of public transport, efficiency of local council offices, and the extent of traffic congestion at the province level, respectively (10 corresponds to the highest level of dissatisfaction). The table reports marginal effects and, in parentheses, bootstrap *p*-values based on 500 replications. For ease of legibility, the variance of consumption growth (consumption uncertainty) has been multiplied by 100.

As predicted by the theory, a higher level of uncertainty reduces the probability of adjusting (a bootstrap *p*-value of 0.8%). The effect is quite substantial. For instance, we calculate that doubling consumption volatility would halve the probability of adjustment.

Our main control for the drift (expected household non-durable consumption growth) is borderline significant (*p*-value 0.056) and its sign agrees with theoretical predictions: households with faster-growing non-durable consumption (a more negative value for ϑ) adjust upward more frequently than those with a flatter profile. Nevertheless, the economic effect is quite small. If the drift is twice the sample average, the probability of adjusting the stock of vehicles increases by

less than 1%. 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 the expectations of changes in the durable stock have a theoretically consistent effect.

The index of inefficiency of the public administration has a negative, statistically significant impact on the probability of adjustment, consistently with theoretical predictions if it does proxy for adjustment costs. The economic effect is also substantial. A one standard deviation increase in the efficiency of the local public administration increases the probability of adjustment by about 8.5 percentage points (about 48% of the sample mean). Bringing all provinces to the level of efficiency of the most efficient province would increase the probability of adjustment by almost 20 percentage points.

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 parameter is estimated with poor precision). 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 increases not only the drift but also uncertainty about the X/C ratio, with counteracting effects on the probability of adjustment.

Overall, the estimated effects on the probability of adjustment of the main variables that theory predicts should affect the adjustment decision (the initial stock, the value of uncertainty, drift, and adjustment cost) lend considerable support to the model in the case of vehicles.

Some of the demographic variables also appear relevant to the likelihood of adjusting, which declines significantly with age, and is higher for families with multiple earners and not living in a metropolitan area. Theory suggests that adjustment should be more likely for individuals whose preferences attach a larger weight to durables, and most of the effects are at least superficially consistent with this interpretation. For instance, cars are likely to be more important for large households than for the elderly, leading the former to adjust more frequently than the latter.

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. Uncertainty reduces the likelihood of upgrading, though standard errors are higher than in the case of vehicles (p -values of 21% and 8%, respectively, for furniture and jewelry). The drift term is correctly signed but not significant. The index of inefficiency of the local public administration is insignificant, and not surprisingly so since this variable is a poor adjustment-cost proxy for durable goods other than vehicles. The own durables stock/non-durable consumption ratio is statistically significant and of the right sign in the furniture equation, but insignificant in the jewelry equation. This indicates that our theoretical and empirical approach cannot capture crucial features of jewelry consumption, which is likely to depend on unobservable taste heterogeneity in ways that are poorly controlled by demographic variables. For example, households with a relatively high jewelry/non-durable consumption desired ratio may also derive lower disutility from visiting jeweller shops, and our result could then be explained by the upward bias of initial stock coefficient estimates in a regression that fails to control for taste heterogeneity that increases both the likelihood of adjustment and the initial jewelry stock.

6.2. *The size of adjustment*

Theory also delivers sign predictions and exclusion restrictions for the size of the adjustment conditional on adjusting (or, equivalently, for the width of the inaction band). Formally,

$$E(\ln S_i \mid Q_i, D_i^* > 0) = Q_i' \pi + E(\varepsilon_i \mid Q_i, D_i^* > 0), \quad (17)$$

where S_i is individual i 's optimal adjustment size, Q_i a vector of explanatory variables, and $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$ a Gaussian error term capturing measurement error and unobserved heterogeneity. Theory predicts that higher uncertainty, adjustment costs, and drift all increase the size of adjustment.

