Adoption Costs of Financial Innovation: Evidence from Italian ATM Cards*

Kim P. Huynh[†] Philipp Schmidt-Dengler[‡] Gregor W. Smith[§] Angelika Welte[¶]

March 19, 2023

Abstract

The discrete choice to adopt a financial innovation affects a household's exposure to inflation and transactions costs. We model this adoption decision as subject to an unobserved cost. Estimating the cost requires a dynamic, structural model, to which we apply a conditional choice simulation estimator. A novel feature is that preference parameters are estimated separately, from the Euler equations of a shopping-time model, to aid statistical efficiency. We apply this method to study ATM card adoption in the Bank of Italy's *Survey of Household Income and Wealth*. There, the implicit adoption cost is large, but varies significantly by age cohort, education, or region.

Keywords: dynamic discrete choice, money demand, financial innovation.
JEL codes: E41, D14, C35.

*Smith acknowledges the Social Sciences and Humanities Research Council of Canada and the Bank of Canada research fellowship programme for support of this research. Schmidt-Dengler acknowledges financial support from the German Science Foundation through Sonderforschungsbereich Transregio 15. Dörte Heger provided excellent research assistance. For constructive comments, we thank Jason Allen, Thomas Lemieux, Héctor Pérez Saiz, Ben Tomlin, and seminar and conference participants. The views in this paper represent those of the authors alone and are not those of the Bank of Canada.

[†]Currency Department, Bank of Canada. kim@huynh.tv
[‡]Department of Economics, University of Vienna. philipp.schmidt-dengler@univie.ac.at
[§]Department of Economics, Queen's University. smithgw@econ.queensu.ca
[¶]Currency Department, Bank of Canada. awelte@bankofcanada.ca

1 Introduction

We study how households adopt a new financial product.[1] Specifically, we provide a measure of households' perceived adoption costs and benefits and whether these perceptions explain observed adoption patterns. Our application is to the adoption of ATM cards in Italy, as tracked in the Bank of Italy's *Survey of Household Income and Wealth*, a rich survey data on household financial decisions.

Adopting a card provides ongoing benefits to consumers, for example by helping them use cash efficiently. Yet adoption appears slow and incomplete historically. Measuring an implicit (unobserved) adoption cost is important to predicting the speed with which financial innovations can spread. By measuring this cost we also provide a framework in which counterfactual scenarios can be simulated. For example, we consider the scenario where all households adopt the financial innovation and estimate what compensation would have to provided to achieve this universal adoption.

Households' adoption patterns have two features that we strive to incorporate into the modelling and estimation strategy. First, adoption is a dynamic, discrete choice where the household weighs the future benefits of the new technology against a one-time adoption cost. Second, households are heterogeneous in how they use the new technology and hence benefit from it. The study contributes to the literature on adoption of financial innovation by proposing a conditional choice probability estimator for the parameters of these benefits and costs. The estimator builds on the insights by Hotz et al. [1994] and is based the methods of Hotz and Miller [1993], Aguirregabiria and Magesan [2013] and Arcidiacono and Miller [2011].

There are three key features of the estimator. First, we combine the simulation estimator with the estimation of preference parameters via Euler equations for households with or without ATM cards. The Euler equations come from a shopping-time model that describes both the intensive margin of money-holding and the additional gains from holding an ATM card. Time effects in the implied money-demand function also allow for the diffusion of ATM machines and bank branches over the historical sample. It is important to control for this diffusion in banking services in estimating the adoption cost. Second, we allow for both observed and unobserved heterogeneity among households, in their cash-holding behavior. Tracking their decisions over time allows us to control for this heterogeneity, which we find to be substantial. Third, we assume that adoption is irreversible in that a household cannot 'un-adopt' the new technology. This assumption implies a finite-dependence property of the decision problem (outlined by Aguirregabiria and Magesan [2013], Arcidiacono and Miller [2011] and Arcidiacono and Ellickson [2011]) which facilitates computation.

[1] KPH: motivate via EU financial inclusion initiative `http://ec.europa.eu/finance/finservices-retail/inclusion/index_en.htm`.

FINDINGS

Section 2 describes the data sources. Section 3 outlines the household decision problem. Section 4 describes the econometric building blocks while Section 5 contains the results. An appendix contains extensive sensitivity analysis.

2 Related Research

This project is related to behavioral perspectives on financial decisions, to econometric work on the participation decision, and to numerical portfolio models. This section briefly sets our work in these contexts.

First, financial decisions, including the decision to adopt a new type of account, may be made infrequently and may be subject to deferred benefits that are hard to measure. Households thus may well make mistakes in these decisions, as a wealth of research in behavioral economics has emphasized. We explore the possibility that some households may rationally not adopt an ATM card, because of the cost involved in doing so. Non-adoption may be a wise choice if fees are high or ATM locations are inconvenient. Vissing-Jørgensen [2003] suggests that to show non-participation to be rational requires investigators to find the participation cost not be implausibly large. We explore whether this conclusion holds for those without ATM cards.

Second, the extensive margin or participation choice is a key issue in empirical work on household portfolios. Most work on participation or account adoption uses limited-dependent-variable econometrics to statistically explain the dichotomous variable I_{it}. Miniaci and Weber [2002] provide an excellent discussion of the econometric issues in estimating these models with data from household surveys. Most empirical work concerns the decision to hold risky assets. Guiso et al. [2003], Perraudin and Sørensen [2000], and Vissing-Jørgensen [2002] study participation in stock markets. The same methods also have been applied to the decision to open a bank account. Attanasio et al. [2002] study the demand for currency allowing for the adoption decision.

A number of studies use panel data that allow researchers to track adoptions and potentially allow for unobserved heterogeneity across households. Alessie et al. [2004], for example, track the ownership of stocks and mutual funds in a panel of Dutch households. These studies typically find that previous participation is significant in statistically modelling current participation. A key issue is whether this pattern reflects true state-dependence or persistent unobserved exogenous variables. This pattern of persistence may be evidence of a fixed cost to adoption and so may help identify that cost.

Third, a number of researchers have studied the extensive margin of stock-holding using numerical portfolio models solved by dynamic programming. For example, Haliassos and Michaelides [2003] introduce a fixed cost into an infinite-horizon consumption-portfolio model. Gomes and Michaelides [2005] also solve and simulate a life-cycle model with a fixed cost of holding stocks. These studies calibrate planning problems and carefully study the outcomes.

Alan [2006] takes the important step from simulation to estimation, by indirect inference. She estimates preference parameters and a fixed cost of entry to a stock market using the solved consumption-portfolio model so that statistics from a participation equation in the simulated data match those from the same equation in historical, panel data. Sanroman [2007] adopts a similar method. She outlines a planning problem that involves both a participation decision and an asset-allocation or portfolio problem. She then solves the dynamic programme by discretization. Finally, she estimates parameters by indirect inference using the Italian SHIW with a logit model as the auxiliary estimating equation. She estimates that the participation cost for holding stocks varies from 0.175 to 6 percent of income, or from 10 to 1126 euros, with households with higher education implicitly facing lower costs.

One model of asset allocation in the case of cash management is the Baumol-Tobin model or a related inventory model. In this framework, Alvarez and Lippi [2009]'s work constitutes the state of the art in theory and empirical evidence on household cash management. They model the household's cash withdrawal conditional on the adoption of an ATM card. Their framework also measures changes over time in withdrawal costs. They find a relatively small benefit to adopting an ATM card, though they note that it is based only on a reduction in withdrawal costs and not on the card's use as a debit card. Yang and Ching [2014] model both the extensive and intensive margins, again using the Baumol-Tobin model to describe the latter. They estimate a significantly larger cost of ATM adoption. We adopt an alternate, shopping-time model of money holding that also is widely used in macroeconomics. A key feature of this model that constrasts with the Baumol-Tobin model is that we do not assume that the cost of a cash withdrawal is proportional to consumption or income. Section 7 discusses the impact of this distinction on the results.

This study differs from the work on dynamic consumption-portfolio models in that we estimate the parameters of the mixed discrete (adoption) continuous (money demand) without solving the dynamic programming problem. We can therefore incorporate heterogeneity in time horizons and adoption costs with very little computational costs. Separately estimating the a subset of parameters of the utility function conditional on the discrete action using Euler equation methods was suggested by Pakes [1994], albeit in a different context. To estimate the cost of adoption we employ a conditional choice simulation estimator in the spirit of Hotz et al. [1994] and Hotz and Miller [1993]. The finite dependence nature of the adoption problem allows us to estimate the dynamic model by simulating only one-period ahead as shown by Arcidiacono and Miller [2011] and illustrated in Arcidiacono and Ellickson [2011]. A particular feature of our setting with finite dependence is that we can identify the household discount factor as shown in Abbring and Daljord [2020].

In sum, we combine some of the economic structure from numerical portfolio models—planning horizons and parameters of discounting and intertemporal substitution—with the ability to accommodate all the household heterogeneity and econometric tractability of the discrete-choice econometric models.

3 Data Sources

3.1 Survey of Household Income and Wealth

Our study relies on household-level data from the Bank of Italy's *Survey of Household Income and Wealth* (*SHIW*), which is the gold standard for panel surveys involving wealth and savings. It has detailed information on account status, wealth, and consumption, and the largest and longest coverage of any such panel. The *SHIW* is the main data source for studies on money demand and financial innovation by Attanasio et al. [2002], Alvarez and Lippi [2009], and Lippi and Secchi [2009], among others.

The *SHIW* is a biennial survey run by the Banca d'Italia. We use the 1991, 1993, 1995, 1998, and 2000, 2002, and 2004 waves. We stop at 2004 as one of our main variables—average currency holdings—is discontinued from 2006 onwards with the exception of 2008. The three year spacing from 1995 to 1998 was a result of the Banca d'Italia switching survey providers. The Banca d'Italia spends considerable resources to ensure that the data is nationally representative, as outlined by Brandolini and Cannari [1994]. The *SHIW* survey is a rotating panel with about 8,000 households per wave. The rotating panel design is incorporated because there is an attrition rate of roughly 50%. Jappelli and Pistaferri [2000] provide an extensive discussion of the quality of the *SHIW* data and also provide a comparison with Italian National Accounts data to address issues of sample representativeness, attrition, and measurement. Details about the variables we used are available in a separate technical appendix.

ATM cards involve a small annual fee, but no additional charges for withdrawals at machines owned by the issuing bank. Their first benefit is that they allow card-holders to withdraw cash rapidly and when banks are closed. Checking accounts bear interest, so the ability to make withdrawals at lower cost can reduce foregone interest earnings from holding cash. A second benefit is that they can be used as point-of-sale debit cards for retail transactions. Despite these benefits, though, the use of cash remained very widespread in Italy throughout this period.

