

Financial Innovation and the Transactions Demand for Cash*

Fernando Alvarez
University of Chicago and NBER[†]

Francesco Lippi
University of Sassari and CEPR[‡]

September 2007

Abstract

We document cash management patterns for households that are at odds with the predictions of deterministic inventory models that abstract from precautionary motives. We extend the Baumol-Tobin cash inventory model to a dynamic environment that allows for the possibility of withdrawing cash at random times at a low cost. This modification introduces a precautionary motive for holding cash and naturally captures developments in withdrawal technology, such as the increasing diffusion of bank branches and ATM terminals. We characterize the solution of the model and show that qualitatively it is able to reproduce the empirical patterns. Estimating the structural parameters we show that the model quantitatively accounts for key features of the data. The estimates are used to quantify the expenditure and interest rate elasticity of money demand, the impact of financial innovation on money demand, the welfare cost of inflation, the gains of disinflation and the benefit of ATM ownership.

JEL Classification Numbers: E5

Key Words: money demand, technological progress, inventory models.

*We thank Alessandro Secchi for his guidance in the construction and analysis of the database. We benefited from the comments of Manuel Arellano, V.V. Chari, Bob Lucas, Greg Mankiw, Rob Shimer, Pedro Teles and seminar participants at the University of Chicago, University of Sassari, Harvard University, FRB of Chicago, FRB of Minneapolis, Bank of Portugal, ECB, Bank of Italy and CEMFI.

[†]University of Chicago, 1126 E. 59th St., Chicago, IL 60637.

[‡]Ente Einaudi, via Due Macelli 73, 00184 Rome, Italy.

1 Introduction

There is a large literature arguing that financial innovation is important for understanding money demand, yet seldom this literature integrates the empirical analysis with an explicit modeling of the financial innovation. In this paper we develop a dynamic inventory model of money demand that explicitly incorporates the effects of financial innovation on cash management. We estimate the structural parameters of the model using detailed micro data from Italian households, and use the estimates to revisit several classic questions on money demand.

As standard in the inventory theory we assume that non-negative cash holdings are needed to pay for consumption purchases. We extend the Baumol-Tobin model to a dynamic environment which allows for the opportunity of withdrawing cash at random times at low or zero cost. Cash withdrawals at any other times involve a fixed cost b . In particular, the expected number of such opportunities per period of time is described by a single parameter p . Examples of such opportunities are finding an ATM that does not charge a fee, or passing by an ATM or bank desk at a time with a low opportunity cost. Another interpretation of p is that it measures the probability that an ATM terminal is working properly or a bank desk is open for business. Financial innovations such as the increase in the number bank branches and ATM terminals can be modeled by increases in p and decreases in b .

Our model changes the predictions of the Baumol-Tobin model (BT henceforth) in ways that are consistent with stylized facts concerning households' cash management behavior. The randomness introduced by p gives rise to a precautionary motive for holding cash: when agents have an opportunity to withdraw cash at zero cost they do so even if they have some cash at hand. Thus, the average cash balances held at the time of a withdrawal relative to the average cash holdings, \underline{M}/M , is a measure of the strength of the precautionary motive. For larger p the model generates larger values of \underline{M}/M , ranging between zero and one. Using household data for Italy and the US we document that \underline{M}/M is about 0.4, instead of being zero as predicted by the BT model. Another property of our model is that the number of withdrawals, n , increases with p , and the average withdrawal size W decreases, with W/M ranging between zero and two. Using data from Italian households we measure values of W/M smaller than two, the value predicted by the BT model.

We organize the analysis as follows. In Section 2 we use a panel data of Italian households to illustrate key cash management patterns, including the strength of precautionary motive, to compare them to the predictions of the BT model, and

motivate the analysis that follows.

