NBER WORKING PAPER SERIES

THE DYNAMICS OF CAR SALES: A DISCRETE CHOICE APPROACH

Jérôme Adda
Russell Cooper

Working Paper 7785
http://www.nber.org/papers/w7785

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2000

We are grateful to Jon Willis for research assistance, and to seminar participants at Boston University, Boston College, Haiffa University, Tel Aviv, Hebrew University, SED 1999, NBER Summer Meetings 1998 and CEPR Macro Conference Hydra 2000. Cooper is grateful to the NSF for financial support. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

The Dynamics of Car Sales: A Discrete Choice Approach
Jérôme Adda and Russell Cooper
NBER Working Paper No. 7785
July 2000
JEL No. D12, D91, E21

ABSTRACT

Mankiw [1982] explores the Permanent Income Hypothesis implication that durable expenditures follow an ARMA(1,1) representation. He finds that durable expenditures are represented by an AR(1) process which implies that the rate of depreciation of durables, under the PIH model, is 100%. This finding presents a puzzle. Our paper builds on earlier work which attempts to explain this puzzle by considering the aggregation of the discrete dynamic choices of heterogeneous households. We implement this approach by estimating a dynamic discrete choice model of car replacement. We find that through aggregation we can explain both the AR and MA components of Mankiw's results. Further we find that our model is able to match a VAR representation of car sales, prices and income. We find that most of the variation in car sales is due to shocks which influence the replacement probability.

Jérôme Adda
Department of Economics
University College London
Gower Street
London WC1E 6BT
England

Russell Cooper
Department of Economics
270 BSR, Boston University
Boston, MA 02215
and NBER
rcooper@bu.edu

The Dynamics of Car Sales: A Discrete Choice Approach

Jérôme Adda and Russell Cooper

1 Introduction

This paper focuses on understanding the behavior of durable consumption expenditures. As is well known, spending on durables is an important component of aggregate spending and one that fluctuates considerably over the business cycle.

From the aggregate perspective, Mankiw [1982] presents evidence that the permanent income hypothesis (PIH) model of durable expenditures is inconsistent with observed data. In particular, he argues that in a single agent choice problem in which utility is a quadratic function of the stock of durables, the optimal choice of the agent implies that expenditures on durables will follow an ARMA(1,1) process, where the MA component is parameterized by the rate of durable goods depreciation. Mankiw estimates an ARMA(1,1) time series representation of quarterly durable goods expenditures for the US. In contrast to the predictions of the theory, he finds that durable goods expenditures follow an AR(1) process. Put differently, Mankiw estimates the rate of depreciation of durable goods to be 100%. We call this finding the "Mankiw puzzle".

From the household perspective, Lam [1991] reports that households only occassionally adjust their stock of durables. Consistent with this finding, Bar-Ilan and Blinder [1988,1992], Bertola and Caballero [1990] and Caballero [1990,1993] view aggregate observations on durable purchases as the outcome of the aggregation over heterogeneous microeconomic agents. Taken together, these papers certainly suggest that a model of heterogeneity and discrete adjustment can qualitatively match relevant parts of the data.

However, there is no characterization of the time series properties of durable purchases offered in these papers and thus the "Mankiw puzzle" remains open.[1] So, the goal of this paper is to study the determinants of the time series representation of durable expenditures in an explicit dynamic, discrete choice framework: can a dynamic discrete choice representation of household durable purchases produce the observed time series behavior of durable expenditures?

We address this question by looking specifically at two distinct features

[1]For example, the final section of Caballero [1993], entitled "ARMA Representation and Impulse Responses" displays impulse response functions for Cars and Furniture and states that "The shapes are broadly consistent with the description given in the paper." Whether or not the estimated model can produce an ARMA representation close to that reported by Mankiw is not specifically addressed.

In a related, independent study, Attanasio [1997] estimates (S,s) rules for automobile purchases using microeconomic data. After estimating the model, he undertakes an evaluation of the aggregate time series implications of the model, as we do in this paper. He finds that if there is more persistence in the shocks to the target relative to the persistence in the shocks to the band, then the model is able to match observed aggregate behavior.

of spending on an important component of consumer durables, aggregate car sales. First, we confirm the ARMA(1,1) representations that underlie the "Mankiw puzzle" for our various measures of automobile sales. Second, we estimate and study a VAR representation of automobile sales, prices and income. Here we find that the impulse response function displays dampenned oscillations in response to an innovation in income. So, besides confronting the Mankiw puzzle for car sales, we ask whether an aggregated discrete choice model can match and explain this rich time series response to an income shock.

This paper builds upon the framework of Adda and Cooper [2000] who investigate the effects of scrapping subsidies on car purchases. An important difference between this paper and the existing literature is that the empirical implications are drawn directly from the dynamic optimization problem without imposing any structure directly on agents' decision rules. In particular, while our model of durable replacement is of the optimal stopping variety, we do not specify *(S,s)* bands directly nor do we find it necessary to specify a "desired" stock of durables in our estimation. We do this for two reasons. First, we find that empirically the PIH assumptions which underlies this "desired stock" approach are not supported by the data.[2] Second, deriving the optimal durable expenditure policy from a dynamic optimization framework and then using this same structure for estimation is more consistent theoretically.[3]

We find that the aggregate model based upon the dynamic optimization of heterogeneous microeconomic units can "explain" both the AR and MA parts of Mankiw's regression results. Further, a comparison of the impulse response functions generated by the models with that obtained through an unrestricted VAR reveals why a PIH model has difficulty matching the data. Suppose that there is an income shock. In the data, the initial burst of sales is followed by a reduction in sales and then dampened oscillations (relative to the initial level). It is precisely these endogenous fluctuations in sales

[2]We argue that the finding, reported in Caballero [1990], that an ARMA(1,q) representation of durable expenditures reconciles the evidence and the PIH model is not robust across samples and the choice of q.

[3]These two points are, of course, related. Since we do not have convincing evidence that the PIH prediction holds even in the "long-run", linking the estimation to a target seems unwarranted. Our approach to estimation through a characterization of the complete household dynamic optimization overcomes this problem as we do not require the specification of a target.

that the estimated PIH model misses. It is captured in our model by the interaction of a state dependent hazard function and the evolution of the cross sectional distribution of car vintages.

We then use our structure to uncover the sources of these dynamics. In general, the dynamic discrete choice structure generates variations in aggregate sales from two sources: fluctuations in the replacement probability and the evolution of the cross sectional distribution of car vintages. We find that most of the variation in car sales is due to shocks which influence the probability of replacement. Put differently, the endogenous evolution of the cross sectional distribution contributes surprisingly little to the time series variation of car sales.[4]

2 Evidence on Aggregate Durables Dynamics

This section presents evidence on the behavior of aggregate durables. We extend the ARMA(1,1) representation stressed by Mankiw in three ways. First our sample period is longer. Second we study both the U.S. and France. Third, we focus on both total durables and cars. In addition, we estimate a VAR for both the U.S. and France and thus characterize the joint behavior of prices, income and car sales. These two pieces of evidence provide empirical motivation for our work and are then used in assessing our model.

2.1 ARMA(1,1) Representation

The starting point of the analysis is the representative agent model in which durables expenditures is a continuous choice variable. This model provided the basis for the initial empirical literature. Following Hall [1978], Mankiw [1982] extended the permanent income hypothesis model to account for durability. In this model, the agent only consumes a durable good and faces an uncertain income. If the utility function is quadratic, then expenditures on durables e_t by the representative household follow:

$$e_{t+1} = \delta\alpha_0 + \alpha_1 e_t + \varepsilon_{t+1} - (1-\delta)\varepsilon_t. \tag{1}$$

[4]This is similar to the findings reported in Cooper, Haltiwanger and Power [1999] which studied the implications of machine replacement for aggregate investment. This finding is not inconsistent with the emphasis placed on the movement of the cross sectional distribution in Adda and Cooper [2000] since in that exercise there was a policy action that had a significant effect on the cross sectional distribution of vintages.

where δ is the depreciation rate of the stock of durables and ε_t is the innovation to income. Using aggregate quarterly U.S. data on durables, Mankiw [1982] shows that the series are better described by an AR(1) process than an ARMA(1,1).

Working with annual series for France and the US, we report very similar results in Table 1. The rows pertain to both aggregated durable expenditures and estimates based on cars. For the latter, we have data on both total expenditures on cars (for France) and new car registrations. The columns refer to estimates with and without the removal of a linear trend.