Table 1 reveals that the fraction of households with an ATM card in 1991 was 29% and that it steadily increased to 58% in 2004. Although the survey has a high attrition rate, many actual ATM card adoptions can be observed. On average, the share of households who did not have an ATM card in the previous wave of the survey, were in both the current and previous waves, and had a card in a given, current wave was 16.7%.

Next, we focus on average currency holdings, consumption, and wealth. All the nominal variables are expressed in 2004 equivalent euros. During this period the average currency holdings fell for both the households with and without an ATM card. However, with the exception of 1991 the average cash holdings of ATM holders were lower than those of non-ATM holders. Not surprisingly, those with ATM cards tended to have higher consumption and financial wealth than those without ATM cards. Notice that the difference in consumption and wealth increased over time as was detailed by Jappelli and Pistaferri [2000].

3.2 Regional Inflation and Interest Rates

We also use data on inflation and interest rates from a variety of sources. The inflation rate, measured as the per-annum change in consumer prices, is taken from the *International Financial Statistics* of the International Monetary Fund. The data are on an annual basis from 1989 to 2010. The Banca d'Italia *Base Informativa Pubblica* online historical database is the source for for regional nominal deposit interest rates. These interest rates are constructed from a variety of historical tables at a quarterly frequency. The quarterly data are then aggregated to an annual frequency using simple sum averaging to derive annual data from 1989 to 2010. We refer to Alvarez and Lippi [2009] for more details on the data sources and institutional details.

4 Household Choice Problem

We propose to model households decision to adopt a financial innovation, specifically an ATM card, as an optimal stopping process. A household will adopt the innovation when they expect the benefits of adoption to outweigh the opportunity costs of not adopting. The model is dynamic since expectations about adoption benefits are computed from summing up (discounted) future per-period utilities. A household in period t is described by

1. ATM card adoption status from the previous period $I_{t-1} \in \{0, 1\}$.

2. Choice specific adoption shocks $\epsilon_t = (\epsilon_t^0, \epsilon_t^1) \in \mathbb{R}^2$ where ϵ_t^I is incurred for choosing $I \in 0, 1$ in period t. These shocks are known to the household before they make their decision, but they are not observed by the econometrician. We assume that ε_t are independently and identically distributed across adoption choices, households, and over time.

3. a vector of state variables z_t that are exogenous, ie do not depend on I_{t-1}. These state variables are for example the size of the household, its wealth, education and age of its members.

A household's utility function depends on the adoption status of financial innovation I_t, its real consumption expenditures c_t as well as its real money holdings m_t, and is given by $u(I_t, c_t, m_t)$. It is increasing in consumption ($u_c > 0$). We consider a shopping-time model of money-holding, as outlined by McCallum [1989] (pp 35–41) and Walsh [2003] (pp 96–100). Holding money also adds to utility because it reduces the time spent shopping and so adds to leisure ($u_m > 0$). We also impose the usual concavity assumption ($u_{cc} < 0, u_{mm} < 0$). Using a financial innovation, or more specifically access to cash through an ATM card increases the utility of consumption for a given real cash balance m_t because the household can shop more efficiently. In the shopping-time utility function, this benefit shows up as $u(1, c, m) > u(0, c, m)$.. We specify the functional form of u in 5.2.

Adoption of an ATM card of financial innovation is costly. It involves a deterministic cost $\bar{\kappa}$ and the choice specific cost shock ε_t^ℓ:

$$\kappa_t(I_t, I_{t-1}, \epsilon_t) = \mathbb{1}\{I_{t-1} = 0\}(I_t - I_{t-1})\bar{\kappa} + \sigma_\kappa \sum_{\ell=0,1} \epsilon_t^\ell \mathbb{1}\{I_t = \ell\} \tag{1}$$

The household's per-period pay-off is thus given by

$$u(I_t, c_t, m_t|z_t)) - \kappa_t(I_t, I_{t-1}, \epsilon_t). \tag{2}$$

With the specification of the cost shock, the period utility can written in terms of an observed and an unobserved component:

$$\underbrace{u(I_t, c_t, m_t|z_t) - \mathbb{1}\{I_{t-1} = 0\}(I_t - I_{t-1})\bar{\kappa}}_{observed} - \underbrace{\sigma_\kappa \sum_{\ell=0,1} \epsilon_t^\ell \mathbb{1}\{I_t = \ell\}}_{unobserved}. \tag{3}$$

In each period t the household decides on real consumption, c_t, and real cash holdings, m_t. These two choice variables are continuous and observed by the econometrician. If the household does not yet have an ATM card ($I_{t-1} = 0$), it decides whether to adopt or not, that is, chooses $I_t \in \{0, 1\}$.

Let A denote the household's life span. The household discounts future payoffs with discount factor $\beta \in [0, 1)$. A household's state is fully described by $(I_{t-1}, z_t, \epsilon_t)$. The household chooses the sequence of ATM adoption decision, consumption, and money holding $\{I_\tau, c_\tau, m_\tau\}, \tau = t, \ldots, A$ to maximize the discounted sum of future pay-offs or the value function:

$$W_t(I_{t-1}, z_t, \epsilon_t,) = \max_{\{I_\tau, c_\tau, m_\tau\}} E_t \left(\sum_{\tau=t}^{A} \beta^{\tau-t} \left[u(I_\tau, c_\tau, m_\tau | z_\tau) - \kappa_\tau(I_\tau, I_{\tau-1}, \epsilon_t) \right] . \right) \tag{4}$$

The data support the assumption that ATM card adoption is irreversible. Of the observed households who appear in the data more than once, only 920 or 11.6% appear to adopt the card in one period and then report not having it in a later period. For 662 of these 920 households, at least one of the following variables is not reported consistently across time periods: the age, gender or education level of the household head, the region where the household resides or the number of adults living in the household. The inconsistent adoption pattern of these households may therefore be explained by their inconsistent overall reporting or a change in the reporting household member from one period to the next.

4.1 Intratemporal Euler equations

We exploit the fact, that conditional adoption choice I_t, consumption choices c_t, and real cash holdings m_t have to satisfy the standard Euler equations. In particular, the opportunity cost of holding real cash balances is the interest that could be earned when holding the deposits on an interest bearing checking account. Let r_t denote the nominal interest rate.The intratemporal Euler condition is given by:

$$u_m(I_t, c_t, m_t) = r_t u_c(I_t, c_t, m_t). \tag{5}$$

In this equation, u_c and u_m are the first derivatives of u with respect to consumption c and cash holdings m and r denotes the nominal net interest rate. Equation (5) can be derived following Carlstrom and Fuerst [2001]. Let the nominal cash balance available for transaction purposes in period t be given $\mathcal{M}_t$, and the price level be P_t. Real cash balances available for consumption purposes are thus given by $M_t = \mathcal{M}_t/P_t$). The household begins the period with M_t nominal balances and D_{t-1} holdings of deposits on which it earned nominal interest rate r_{t-1}. It receives nominal income Y_t (including income from wealth). Before making consumption choices, the household decides how much cash to keep for consumption purposes, such that $\mathcal{M}_t = M_t + (1+r_{t-1})D_{t-1} - D_t + Y_t$. The household ends the period with balances given by the intertemporal budget constraint $M_{t+1} = M_t + (1 + r_{t-1})D_{t-1} - D_t + Y_t - P_t c_t$. By substituting $\mathcal{M}_t$ into (4) and maximizing subject to the intertemporal budget constraint with respect to c_t, D_t, and M_{t+1} we obtain (5). Through this intratemporal Euler condition, the dimension of the per-period decision problem condition on ATM adoption status has been reduced to one from two. Given c and I, we can solve for the optimal m

using equation 5.

4.2 The adoption decision

To study the optimal dynamic decision of adopting an ATM card, it useful to define the conditional choice value function (see for example Hotz and Miller [1993] as the value of choosing I_t net of the choice specific shock ε_t^ℓ

$$v(I_t, z_t) = u(I_t, c_t, m_t|z_t) - \mathbb{1}\{I_{t-1} = 0\}(I_t - I_{t-1})\bar{\kappa} + \beta\left(E(W(I_t, \epsilon_{t+1}, z_{t+1}|z_t)\right). \tag{6}$$

where the expectation is taken by integrating out the stochastic components of the state variable z_{t+1} and the future adoption shocks, conditional on the state variable z_t and the choice I_t..

Recall that after households adopt the technology, they keep it forever or until the end of the decision problem's time horizon. Let $p(I_t|I_{t-1}, z_t)$ denote probability of adopting conditional and adoption status I_{t-1} and state z_t.

By irreversibility, the probability of "unadopting" is zero ($p(I_t = 0|I_t = 1, z_t) = 0$) and the probability of keeping the ATM card is one ($p(1|z_t, 1) = 1$). For this reason, the probability of interest is $p(I_t = 1|I_t = 0, z_t)$, which corresponds to $1 - p(I_t = 0|I_t = 0, z_t)$. To economize on notation we will from now abbreviate $p_t = p(I_t = 1|I_t = 0, z_t)$. Assume that the household has not adopted the ATM card prior to time t, thus $I_{t-1} = 0$. At time t, the household makes the adoption decision based on maximizing $W(0, \epsilon_0, z_t)$ for the revealed choice specific error terms ϵ_t^0 and ϵ_t^1, that is by comparing:

$$V^1(\epsilon_t, z_t) := v(1, z_t) - \bar{\kappa} - \sigma_\kappa \epsilon_t^1 \tag{7}$$

$$V^0(\epsilon_t, z_t) := v(0, z_t) - \sigma_\kappa \epsilon_t^0 \tag{8}$$

The adoption rule $V^1(\epsilon_t, z_t) > V^0(\epsilon_t, z_t)$ can be re-written in terms of the adoption cost and the conditional value functions:

$$\begin{aligned} & V^1(\epsilon_t, z_t) > V^0(\epsilon_t, z_t) \\ \Leftrightarrow\quad & v(z_t, 1) - \bar{\kappa} - v(z_t, 0) > \sigma_\kappa\big(\epsilon_t^1 - \epsilon_t^0\big). \end{aligned} \tag{9}$$

Keeping in mind that the right hand side of (9) is a random variable whose distribution is known up to its scale σ_κ, the probability of adoption is given by:

$$p_t = \mathrm{prob}\left[\frac{v(1, z_t) - \bar{\kappa} - v(0, z_t)}{\sigma_\kappa} \geq \epsilon_t^1 - \epsilon_t^0\right]. \tag{10}$$

After setting up the household's decision problem, we will now work through the econometric building blocks and the estimation procedure.