Sections 3, 4 and 5 present the theory. Section 3 analyzes the effect of financial diffusion using a version of the BT model where agents have a deterministic number of free withdrawals per period. This model provides a simple illustration of how technology affects the level and the shape of the money demand (i.e. its interest and expenditure elasticities). Section 4 introduces our benchmark stochastic dynamic inventory model. In this model agents have random meetings with a financial intermediary in which they can withdraw money at no cost, a stochastic version of the model of Section 3. We solve analytically for the Bellman equation and characterize its optimal decision rule. We derive the distribution of currency holdings, the aggregate money demand, the average number of withdrawals, the average size of withdrawals, and the average cash balances at the time of a withdrawal. We show that a single index of technology, $b \cdot p^2$, determines both the shape of the money demand and the strength its precautionary component. While technological improvements (higher p and lower b) unambiguously decrease the level of money demand, their effect on this index –and hence on the shape and the precautionary component of money demand– is ambiguous. We conclude the section with the analysis of the welfare implications of our model and a comparison with the standard analysis as reviewed in Lucas (2000). Section 5 generalizes the model to one where withdrawals upon random meetings involve a small fixed cost f , with $0 < f < b$, which implies a more realistic distribution of withdrawals.

Sections 6, 7 and 8 contain the empirical analysis. In Section 6 we estimate the model using the panel data for Italian households. The two parameters p and b are overidentified because we observe four dimensions of household behavior: M , W , \underline{M} and n . We argue that the model has a satisfactory statistical fit and that the patterns of the estimates are reasonable. For instance, we find that the parameters for the households with an ATM card indicate their access to a better technology (higher p and lower b). The estimates also indicate that technology is better in geographic locations with higher density of ATM terminals and bank branches. Section 7 studies the implications of our findings for the time pattern of technology and for the expenditure and interest elasticity of the demand for currency. The estimated parameters reproduce the sizeable precautionary holdings present in the data, a feature absent in the BT model. Even though our model can generate interest rate elasticities between zero and $1/2$, and expenditure elasticities between $1/2$ and one, the values implied by our estimates are close to $1/2$ for both, the values of the BT model. We discuss how to reconcile our estimates of the interest rate elasticity

with the smaller values typically found in the literature. In Section 8 we use the estimates to quantify the welfare cost of inflation—in particular the gains from the Italian disinflation in the 1990s—and the benefits of ATM card ownership.

In the paper we abstract from the cash/credit choice. That is, we abstract from the choice of whether to have a credit card or not, and for those that have a credit card, whether a particular purchase is done using cash or credit. Our model studies how to finance a constant flow of cash expenditures, the value of which is taken as given both in the theory and in the empirical implementation. Formally, we are assuming separability between cash vs. credit purchases. We are able to study this problem for Italian households because we have a measure of the consumption purchases done with cash. We view our paper as an input on the study of cash/credit decisions, an important topic that we plan to address in the future.

2 Cash Holdings Patterns of Italian Households

Table 1 presents some statistics on the cash holdings patterns by Italian households based on the *Survey of Household Income and Wealth*.¹ For each year we report cross section means of statistics where the unit of analysis is the household. We report statistics separately for households with and without ATM cards. All these households have checking accounts that pay interests at rates documented below. The survey records the household expenditure paid in cash during the year (we use cash and currency interchangeably to denote the value of coins and banknotes). The table displays these expenditures as a fraction of total consumption expenditure. The fraction paid with cash is smaller for households with an ATM card, it displays a downward trend for both type of households, though its value remains sizeable as of 2004. These percentages are comparable to those for the US between 1984 and 1995.² The table reports the sample mean of the ratio M/c , where M is the average currency held by the household during a year and c is the daily expenditure paid

¹This is a periodic survey of the Bank of Italy that collects information on several social and economic characteristics. The cash management information that we are interested in is only available since 1993.

²Humphrey (2004) estimates that the mean share of total expenditures paid with currency in the US is 36% and 28% in 1984 and 1995, respectively. If expenditures paid with checks are added to those paid with currency, the resulting statistics is about 85% and 75% in 1984 and 1995, respectively. The measure including checks is used by Cooley and Hansen (1991) to compute the share of cash expenditures for households in the US where, contrary to the practice in Italy, checking accounts did not pay an interest. For comparison, the mean share of total expenditures paid with currency by all Italian households is 65% in 1995.

with currency. We notice that relative to c Italian households hold about twice as much cash than US households between 1984 and 1995.³