Table 1: ARMA(1,1) Estimates on US and French Data

Specification	No trend		Linear trend	
	α_1	δ	α_1	δ
US durable expenditures	1.00(.03)	1.5 (.15)	0.76 (0.12)	1.42 (0.17)
US car registration	0.36(.29)	1.34 (.30)	0.33 (0.30)	1.35(0.31)
France durable expenditures	0.98 (0.04)	1.20 (0.2)	0.56 (0.24)	1.2 (0.36)
France car expenditures	0.97(0.06)	1.3 (0.2)	0.49 (0.28)	1.20 (0.32)
France car registrations	0.85 (0.13)	1.00 (0.26)	0.41 (0.4)	1.20 (0.41)

Notes: Annual data. For the US, source FRED database, 1959:1-1997:3. French data: source INSEE, 1970:1-1997:2. US registration: 1968-1995.

For both countries, the estimated rate of depreciation is quite high. Clearly the hypothesis that the rate of depreciation is close to 100% per year would not be rejected for most of the specifications. Further, the results are robust to the detrending method (we have also tried exponential trends and obtained very similar results). For France, the data exhibit a trend so we get different values for the AR coefficient.[5] Thus Mankiw's "puzzle" seems to be robust across categories of durables, countries, time periods and the method of detrending. That is, under the null hypothesis of the PIH model, the estimated rate of depreciation is quite high.

[5]Note that contrary to the argument in Bar-Ilan and Blinder [1988,1992], sales and expenditures have very similar time series properties.This may reflect the fact that variations in expenditures are largely a consequence of the extensive margin (to buy a car or not) rather than the intensive margin (how much to spend on a given car). Leahy and Zeira [1999] present a model in which income variations lead solely to changes in decisions on the extensive margin.

2.2 Unrestricted VAR

A second and more general way of representing durable expenditures is through a VAR. While focusing on ARMA representations, much of the literature on durable expenditures has ignored the joint dynamics of durables, income and prices over time. Here we present results using a VAR composed of automobile sales, automobile prices and income. While this representation has no structural interpretation, it provides a better characterization of the dynamics of sales. We estimate VAR models for sales for both the US and France. The optimal number of lags were chosen with an Akaike criterion. The optimal lag lengths are one for the US and three for France. All series were detrended by filtering logged variables with a Hodrick-Prescott filter. For the US, the sample period is 1960-1994, and for France, 1968-96. From the estimated VAR, we compute the impulse response functions of sales, prices and income to a shock to income. The variables have been ordered in the following way: income, prices and sales.

The impulse responses for new car registrations, prices and income, both for France and the US are reported in Figure 1. The first two graphs displays the response of sales to a shock on income. In both countries, after an initial increase, the sales are characterized by dampened oscillation around the baseline. These oscillations could arise for two reasons. First, as emphasized in the literature on non-convex adjustment costs with heterogenous agents, the endogenous evolution of the stock of cars can potentially produce replacement cycles and thus oscillations in sales. A second explanation is that income and prices are serially correlated and have some cross dynamics. Indeed, the next four graphs of Figure 1 show that prices and income also oscillate around baseline. We return to an evaluation of the relative importance of these two sources of dynamics later.

2.3 Summary of Evidence

We use the facts reported thus far in two ways. First, in the estimation of our model, we use the ARMA(1,1) representation as one means of characterizing the aggregate behavior of durables. Second, we use the VAR in an informal overidentification test of our model: can our estimated model produce impulse response functions that are similar to those from an estimated VAR?

There are other attempts to analyze the time series of durables. Bernanke

[1985] examines a model which includes nondurable goods, convex adjustment costs for the stock of durables and price variations. For his specification, Bernanke assumes that utility is quadratic in both nondurables and durables. Bernanke argues that in the presence of adjustment costs, the stock of durables will follow an AR(2) process.[6] This implies that durable expenditures are given by an ARMA(2,1) process. As reported in the Data Appendix, we again find that the rate of depreciation is generally quite high.

A final representation of durable expenditures follows from Caballero [1990], who extends Mankiw's analysis by considering a model with additional MA components. The idea is that perhaps the PIH is a valid representation of the data in the long-run. In the short-run, because agents are characterized by inertia, the response of sales to income innovations is slow. By allowing additional moving average terms the PIH might be reconciled with the data. Inference about the underlying rate of depreciation from the MA coefficients is made possible by a simple model of sluggish adjustment that Caballero [1990] postulates. This approach then forms the basis of Caballero [1993] where the idea that agents eventually abide by the PIH prediction is central to the formulation of the nonlinear target adjustment structure.

Our attempts at fitting an ARMA(1,q) model to our series are reported in the Data Appendix. In constrast to ARMA(1,1) representation, the results are quite sensitive to the nature of detrending, the choice of series, the frequency of observation, etc. Owing to the lack of robustness and precision in the ARMA(1,q) representations, we cannot use these results in the estimation or evaluation of our model.

3 Dynamic Discrete Choice Model

The representative agent models discussed in the previous section must deal with the aggregation of durable goods of different vintages. For these specifications, the aggregation is achieved by the creation of a durable aggregate stock that is diminished by depreciation and increased by expenditures. One possible extension of this representative agent structure would be to view expenditures on goods in different periods simply as different goods. In this

[6] His AR(2) representation, equation 4.2, is a quasi-differenced version of the optimal decision rule.

way, the entire spectrum of different vintages would be reflected in the representative agent's utility function.

However, this specification goes only part way to fully modeling the implications of heterogeneity. The model we propose in this section goes a step farther as it allows for heterogeneity across the decision making units due to taste and income shocks and also recognizes that households may have durables of different vintages.

The next section discusses the theoretical specification of the discrete choice model. We then present our estimation results for this model. The final sections return to the results reported for the representative agent model to discuss the extent to which the dynamic discrete choice model can capture the aggregate behavior of durable expenditures.

3.1 Theory

3.1.1 Household Behavior

The starting point for this analysis is the dynamic programming problem explored in Adda and Cooper [2000]. Consider an agent with a car of age i in state z=(p,Y,y,ε). As above, p is the price of the durable good. Current income is given by the sum $Y+y$ where Y represents aggregate income and y represents idiosyncratic shocks to nondurable consumption that could reflect variations in household income, tastes or expenditures on car maintenance and other necessities.[7] Finally, every household experiences a taste shock which is represented by ε. At every point in time, the household decides whether to retain a car of age i or scrap it. If the household decides to scrap the car, then it receives the scrap value of π and has the option to purchase a new car. If the household retains the car, then it receives the flow of services from that car and cannot, by assumption, purchase another car. Thus the household is constrained to own at most a single car.

For tractability, we place a number of restrictions on this household optimization problem. First, we do not explore the operation of a second-hand market: cars are either kept or scrapped.[8] Second, the household is forced

[7] Adda and Cooper [2000] explicitly views this as a household specific income shock but a broader interpretation is acceptable, particularly in light of the iid assumption associated with this source of variation.

[8] Adda-Cooper [2000] compute a no trade equilibrium for the case in which households are homogenous. Computing an equilibrium with a second hand market as part of the

to finance the durable purchase from current income: there is no borrowing or lending in the model.[9]

Formally, let $V_i(z)$ represent the value of having a car of age i to a household in state z. Further, let $V_i^k(z)$ and $V_i^r(z)$ represent the values from keeping and scrapping an age i car in state z. Then,

$$V_i(z) = \max[V_i^k(z), V_i^r(z)]$$

where

$$\begin{aligned} V_i^k(z) &= u(s_i, y+Y, \varepsilon) + \beta(1-\delta)EV_{i+1}(z') + \\ &\quad \beta\delta\{EV_1(z') - u(s_1, y'+Y', \varepsilon') + u(s_1, y'+Y'-p'+\pi, \varepsilon')\} \end{aligned} \tag{2}$$

and

$$\begin{aligned} V_i^r(z) &= u(s_1, y+Y-p+\pi, \varepsilon) + \beta(1-\delta)EV_2(z') + \\ &\quad \beta\delta\{EV_1(z') - u(s_1, y'+Y', \varepsilon') + u(s_1, y'+Y'-p'+\pi, \varepsilon')\}. \end{aligned}$$

In the definition of $V_i^k(z)$, the car is assumed to be destroyed (from accidents and breakdowns) with probability δ leading the agent to purchase a new car in the next period. The cost of a new car in numeraire terms is $p'-\pi$, which is stochastic since the price of a new car in the next period is random. Further, since we assume that there is no borrowing and lending, the utility cost of the new car is given by $u(s_1, y'+Y', \varepsilon') - u(s_1, y'+Y'-p'+\pi, \varepsilon')$ which exceeds $p'-\pi$ as long as $u(\cdot)$ is strictly concave in nondurable consumption. It is precisely at this point that our borrowing restriction appears as an additional transactions cost.

estimation of a discrete choice non-representative household model is beyond the scope of this paper as this would require the characterization of an equilibrium as part of our estimation. However, given that our focus is on the decision to scrap a car and purchase a new one, the basis of new car sales, perhaps our omission of resale opportunities is not too unfortunate.