5 Econometric Building Blocks

Our simulation estimator is constructed from five building blocks. We first show how a static optimality condition of the household problem in 4 yields the demand for real cash balances which allows us to recover the parameters of the period utility function. Second, we describe how we estimate the transitions for exogenous variables such as the regional interest rate and the inflation rate as well as the consumption process. We then state key assumptions on the dynamic choice problem for the household's ATM adoption and the distribution how we estimate the adoption costs. Last, we estimate transitions for endogenous variables—consumption, wealth, and ATM-card adoption—denoting these transitions f. The next four sub-sections describe these steps in turn. The last and fifth section describes the simulator and its implementation.

To motivate the per-period utility function presented in the next sections, we provide evidence that adoption is directly associated with changes in money holding. Figure 1 plots the money-consumption (mr/c) and wealth-consumption (w/c) ratios over sequences of three waves of the SHIW for the adopters, denoted by (0,1,1), the always adopters, denoted by (1,1,1), and the never-adopters, denoted by (0,0,0). The plots apply to three time windows: 1991–1993–1995 (denoted W1), 1998–2000–2002 (denoted W2), and 2000–2002–2004 (denoted W3). The top panel shows the ratio mr/c. It illustrates that the never-adopters have the highest ratios, followed by adopters, and then the always-adopters. In each window, the ratio mr/c is decreasing, consistent with earlier analysis that the overall mr/c ratio is falling over time. The bottom panel shows the w/c ratio for the same households. Comparing the three groups in the top panel suggests that adoption per se is associated with a fall in money holding relative to consumption (in time periods W1 and W2), and economizing on money balances, which will raise utility. In addition, future adopters already hold less money compared to consumption than those who will not adopt in the next period, as can be seen by comparing (0, 1,1) to (0,0,0).

In our shopping time model, we specifiy the per-period utility function $u(c, m, I)$ of each household i in period t as follows:

$$u(c_{i,t}, m_{i,t}, I_{i,t}) = (1 + \gamma I_{it})^{\omega} \frac{c_{i,t}^{1-\alpha} - 1}{1-\alpha} + e^{\omega \cdot (\eta_i + \tau_t)} \frac{m_{it}^{1-\omega} - 1}{1-\omega} \tag{11}$$

One verifies immediately that for $\alpha \geq 1$ and $\omega \geq 1$, it is true that $u_c > 0, u_m > 0$ and $u_{cc} < 0, u_{mm} < 0$ which are the usual montonicity and concavity conditions. With the specification in (11), u is additively separable in consumption c and cash holdings m since it consists of one summand that varies in c and one summand that varies in m. Each of those summands has the constant relative risk aversion (CRRA) shape with risk aversion parameters α for the consumption part and ω for the cash holdings part. The utility function also depends on household and period specific parameters η_i and τ_t. The variable I_{it} equals 1 if the household has adopted an ATM card and 0, if not. ATM card adoption leads to increased utility from consumption due to the "technology parameter" $\gamma > 0$. In the specification (11), we allow for individual and time varying heterogeneity in the parameters. The inclusion of individual level fixed-effects η_i is motivated by the observation that future adoption is correlated with lower cash balances relative to consumption, as documented in Figure 1.

The estimation of the parameters of the utility function begins with the intra-temporal Euler conditions:

$$\begin{aligned} r u_c &= u_m \\ r_{it} \cdot \left((1 + \gamma_i I_{it})^{\omega} \cdot c_{it}^{-\alpha} \right) &= e^{\omega(\eta_i + \tau_t)} m_{it}^{-\omega} \end{aligned}$$

Recall that the adoption costs do not interact with c and m in the per-period utility and hence would disappear in the partial derivatives u_c and u_m.

Taking logarithms on both sides then gives

$$\ln(r_{it}) + \omega \cdot \ln(1 + \gamma_i) I_{it} - \alpha \ln(c_{it}) = \omega \cdot (\eta_i + \tau_t) - \omega \ln(m_{it}). \tag{12}$$

For general $\alpha, \omega \in (1, \infty]$, isolating $\ln(m)$ on the right hand side yields

$$\ln(m_{it}) = -\frac{\ln(r_{it})}{\omega} + \frac{\alpha \ln(c_{it})}{\omega} + \eta_i + \tau_t - \ln(1 + \gamma_i) I_{it} \tag{13}$$

In the limiting case $\alpha = \omega = 1$, the utility function is still well-defined since $\lim_{\alpha \to 1} \frac{x^{1-\alpha} - 1}{1 - \alpha} = \ln(x)$ for all positive values of x.The intratemporal Euler equation (12) simplifies to

$$\ln \left(\frac{m_{it} r_{it}}{c_{it}} \right) = \eta_i + \tau_t - \ln(1 + \gamma_i) I_{it} \tag{14}$$

As discussed in Section 4, the Euler equations reduce the dimension of per-period decision problem. Given consumption c and ATM card adoption I, equations (14) and (13) implicitly define the optimal level of cash holdings m.

5.1 State variables and transition functions

The model three types of observable state variables that make up the vector (I_{t-1}, z_t). The first group are time-invariant or static state variables which describe the household composition, education level, and place of residence and which do not change from period to period. The second group is deterministic and time-varying and includes the age of the household head and their employment status. Age advances by the length of the time period between t and $t+1$, in the case of the SHIW by 2 years. Employment status is maintained until age 65 when the household head retires. The third group includes time-varying stochastic state variables, namely interest rates and inflation, wealth w_t and adoption status I_t.

For all transitions, time is measured in two-year steps and is indexed by t, since the survey is sampled every two years (three years between 1995 and 1998). We need to specify transition functions for the third group, the time-varying stochastic state variables.

5.1.1 Interest-Rate and Inflation Process

The inflation rate π_t is the year-to-year growth rate of the consumer price index, from 1989 to 2010. Interest rates r_t are regional nominal deposit rates for each of the twenty administrative regions of Italy.

To parametrize the transition function for $\{\pi_t, r_t\}$ we use ordered VARs and test the lag length with standard information criteria. It makes sense to penalize models with large numbers of parameters given the short time-series sample. Since we have 2-year household survey data, it makes sense to estimate to estimate transitions in two year steps. We work with natural logarithms to guarantee positive, simulated interest rates and inflation rates.

We find that inflation can be described autonomously:

$$\ln \pi_t = a_0 + a_1 \ln \pi_{t-2} + \epsilon_{\pi t}, \tag{15}$$

with $\epsilon_{\pi_t} \sim IID(0, \sigma_\pi^2)$. In each region the deposit rate is well-described by:

$$\ln r_t = b_0 + b_1 \ln r_{t-2} + b_2 \ln \pi_t + \epsilon_{rt}, \tag{16}$$

with $\epsilon_{rt} \sim IID(0, \sigma_r^2)$. This setup ensures that $cov(\epsilon_{\pi t}, \epsilon_{rt}) = 0$ (which simplifies simulations). We use this specific ordering because it fits with the difference in the time periods to which the inflation rate and interest rate in a given year apply.

We estimate the r-equation for each of the twenty regions and report the average estimates over this set (rather than averaging the interest rates, which would lead to an understatement of uncertainty in a typical region). In practice, though, the variation in estimates across regions is quite small.

Table 2 contains the estimates for the parameters, their standard errors, and the two residual variances. Later, we will assume that $\{\epsilon_{\pi t}, \epsilon_{rt}\}$ jointly normal and, with this ordering, the two shocks are uncorrelated.

5.1.2 Wealth and the consumption decision

We estimate reduced form equations for the wealth transition $(w_{t+1}) = f^w(w_t, z_t)$ and the consumption decision $c_t = f^c(z_t, I_t)$. In our model, wealth will depend on previous period wealth, but the consumption decision does not depend on lagged variables. Graphically, the bottom panel of 1 the same three groups in the bottom panel suggests that adoption is associated with a rise in financial wealth relative to consumption in W1 and W3 but not W2. We do observe less variation over time in financial wealth relative to consumption for non-adopters. We account for these observations in two ways: first, we include the fixed effects from equations 13 and 14 in the consumptions and wealth transitiuon, and second, we assess whether the transitions for adopters and non-adopters are statistically different.

We use ordinary least squares (OLS) regression of the logarithm of current period wealth on the observable state variables z_t and previous period wealth. We show that the estimated transitions for adopters and non-adopters are not statistically different. We thus pool the sample to estimate the transition function f^w which this does not depend on the ATM adoption status.[2]

Using the same wealth transition for adopters and non-adopter ensures that the discrete choice process has the "renewal property" (see Rust [1987]). If the state variables evolve independently of the ATM card decision, then $E(u(I_t, c_t, m_t))$ depends on I_t, but not on the adoption history $I_1, \ldots, I_{t-1}$. The "renewal property" could potentially be used to generalize to a setting where adoption is not a terminal state.

Given their state z_t, households decide how much to consume per period. For flexibility, we estimate separate decision functions for adopters on non adopters. The decision function $f^{c,I}$ is computed using OLS regression of the logarithm of consumption c_t on observable state variables z_t, including wealth w_t. Furthermore, the fixed effects η_i from Section 5.2 also enter the consumption equation as independent variables. These reduced form equations allow to re-parametrize the utility function as $u(I_t, z_t)$ and also compute $u(0, z_t) - u(1, z_t)$ in (18).

5.2 The dynamic discrete choice process

The dynamic discrete choice (DDC) process optimal stopping problem is tractable under certain assumptions. This section states these assumptions and then specifies the functional

[2]Fractional polynomials did not yield a large improvement in fit relative to the added complexity. The results are available upon request.

form of the model components and the parameters to be estimated.

The first assumption is additive separability of the per-period utility in the observables and unobservables, see (2). The second assumption is that the random variables $\kappa(I_t|I_{t-1})$ are independently and identically distributed over time with probability density function g. The third assumption is condition independence which means that the state variables z_t follow a Markov process that is not affected by the unobservable adoption cost κ_t.To fulfill this assumption, it suffices that the probability density function of the state variable z_t has the property $f(z_{t+1}|I_t, \epsilon_{t+1}, z_t) = f(z_{t+1}|z_t)$.

Typically, κ is assumed to follow a parametric distribution and the structural parameters to be estimated are the parameters characterizing this distribution. Since adoption probabilities depend on the difference $\epsilon_t^1 - \epsilon_t^0$, but not on the individual error terms, there is no loss of generality in normalizing $\epsilon_t^0 = 0$. These normalization mean that only one choice specific error term, namely ϵ_t^1, is left. We will therefore denote this error by $\sigma_\kappa \epsilon$ where σ_κ is the variance of κ relative to a chosen distribution for ϵ.