Table 1: Households' currency management

Variable	1993	1995	1998	2000	2002	2004
Expenditure share paid w/ currency ^a						
w/o ATM	0.68	0.67	0.63	0.66	0.65	0.63
w. ATM	0.62	0.59	0.56	0.55	0.52	0.47
Currency ^b : M/c (c per day)						
w/o ATM	15	17	19	18	17	18
w. ATM	10	11	13	12	13	14
M per Household, in 2004 euros ^c						
w/o ATM	430	490	440	440	410	410
w. ATM	370	410	370	340	330	350
Currency at withdrawals ^d : \underline{M}/M						
w/o ATM	0.41	0.31	0.47	0.46	0.46	na
w. ATM	0.42	0.30	0.39	0.45	0.41	na
Withdrawal ^e : W/M						
w/o ATM	2.3	1.7	1.9	2.0	2.0	1.9
w. ATM	1.5	1.2	1.3	1.4	1.3	1.4
# of withdrawals: n (per year) ^f						
w/o ATM	16	17	25	24	23	23
w. ATM	50	51	59	64	58	63
Normalized: $\frac{n}{c/(2M)}$ (c per year) ^f						
w/o ATM	1.2	1.4	2.6	2.0	1.7	2.0
w. ATM	2.4	2.7	3.8	3.8	3.9	4.1
# of observations ^g	6,938	6,970	6,089	7,005	7,112	7,159

The unit of observation is the household. Entries are sample means computed using sample weights. Only households with a checking account and whose head is not self-employed are included, which accounts for about 85% of the sample observations.

Notes: - ^aRatio of expenditures paid with cash to total expenditures (durables, non-durables and services). - ^bAverage currency during the year divided by daily expenditures paid with cash. - ^cThe average number of adults per household is 2.3. In 2004 one euro in Italy was equivalent to 1.25 USD in USA, PPP adjusted (Source: the World Bank ICP tables). - ^dAverage currency at the time of withdrawal as a ratio to average currency. - ^eAverage withdrawal during the year as a ratio to average currency. - ^fThe entries with $n = 0$ are coded as missing values. - ^gNumber of respondents for whom the currency and the cash consumption data are available in each survey. Data on withdrawals are supplied by a smaller number of respondents. Source: Bank of Italy - *Survey of Household Income and Wealth*.

Table 1 reports three statistics which are useful to assess the empirical performance of deterministic inventory models, such as the classic one by Baumol and Tobin. Similar information can be drawn from Figures 3, 4, 5 where each circle represents the average for households with and without ATM in a given year and

³Porter and Judson (1996), using currency and expenditure paid with currency, estimate that M/c is about 7 days both in 1984 and in 1986, and 10 in 1995. A calculation for Italy following the same methodology yields about 20 and 17 days in 1993 and 1995, respectively.

province (the size of the dot is proportional to the number of household observations). There are 103 provinces in Italy (the size of a province is similar to that of a U.S. county).

The first statistic is the ratio between currency holdings at the time of a withdrawal (\underline{M}) and average currency holdings in each year (M). While this ratio is zero in deterministic inventory theoretical model, its sample mean in the data is about 0.4. A comparable statistic for US households is about 0.3 in 1984, 1986 and 1995 (see Table 1 in Porter and Judson, 1996). The second one is the ratio between the withdrawal amount (W) and average currency holdings. While this ratio is 2 in the BT model, it is smaller in the data. The sample mean of this ratio for households with an ATM card is below 1.4, and for those without ATM is slightly below 2. Figure 4 shows that there is substantial variation across provinces and indeed the median across households (not reported in the table) is about 1.0 for households with and without ATM.⁴ The third statistic is the normalized number of withdrawals per year. The normalization is chosen so that in BT this statistic is equal to 1. In particular, in the BT model the following accounting identity holds, $nW = c$, and since withdrawals only happen when cash balances reach zero, then $M = W/2$. As the table shows the sample mean of this statistic is well above 1, especially so for households with ATM.

The second statistic, $\frac{W}{M}$, and the third, $\frac{n}{c/(2M)}$, are related through the accounting identity $c = nW$. In particular, if W/M is smaller than 2 and the identity holds then the third statistic must be above 1. Yet we present separate sample means for these statistics because of the large measurement error in all these variables. This is informative because W enters in the first statistic but not in the second and c enters in the third but not in the second. In the estimation section of the paper we document and consider the effect of measurement error systematically, without altering the conclusion about the drawbacks of deterministic inventory theoretical models.