[9]This then implies that the cost of buying a durable good cannot be spread over time, thus implicitly increasing the cost of such expenditures. To control for these effects, we also estimate a model in which utility is assumed to be a linear function of nondurable consumption. See Campillo [2000] for simulation results of a model with durables and borrowing restrictions. We omit savings in our approach for two reasons. First, solving the model with savings is very demanding computationally as it requires an additional continuous state variables. Second, we would need data on the joint distribution of car vintages and savings, which are difficult or even impossible to get for all years and countries under study.

For the application we define the utility function to be additively separable between durables and nondurables:

$$u(s_i, c) = \left[i^{-\gamma} + \frac{\varepsilon (c/\lambda)^{1-\xi}}{1-\xi} \right]$$

where c is the consumption of non-durable goods, γ is the curvature for the service flow of car ownership, ξ the curvature for consumption and λ is a scale factor. In this specification, the taste shock (ε) influences the contemporaneous marginal rate of substitution between car services and nondurables.

In order for the agent's optimization problem to be solved, a stochastic process for income, prices and the aggregate taste shocks must be specified. We assume that aggregate income, prices and the unobserved preference shock follow a VAR(1) process given by:[10]

$$\begin{aligned} Y_t &= \mu_Y + \rho_{YY} Y_{t-1} + \rho_{Yp} p_{t-1} + u_{Yt} \\ p_t &= \mu_p + \rho_{pY} Y_{t-1} + \rho_{pp} p_{t-1} + u_{pt} \\ \varepsilon_t &= \mu_\varepsilon + \rho_{\varepsilon Y} Y_{t-1} + \rho_{\varepsilon p} p_{t-1} + u_{\varepsilon t} \end{aligned}$$

The covariance matrix of the innovations $u = \{u_{Yt}, u_{pt}, u_{\varepsilon t}\}$ is

$$\Omega = \begin{bmatrix} \omega_Y & \omega_{Yp} & 0 \\ \omega_{pY} & \omega_p & 0 \\ 0 & 0 & \omega_\varepsilon \end{bmatrix}$$

As the aggregate taste shock is unobserved, we impose a block diagonal structure on the VAR, which enables us to identify all the parameters involving prices and aggregate income in a simple first step regression. This considerably reduces the number of parameters to be estimated in the structural model. We allow prices and income to depend on lagged income and lagged prices.

The aggregate taste shock potentially depends on lagged prices and income. The coefficients of this process along with ω_ε are estimated within the structural model. By allowing a positive correlation between the aggregate taste shock and lagged prices, given that prices are serially correlated, we can reconcile the model with the fact that sales and prices are positively correlated in the data. This allows us to better capture some additional dynamics of sales and prices in the structural estimation. An alternative way would

[10]Here we have only a single lag to economize on the state space of the agents' problem.

be to model jointly the producer and consumer side of the economy, to get an upward slopping supply curve. However, solving for the equilibrium is computationally very demanding.

The policy functions that are generated from this optimization problem are of an optimal stopping variety. That is, given the state of the household, the car is scrapped and replaced iff the car is older than a critical age. Letting $H_k(z_t;\theta)$ represent the probability that a car of age k is scrapped, the policy functions imply that $H_k(z_t;\theta)=\delta$ if $k < J(z_t;\theta)$ and $H_k(z_t;\theta) = 1$ otherwise. Here $J(z_t;\theta)$ is the optimal scrapping age in state z_t when θ is the vector of parameters describing the economic environment.

The remaining part of the model is firm behavior. As in Adda and Cooper [2000], we assume that the costs of production are independent of the level of production. Combined with an assumption of constant mark-ups, this implies that the product price is independent of the cross sectional distribution of car vintages.

This assumption of an exogenous price process greatly simplifies the empirical implementation of the model since we do not have to solve an equilibrium problem. In fact, we have found that adding information on the moments of the cross sectional distribution of car vintages has no explanatory power in forecasting car prices in the French case. Results are mixed for the US case, as the average age of cars significantly predicts future prices.

Before proceeding further, note that the underlying model stresses the replacement of older cars with new ones. Actual data presumably includes a component of car sales to agents who do not scrap a car before buying a new one. Further, there are surely instances where an agent scraps a car but does not buy another. Clearly movement of this type on the "extensive margin" creates a variation in sales not included in our model. To deal with this issue, we have detrended the data to remove the effects of population growth on sales. Further, from our investigation of some additional panel data on French households, we find that less than 2% of the sales are by households which have no cars.

3.1.2 Aggregate Sales

Aggregating over all possible types of individual leads to a prediction of the aggregate demand for cars and a prediction of the cross section distribution of car vintages. Letting $f_t(k)$ the period t cross sectional distribution of k, aggregate sales are given by

$$S_t = \sum_k H_k(z_t; \theta) f_t(k) \tag{3}$$

where θ is a vector of parameters. From an initial condition on the cross sectional distribution, it is possible to generate a time series for the cross sectional distribution given a particular parameterization of the hazard function. The evolution of $f_t(k)$ is given by:

$$f_{t+1}(k) = [1 - H_k(z_t; \theta)] f_t(k-1) \text{ for } k>1 \tag{4}$$

and

$$f_{t+1}(1) = S_t$$

Thus for a given θ we can simulate both sales and the cross sectional distribution.

3.2 Estimation

The parameters describing the joint process of aggregate income and prices are estimated in a first step, to reduce the number of parameters in the structural estimation. The estimation results are displayed in appendix A.2. After several trials, we imposed $\rho_{\varepsilon Y} = 0$ as the results were not sensitive to this parameter. The remaining parameters, $\theta = \{\gamma, \delta, \lambda, \xi, \sigma_y, \rho_{\varepsilon p}, \omega_\varepsilon\}$, are estimated from the policy functions generated by the solution of the households' optimization problem.

A natural estimation strategy is to find the parameters that bring the simulated model as close as possible to the data. In our estimation, we make use of different types of observations. First, we use time series observations on sales, prices and income to match the sales predicted by our model. Second, we use moments of the cross sectional distribution.[11] Third, we capture the dynamics of sales using the autoregressive and moving average coefficients from an ARMA(1,1) model to match up with the one generated on predicted data.[12] These additional moments are used to help the identification

[11] For the U.S. data on the cross sectional distribution comes from issues of Ward's Automotive. For France, this information comes from the CCFA.

[12] Here we focus on the ARMA (1,1) rather than the VAR representation for two reasons. First, the ARMA(1,1) representation has been the focus of the literature since Mankiw [1982]. Second, we can then investigate how well our estimated model matches the VAR representation as part of an informal evaluation of our model.

of the parameters characterizing the dynamic programming problem. The price, income and sales series were linearly detrended prior to the structural estimation since the model itself has no trends.

Given a vector of parameters θ and a realized path of preference shocks ε, the model predicts aggregate sales and the evolution of the cross sectional distribution. These simulated series can be compared to their empirical counterparts. The estimation strategy is to find θ to minimize the "distance" between the actual and simulated data.

Formally, the criterion we minimized, via the simplex algorithm, was a weighted average of the difference between actual and predicted sales and between actual and simulated moments characterizing the cross sectional distribution of car ages and ARMA coefficients, weighted by their actual variance. Concentrating on sales exclusively does not identify the parameters of the model such as the depreciation rate. Sales are the results of the interaction of a hazard function and the cross section distribution, both to be estimated. If the shape of the cross section distribution is not pinned down, there are several sets of parameters that would produce the same level of sales. Therefore, we also match three moments characterizing the cross sectional distribution and the ARMA coefficients described above.

So, the parameters were estimated by minimizing the overall criterion:

$$\mathcal{L}(\theta) = \alpha\mathcal{L}^1(\theta) + \mathcal{L}^2(\theta)$$

where α is a scale parameters defined to be equal to the inverse of the variance of sales. The minimization of the overall criterion yields a root-T consistent estimate for any fixed number of simulations. We discuss the two components of this objective function in turn.

The first component of the criterion is:

$$\mathcal{L}^1(\theta) = \frac{1}{T}\sum_{t=1}^{T}\left[(S_t - \bar{S}_t(\theta))^2 - \frac{1}{N(N-1)}\sum_{n=1}^{N}(S_{tn}(\theta) - \bar{S}_t(\theta))^2\right].$$

Thus $\mathcal{L}^1(\theta)$ contains two terms. The first one is the standard nonlinear least square criterion which measures the squared distance between observed and average predicted values of the variables. For the sales series, the estimation is conditional on the realization of aggregated income and prices at each date, as well as on the initial cross sectional distribution. Given these realizations and the initial condition, for each value of the parameters, we can see how close our simulated sales is to observed sales.

Specifically, Let S_t be the observed aggregate sales for the year t. Let $S_{tn}(\theta) = S(Y_t, p_t, \varepsilon_{t,n}, \theta)$ be the predicted sales for year t and for the draw n of the unobserved aggregate shock, $(n = \{1, \dots, N\})$. The first component of the objective is essentially the squared distance between S_t and an average measure (over the taste shocks) of $\bar{S}_t(\theta) = \sum_n S_{tn}(\theta)/N$. However, such a criterion produces an inconsistent estimator for a fixed number of simulation.