In this paper, ϵ is a normal random variable. We conducted robustness checks for a standard logistic random variable and found that the results were quantitatively and qualitatively very similar.

If $v(1, z_t) - v(0, z_t)$ is known and adoption choice at t are observed then the structural parameters can be estimated from the condition (10). Note that $\kappa_\tau, \tau > t$ enters $v(z_t, 0)$ since the household may adopt later and then pay the adoption cost. In the following section, we show that the terminal choice assumption addresses this recursivity problem.

We now describe how the conditional value functions $v(I, z)$ are estimated from the data. The key points are first, to exploit the assumption that ATM card adoption is a terminal choice and second, to show that the structural parameters are identified.

We use the Euler equations in Aguirregabiria and Magesan [2013] for the ATM card adoption decision problem. From these Euler equations we retrieve moment conditions and the Likelihood. Detailed proofs are available as an appendix. For the remainder of this section, we will re-parametrize the utility function in terms of the exogeneous state vector $u(I, c, m) = u(I, c(z, I), m(c(z), I)) = u(I, z)$. The functional forms of $c(z)$ and $m(c, I)$ will be derived further below.

As before ϵ follows a normal distribution with mean 0 and variance 1, corresponding to a dynamic probit model. We noe that the Euler equation for a terminal choice in the logit case also follows from Arcidiacono and Miller [2011]

$$\begin{aligned} 0 = & \big(u(1, z_t) - (1-\beta^2)\bar{\kappa} - u(0, z_t)\big) - \sigma_\kappa \Phi^{-1}(p_t) \\ & - \beta^2 \sigma_\kappa \int \big(\phi(\Phi^{-1}(p_{t+1})) - (1-p_{t+1})\Phi^{-1}(p_{t+1})\big) f(z_{t+1}|z_t) \mathrm{d}z_{t+1} \end{aligned} \tag{17}$$

From this Euler equation, $\Phi^{-1}(p_t)$ can be expressed as:

$$\Phi^{-1}(p_t) = \frac{u(1, z_{t+1}) - u(0, z_{t+1}) - (1-\beta^2)\bar{\kappa}}{\sigma_\kappa} + \beta^2 E\left((1-p_{t+1})\Phi^{-1}(p_{t+1}) - \phi\left(\Phi^{-1}(p_{t+1})\right)\right), \tag{18}$$

and if we apply the injective function $\Phi(\cdot)$ to both sides:

$$p_t = \Phi\left(\frac{u(1, z_t) - u(0, z_t) - (1-\beta^2)\bar{\kappa}}{\sigma_\kappa} + \beta^2 E\left((1-p_{t+1})\Phi^{-1}(p_{t+1}) - \phi\left(\Phi^{-1}(p_{t+1})\right)\right)\right). \tag{19}$$

In what follows, we will refer to the equation (18) as "**Lin**ear specification" and the equation (19) as "**MLE** specification."

Now, if we are able to compute $u(1,z)-u(0,z)$, p, and $E((1-p_{t+1})\Phi^{-1}(p_t)-\phi(\Phi^{-1}(p_{t+1})))$ we can fit a linear model or a fractional probit model to obtain estimates of the structure parameters κ, σ_κ and also the discount factor β^2. Specifically, our estimation strategy will be as follows: First, we attach an error term to the equations for "Lin" namely (18).[3] We then use a least squares estimator. Second, we maximize the pseudolikelihood corresponding to equation (19), the "MLE specification." This is the same maximization problem as for a fractional logit or probit regression.

The next section will provide the building blocks for for $u(1,z)-u(0,z)$, p, and $E((1-p_{t+1})\Phi^{-1}(p_t) - \phi(\Phi^{-1}(p_{t+1})))$. [4]

The parameters of interest are thus $\alpha, \omega, \eta_i, \tau_t$ and γ_i. In the data, c, m, r and I are observed. To estimate the parameters of interest, we attach an error term to the money demand equations (13) and (14). Since the money demand equation contains a household level fixed effect on the right hand side, an linear model with within-household fixed effects is appropriate. Since the same household is observed multiple times and some households are observed before and after adoption, the coefficient of I_{it} is identified. The results of the regressions for the intratemporal Euler equations (14) and (13) are summarized in Table 3.

We find that allowing γ to vary with observable household characteristics does not improve model fit as measured by R^2. In what follows, we will always use the specification with

[3]We use the linear specification with the log odds ratio (resp. the inverse cumulative normal of the adoption probability) on the left hand side which becomes the dependent variable in the linear model. As discussed in Aguirregabiria and Magesan [2013], it is also possible to re-arrange this equation so that U appears on the left hand side. We conducted a detailed simulation study based in which the former specification converges faster. Details are available upon request.

[4]To highlight the computational gains, $\frac{u(z,1)-u(z,0)-(1-\beta^2)\bar{\kappa}}{\sigma_\kappa} + \beta^2 E(\cdot)$ is equal to $EV_{(\bar{\kappa},\sigma_\kappa)}$. Thus, with the notation in Rust [2000], we can compute, up to simulating p', $EV_{(\bar{\kappa},\sigma_\kappa)}$. Indeed, the algorithm has been reduced to estimating the structural parameters only using the "outer" steps of the fixed point algorithm which maximizes the partial maximum likelihood for MLE in Table 5 (ie solving a fractional probit or logit model).

constant γ as reported in the first row of Table 3.

5.3 ATM card adoption

Finally, the probability to adopt an ATM card in period t comes from a static binary probit model of the dichotomous variable I_t on the deterministic state variables and real regional interest rates. Thus we can compute p_t. Since the right hand side of the reduced form for p_t contains stochastic state variables (inflation and regional interest rates), we need to integrate them out to obtain the offset term $E((1-p_{t+1})\Phi^{-1}(p_{t+1}) - \phi(\Phi^{-1}(p_{t+1})))$. We use the following Monte Carlo algorithm:

First, S denotes number of Monte Carlo draws, Second, N is the number of households in the sample for which $I = 0$ and I' is observed, and L is the number of distinct regions that these households live in. Third, as before z and z_{t+1} are the state variables. To be clear on the timing, the households are observed at two time intervals $t-1$ and t.

For the algorithm, the vector z of state variables is split into three vectors $z = (z_1, \pi, r, w)$ where z_1 consists of all deterministic and static state variables, and the remaining state variables are Markovian, in particular, π is the inflation rate, r the regional deposit rate and w is household wealth.

Algorithm The expected value $E((1-p_{t+1})\Phi^{-1}(p_{t+1}) - \phi(\Phi^{-1}(p_{t+1})))$ is approximated in three steps:

1. Obtain S draws of the state variable z_{t+1}:
   (a) Update the deterministic components of $(z_1)_{t+1}$ according to the rules for age and employment.
   (b) Draw S shocks ϵ_π^s for the inflation process. Simulate S paths for inflation π_{t+1} as $(\ln \pi_{t+1})^s = a_0 + a_1\pi + \sigma_\pi \epsilon_\pi^s$.
   (c) Draw S shocks $\epsilon_{r,l}^s$ for each of the L regional deposit rates. and simulate $\ln(r_{l,t+1})^s = b_0 + b_1 \ln r_{l,t} + b_2 \ln \pi_{t+1}^s + \sigma_r \epsilon_{r,l}^s$
   (d) Draw S shocks $\epsilon_{w,i}^s$ for each for the wealth processes of the N households. Simulate $S \times N$ paths of the wealth process $w_{i,t+1} = f^w(z_i) + \sigma_w \epsilon_{i,w}^s$.
   (e) Define $(z_{t+1})^s = \big((z_{t+1})_1, \pi_{t+1}^s, r_{t+1}^s, w_{t+1}^s\big)$
2. Compute $p_{t+1}^s = p(z_{t+1}^s)$ using the coefficients of the static binary probit model.

3. Output:

$$E((1-p_{t+1})\Phi^{-1}(p_t)-\phi(\Phi^{-1}(p_{t+1}))) \cong \frac{1}{S}\sum((1-p^s_{t+1})\Phi^{-1}(p^s_{t+1})-\phi(\Phi^{-1}(p^s_{t+1}))). \tag{20}$$

6 Results

In this section, we present the estimation result for the structural parameters and the discount factor β^2. We also provide an estimate of the monetary compensation that a household would require to adopt. We use the estimation strategy outlined at the end of Section 5.3 and obtain four sets of results, corresponding to two different functional forms of the utility (CARA and CRRA) as well as two distinct estimators (LIN and MLE).

6.1 Parameter estimates

Table 5 summarizes results for Lin and MLE, and CARA and CRRA utilities, for a total of 4 different specifications. We first obtain the coefficient vector $(1/\sigma_\kappa, (1-\beta^2)\kappa/\sigma_\kappa, \beta^2)$, subject to $\sigma_\kappa > 0$ and $\beta^2 \in [0,1)$. We transform these estimates into $\kappa, \sigma_\kappa, \beta^2$, provided $\beta^2 \neq 1$. Confidence intervals are obtained from a parametric bootstrap. Note that, in order to present confidence intervals, we include estimates of the discounted adoption cost $\bar{\kappa}(1-\beta^2)$ instead of $\bar{\kappa}$ in the table. The reason is that the upper bound of the confidence interval for $1-\beta^2$ is close to 0 while the estimate $\bar{\kappa}(1-\beta^2)$ is bounded away from 0.[5]. In addition to estimating $\bar{\kappa}, \sigma_\kappa$ and β^2,, we also estimate $\bar{\kappa}$ and σ_κ for fixed $\beta^2 \in [0,1)$. We find that the estimates in Table 5 differ somewhat across the specification choices for teh estimators (Lin or MLE), but that different specifications of the per-period utility functions lead to even larger differences in estimated parameters.

For each parameter, we discuss CARA utilities, then the more general CRRA utilities, followed by a comparison between the two.

Starting with the point estimates for $(1-\beta^2)\bar{\kappa}$, CARA utility estimates are around 2.9 for Lin and MLE, and the 95% bootstrap confidence intervals have approximately the same width (around 1.35) and shape. In the CRRA case, Lin and MLE estimates are very close (7.65) and differ only after the second decimal. The width of the confidence interval orders of magnitude larger than point estimate (over 200 units). The confidence intervals are also not symmetric, with the right end point being much further away from the point estimate than the left end point. Comparing the CARA and CARRA specifications, CARA gives

[5]It is possible that $\hat{\bar{\kappa}}$ does not have well-defined second order moment, similar to a Cauchy distribution or the certain ratios of two normally distributed random variables

lower estimates that also appear to be more precise. We observe that $(1-\beta^2)\bar{\kappa}$ and $\bar{\kappa}$ are senstive to assumptions about β^2.