In the equation for the size of the adjustment, however, the disturbance depends upon unobserved heterogeneity as well as measurement error. Hence, $E(\varepsilon_i \mid Q_i, D_i^* > 0)$ is not zero in general, and simple OLS regressions conducted on the sample of those who adjust will provide inconsistent estimates of the parameters of interest. This is a standard problem in microeconometrics, which we treat in what follows with Heckman (1979) selectivity corrections in regressions for the size of adjustment. This approach is most suitable in our setting. First, since unobserved heterogeneity in the two margins is likely to exist along a variety of dimensions (taste for durables, transaction costs, etc.) that affect both the likelihood of observation in the neighbourhood of trigger points and the width of the adjustment bands, the relevant self-selection mechanism implies that unobservable heterogeneity in the extensive margin is correlated with unobserved heterogeneity in the intensive margin (*i.e.* $\text{cov}(u_i, \varepsilon_i) = \rho \neq 0$), thus precluding use of a Cragg (1971) model (see Lee and Maddala, 1985). Second, theory suggests that the decision to adjust does not depend on the same variables affecting the decision about how much to adjust. Hence, a simple Tobit model would not be appropriate, and suitable identifying restrictions are available for a Heckman correction procedure.

As mentioned, theory predicts that the value of X/C prior to adjustment affects the likelihood of adjusting, but not the size of the adjustment if it occurs. This is an important, theory-based exclusion restriction, and the evidence presented above suggests that it is powerful enough for identification purposes. The results of the second stage of the Heckman selectivity regressions are reported in Table 5. In the first column we see that income uncertainty increases the size of vehicles adjustment, as predicted by the theoretical model. The effect is statistically significant (a p -value of 4.4%) and economically sizable: a 10% increase in consumption uncertainty increases the size of vehicle adjustment by 5.5% conditioning on adjustment. The effect of the drift is poorly measured and thus no inference can be made from the fact that the point estimate has the wrong sign. Our proxy for adjustment costs has the expected positive sign and the point estimate (though significant only at the 16% level) implies a strong effect: raising adjustment costs by one standard deviation increases the size of vehicles' adjustment by about 5% conditional on adjusting. Among demographic variables, only age and the number of teenage family members are significant (young people buy smaller cars, and the size of adjustment is also smaller for families with teenage members). There is evidence of self-selection: the Wald test for independent equations has a p -value below 1%, and strongly supports our specification's allowance for unobserved taste factors.

For furniture, where adjustment is measured on the basis of purchases, results are less in line with theoretical predictions. More uncertainty actually decreases the size of the adjustment (the effect is borderline but still significant at conventional levels). The drift has the expected sign and is strongly significant. People living in large cities and households with young children adjust by smaller amounts, while households living in the South and in provinces with good public transports and less traffic congestion make larger purchases. We again find evidence for selection and negative correlation between unobservables in the adjustment equation and the selection equation.

For jewelry, results are broadly in agreement with the theory, but poorly measured. The only variables to have a statistically reliable effect on the size of adjustment of the stock of jewelry are age and residence in the South (the elderly and those living in the South buy larger amounts of precious objects). Finally, we find again that the estimate of ρ is negative and statistically significantly different from zero. In summary, regression evidence shows that the model seems