To overcome this problem, we follow Laffont, Ossard and Vuong [1995] by including the second term which is a second order correction for the inconsistency bias introduced by the random draws of the preference shocks. Under standard regularity conditions, the asymptotic distribution of the estimators is normal and root-T consistent, for any fixed number of simulations N, (see Laffont, Ossard and Vuong [1995]). In practice, we fix the number of random draws to 50. We find that the correction for simulation error is then negligible.

The second piece of the objective contained additional moments and is specified as: :

$$\mathcal{L}^2(\theta) = \sum_{i=\{5,10,15,AR,MA\}} \alpha_i(\bar{F}^i - \bar{F}^i(\theta))^2$$

where α_i is a weight equal to the empirical inverse of the variance of each moment.[13] The moments of the cross sectional distribution we match are the average (over the 1981-95 sample for the US, over the 1972-92 sample period for France) fraction of cars of ages 5, 10 and 15. The idea is to use these critical ages to characterize the average cross sectional distribution. The ARMA(1,1) model was run on observed annual sales for both countries. The predicted counterparts were obtained on simulated data from the model, for a similar sample size.

Formally, let $\bar{F}^i = (1/T)\sum_t F_t^i, i = 5, 10, 15$ be the average fraction of cars of age 5, 10 or 15 during the sample period. Let $\bar{F}^{AR}$ and $\bar{F}^{MA}$ be the AR and MA coefficients in an ARMA(1,1) regression of sales of new cars. Similarly let $F_{t,n}^i(\theta), i = 5, 10, 15$ be the predicted fraction of cars of age 5, 10 or 15, in period t and given draw n of the unobserved taste shock. Let $\bar{F}^i = \sum_{T,n} F_{t,n}^i(\theta)/(TN), i = 5, 10, 15$ be the average predicted fraction of cars of age 5, 10 or 15. Let $\bar{F}^{AR}(\theta)$ and $\bar{F}^{MA}(\theta)$ be the average AR and MA coefficients from an ARMA(1,1) regression on simulated sales.

[13]Note that the use of an ARMA representation as an auxiliary model can be seen as an indirect inference method as in Gourieroux et al [1993].

3.3 Estimation Results

Table 2 provides a summary of our estimated parameters, for both countries and for two different specifications of the utility function. The linear case refers to a model where the utility is linear (i.e. $\xi = 0$). The linear model provides a framework where income fluctuations do not matter, but also removes the effect of idiosyncratic shocks to income. However, idiosyncratic shocks are needed to provide a smooth upward sloping hazard function. To keep idiosyncratic shocks in the model, we specify a taste shock which affect the contemporaneous marginal rate of substitution between car services and nondurables. This shock is individual specific, with mean one and variance σ_y.

From this table, note that the rate of physical depreciation of cars is about 6 percent for France and 4 percent for the US. [14] This might reflect the greater number of car accidents in France. The rate of depreciation of the service flow is lower, about 2 percent on an annual basis for both countries. Further, we find that there is some curvature in the utility function for both countries, with a parameter around 1.5, which is in the usual range for this kind of utility functions. The variance of the idiosyncratic shock seems to be higher in France than in the US. The distribution of the idiosyncratic shock is important for the slope of the aggregate hazard function, as we aggregate out the individual component. A high variance means that the shape of the hazard function is flatter. This is needed for France as, in the data, the stock of cars is older. [15]

All models are able to match rather closely the aggregate sales as the R^2 vary from 0.52 to 0.69. By comparison, the R^2 obtained from an OLS regression of sales on lag sales, prices and income is 0.46 for France and 0.6 for the US. Our model also predicts rather well the cross sectional distribution of car vintages, as shown in Table 3. Both the linear models and the nonlinear ones do a good job of predicting the fraction of cars of age 5 and 10, but underpredict the fraction of age 15 cars. Anticipating the results of next section, the models do a relatively good job as well in matching the AR and MA parameters for aggregate sales for both countries.

[14] The results for France are close to those reported in Adda and Cooper [2000].

[15] The standard errors were computed using a bootstrapping method. Given a value of ω_ε, we simulated 100 series and reestimated the model on the simulated data. This gives us a distribution for each parameter. Given that the fit of the model is relatively good, the standard errors of the parameters are small.

However, when testing the over-identifying restrictions of the model, we reject them at the 5% level. At the 10% level, the restrictions for the nonlinear model for France are accepted. For the US, the test is rejected mainly because the predicted fraction of cars of age 15 is too low. Without this moment, the restrictions for the nonlinear model would be valid. The failure to accept the over identifying restrictions, which is a common feature of structural models, also reflects the inability of the model to match the observed (positive) correlations between car prices and sales without an upward sloping supply curve, as discussed in Adda and Cooper [2000].

In terms of these chi-square tests, the nonlinear models seem to better match the data than the linear ones. Consequently, we concentrate on the nonlinear specifications from now on. Given the no borrowing assumption, the nonlinear model provides a framework with liquidity constraints where purchase decisions are sensitive to income shocks. Given that aggregate sales appear to be correlated with income, the nonlinear models are preferred by the data. On micro data, Eberly [1994] and Attanasio et al. [2000] also provide empirical evidence of borrowing constraints using data for the US.

4 Time Series of Car Expenditures

Our motivation in this section is to compare the aggregate time series implications of the model estimated in section 3 with observed car sales in France and the US. More specifically, we compare the ARMA representations of the actual sales series with those created by simulating our estimated discrete choice model. We find that our estimated model is able to reproduce the observed time series behavior of car expenditures.[16]

The second part of this section provides an interpretation of these findings by looking at the impulse response functions produced by the model and investigating a decomposition of car expenditures into two components: variations in the cross sectional distribution and shocks to the hazard function. Here we find that most of the variations in car sales reflects shifts in the hazard function (representating time series variations in the probability of car replacement) rather than movements in the cross sectional distribution.

[16]Note that while the ARMA coefficients were used in the estimation procedure, the model was overidentified. Thus it is by no means "automatic" that we reproduce the ARMA representation. Removing the ARMA coefficients from the objective function did not change the results.

Table 2: Estimated Parameters for Discrete Choice Model

	U.S.			
	Non Linear Case		Linear Case	
Parameters	Estimates	S.E.	Estimates	S.E.
γ	0.02	0.0015	0.02	0.002
δ	0.038	0.0025	0.03	0.003
ζ	1.49	0.0037	0	-
λ	483.5	10.5	1.6e-5	1.4e-06
ω_ε	0.085	0.015	0.03	0.097
σ_y	0.52	0.016	0.26	0.013
$\rho_{\varepsilon,p}$	1.14e-4	4.2e-05	-3.9e-5	1.9e-05
Pseudo-R^2:	0.69		0.52	
P(Overidentification test)	0.84		0.60	
	FRANCE			
	Non Linear Case		Linear Case	
Parameters	Estimates	S.E.	Estimates	S.E.
γ	0.017	0.003	0.028	0.014
δ	0.061	0.016	0.06	0.03
ζ	1.46	0.007	0	-
λ	108.0	12.5	6e-6	2.8e-06
ω_ε	0.215	0.09	0.06	0.037
σ_y	0.99	0.003	0.44	0.22
$\rho_{\varepsilon,p}$	1.2e-4	4.3e-05	-2.9e-5	1.5e-05
Pseudo-R^2:	0.61		0.52	
P(Overidentification test)	0.91		0.67	

Note: Estimates obtained on annual data, France: 1972:1994, US: 1968:1995. The Pseudo R^2 measure the fraction of the variance of aggregate sales explained by the model. Probability values are displayed for the over identification tests.

4.1 ARMA Representation

Our interest is in the aggregate time series of new car sales produced by our estimated model. ARMA(1,1) representations of the sales series are reported in Table 4. We use 100 series simulated over 21 periods (years) for France and 28 years for US, so that the results are comparable with the observed data. We estimate ARMA representations for new car sales. The table reports the

Table 3: Observed and Predicted Moments from CDF

	US		
	Observed	Non Linear	Linear
Fraction of cars of age 5	8.6	8.5	8.4
	(1.1)	(1.4)	(1.3)
Fraction of cars of age 10	6.8	6.0	5.9
	(1.0)	(1.1)	(1.1)
Fraction of cars of age 15	2.7	1.7	1.8
	(0.6)	(0.3)	(0.3)
	FRANCE		
	Observed	Non Linear	Linear
Fraction of cars of age 5	8.9	9.0	9.7
	(0.8)	(0.8)	(0.6)
Fraction of cars of age 10	5.0	4.7	4.2
	(0.7)	(0.5)	(0.3)
Fraction of cars of age 15	1.2	1.4	0.9
	(0.3)	(0.2)	(0.1)

averages and the standard deviations of the coefficients.