The parameter σ_κ is estimated at 0.82 (0.86) for Lin and MLE . For the CRRA utility, they are very close,and only differ in the third decimal. The bootstrap confidence intervals for the estimate σ_κ are large compared to estimate itself. While they are bounded away from zero at the lower end, the width of the 95% bootstrap confidence interval is around 3.5 the size of parameter estimate for CARA utilities and 500 times the size of parameter estimate for CRRA utilities. A likely explanation is that our estimators involve transforming an estimate by inversion $(1/\sigma_\kappa)$. Thus, small ranges of the untransformed estimate can translate into large ranges when the untransformed estimates are close to zero.

With regards to β^2, it appears that the households are less myopic (larger β^2) under constant absolute risk aversion (CARA) than under constant relative risk aversion (CRRA). Specifically, the squared annual discount factor β^2 is around $0.97-0.98$ for CARA and around $0.78-0.79$ in the CRRA case. We note however that our estimator only use one period ahead and thus cannot separate discounting from present-bias.

Since most previous literature has relied on calibrated values of β^2 or investigated a range of values for β^2 (as in Yang and Ching [2014]) it is instructive to estimate the structural parameters for fixed β^2. These estimates are shown in Figure 2.

We note that the estimates for $\bar{\kappa}$ are large compared to the utility from consumption. For example, with CARA utility and a typical consumption of €20,000, the utility of a non-adopter is around 10 units while the estimate for $\bar{\kappa}$ is close to 100 units.

That $\bar{\kappa}$ increases with β^2 could be driven by the observation that $\bar{\kappa}$ enters as $(1-\beta^2)\bar{\kappa}$. To illustrate this dependency, we plot $(1-\beta^2)\bar{\kappa}$ in Figure 3. We see that this discounted value of $\bar{\kappa}$ decreases in β^2 for all eight model specifications. Furthermore, the rate of decrease is constant in β^2 for the CARA utility, but accelerates in β^2 for the CRRA utility.

Myopic consumers might perceive adoption as more costly. Adopting earlier would allow the household to reap greater benefits in the future which become more important the larger the discount factor β is. This is supported by the observation that the CRRA specifications show a stronger increase of the discounted adoption cost with respect to β^2.

6.2 Compensating variation

In this section, we use the structural estimates from previous period to estimate the required on-time per-household subsidy to encourage adoption in the current period. This exercise can be viewed as a counterfactual were the government provides We follow the approach of Cooley and Hansen [1989] and Goolsbee and Klenow [2006] in obtaining a measure in monetary units (here: Euros) to answer the following question

How much do we have to compensate Italian households in consumption so that they are indifferent between adoption and not adopting an ATM card in the current period?

The compensating consumption CV is defined implicitly as

$$u(1, c^1 + CV, m^1(c)) - \bar{\kappa} + \beta EW(z_{t+1}|I_t = 1) \quad = \quad u(0, c^0, m^0(c)) + \beta EW(z_{t+1}|I_t = 0)$$

Here the left hand side represents the expected future value of adopting if the household's consumption were increased by CV in the initial period and the right hand side represents the future value of not adopting. Note that right hand side is $v(0, z_t)$

From the definition of $v(z, 1)$ and $v(z, 0)$, the equation can be re-written as

$$\begin{aligned}
u(1, c^1 + CV, m^1(c^1)) - \bar{\kappa} + \beta EW(z_{t+1}|I_t = 1) &= v(0, z_t) \\
u(1, c^1, m^1(c^1)) - \bar{\kappa} + \beta EW(z_{t+1}|I_t = 1) - (1, u(c^1, m^1) - u(1, c^1 + CV, m^1)) &= v(0, z_t) \\
v(1, z_t) - \bar{\kappa} - u(1, c^1, m^1) - u(1, c^1, m^1)) &= v(0, z_t) \\
v(1, z_t) - \bar{\kappa} - v(0, z_t) = u(1, c^1, m^1) &- u(1, c^1 + CV, m^1)
\end{aligned}$$

Here, $c^I = f^c(I, z_t)$ and $m^I = m(I, c^I)$ Recall that $I_t = 1$ is the optimal choice if and only if $v(1, z_t) - \kappa - \sigma_\kappa \epsilon^1 > v(0, z_t) - \sigma_\kappa \epsilon^0$; thus there exists a unique $\bar{\epsilon}_\kappa$ such that

$$\begin{aligned}
v(1, z_t) - \bar{\kappa} - \sigma_\kappa \bar{\epsilon}_\kappa &= v(0, z_t) \\
\iff v(1, z_t) - \bar{\kappa} - v(0, z_t) &= \sigma_\kappa \bar{\epsilon}_\kappa \\
\iff v(1, z_t) - \bar{\kappa} - v(0, z_t) &= \sigma_\kappa F^{-1}(p_t)
\end{aligned}$$

where F^{-1} is the mapping from probabilities to differences in expected future values, induced by (10) Thus CV is implicitly defined by

$$u(1, c^1 + CV, m^1) - u(1, c^1, m^1) \quad = \quad -\sigma_\kappa F^{-1}(p_t) \tag{21}$$

(22)

To be able to derive meaningful results for variety of functional forms of u, we linearize u and compute the compensating variation as

$$LCV = \frac{-\sigma_\kappa F^{-1}(p_t)}{u(1, c^1, 0)} c^1. \tag{23}$$

For CRRA utilties, we thus obtain

$$LCV = \frac{(1-\alpha)\sigma_\kappa F^{-1}(p)}{(1+\gamma)^\omega(c^{1-\alpha}-1)}c \tag{24}$$

Our choice of linearizing u in this fashion is motivated by Goolsbee and Klenow [2006] who show that logarithmic demand functions can lead to large estimates for the consumer surplus of technological adoption when compared to e.g. linear demand functions. The reason is that logarithmic demand functions, CARA and CRRA utility functions are steep at small values and flat at large values. [6] Note that in the special case of a linear utility function, $u = \alpha c$, $CV = LCV = \alpha\sigma_\kappa F^{-1}(p)$. Figure 4 illustrates the compensating variation measure in both cases for a convex utility function.

We now provide numerical estimates of the compensating variation. First, the two bottom panels of Table 5 show the average value of LCV for different age groups, educational attainment and regions. We also provide confidence intervals from a bootstrap procedure. For CARA utility, point estimates for LCV range from €786 (highly educated) to €1810 (basic education), when taken across all specifications. For CARA utility, they range from €52 (highly educated) to €157 (basic education). In general, LCV increases with age, decreases with educational attainment and is lowest in the North of Italy and highest in the South. We also observe that confidence intervals are narrower for CARA than for CRRA.

In addition to these point estimates and their confidence intervals, we also illustrate how LCV varies within the population. Violin plots, which are an extension of boxplot, are a useful tool to visualize the distribution of LCV. The violinplots in Figure 5, 6 and 7 are computed from the point estimates in Table 5. They show that the distribution of LCV is qualitative similar for all specifications. In particular, we observe that oldest age group has a bimodal distribution of LCV, hinting at heterogeneity within this group.

We use the estimates of LCV from the MLE specification to compute the counterfactual adoption rates in the case of subsidies of 10, 50, 100 and 200 Euros for the CRRA preferences. We assume that the household will adopt when the subsidy exceeds LCV. As shown in Table 6, most of the older, less educated and those living in the South require subsidies exceeding €100 to adopt immediately. On the other hand, those with high education require lower incentives. Overall, incentives of around €100 would be required to get about half of

[6] This linearization strategy also overcomes the following challenge: Because u is monotonically increasing in c, there exists at most one solution for CV 21. However, since u is not surjective, the real solution set could be empty. For the CARA utility, we can obtain exact real solutions of (21), that is a number in $\mathbb{R}$. For CRRA utilities, this inversion may lead to complex solutions. We thus rely on the first order approximation.To find CV for a variety of functional forms of u, we could use a first order Taylor approximation at c^1 to solve for CV. That is we approximate $u(1, c^1 + CV, m^1) \equiv u(1, c^1, m^1) + (\partial u/\partial c)|_{c=c^1} CV$. Since the derivative of the utility function with respect to c is $\partial u/\partial c = (1+\gamma)^\omega c^{-\alpha}$, approximately $u(1, c+CV, m) - u(c, m, 1) \equiv (1+\gamma)^\omega c^{-\alpha} CV$, and thus $CV \cong \frac{-\sigma_\kappa F^{-1}(p) c^\alpha}{(1+\gamma)^\omega}$.

the non-ATM card adopters to adopt the cards immediately.

7 Conclusion

This paper provides an application of discrete dynamic choice (DDC) models to the adoption of financial innovation, contributing insights both to the literature on the identification of DDC models and the technology adoption in the banking sector. Our method is applicable to a range of additional financial adoption decisions.

Our paper uses several recent advances to modify the conditional choice simulation estimator of Hotz et al. [1994]. In addition of exploiting the finite dependence property of the technology adoption problem(Arcidiacono and Ellickson [2011]), we also implement the Euler equations in Arcidiacono and Miller [2011] and estimation of the discount parameter from Abbring and Daljord [2020]. In doing so, we reduce the computational complexity of the DDC model. As a byproduct, we provide a closed form for dynamic probit model similar to the often-cited formula for the dynamic logit model.

A key feature of the economic environment is the return or utility function. That is based on a shopping-time model of money demand, with two distinctive features. First, it allows for a gradual diffusion of bank branches and ATM machines between 1989 and 2004, which enhanced the efficiency of money holding (and so reduced the ratio of money to consumption) for both card-holders and non-card-holders. Second, it includes a parameter (γ) that isolates the additional degree to which card-holders economized on money holding. We estimate these features of money demand via the Euler equations in a first step, using data from more than 52,000 household-year observations. We also estimate transitions for consumption and wealth for both groups. We can therefore compute per-period utilities up to adoption cost. The adoption cost is then estimated from the dynamic choice model.

While we find that the discounted adoption cost is larger for myopic households, we provide empirical evidence that households have discount factors between 0.88 and 0.99, hence are not myopic when deciding whether to adopt financial innovation. Furthermore, if household are have constant relative risk aversion, the data suggest that they are also somewhat more myopic when compared to the case of constant absolute risk aversion. This finding is important since many dynamic decision problem take the discount factor, a measure of myopia, as given. It is also consistent with the experimental finding of Andersen et al. [2008].