TABLE 5
Heckman selectivity model

	Vehicles	Furniture	Jewelry
Consumption uncertainty	0.2440 (0.0439)	-0.1658 (0.0479)	0.0724 (0.3433)
Drift	-3.1898 (0.4870)	8.8917 (0.0080)	6.5912 (0.2475)
Adjustment costs	0.3780 (0.1637)	-0.0455 (0.4870)	-0.4814 (0.1876)
Education	-0.0172 (0.8224)	-0.0038 (0.6826)	-0.0298 (0.4750)
Age	0.0322 (0.0599)	-0.0047 (0.8822)	0.0505 (0.0120)
Family size	-0.0879 (0.7026)	-0.0405 (0.4112)	0.0452 (0.6467)
Number of children aged 0-5	0.2757 (0.2834)	-0.0216 (0.7345)	0.2523 (0.3832)
Number of children aged 6-13	0.0382 (0.8064)	-0.2442 (0.0679)	-0.0940 (0.5988)
Number of children aged 14-17	-0.5644 (0.0200)	0.0746 (0.3633)	0.1319 (0.5749)
Number of earners	-0.1723 (0.3313)	0.1132 (0.1637)	-0.0579 (0.7066)
Small town	0.1243 (0.9022)	0.5712 (0.0479)	-0.4232 (0.5150)
Medium town	0.3396 (0.4990)	0.2109 (0.2036)	-0.3003 (0.6028)
Large town	0.2176 (0.6986)	0.2924 (0.1238)	-0.3420 (0.5190)
Car accidents	0.0940 (0.2635)	-0.0522 (0.4112)	0.0283 (0.7066)
Public transport	-0.2060 (0.3313)	-0.3209 (0.0798)	0.2268 (0.2036)
Traffic congestion	-0.0255 (0.7944)	0.2366 (0.0279)	0.1021 (0.5269)
South	-0.1068 (0.5908)	0.5737 (0.0719)	1.0740 (0.0599)
Centre	0.2340 (0.5110)	0.3320 (0.0838)	0.4991 (0.1717)
ρ	-0.9745 (0.0000)	-0.6188 (0.0000)	-0.9529 (0.0000)

For each durable considered, the L.H.S. variable is the log of the value of purchases net of sales. For furniture, only purchases are considered. Our measures of consumption uncertainty, drift, and adjustment costs are described in the text. The variable "Car accidents" is the frequency of car accidents per 1000 cars at province level. "Public transport", "Local council offices", and "Traffic congestion" are 1-10 indices of dissatisfaction with the efficiency of public transport, efficiency of local council offices, and the extent of traffic congestion at the province level, respectively (10 corresponds to the highest level of dissatisfaction). Bootstrap *p*-values based on 500 replications are reported in parentheses. For ease of legibility, the variance of consumption growth (consumption uncertainty) has been multiplied by 100.

to be appropriate for modelling the behaviour of car adjustment. The evidence is more mixed for furniture, and the results for jewelry conform quite poorly to theoretical predictions. In shaping the size and timing of jewelry purchases, portfolio considerations may well be more important than utility flows. Thus, finding that our consumption-based theory fails to explain them may actually foster confidence in our theoretical framework's real-life relevance. Like a dog that does not bark, negative evidence for jewelry indicates that the theory's good fit in cases it is designed to model (like that of cars) does not mechanically reflect spurious features of the data.

It is interesting, from our theoretical perspective, to find that variables that have positive coefficients in the selection equation also tend to have negative coefficients in the adjustment-size equation, and *vice versa*. This is quite consistent with the infrequent-adjustment

perspective of our theoretical approach: the likelihood of adjustment (conditional on the initial durables/non-durable consumption ratio) is lower for wider inaction bands, which in turn imply larger adjustment sizes. Hence, sources of heterogeneity that make action less likely should indeed also imply larger adjustment upon action. From the viewpoint of the dynamic model outlined in Section 3, it is particularly intriguing and quite interesting to find that not only the effects of observable variables, but also those of unobservable heterogeneity, are negatively related in selection and outcome controlled regressions estimated on cross-sectional data. 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 for each type of durable good considered, indicating that 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. Consider individuals with similar target levels but different unobserved tastes for vehicles. Individuals whose utility attaches more weight to vehicles face steeper flow utility losses when deviating from the frictionless optimal stock, and hence have narrower inaction bands (see Figure 1) and adjust more frequently and by smaller amounts than individuals with weaker tastes for vehicles. Thus, unobserved heterogeneity in the intensive margin equation should indeed be negatively correlated with unobserved heterogeneity in the extensive margin equation. In this respect the evidence in support of the infrequent adjustment model is consistent across the three types of durables.