Our principle finding is that our model is able to reproduce the abnormally high value for the "rate of depreciation" inferred from the MA(1) coefficient, when the estimation is done through an ARMA(1,1) representation as in Mankiw [1982]. To this extent, our model is able to reconcile a low depreciation rate at the micro level (7%) with a coefficient close to one at the macro level, as viewed through the PIH model. The autoregressive coefficient is also estimated quite close to its value in the annual data once the time series is detrended.

Table 4: ARMA Coefficient Observed and Predicted.

US	Observed	0.33	-0.35
		(0.30)	(0.30)
US	Predicted	0.23	-0.30
		(0.36)	(0.38)
		AR_1	MA_1
France	Observed	0.51	-0.16
		(0.27)	(0.31)
France	Predicted	0.06	-0.05
		(0.63)	(0.72)

Note: Monte Carlo results obtained over 100 replications with sample length 21 (France), 28 (US). French estimates on annual car registration, 1972-1994. US estimates on annual car registrations 1968-1995. All ARMA models included a linear trend.

4.2 VAR Representations

In this subsection, we return to the VAR representations reported in Section 2.2. Our purpose is first to explore how well our model fits the impulse response functions estimated for our data and then to use these impulse response functions to explain our results.

4.2.1 Matching Impulse Response Functions

A final check of our model is its ability to reproduce the impulse response functions presented in section 2.2. Given the estimated models for the two countries, we simulate data on prices, income and compute the predicted sales. We simulate 100 series of length 30 years to be comparable to the observed data. For each simulated series, we compute an impulse response by estimating an unrestricted VAR and these are then averaged. Figure 2 displays the results of a one standard deviation shock to income, with the variables ordered as income, prices and sales. The first graph is the response of US sales to a shock to income. The predicted response is plotted against the one estimated from the observed data between 1960 and 1995 (and displayed in Figure 1). For the US, the model predicts an increase in sales in the first period, followed by a fall below baseline. The predicted response differ in two ways from the observed one: the initial response is much higher and the dampened oscillations do not have the same periodicity. However, when we restrict the sample period to be the one used in the structural estimation (1975-1995), the observed and predicted responses are very close. This indicate that the match of the estimated model is good, but that the impulse responses are sensitive to the sample period, which is not very surprising. For France, the predicted response is close to the observed one. Here the sample periods are roughly the same.

4.2.2 Explaining the mechanism through IRFs

So, the estimated model has the ability to match an ARMA(1,1) time series representation of car expenditures, as well as impulse response functions of sales. Given this "empirical success", we now turn to a more intuitive discussion/evaluation of the model.

Using our estimates of the parameters of the agent's optimization problem, we can simulate the effect on aggregate car sales of an income shock. For example, consider the model estimated off the U.S. data which allowed for curvature in the utility function. We start with an initial cross section distribution in year 1981. We then introduce a shock to income of one standard deviation, and simulate aggregate sales over time (we assume here for simplicity that the covariance matrix for the innovations is diagonal). Figure 3 shows the response in new car sales.

The picture also displays two more impulse responses. First we graph the impulse response obtained by fitting an ARMA(1,1) to the data. This simple PIH model is not able to pick up the dynamics of car sales over time. Using an ARMA(1,1), the impact of a shock on sales is $\rho^{t-1}(\rho-\alpha)$, t periods after the shock, where ρ is the AR coefficient and α is the MA one. First, as ρ is positive in the estimation, there is no way the ARMA(1,1) can reproduce the oscillations in the impulse response function. Second, in order to match the patterns of response over time, $(\rho-\alpha)$ has to be positive. As ρ is estimated between 0.4 and 0.8 on aggregate data (see Table 1), this means that the implied depreciation rate δ cannot be lower than 0.2 to 0.6. We see here an important point about the ARMA(1,1) model: it is structurally unable to deliver a depreciation rate low enough to be credible.

The last impulse response displayed in Figure 3 has been obtained by fitting an unrestricted linear VAR to the observed data (sample period 1975-1996) as in Figure 2. When using a linear VAR, the impulse response is qualitatively closer to the one obtained from our structural model. However, the cycles do not have the same frequency. This can arise for two reasons, either because our model better picks up the dynamics of sale through the modelling of the cross-section distribution of cars, or because the VAR is less restricted in its price and income effects and is thus a better representation of the data.

However, given the exercise in the previous section, our model is able to reproduce the results of a linear VAR, for that particular sample period. This suggests that our model is correctly specified, and that an unrestricted VAR model might be too linear to pick up the dynamics of car sales.

The results of the impulse responses, of course, may depend on the cross sectional distribution of car ages at the time of the shock. To illustrate, Figure 4 shows the impulse response computed using the same shocks but comparing 1981 and 1995 as initial cross section distributions. In 1981, the average age of cars is higher, thus an income shock has a bigger instantaneous effect as more household replace their cars. This effect is then propagated through time.

4.3 Decompositions

Caballero [1993] explains why a dynamic discrete choice model might explain the response of durable expenditures to an income shock. The key point is that a shock to income produces a dynamic in durable expenditures as agents respond differentially. While agents may differ along a number of dimensions, our analysis focuses on the cross sectional distribution of car vintages. Essentially agents with relatively old cars will respond to the income shock by replacing their car first and then agents with younger cars will respond later. The delayed response simply reflects the upward sloping adjustment hazard: all else the same, agents with younger cars are less likely to respond to income variations than are agents with older cars. The evolution of the cross section distribution through time can be a source of fluctuations, which are picked up by the impulse responses. As the distribution evolves, following (4), sales will respond. In fact, the magnitude of this response depends on the slope of the hazard function: the flatter is the hazard, the less responsive will sales be to the evolution of the cross sectional distribution.

A second source of movement is the dynamics induced by prices and income as these processes are serially correlated. Movements in these variables are represented by shifts in the probability of adjustment (hazard).

We study the relative importance of these two influences (hazard and CDF shifts) in two ways. First, we recompute the impulse response functions either by holding the CDF fixed or by limiting the shifts in the hazard function. Second, we decompose the time series of sales into these two influences.

4.3.1 Impulse Response Functions

Figure 5 shows the impulse response functions from two exercises. In the first, we hold the CDF fixed at its value in 1981 and consider the effects of an innovation to income. In the second case, we allow the CDF to evolve but impose that the income variation be temporary. Thus the hazard function shifts out for one period only.

We find that with a fixed CDF, the impulse response is close to the global impulse response reported in Figure 3 and reproduced in Figure 5. Thus the dynamic is mainly due to the evolution of the prices and income. The dynamics induced by the evolution of the cross section distribution contributes surprisingly little. Evidently, the depreciation of cars along with the household specific shocks are significant enough to eliminate replacement cycles.

Narrowing down our search for an explanation of these oscillations, we simulate the model with a fixed CDF, but we eliminate the cross effect of prices and income (we set $\rho_{Yp} = \rho_{pY} = 0$). Figure 6 displays the impulse response functions. We then find that all the oscillations are gone. From this we conclude that the oscillations in sales are mainly due to the dynamics of prices and income, and more particularly to the cross effects.

4.3.2 Time Series Decompositions

For the time series of sales, using (3), the change in aggregate sales can be decomposed into two terms. The first term is the change due to shifts in the hazard functions, such as price or income movements against a fixed cross-section distribution. The second term is the contribution of the shifts in the cross-section distribution, holding the hazard function fixed:

$$S_t - S_{t-1} = \sum_k [H_k(z_t;\theta) - H_k(z_{t-1};\theta)] f_{t-1}(k) + \sum_k H_k(z_t;\theta)[f_t(k) - f_{t-1}(k)] + u_t \tag{5}$$

For our simulated data, this decomposition is exact. The error term reflects the fact that, in the actual data, there are measurement problems and not all sales variations are a consequence of replacement. Inspection of the time series of the error process indicates little structure to this error supporting the view that it is mainly due to measurement problems.

Given data on the cross section distribution, we can compute the contributions to the change in sales of the fixed hazard and fixed cross-section distribution components. The fixed cdf component tracks the change in sales very closely both for the US and for France, whereas the other component has a much smaller variance and have a low correlation with the change in sales. In particular, the R^2 associated with shifts in hazards is equal to 0.93 both for the US and for France and the one associated with shifts in the cross section distribution is only 0.25 for the US and 0.16 for France. **Thus in the actual data hazard shifts are the main source of fluctuations.**

From the simulated data of our estimated model, we can also evaluate the contribution of each term to the variability in aggregate sales. From a simulated sample of length 400, we find very similar results to the real data: shifts in hazards are the most important determinant of sales. The R^2 associated with shifts in hazard is equal to 0.97, whereas the R^2 associated with shifts in the cross-section distribution is only 0.1.