We compute the compensating variation, that is the amount of consumption units that corresponds to the disutility of adopting, after that adoption costs are balanced against the enhanced efficiency of money holding. In this context, the compensating variation could be interpreted as the size of financial incentive or subsidy to encourage adoption of the financial innovation. The average compensating variation is €90 (CRRA) respectively €1400 (CARA). Older, less educated household in less prosperous regions would need to receive larger financial incentives. These non-trivial costs could explain some of the slow uptake of financial innovations, especially among the elderly, rural or less educated.

References

J. H. Abbring and O. Daljord. Identifying the discount factor in dynamic discrete choice models. *Quantitative Economics*, 11(2):471–501, May 2020. doi: 10.3982/QE1352.

V. Aguirregabiria and A. Magesan. Euler Equations for the Estimation of Dynamic Discrete Choice Structural Models. In *Structural Econometric Models*, volume 31 of *Advances in Econometrics*, pages 3–44. Emerald Group Publishing Limited, April 2013. doi: 10.1108/S0731-9053(2013)0. URL https://ideas.repec.org/h/eme/aecoll/s0731-9053(2013)0000032001.html.

S. Alan. Entry costs and stock market participation over the life cycle. *Review Economic Dynamics*, 9(4):588–611, 2006.

R. Alessie, S. Hochguertel, and A. van Soest. Ownership of stocks and mutual funds: A panel data analysis. *The Review of Economics and Statistics*, 86(3):783–796, August 2004.

F. Alvarez and F. Lippi. Financial innovation and the transactions demand for cash. *Econometrica*, 77(2):363–402, 03 2009.

S. Andersen, G. W. Harrison, M. I. Lau, and E. E. Rutström. Eliciting risk and time preferences. *Econometrica*, 76(3):583–618, 2008.

P. Arcidiacono and P. B. Ellickson. Practical methods for estimation of dynamic discrete choice models. *Annual Review of Economics*, 3(1):363–394, 2011. URL https://EconPapers.repec.org/RePEc:anr:reveco:v:3:y:2011:p:363-394.

P. Arcidiacono and R. A. Miller. Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79(6):1823–1867, 2011. ISSN 1468-0262. doi: 10.3982/ECTA7743. URL http://dx.doi.org/10.3982/ECTA7743.

O. R. Attanasio, L. Guiso, and T. Jappelli. The demand for money, financial innovation, and the welfare cost of inflation: An analysis with household data. *Journal of Political Economy*, 110(2):317–351, 2002.

A. Brandolini and L. Cannari. *Saving and the accumulation of wealth: Essays on Italian household and government saving behavior*, chapter Appendix: The Bank of Italy's Survey of Household Income and Wealth. Cambridge University Press, 1994.

C. T. Carlstrom and T. S. Fuerst. Timing and real indeterminacy in monetary models. *Journal of Monetary Economics*, 47(2):285–298, April 2001. URL https://ideas.repec.org/a/eee/moneco/v47y2001i2p285-298.html.

T. F. Cooley and G. D. Hansen. The Inflation Tax in a Real Business Cycle Model. *American Economic Review*, 79(4):733–748, September 1989. URL https://ideas.repec.org/a/aea/aecrev/v79y1989i4p733-48.html.

F. Gomes and A. Michaelides. Optimal life-cycle asset allocation: Understanding the empirical evidence. *Journal of Finance*, 60(2):869–904, 04 2005.

A. Goolsbee and P. J. Klenow. Valuing consumer products by the time spent using them: An application to the internet. *American Economic Review*, 96(2):108–113, May 2006. doi: 10.1257/000282806777212521. URL https://www.aeaweb.org/articles?id=10.1257/000282806777212521.

L. Guiso, M. Haliassos, and T. Jappelli. Household stockholding in europe: Where do we stand and where do we go? *Economic Policy*, 18(36):123–170, 04 2003. URL http://ideas.repec.org/a/bla/ecpoli/v18y2003i36p123-170.html.

M. Haliassos and A. Michaelides. Portfolio choice and liquidity constraints. *International Economic Review*, 44(1):143–177, February 2003.

V. Hotz and R. Miller. Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies*, 60(3):497–531, 1993.

V. J. Hotz, R. A. Miller, S. Sanders, and J. Smith. A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies*, 61(2):265–89, April 1994. URL http://ideas.repec.org/a/bla/restud/v61y1994i2p265-89.html.

T. Jappelli and L. Pistaferri. The dynamics of household wealth accumulation in italy. *Fiscal Studies*, 21(2):269–295, June 2000. URL http://ideas.repec.org/a/ifs/fistud/v21y2000i2p269-295.html.

F. Lippi and A. Secchi. Technological change and the households' demand for currency. *Journal of Monetary Economics*, 56(2):222–230, March 2009. URL http://ideas.repec.org/a/eee/moneco/v56y2009i2p222-230.html.

B. T. McCallum. *Monetary Economics*. Macmillan, 1989.

R. Miniaci and G. Weber. *Household Portfolios*, chapter Four: Econometric Issues in the Estimation of Household Portfolios. MIT Press, 2002.

A. Pakes. Dynamic structural models, problems and prospects: mixed continuous discrete controls and market interactions. In C. A. Sims, editor, *Advances in Econometrics*, volume 2, pages 171–274. Cambridge University Press, 1994. ISBN 9781139052047. URL http://dx.doi.org/10.1017/CCOL0521444608.005. Cambridge Books Online.

W. Perraudin and B. E. Sørensen. The demand for risky assets: Sample selection and household portfolios. *Journal of Econometrics*, 97:117–144, 2000.

J. Rust. Optimal replacment of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033, 1987.

J. Rust. Nested fixed point algorithm documentation manual, October 2000. URL https://editorialexpress.com/jrust/nfxp.pdf.

G. Sanroman. Estimating heterogeneous costs of participation in the risky asset markets. CEMFI, 2007.

A. Vissing-Jørgensen. Towards an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures. mimeo, University of Chicago, 2002.

A. Vissing-Jørgensen. Perspectives on behavioral finance: Does 'irationality' disappear with wealth? evidence from expectations and actions. *NBER Macroeconomics Annual*, 18:pp. 139–194, 2003. URL http://www.jstor.org/stable/3585252.

C. E. Walsh. *Monetary Theory and Policy*. MIT Press, second edition, 2003.

B. Yang and A. T. Ching. Dynamics of consumer adoption of financial innovation: The case of atm cards. *Management Science*, 60(4):903–922, April 2014.

Tables

Table 1: Descriptive Statistics

Average	1989	1991	1993	1995	1998	2000	2002	2004
Percent with ATM card	15%	29%	34%	40%	49%	52%	56%	58%
Currency holdings(m)								
with ATM card	741	527	398	421	374	349	341	352
without ATM card	696	607	457	498	438	443	458	468
Non-durable consumption(c)								
with ATM card	28395	26517	26430	27242	24463	25796	24941	25343
without ATM card	20585	17757	17237	17299	15341	15405	15138	15899
Financial wealth (w)								
with ATM card	161293	190016	210880	203246	190581	194804	201501	208796
without ATM card	115507	119261	126080	120216	111700	112830	110951	129692
Interest rate (r)								
with ATM card	8.3%	8.7%	9.0%	7.1%	3.2%	2.2%	1.6%	0.4%
without ATM card	8.2%	8.5%	8.9%	7.0%	3.2%	2.2%	1.6%	0.4%
mr/c								
with ATM card	0.23	0.19	0.15	0.12	0.06	0.03	0.03	0.006
without ATM card	0.32	0.33	0.26	0.22	0.10	0.07	0.05	0.013
Observations	8038	7951	7799	7844	6801	7641	7660	7639

Note: Currency holdings, consumption, and wealth are expressed in terms of 2004 Euros. The source is the Bank of Italy's Survey of Household Income and Wealth.

Table 2: Interest-Rate and Inflation Process

$$\begin{aligned} \ln \pi_t &= a_0 + a_1 \ln \pi_{t-2} + \epsilon_{\pi t} \\ \ln r_t &= b_0 + b_1 \ln r_{t-2} + b_2 \ln \pi_t + \epsilon_{rt} \\ & \epsilon_{\pi_t} \sim IIN(0, \sigma_\pi^2) \\ & \epsilon_{rt} \sim IIN(0, \sigma_r^2) \end{aligned}$$

Parameter	Estimate	Standard Error
a_0	0.182	0.256
a_1	0.694	0.208
σ_π^2	0.162	0.054
b_0	-1.185	0.288
b_1	0.826	0.127
b_2	0.947	0.312
σ_r^2	0.314	[0.261, 0.372]

Notes: Estimation uses annual observations from 1989–2010 on CPI inflation and regional deposit rates. The interval in the Standard Error column for σ_r shows the second lowest and second highest value among the twenty per-region regressions.

Table 3: Intratemporal Euler equation for CARA and CRRA preferences

	CARA		CRRA			
	$\bar{\gamma}$	R^2	α	ω	$\bar{\gamma}$	R^2
γ constant	0.27 (0.04)	0.78	1.52(0.87)	6.68 (3.8)	0.23 (0.03)	0.35
γ varies	0.27 (0.03)	0.78	1.76(1.17)	7.70 (5.1)	0.23 (0.06)	0.35
across demographics	p = 0.3				p = 0.3	

Note: This table is derived from the fixed effects regression for the money demand equations (14) and (13). The p-value reports the test that the interaction effects of demographic variables and the ATM adoption status are significant. Standard errors in parentheses. The Stata module -fese- gives standard errors for the fixed effects. The reported standard errors are obtained from Delta method,

Table 4: Marginal effects of the ATM static adoption regression

	Marginal effects	se
Log real rate	6.6685**	2.3983
Age	0.0067*	0.0037
Age Squared	-0.0001**	3.20E-05
Male	0.0174	0.0159
Employed	0.0105	0.0201
Self-employed	-0.0473**	0.0228
No Schooling	-0.2671***	0.0394
Elementary School	-0.1999***	0.0256
Middle School	-0.1306***	0.0248
High School	-0.0207	0.0253
Number of Adults	0.0333***	0.0066
Number of Children	-0.0057	0.0087
Countryside	-0.0680**	0.0289
Town outskirts	0.0038	0.0157
Suburbs	0.0172	0.0153
North-West	0.0555**	0.019
North-East	0.0501**	0.0187
South	-0.0516**	0.018
Islands	-0.0563*	0.0251
Year 1991	-0.3719***	0.0994
Year 1993	-0.5515***	0.1456
Year 1995	-0.3573***	0.0861
Year 2000	-0.1731***	0.044
Year 2002	-0.0823*	0.0327
N	6402	

Note: The outcome variable is ATM card adoption. Based on the sample of households who had not adopted the card in the previous period. The year category for 2004 is the baseline while the year 1998 is dropped since data for 1996 were not collected.