6.3. Aggregate effects

Much recent theoretical and empirical work has focused on aggregate implications of infrequent adjustment at the microeconomic level. In the dynamic aggregation studies proposed by Bertola and Caballero (1990), Caballero (1993), and Eberly (1994), a crucial role is played by the aggregate component of individual uncertainty, whose empirical relevance is analysed by Attanasio (2000) and Hassler (2001) but cannot be analysed in the context of our essentially cross-sectional empirical work. Since at any given point in time aggregate *per capita* durable purchases are given by the product of the fraction of households adjusting and the size of adjustment for those adjusting, however, our separate estimation of the effects of uncertainty on the two margins is of considerable interest, and our estimates offer useful information regarding the role of uncertainty (and of other variables) in shaping 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: since households that do not adjust make no purchase,

$$E(S_i) = E(S_i | D_i^* > 0) \Pr(D_i^* > 0),$$

where $E(S_i | D_i^* > 0)$ is average expenditure in the sub-sample of those who adjust and $\Pr(D_i^* > 0)$ the probability of adjusting. The marginal effect of a given variable q_{ij} (affecting both the size of adjustment and the probability of adjustment) on average household expenditure is then:

$$\frac{\partial E(S_i)}{\partial q_{ij}} = \underbrace{\frac{\partial E(S_i | D_i^* > 0)}{\partial q_{ij}} \Pr(D_i^* > 0)}_{\text{intensive margin}} + \underbrace{\frac{\partial \Pr(D_i^* > 0)}{\partial q_{ij}} E(S_i | D_i^* > 0)}_{\text{extensive margin}}. \quad (18)$$

The marginal effect is given by the sum of two components. The first is the change in buyers' purchases induced by a change in q_{ij} (the change in the intensive margin). The second is the change in purchases of those who change adjustment decision (from no-adjustment to adjustment

TABLE 6
Unconditional marginal effects and elasticities

	Marginal effect	Intensive margin	Extensive margin	Elasticity
Consumption uncertainty	-28.14	-6.01	-22.13	-0.095
Drift	11.05	4.79	6.27	0.001
Adjustment costs	-0.81	-0.29	-0.51	-0.528
Education	0.03	0.01	0.02	0.046
Age	-0.02	0.00	-0.02	-0.124
Family size	0.10	0.02	0.08	0.052
Number of children aged 0-5	-0.30	-0.06	-0.24	-0.013
Number of children aged 6-13	-0.30	-0.15	-0.15	-0.019
Number of children aged 14-17	-0.01	-0.22	0.21	-0.000
Number of earners	0.49	0.20	0.29	0.130
Small town	1.17	0.68	0.49	0.044
Medium town	0.98	0.66	0.32	0.033
Large town	1.00	0.63	0.38	0.060
Car accidents	0.02	0.05	-0.03	0.030
Public transport	0.40	0.14	0.26	0.272
Traffic congestion	0.07	0.03	0.05	0.063
South	-0.45	-0.28	-0.16	-0.025
Centre	-0.04	0.07	-0.10	-0.001

Both the marginal effects $\partial E(\cdot)/\partial q_{ij}$ and the elasticities $(\partial E(\cdot)/\partial q_{ij}) \times (q_{ij}/E(\cdot))$ are evaluated at the mean of the variables and refer to car purchases. Marginal effects (and their breakdown in intensive and extensive effects) are in 1000 euro. Our measures of consumption uncertainty, drift, and adjustment costs are described in the text. The variable “Car accidents” is the frequency of car accidents per 1000 cars at province level. “Public transport”, “Local council offices”, and “Traffic congestion” are 1–10 indices of dissatisfaction with the efficiency of public transport, efficiency of local council offices, and the extent of traffic congestion at the province level, respectively (10 corresponds to the highest level of dissatisfaction).

or *vice versa*) because of a shift in q_{ij} (the change in the extensive margin).¹⁶ Using (16) and (17), it is easy to recover the expression for (18). We calculate the marginal effect (the L.H.S. of (18)), its breakdown in intensive and extensive margin effects, and the corresponding elasticity $\frac{\partial E(S_i)}{\partial q_{ij}} \frac{q_{ij}}{E(S_i)}$ (evaluated at the sample mean of the variables). The results are reported in Table 6 for the case of vehicles.¹⁷

Consider the effect of consumption uncertainty. The extensive margin effect is negative because, from Table 4, an increase in uncertainty reduces the probability of adjusting the stock of vehicles, for given $E(S_i | D_i^* > 0)$. The sign of the effect on the intensive margin depends on two influences: for given probability of adjustment $\Pr(D_i^* > 0)$, there is the direct effect on the size of adjustment which, from Table 5, is positive; but there is also an indirect offsetting effect, namely the negative change in the mean of the heterogeneity term conditioning on adjusting. Which of the two effects dominates determines the sign of the intensive margin.