Spectral analysis using the series for the fixed cdf and the fixed hazard gives more insights on the relative contributions of both series. The shifts in the cross section distribution have an overall low contribution, but has more long run effects at a frequency corresponding to a period of approximately 8 years for the US and 13 years for France. The spectrum for the shifts in the hazard function peaks at a period of about 5 years for the US and 2 years for France. Here again, the results for the simulated data are very similar to the results obtained on real data.

5 Conclusions

We have found that a model which stresses the dynamic optimizing behavior of households over the choice of purchasing a new car goes a considerable way towards solves the "durables puzzle" of Mankiw [1982]. The point that this type of model might explain this puzzle is certainly not new.[17]

Our approach to the problem follows a methodology that is quite different from that put forth by Bar-Ilan and Blinder and utilized in much of the subsequent work. We specifically avoid the specification of individual optimization in terms of (S,s) bands and instead focus on the underlying parameters of the individual's dynamic discrete choice problem. Further, our model was estimated using data which emphasized the cross sectional variations in the data. Still, we find that this modelling approach delivers time series implications that match certain features of the data. In contrast to other studies, are results come from a regression on simulated data that essentially mimics Mankiw [1982].

In trying to understand our finding, we are naturally led to a decomposition of the movements in sales into two components: shifts in the hazard function and the evolution of the cross sectional distribution. We report that most of the variation in the change in sales can be attributed to shifts in the hazard function though the evolution of the cross sectional distribution is certainly present. This leads us to one of our main findings: **the ARMA representation can be reproduced by a dynamic discrete choice model.**

In terms of further work, there are two elements of our basic model that deserve additional attention. First, as in Caplin-Leahy [1998] there is undoubtedly some room for endogenous price variations due to upward sloping supply to explain some of the results. This assumption dramatically simplifies our numerical analysis since the cross sectional distribution of car vintages would then be an element in households' state vector. Second, the model we have studied imposes borrowing restrictions. This assumption certainly simplifies our analysis by reducing the dimensionality of our state space. Further, we found that a model with linear utility where borrowing is irrelevant does not fit the data as well. Understanding the robustness of our findings to relaxing these assumptions is of considerable interest.

[17] To our knowledge, Bar-Ilan and Blinder [1988] deserve credit for drawing the attention of the profession to this point.

A Data Appendix

In this appendix, we report additional time series represenations of the data that are useful in evaluating our results and models.

A.1 ARMA(2,1) Representation

Bernanke [1985] argues that durable expenditures are given by an ARMA(2,1) process:

$$e_{t+1} = \alpha_1 e_t + \alpha_2 e_{t-1} + \varepsilon_{t+1} - (1-\delta)\varepsilon_t$$

where the coefficient α_1 and α_2 on the lagged expenditures sum to 1, and are functions of the discount factor and parameters of the utility function (see the appendix in Bernanke[1985] for more details). [18] As in Mankiw's specification, the MA coefficient implies an estimate of the rate of depreciation.

Motivated by this analysis, Table 5 presents ARMA(2,1) representations for the same series analyzed in Table 1.

Table 5: ARMA(2,1) Estimates on US and French Data

Specification	No trend			Linear trend		
	α_1	α_2	δ	α_1	α_2	δ
US durable expenditures	1.95 (0.07)	-0.95 (0.07)	0.13 (0.1)	0.62(0.7)	0.33(0.7)	1.24 (0.7)
US car registration	0.89 (0.31)	-0.04 (0.21)	0.5 (0.3)	0.89(0.31)	-0.04(0.21)	0.5 (0.3)
France durable expenditures	0.7(0.4)	0.3(0.4)	1.03 (0.4)	0.7(0.4)	0.2 (0.4)	1.0 (0.45)
France car expenditures	0.8 (0.3)	0.2(0.3)	0.9(0.3)	0.65 (0.2)	0.2 (0.3)	0.96 (0.4)
France car registrations	0.85 (0.24)	0.09 (0.22)	1.43(0.23)	0.72 (0.36)	0.09 (0.23)	1.35 (0.35)

Notes: Estimation done on quarterly data. For the US, source FRED database, 1959:1-1997:3. French data: source INSEE, 1970:1-1997:2. US Registration: 1968-1995.

Here we see that again the implied values of the depreciation rate are close to 100%. Further, the coefficient on the second lag of expenditures should be negative, according to the theory, while in the estimation this is always positive. Finally, note that the sum of the estimated AR coefficients is close to one, as predicted by the theory.

[18] If there are no adjustment costs, then α_2 is equal to zero.

A.2 Estimation Results for joint Process of Income and Prices

Table 6 displays the estimation results for the joint process of income and prices, used in the structural model.

Table 6: VAR for prices and Income

	US		France	
Parameter	Estimate	standard error	Estimate	standard error
ρ_{YY}	0.75	0.12	0.67	0.12
ρ_{Yp}	1.25	0.47	0.14	0.22
ρ_{pp}	0.68	0.14	0.65	0.17
ρ_{pY}	-0.10	0.03	-0.04	0.09
ω_Y	3.0e6	-	2.6e6	-
ω_p	2.7e5	-	1.5e6	-
ω_{Yp}	2.6e5	-	-1.6e5	-

Note: Regression done on detrended ***annual*** *series.*

A.3 ARMA(1,q) Representations

Following Caballero [1990], we present estimates of the implied depreciation rates using higher order moving average representations. Caballero [1990], Table II, page 734, using US annual aggregate durable expenditures, finds that the MA coefficients sums to -.95 implying a rate of depreciation of 5%. Using the same method and series, but with a sample period 1959-97 instead, we find an implied rate of depreciation of 10%. However, as we show below, this result is not robust across series and countries, to the specification of the ARMA model or to the detrending method.

Table 7 panel A displays the results for the US durables as well as for four other series, US new car registrations, French durable expenditures, French car expenditures and French new car registrations. As in Caballero [1990], all series are annual and exponentially detrended. An MA(5) was fitted to the changes in the series. Apart from the US durables, none of the series seem consistent with a long run PIH model. Increasing the number of lags beyond five did not produce any smaller implied depreciation rates.

The model fitted in Panel A imposes a coefficient of one on the lagged series and leads to an over-differencing if the process is not a unit root. This over differencing could bias the results. Given that the series are exponentially detrended, at the 5% level, a unit root is rejected for the following detrended series: US durable expenditures and US car registrations. Thus a preferred specification for both series should be an ARMA(1,5) model. Table 7, panel B displays the sum of the MA coefficients for such a specification. The implied depreciation rate for the US durables is now 96%. The results for the other series are also sensitive to the model specification. None of the point estimates are close to a reasonable annual rate of depreciation.

The results could also depend on the detrending method. Table 7, panel C displays results for the series fitting an MA(5) on changes. An augmented Dickey-Fuller test reveals that all the changes in these series are stationary, so detrending is not necessarily . Here we find a depreciation rate for the US durables of 51%. Two series have a low implied rate of depreciation, French durable expenditures and French new car registrations, close to 20%. However, when extending the number of lags in the moving average structure to seven, the implied depreciation rates are 74% and 58%. In general, extending the number of lags in the moving average structure does not produce a lower rate of depreciation, except for the US durable expenditures, when changes in the series are considered.

As pointed out in Caballero [1990], a completely equivalent way of testing whether the moving average coefficients sum to one is to test whether the sum of autocorrelations of the changes of the series converges to -0.5 as the number of lags increases. The cumulative sum of the autocorrelations are displayed in Figure 7 for four annual series: US durable expenditures, US car registrations, French durable expenditures and French car registrations. For US durables, when we consider lags up to 8 years, we obtain comparable results with those depicted in Caballero [1990, Figure I], as the cumulative sum seems to converge to -.5 after 7 to 8 years (30 quarters). [19] However, when we extend the cumulative sum beyond this point, there is no longer any evidence of convergence as the cumulative sum increases to around -0.1, a value which is not consistent with a model of sluggish adjustment to the PIH, even with slow adjustment. Similar results, although more mixed, are found for the other series. Even after 8 years, the cumulative sums of

[19] Caballero [1990] reports evidence for both quarterly and annual series. His Figure 1 is for quarterly observations of durable expenditures. The sum of the autocorrelation coefficients is near -.5 after 17 quarters, rises to nearly -.3 and then returns to -.5 by about 38 quarters.