For each year Table 2 reports the mean and standard deviation across provinces for the diffusion of bank branches and ATM terminals, and for two components of the opportunity cost of holding cash: interest rate paid on deposits and the probability of cash being stolen. The diffusion of Bank branches and ATM terminals varies significantly across provinces and is increasing through time. Differences in the

⁴An alternative source for the average ATM withdrawal, based on banks' reports, can be computed using Tables 12.1 and 13.1 in the ECB Blue Book (2006). These values are similar, indeed somewhat smaller, than the corresponding values from the household data (see the working paper version for details).

Table 2: Financial innovation and the opportunity cost of cash

Variable	1993	1995	1998	2000	2002	2004
Bank branches ^a	0.38 (0.13)	0.42 (0.14)	0.47 (0.16)	0.50 (0.17)	0.53 (0.18)	0.55 (0.18)
ATM terminals ^a	0.31 (0.18)	0.39 (0.19)	0.50 (0.22)	0.57 (0.22)	0.65 (0.23)	0.65 (0.22)
Interest rate on deposits ^b	6.1 (0.4)	5.4 (0.3)	2.2 (0.2)	1.7 (0.2)	1.1 (0.2)	0.7 (0.1)
Probability of cash being stolen ^c	2.2 (2.6)	1.8 (2.1)	2.1 (2.4)	2.2 (2.5)	2.1 (2.4)	2.2 (2.6)
CPI Inflation	4.6	5.2	2.0	2.6	2.6	2.3

Notes: Mean (standard deviation in parenthesis) across provinces. - ^a Per thousand residents (Source: the Supervisory Reports to the Bank of Italy and the Italian Central Credit Register). - ^b Net nominal interest rates in per cent. Arithmetic average between the self-reported interest on deposit account (Source: Survey of Household Income and Wealth) and the average deposit interest rate reported by banks in the province (Source: Central credit register). - ^c We estimate this probability using the time and province variation from statistics on reported crimes on Purse snatching and pickpocketing. The level is adjusted to take into account both the fraction of unreported crimes as well as the fraction of money stolen for different types of crimes using survey data on victimization rates (Source: Istat and authors' computations).

nominal interest rate across time are due mainly to the disinflation. The variation nominal interest rates across provinces mostly reflects the segmentation of banking markets. The large differences in the probability of cash being stolen across provinces reflect variation in crime rates across rural vs. urban areas, and a higher incidence of such crimes in the North.

Lippi and Secchi (2007) report that the household data display patterns which are in line with previous empirical studies showing that the demand for currency decreases with financial development and that its interest elasticity is below one-half.⁵ From Table 2 we observe that the opportunity cost of cash in 2004 is about 1/3 of the value in 1993 (the corresponding ratio for the nominal interest rate is about 1/9), and that the average of M/c shows an upward trend. Indeed the average of M/c across households of a given type (with and without ATM cards) is negatively correlated with the opportunity cost R in the cross section, in the time series, and the

⁵They estimate that the elasticity of cash holdings with respect to the interest rate is about zero for agents who hold an ATM card and -0.2 for agents without ATM card. See their Section 5 for a comparison with the findings of other papers.

pool time series-cross section. Yet the largest estimate for the interest rate elasticity are smaller than 0.25 and in most cases about 0.05 (in absolute values). At the same time, Table 2 shows large increases in bank branches and ATM terminals per person. Such patterns are consistent with both shifts of the money demand and movements along it. Our model and estimation strategy allows us to quantify each of them.

Another classic model of money demand is Miller and Orr (1966) who study the optimal inventory policy for an agent subject to stochastic cash inflows and outflows. Despite the presence of uncertainty, their model, as the one by BT, does not feature a precautionary motive in the sense that $\underline{M} = 0$. Unlike in the BT model, they find that the interest rate elasticity is $1/3$ and the average withdrawal size W/M is $3/4$. In this paper we keep the BT model as a theoretical benchmark because the Miller and Orr model is more suitable for the problem faced by firms, given the nature of stochastic cash *inflows* and outflows. Our paper studies currency demand by households: the theory studies the optimal inventory policy for an agent that faces deterministic cash outflows (consumption expenditure) and no cash inflows and the empirical analysis uses the household survey data (excluding entrepreneurs).

3 A model with deterministic free withdrawals

This section presents a modified version of the BT model to illustrate how technological progress affects the level and interest elasticity of the demand for currency.