We find that the aggregate effect of uncertainty on vehicle purchases is substantial, because its impact on the frequency and size of adjustment have the same (negative) signs.

16. Recall that we estimate a model for the log of the size of the adjustment, $\ln S_i = Q_i' \pi + \varepsilon_i$.

Noting that $E(S_i | D_i^* > 0) = e^{Q_i' \pi} E(e^{\varepsilon_i} | D_i^* > 0) = e^{Q_i' \pi} e^{E(\varepsilon_i | D_i^* > 0) + \frac{\text{var}(\varepsilon_i | D_i^* > 0)}{2}}$, formulae for the truncated normal can be used to compute the moments of interest. In particular, $E(S_i | D_i^* > 0) = e^{Q_i' \pi} e^{\rho \sigma_\varepsilon \lambda(H_i' \theta) + \frac{\rho^2 \sigma_\varepsilon^2}{2} \left\{ 1 - \lambda(H_i' \theta) [\lambda(H_i' \theta) + H_i' \theta] \right\} + \frac{(1 - \rho^2) \sigma_\varepsilon^2}{2}}$, where $\lambda(H_i' \theta) = \frac{\phi(H_i' \theta)}{\Phi(H_i' \theta)}$. It is also straightforward, if

tedious, to compute $\frac{\partial E(S_i | D_i^* > 0)}{\partial q_{ij}}$ from the expression above.

17. Results for the other two durables categories, not reported, are available on request. Since the estimated marginal effect and elasticity we report lack standard errors, they do not have a rigorous statistical interpretation but provide useful information on the aggregate effects of interest.

The net effect is negative: according to the elasticity reported in the table, a 10% increase in consumption uncertainty decreases household car expenditure by about 1% on average. Estimation-based measures of aggregate effects of marginal uncertainty changes arguably offer a better indication than Hassler's (2001) numerical experiments, based on large, infrequent switches in uncertainty. Further work may fruitfully take an explicitly dynamic perspective on such phenomena, recognizing in particular that rational optimal-adjustment policies would not in general remain invariant in the face of dynamic (rather than comparative-static) uncertainty variation.

Adjustment costs also have an overall depressing effect on aggregate expenditure. A 10% increase in inefficiency (*i.e.* higher adjustment costs) brings about a decrease in average household expenditure of approximately 5%. While the administrative inconvenience of new car purchases may be a small component of the relevant overall adjustment cost, it is the only one observable in our data, and the results indicate that decreasing adjustment costs (for example, through a more liquid market for second-hand cars) would have large effects on steady-state adjustment intensity. The drift in durable expenditure has an overall positive effect, but its impact is very modest: a 10% increase in the drift increases average household expenditure by less than 0.1%.

Finally, more inefficient alternative public transport increases expenditure substantially: a 10% decline in inefficiency induces a decline in average household expenditure of approximately 3%.

7. CONCLUDING REMARKS

We have outlined and tested a set of theoretical predictions concerning optimal infrequent adjustment of durable goods stocks. In our theoretical framework, homotheticity supports a dynamic interpretation of cross-sectional information, and unitary elasticity of intertemporal substitution makes it possible to characterize adjustment-related issues separately from other aspects of consumers' intertemporal problems. Our empirical strategy and results hinge on the use of indicators of subjective labour income variability to control for key aspects of consumers' dynamic environment. In particular, we use subjective income risk as an instrumental variable in the estimation of utility curvature parameters: this is methodologically innovative and empirically powerful in our data, where a single pair of panel periods yields results that conform nicely to theoretical relationships between income variability and consumption risk. Theory offers a clean set of exclusion restrictions and sign predictions to regression estimates of the probability and size of adjustment. Theoretical predictions as to the implications of different degrees of uncertainty across consumers for durable goods replacement patterns are borne out by the data in the case of vehicles, while furniture and jewelry purchase patterns are harder to interpret from the standpoint of the preference specification that supports our empirical approach.