Table 7: ARMA(1,q) Representations for Annual Data

Series	Sample Period	Implied δ
A. MA(5) on Changes. Exponentially Detrended Series.		
US, Durable Expenditures	1959-97	0.1 (0.25)
US, New Car Registrations	1968-95	-0.8 (0.76)
France, Durable Expenditures	1970-97	1.07 (0.27)
France, Car Expenditures	1970-97	1.07 (0.25)
France, New Car Registrations	1968-97	0.44 (0.18)
B. ARMA(1,5) on Levels. Exponentially Detrended Series.		
US, Durable Expenditures	1959-97	0.96 (0.17)
US, New Car Registrations	1968-95	0.54 (1.08)
France, Durable Expenditures	1970-97	1.58 (0.89)
France, Car Expenditures	1970-97	0.99 (1.01)
France, New Car Registrations	1968-97	0.71 (0.67)
C. MA(5) on Changes. No Detrending.		
US, Durable Expenditures	1959-97	0.51 (0.11)
US, New Car Registrations	1968-95	-0.78 (0.73)
France, Durable Expenditures	1970-97	0.23 (0.32)
France, Car Expenditures	1970-97	-1.82 (1.32)
France, New Car Registrations	1968-97	0.21 (0.21)

Note: Estimation done on annual data.

autocorrelations displays variations and does not seem to converge, even if they are closer to -.5. Similar results can be found on quarterly data.

References

Adda, J. and R. Cooper (2000). "Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies." *Journal of Political Economy*, 108, 4.

Attanasio, O. (1997). "Consumer Durables and Inertial Behaviour: Estimation and Aggregation of (S,s) Rules for Automobile Purchases." *mimeo, University College London*, forthcoming *Review of Economic Studies*.

Attanasio, O. P. K. Goldberg and E. Kyriazidou (2000) "Credit Constraints in the Market for Consumer Durables: Evidence from Micro Data on Car Loans", NBER WP 7694.

Bar-Ilan, A.and A. Blinder (1988), "The Life Cycle Permanent-Income Model and Consumer Durables," *Annales D'Economie et de Statistique*, No. 9 (1988), 71-91.

Bar-Ilan, A. and A. Blinder (1992). "Consumer Durables : Evidence on the Optimality of usualy doing Nothing." *Journal of Monetary Economics*, 15, 41–68.

Bernanke, B. (1985). "Adjustment Costs, Durables and Aggregate Consumption." *Journal of Money, Credit and Banking*, 24, 258–272.

Bertola, G.and R. Caballero (1990),"Kinded Adjustment Costs and Aggregate Dynamics," in *NBER Macroeconomics Annual*, vol. 5, ed. by O. Blanchard and S. Fischer, Cambridge, Mass.:MIT Press, 1990.

Caballero, R. J. (1990). "Expenditure on Durable Goods: a Case for slow Adjustment." *Quarterly Journal of Economics*, 105, 727–743.

Caballero, R. J. (1993). "Durable Goods: An Explanation for their Slow Adjustment." *Journal of Political Economy*, 101, 351–384.

Campillo, M. (1999) "Essays on Consumption and Inflation," PhD Thesis, Boston University, 2000.

Caplin, A. and J. Leahy (1997), "Durable Goods Cycles," mimeo, Boston University, October 1997.

Cooper, R. J. Haltiwanger and L. Power (1999),"Machine Replacement and the Business Cycle: Lumps and Bumps" *American Economic Review*, 89, 4, 921–946.

Eberly, J. (1994). "Adjustment of Consumers' Durables Stocks : Evidence from Automobile Purchases." *Journal of Political Economy*, 102, 403–436.

Gourieroux, C., Monfort, A. and Renault, E. (1993), "Indirect Inference" *Journal of Applied Econometrics*, 8, 85–118.

Laffont, P H. Ossard and Q. Vuong. (1995). "Econometrics of First-Price Auctions." *Econometrica*, 63, 953–980.

Lam, P. (1991). "Permanent Income, Liquidity and Adjustments of Automobile Stocks: Evidence form Panel Data." *Quarterly Journal of Economics*, 106, 203–230.

Mankiw, G. N. (1982). "Hall's Consumption Hypothesis and Durable Goods." *Journal of Monetary Economics*, 10, 417–425.

Figure 1: Impulse Response Functions To an Orthogonal Shock to Income

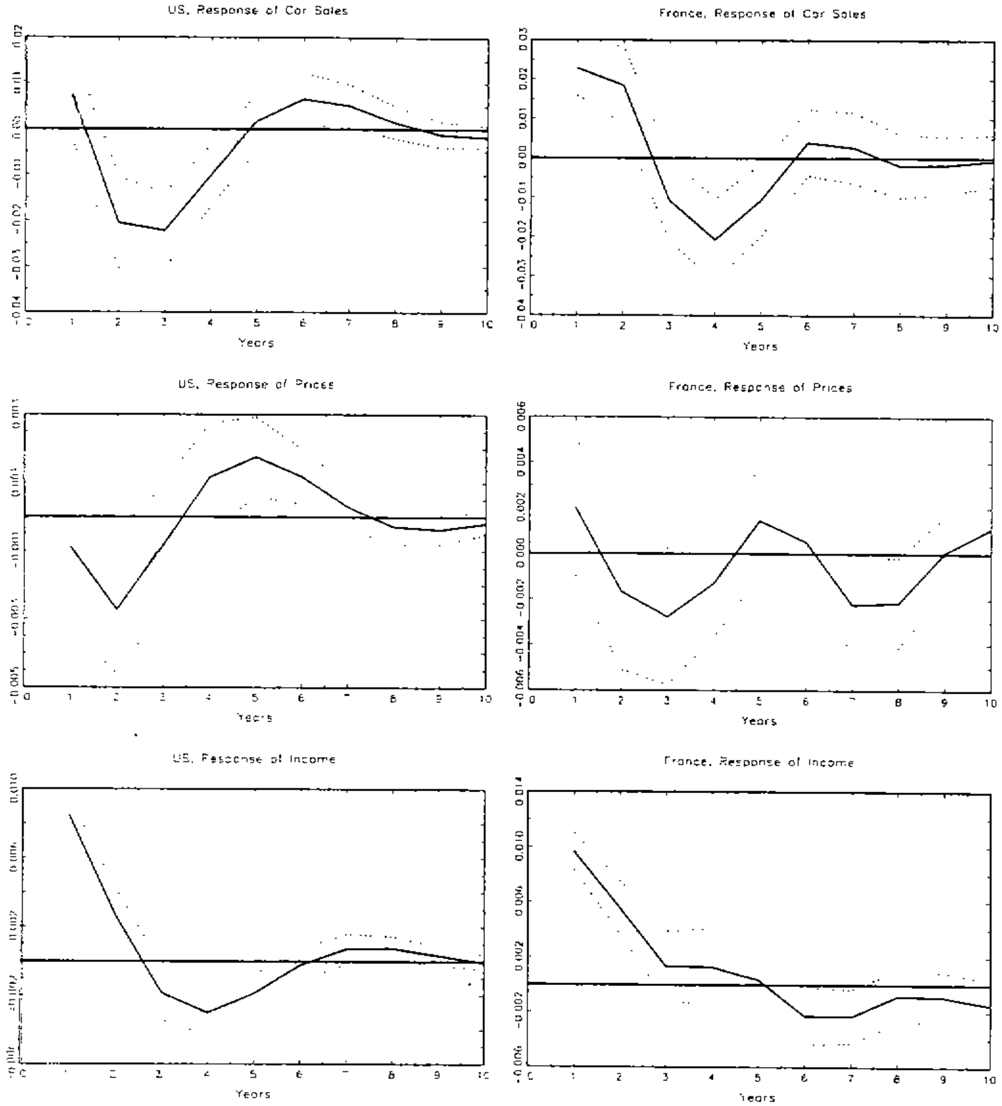

Note: **Impulse response functions** computed from a VAR(1) for the US and a VAR(3) for France (the optimal number of lags was selected with an Akaike criterion). VAR included income, relative price and new car registrations. All series were taken in logs and filtered using an HP filter. One standard deviation to income innovation, with variables ordered as income, prices and expenditures. One standard deviation confidence bands are displayed.