Table 5: Estimators and results

	Linear specification $\Phi^{-1}(p)=\frac{U-(1-\beta)\bar{\kappa}}{\sigma_\kappa}+\beta E\left(S\right)$		**MLE specification** $p=\Phi\left(\frac{U-(1-\beta)\bar{\kappa}}{\sigma_\kappa}+\beta E\left(S\right)\right)$	
	est.	95% bootstrap confidence interval	est.	95% bootstrap confidence interval
CARA				
$\kappa(1-\beta^2)$	2.93	[2.01, 3.38]	2.89	[1.96, 3.30]
σ	0.82	[0.61 , 2.17]	0.86	[0.65, 2.04]
β^2	0.97	[0.57, 1.00)	0.98	[0.51 , 1.00]
CRRA				
$\kappa(1-\beta^2)$	7.65	[4.28 , 215.95]	7.65	[4.05 , 231.71]
σ	0.03	[0.03 , 14.57]	0.03	[0.03 , 17.00]
β^2	0.78	[0.37, 1.00)	0.79	[0.34 , 1.00)
LCV CARA				
Age cohort				
young	952	[594, 2980]	1001	[634, 2702]
middle	1239	[800, 4039]	1303	[843 , 3428]
old	1735	[1180, 5343]	1824	[1222 , 4766]
Education level				
basic	1810	[1241, 5588]	1903	[1288 , 4829]
intermediate	1312	[823 , 4366]	1380	[845 , 3860]
high	786	[360, 2524]	827	[383 , 2388]
Region				
North	1252	[828 , 3884]	1317	[873, 3533]
Centre	1463	[892 , 4753]	1539	[1003, 4296]
South	1601	[1088 , 4792]	1684	[1113, 4400]
LCV CRRA				
Age cohort				
young	62	[36, 2496]	63	[26, 1558]
middle	80	[46, 3005]	82	[37, 2263]
old	110	[61, 3701]	112	[56, 3262]
Education level				
basic	115	[64, 4115]	117	[57, 3143]
intermediate	85	[48 , 2878]	86	[34 , 2320]
high	52	[26 , 1888]	52	[23, 1757]
Region				
North	81	[40, 2713]	82	[39, 2147]
Centre	95	[55, 3897]	97	[40, 2992]
South	101	[58, 3431]	103	[48, 2898]

Note: $U := u(z,1)-u(z,0)$. $S=(1-p_{t+1})\Phi^{-1}(p_{t+1})-\phi\left(\Phi^{-1}(p_{t+1})\right)$. Confidence intervals are obtained from a parametric bootstrap procedure, using the covariance matrices of the coefficient estimates for money demand function, the interest rate and inflation process and the transition functions. Cite Kasahara and Shimotsu?

Table 6: Counterfactual adoption by cohort and subsidy amount

	€10	€20	€50	€100	€200
Age cohort					
young	0.14	0.19	0.39	0.82	1.00
middle	0.08	0.11	0.25	0.70	0.99
old	0.01	0.02	0.06	0.36	0.99
Education level					
basic	0.01	0.01	0.03	0.33	0.98
intermediate	0.03	0.05	0.19	0.67	1.00
high	0.19	0.26	0.48	0.87	1.00
Region					
North	0.10	0.13	0.24	0.64	1.00
Centre	0.06	0.09	0.19	0.51	0.99
South	0.02	0.03	0.11	0.48	0.99

Note: Proportion of sample where LCV is lower than the subsidy amount in the respective column. LCV computed from CRRA and MLE specification.

Figures

Figure 1: ATM Card Adoption, Money-Consumption and Wealth-Consumption Ratio

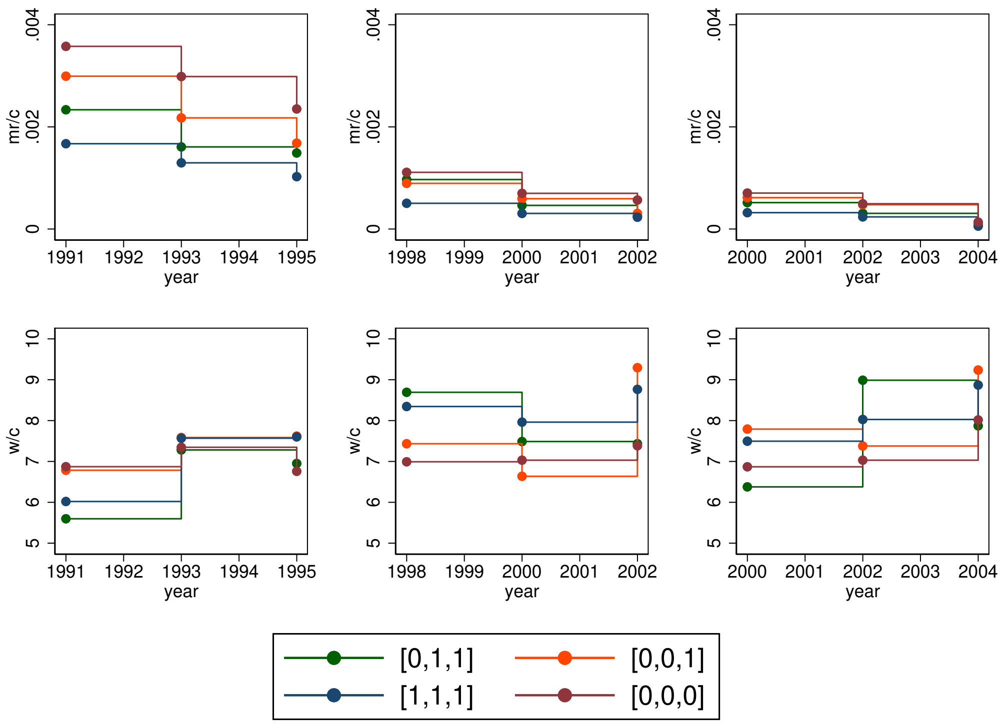

Note: The Figure plots the money-consumption (mr/c) and wealth-consumption (w/c) ratios over sequences of three waves of the SHIW for the adopters, denoted by (0,1,1), the always adopters, denoted by (1,1,1), and the never-adopters, denoted by (0,0,0). The plots apply to three time windows: 1991–1993–1995 (denoted W1), 1998–2000–2002 (denoted W2), and 2000–2002–2004 (denoted W3). The top panel shows the ratio mr/c. The bottom panel shows the w/c ratio for the same households.

Figure 2: Estimates of structural parameters

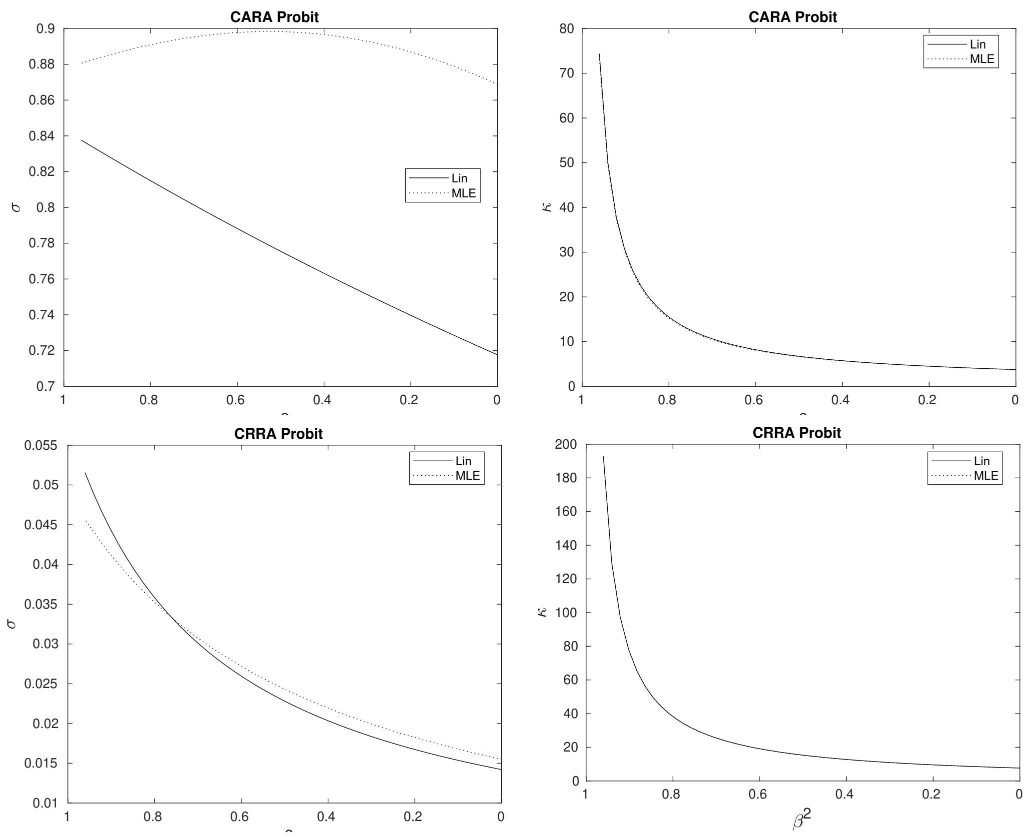

Note: The charts show the estimates of the parameters σ_κ and κ obtained from minimizing the linear least squares conditions corresponding (18) or the maximizing the likelihood functions corresponding to (19) when β^2 is fixed. CARA and CARRA indicate the functional form of the per-period utility function.

Figure 3: Estimates of discounted adoption cost

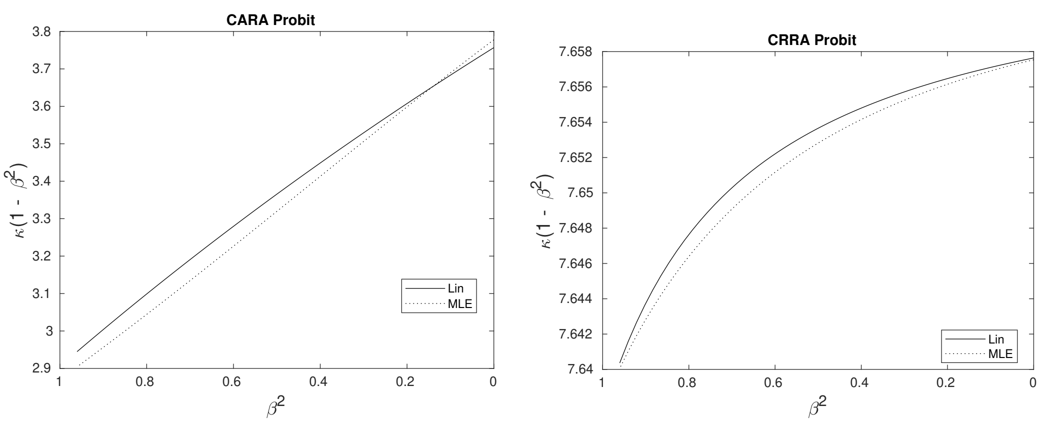

Note: The charts show the estimates of $\kappa(1-\beta^2)$ obtained from minimizing the linear least squares conditions corresponding to (18) or the maximizing the likelihood functions corresponding to (19) when β^2 is fixed. CARA and CARRA indicate the functional form of the per-period utility function.