Our data and empirical approach did not allow us to characterize the effects of adjustment costs proportional to transaction size which, unlike per-adjustment lump sum costs, can in principle allow the width of the inaction band and the size of transactions to vary independently from each other across individuals. The significant negative correlation between the (observable and unobservable) determinants of the likelihood and size of adjustment, however, suggest that lump-sum adjustment costs are the predominant source of "optimal inaction" in our data-set. In principle, the data make it possible to study durable goods other than 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. In practice, the empirical results are uniformly favourable for cars, while empirical support is more mixed for furniture and (especially) jewelry. This is far from surprising, because our theoretical framework's

auxiliary assumptions (such as homotheticity of demand) and our empirical approach (based on demographic controls for taste heterogeneity) are less likely to be suitable for the analysis of jewelry. Hence, lack of uniform support may indicate that the results are not driven by spurious mechanisms, and strengthens our confidence in the explanatory power of the theory when applied to suitable microeconomic problems and data.

APPENDIX A. 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 since the mid-1960s. Over time, it has gone through a number of changes in sample size and design, sampling methodology and questionnaire. However, the sampling methodology, sample size and the broad contents of the information collected are unchanged since 1989. Recent waves of the SHIW have been conducted in 1989, 1991, 1993, 1995, 1998, and 2000. Each wave surveys a representative sample of the Italian resident population and covers about 8000 households, but some portions 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 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% 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.

APPENDIX B. DEFINITIONS OF THE VARIABLES

All demographic variables refer to the household head.

Non-durable consumption: the sum of the expenditure on food, entertainment, education, clothing, medical expenses, housing repairs and additions, and imputed rents.

Net disposable income: the sum of wages and salaries, self-employment income, and income from financial and real assets, less income taxes and social security contributions. Wages and salaries include overtime bonuses, fringe benefits and payments in kind and exclude withholding taxes. Self-employment income is net of taxes and includes income from unincorporated businesses, net of depreciation of physical assets.

Net financial assets: these are imputed from the flow of financial income (interest on checking accounts, saving accounts, money market accounts, certificates of deposit, stock, government and other bonds plus dividends, less interest on household liabilities).

Durables flows: 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 equipment, etc.); "Precious objects" (including jewelry, old and gold coins, works of arts, antiques and antique furniture).

Durable stock: for each of the three categories, the survey reports the end-of-period stock of durables as well as purchases and sales during the period. The value prior to adjustment is computed subtracting purchases and adding sales to the end of period stock.

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

Quality of life indicators: The 1993 SHIW asked each household head to report (on a 0–10 scale) their satisfaction with the quality of various public and private services in the province or neighbourhood of residence (there are 105 provinces in Italy, usually centred around medium- to large-sized cities). We construct the indices reported in the text by taking a provincial average of assessments, ranging from 1 (best) to 10 (worst), of 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 cleanliness and noise pollution.

City size: is coded as follows: 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).

Earnings uncertainty: we use the variance of expected earnings at the individual level. This is computed directly from survey questions asking the employed and job seekers to report: (a) on a scale from 0 to 100, their chances of keeping their job or finding one in the next 12 months; (b) the minimum and the maximum income 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. Assuming a triangular distribution for the probability distribution of earnings and imputing a value for unemployment compensation to each individual in the sample using current legislation, Guiso *et al.* (2002) use this information to recover measures of expected earnings and their dispersion.