Figure 2: Impulse Response of Sales, from Observed and Simulated Data

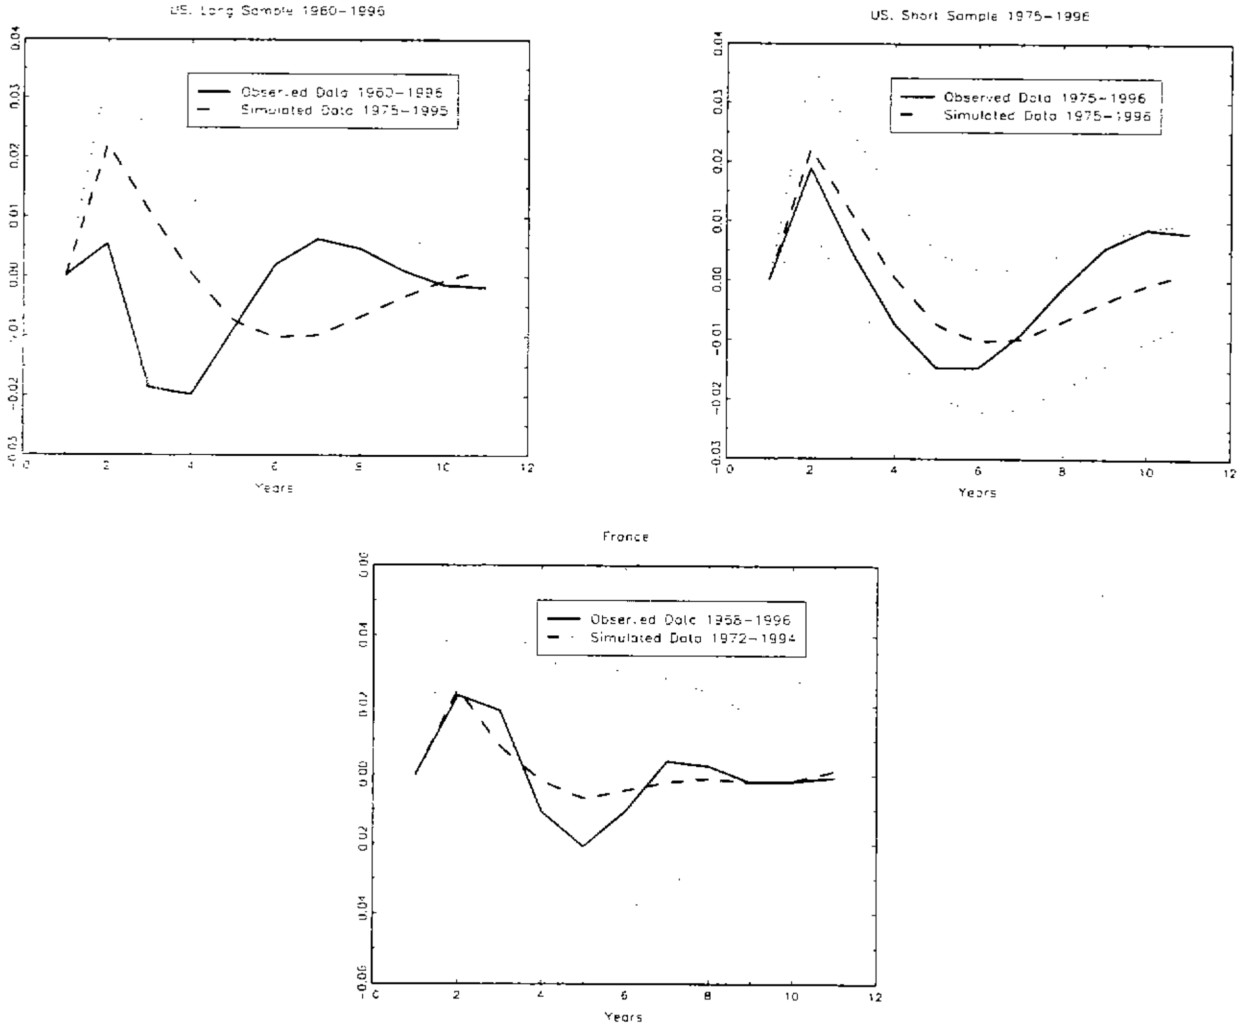

Figure 3:

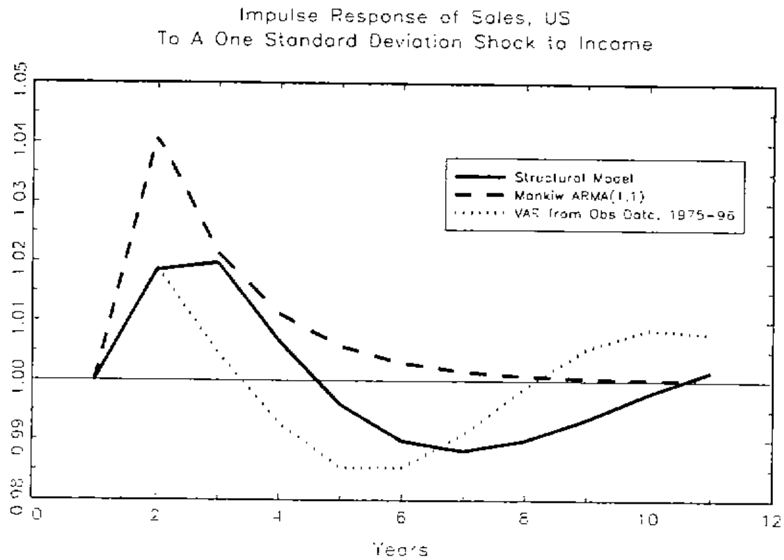

Figure 4:

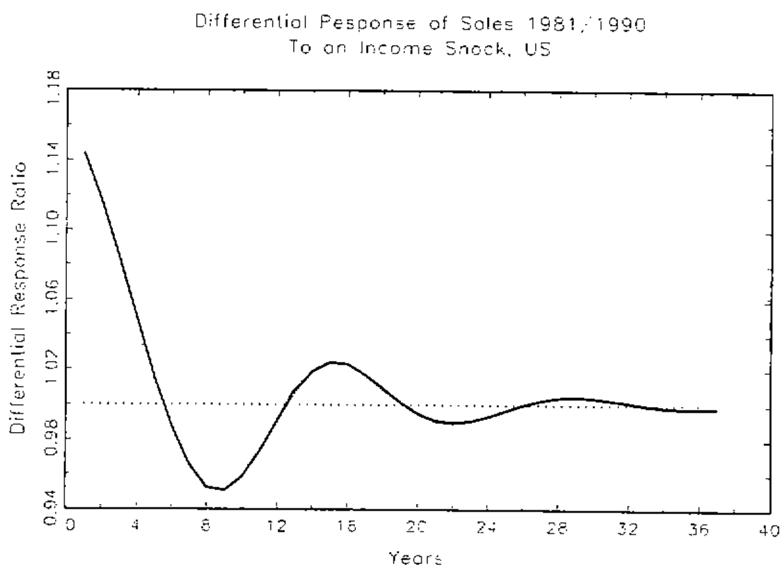

Figure 5:

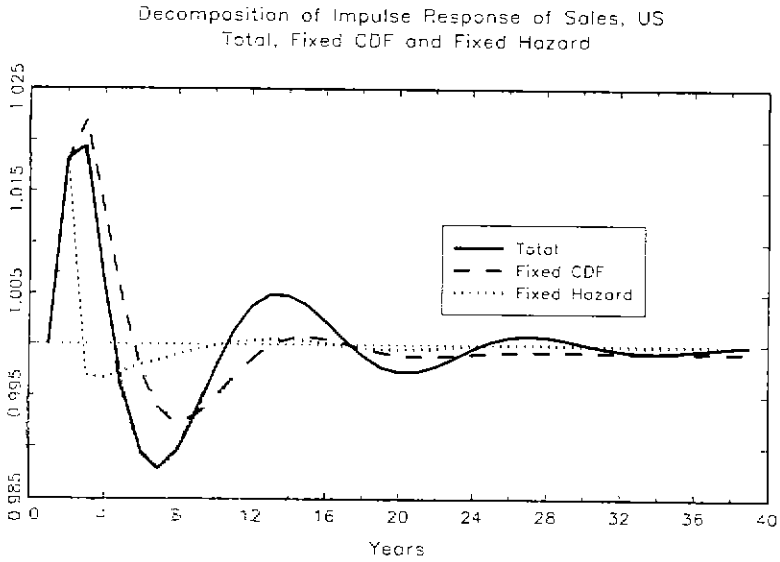

Figure 6:

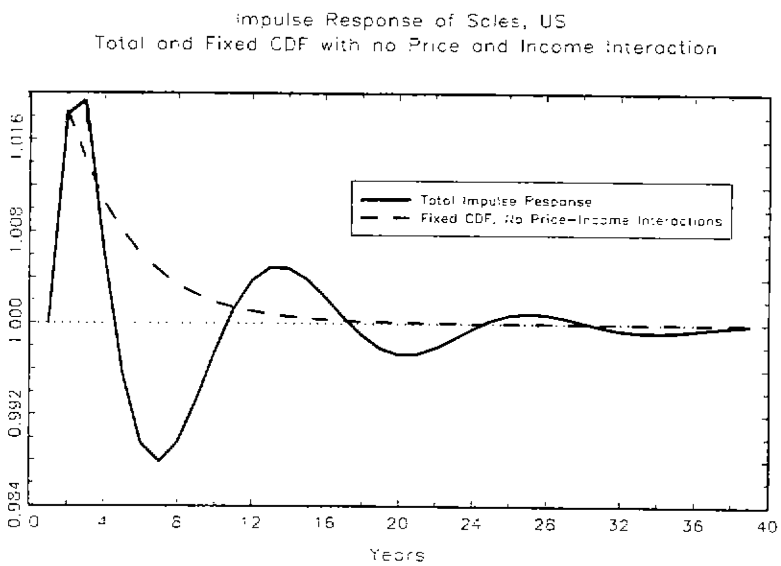

Figure 7: Cumulative Sum of Autocorrelations