Figure 4: Consumption, utility and compensating variation

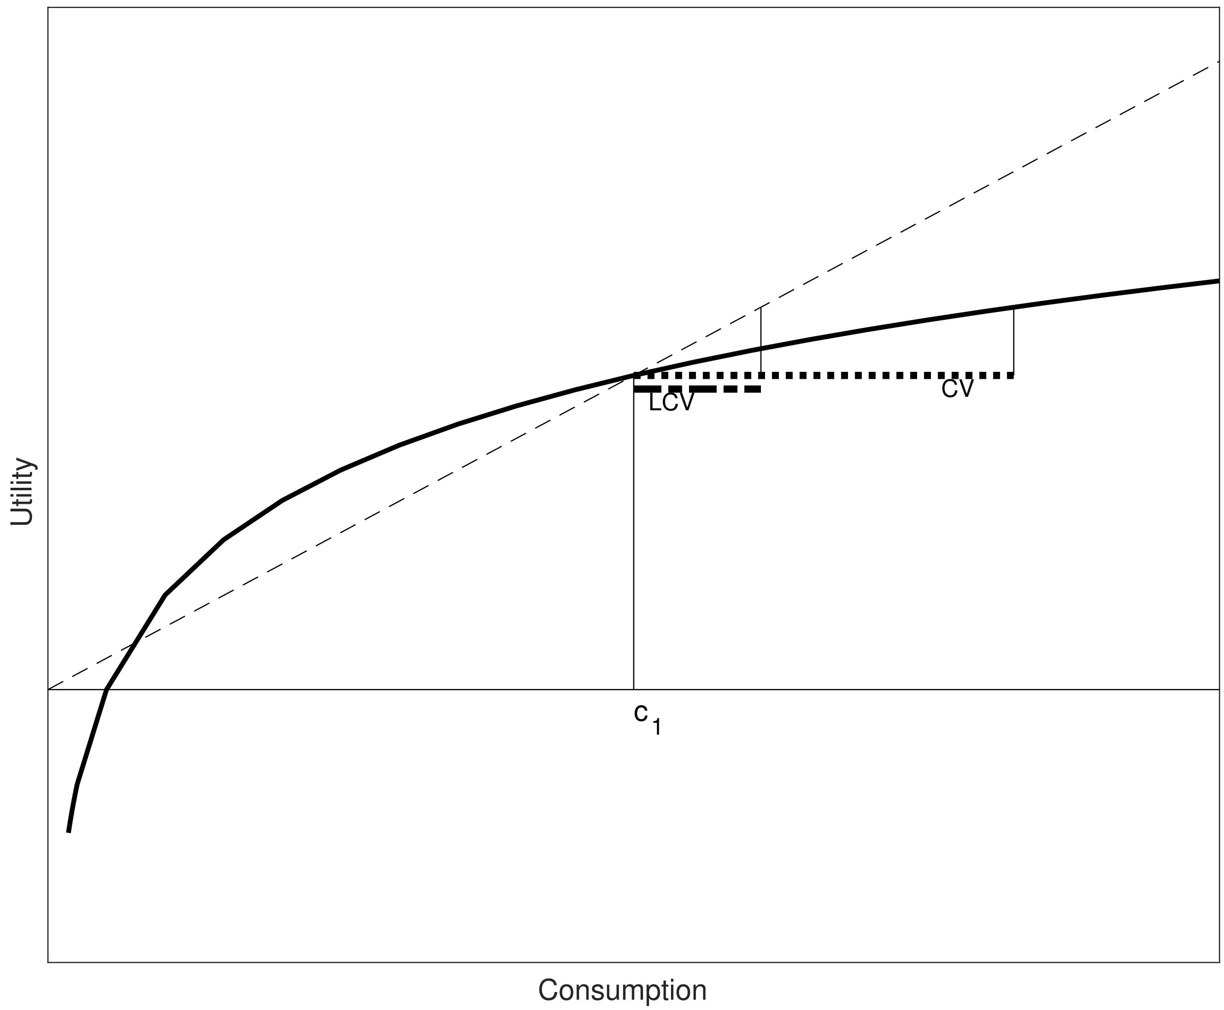

Note: The chart shows how to graphically obtain CV and LCV for a utility shock at a level of consumption c_1, see Goolsbee and Klenow [2006] for further details. Consumption is plotted on the x-axis and the y-axis shows the level of utility corresponding to a certain level of consumption. If utility changes by the amount x represented by the short vertical line, CV represents the solution to $u(c_1 + CV) = u(c_1) + x$ and $LCV = \frac{x}{u(c_1)} c_1$ is the proportional, or linear, compensation.

Figure 5: Estimates of LCV by age cohort

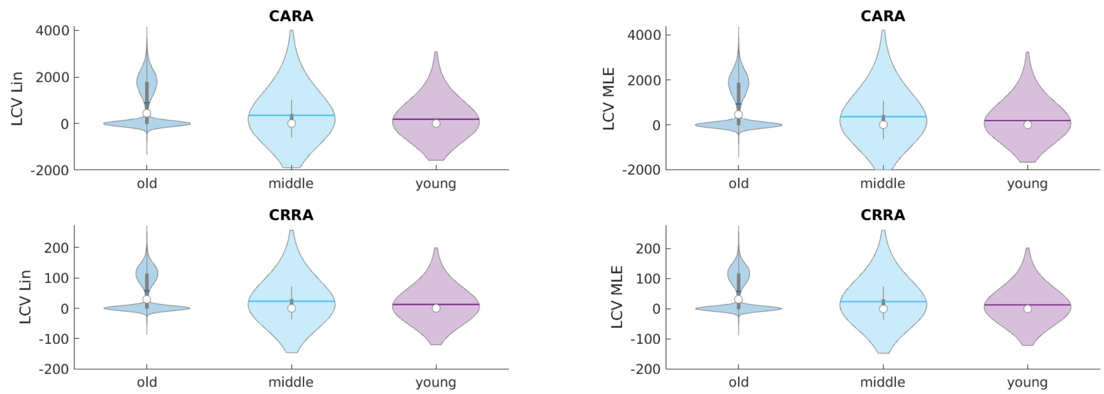

Note: The charts show the distribution of LCV by age cohort, using the parameters obtained from minimizing the linear least squares conditions corresponding to (18) or the maximizing the likelihood functions corresponding to (19). CARA and CARRA indicate the functional form of the per-period utility function.

Figure 6: Estimates of LCV by education

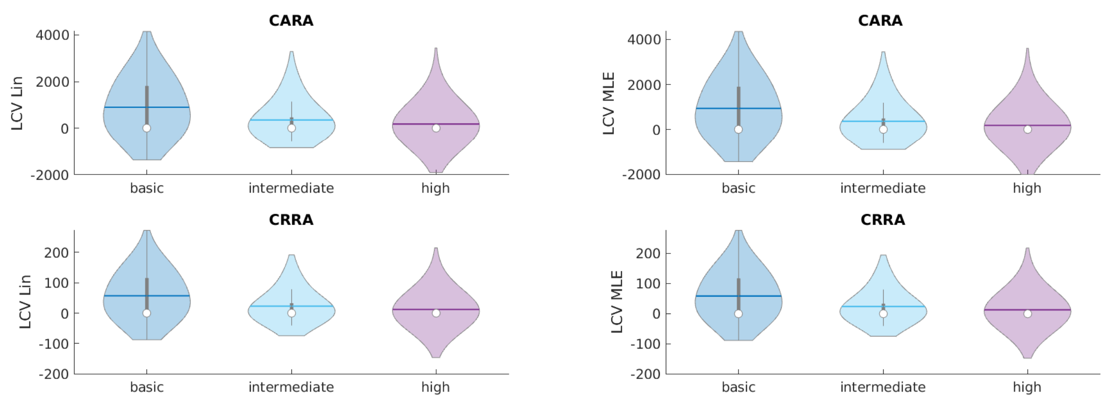

Note: The charts show the distribution of LCV by education level, using the parameters obtained from minimizing the linear least squares conditions corresponding to (18) or the maximizing the likelihood functions corresponding to (19). CARA and CARRA indicate the functional form of the per-period utility function.

Figure 7: Estimates of LCV by region

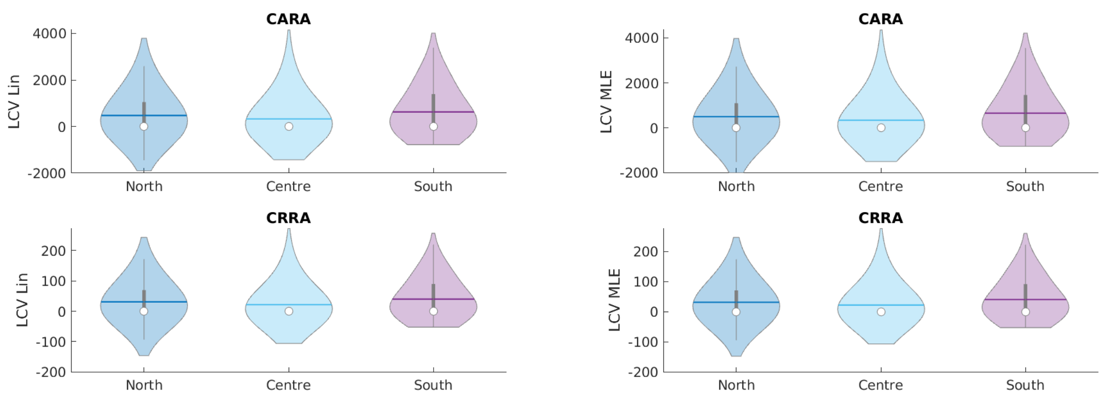

Note: The charts show the distribution of LCV by region, using the parameters obtained from minimizing the linear least squares conditions corresponding to (18) or the maximizing the likelihood functions corresponding to (19). CARA and CARRA indicate the functional form of the per-period utility function.