Liquidity constraints: the SHIW first asks if any member of the household has applied for a loan or mortgage to a bank or other financial intermediary in the previous calendar year. For those who answer yes, it asks whether the application was accepted, partially rejected, or turned down. For those who answer no, it asks whether any member of the household had considered applying for a loan or mortgage, but changed their mind on the expectation that it would be turned down. We define as liquidity constrained those who had their application turned down or partially rejected, and the discouraged borrowers.

APPENDIX C. INFREQUENT ADJUSTMENT

C.1. Optimality conditions

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

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

with solution

$$V(z_t) = -\frac{b}{2} \left(\frac{z_t^2}{\lambda} + \frac{\sigma^2 + 2z_t\vartheta}{\lambda^2} + \frac{\vartheta^2}{\lambda^3} \right) + K_1 e^{\alpha_1 z_t} + K_2 e^{\alpha_2 z_t}, \quad (\text{C.1})$$

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

The constants of integration in $V(\cdot)$ and (L, s, U) must be such as to imply that $V'(x)$ equals the marginal cost of action whenever action is in fact undertaken ("smooth pasting"), and that the value function at the trigger and return points differ by the total cost A of adjusting between the two points ("value matching"):

$$\begin{aligned} V(s) - V(L) &= V(s) - V(U) = A, \\ V'(L) &= V'(s) = V'(U) = 0. \end{aligned}$$

Inserting the functional form (C.1) in these conditions forms a system of equation to be solved for the constants of integration and the action and return points.

C.2. Stable distribution

Our empirical exercise analyses a set of cross-sectional observations, each of which may be viewed 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 behaviour, the cross-sectional information available can be interpreted in terms of the long-run distribution of the controlled variable, z , within the $[L, U]$ optimal inaction interval.

The Kolmogorov equation for the steady-state density reads

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

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

$$f(z) = \begin{cases} \tilde{k}z + \bar{k} & \text{if } \vartheta = 0, \\ \tilde{k}e^{\frac{2\vartheta}{\sigma^2}z} + \bar{k} & \text{otherwise.} \end{cases}$$

The constants of integration \tilde{k} and \bar{k} in each of the state-space segments are determined by continuity of the stable density at the trigger and return points, and by the adding-up constraint $\int_L^U f(z) dz = 1$.

In the long run, 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 s . The same derivations that lead to the stable density—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

$$\frac{\sigma^2}{2} f'_{(+)}(L) = \frac{\sigma^2}{2} \frac{d}{dz} \left(\tilde{k} e^{\frac{2\theta}{\sigma^2} z} + \bar{k} \right)_{z=L} = \tilde{k} \theta e^{\frac{2\theta}{\sigma^2} L},$$

where \tilde{k} 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 neighbourhood of the trigger point, $f'_{(-)}(U)$, and of the intensity of Brownian movements that may push the process towards that point, $\sigma^2/2$.

Acknowledgements. For helpful comments and discussions we thank two anonymous referees, Joao Ejarque, Fumio Hayashi, Ed Vytlačil, Jonathan Willis, and seminar and conference participants at European University Institute, NBER Summer Institute, Tilburg University, Università di Bergamo, CEU Budapest, Università di Pisa, Bureau of Labour Statistics. Mirko Wiederholt and Winfried Koeniger provided competent research assistance and very useful suggestions. This research was conducted in the framework of the “Finance and Consumption in the EU Chair”, sponsored by Findomestic Banca and CETELEM at the EUI Department of Economics. Luigi Guiso also acknowledges financial support from MURST and the European Commission; Luigi Pistaferri acknowledges financial support from the “Taube Faculty Research Fund” at SIEPR.

REFERENCES

- ADDA, J. and COOPER, R. (2000), “Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies”, *Journal of Political Economy*, **108** (4), 778–806.
- ATTANASIO, O. (2000), “Consumer Durables and Inertial Behaviour: Estimation and Aggregation of (S, s) Rules for Automobile Purchases”, *Review of Economic Studies*, **67**, 667–696.
- ATTANASIO, O. and LOW, H. (2004), “Estimating Euler Equations”, *Review of Economic Dynamics* (forthcoming).
- BANKS, J., BLUNDELL, R. and BRUGIAVINI, A. (2001), “Risk Pooling, Precautionary Saving and Consumption Growth”, *Review of Economic Studies*, **68**, 757–779.
- BAR-ILAN, A. and BLINDER, A. (1992), “Consumer Durables: Evidence on the Optimality of Usually Doing Nothing”, *Journal of Money, Credit, and Banking*, **24**, 258–272.
- BERTOLA, G. and CABALLERO, R. (1990), “Kinked Adjustment Costs and Aggregate Dynamics”, *NBER Macroeconomics Annual*, 237–288.
- BRANDOLINI, A. and CANNARI, L. (1994), “Methodological Appendix: The Bank of Italy Survey of Household Income and Wealth”, in A. Ando, L. Guiso and I. Visco (eds.) *Saving and the Accumulation of Wealth: Essays on Italian Households and Government Behavior* (Cambridge: Cambridge University Press).
- CABALLERO, R. (1993), “Durable Goods: An Explanation for their Slow Adjustment”, *Journal of Political Economy*, **101** (2), 351–384.
- CARROLL, C. D. (2001), “Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)”, *Advances in Macroeconomics*, **1** (1), no page numbering.
- CRAGG, J. (1971), “Some Statistical Models for Limited Dependent Variables with Application to the Demand for Durable Goods”, *Econometrica*, **39**, 829–844.
- DYNAN, K. E. (1993), “How Prudent are Consumers?”, *Journal of Political Economy*, **101**, 1104–1113.
- EBERLY, J. (1994), “Adjustment of Consumers’ Durables Stocks: Evidence from Automobile Purchases”, *Journal of Political Economy*, **102**, 403–436.
- FOOTE, C., HURST, E. and LEAHY, J. V. (2000), “Testing the (S, s) Model”, *American Economic Review Papers and Proceedings*, **90**, 116–119.
- GROSSMAN, S. and LAROQUE, G. (1990), “Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods”, *Econometrica*, **58**, 25–51.
- GUISO, L., JAPPELLI, T. and PISTAFERRI, L. (2002), “An Empirical Analysis of Earnings and Employment Risk”, *Journal of Business and Economic Statistics*, **20**, 241–253.
- HASSLER, J. (2001), “Uncertainty and the Timing of Automobile Purchases”, *Scandinavian Journal of Economics*, **103** (2).
- HECKMAN, J. (1979), “Sample Selection Bias as a Specification Error”, *Econometrica*, **47**, 153–162.
- JAPPELLI, T. and PISTAFERRI, L. (2001), “Using Subjective Income Expectations to Test the Excess Sensitivity of Consumption to Predicted Income Changes”, *European Economic Review*, **44**, 337–358.
- KIMBALL, M. S. (1990), “Precautionary Saving in the Small and in the Large”, *Econometrica*, **58**, 53–73.
- LAM, P.-S. (1991), “Permanent Income, Liquidity, and Adjustment of Automobile Stocks: Evidence from Panel Data”, *Quarterly Journal of Economics*, **106**, 203–230.
- LEE, L.-F. and MADDALA, G. S. (1985), “The Common Structure of Tests for Selectivity Bias, Serial Correlation, Heteroscedasticity and Non-normality in the Tobit Model”, *International Economic Review*, **26**, 1–20.
- MANSKI, C. F. (2003), “Inference on Expectations and Decisions” (Unpublished Manuscript, Northwestern University).
- MARTIN, R. (2003), “Consumption, Durable Goods, and Transaction Costs” (FRBG International Finance Discussion Paper #756).
- MEGHIR, C. and WEBER, G. (1996), “Intertemporal Non-separability or Borrowing Restrictions? A Disaggregate Analysis Using a U.S. Consumption Panel”, *Econometrica*, **64**, 1151–1182.
- PADULA, M. (2000), “Durable Goods and Intertemporal Choices: A Survey”, *Giornale degli Economisti ed Annali di Economia*, **60**, 245–269.