Consider an agent who finances a consumption flow c by making n withdrawals from a deposit account. We let R be the net nominal interest rate paid on deposits. In a deterministic setting agents cash balances decrease until they hit zero, when a new withdrawal must take place. Hence the size of each withdrawal is $W = c/n$ and the average cash balance $M = W/2$. In the BT model agents pay a fixed cost b for each withdrawal. We modify the latter by assuming that the agent has p free withdrawals, so that if the total number of withdrawals is n then she pays only for the excess of n over p . Setting $p = 0$ yields the BT case. Technology is thus represented by the parameters b and p .

For example, assume that the cost of a withdrawal is proportional to the distance to an ATM or bank branch. In a given period the agent is moving across locations, for reason unrelated to her cash management, so that p is the number of times that she is in a location with an ATM or bank branch. At any other time, b is the distance that the agent must travel to withdraw. In this setup an increase in the

density of bank branches or ATMs increases p and decreases b .

The optimal number of withdrawals solves the minimization problem

$$\min_n \left[R \frac{c}{2n} + b \max(n - p, 0) \right] . \quad (1)$$

By examining the objective function, it is immediate that the value of n that solves the problem, and its associated M/c , depends only on $\beta \equiv b/(cR)$, the ratio of the two costs, and p . The money demand for a technology with $p \geq 0$ is given by

$$\frac{M}{c} = \frac{1}{2p} \sqrt{\min \left(2 \frac{\hat{b}}{R}, 1 \right)} \quad \text{where} \quad \hat{b} \equiv \frac{b p^2}{c} . \quad (2)$$

To understand the workings of the model, fix b and consider the effect of increasing p (so that \hat{b} increases). For $p = 0$ we have the BT setup, so that when R is small the agent decides to economize on withdrawals and choose a large value of M . Now consider the case of $p > 0$. In this case there is no reason to have less than p withdrawals, since these are free by assumption. Hence, for all $R \leq 2\hat{b}$ the agent will choose the same level money holdings, namely, $M = c/(2p)$, since she is not paying for any withdrawal but is subject to a positive opportunity cost. Note that the interest elasticity is zero for $R \leq 2\hat{b}$. Thus as p (hence \hat{b}) increases, then the money demand has a lower level and a lower interest rate elasticity than the money demand from the BT model. Indeed (2) implies that the range of interest rates R for which the money demand is smaller and has lower interest rate elasticity is increasing in p . On the other hand, if we fix \hat{b} and increase p the only effect is to lower the level of the money demand. The previous discussion makes clear that for fixed p , \hat{b} controls the “shape” of the money demand, and for fixed \hat{b} , p controls its level. We think of technological improvements as both increasing p and lowering b : the net effect on \hat{b} , hence on the slope of the money demand, is in principle ambiguous. The empirical analysis below allows us to sign and quantify this effect.

4 A model with random free withdrawals

This section presents the main model which generalizes the example of the previous section in several dimensions. It takes an explicit account of the dynamic nature of the cash inventory problem, as opposed to minimizing the average steady state cost. It distinguishes between real and nominal variables, as opposed to fi-

nancing a constant nominal expenditure, or alternatively assuming zero inflation. Most importantly, we consider the case where the agent has a Poisson arrival of free opportunities to withdraw cash at a rate p . Relative to the deterministic model this assumption produces cash management behavior that is closer to the one documented in Section 2. The randomness thus introduced gives rise to a precautionary motive, so that some withdrawals occur when the agent still has a positive cash balance and the (average) W/M ratio is smaller than two. The model retains the feature (discussed in Section 3) that the interest rate elasticity is smaller than $1/2$ and is decreasing in the parameter p . It also generalizes the sense in which the “shape” of the money demand depends on the parameter $\hat{b} = p^2 b/c$.

4.1 The agent’s problem

The model we consider solves the problem of minimizing the cost of financing a given constant flow of cash consumption, denoted by c . We assume that agents are subject to a cash-in-advance constraint. We use m to denote the non-negative real cash balances of an agent which decrease due to consumption and inflation:

$$\frac{dm(t)}{dt} = -c - m(t)\pi \quad (3)$$

for almost all $t \geq 0$.

Agents can withdraw or deposit at any time from an account that yields real interest r . Transfers from the interest bearing account to cash balances are indicated by discontinuities in m : a withdrawal is a jump up on the cash balances, i.e. $m(t^+) - m(t^-) > 0$, and likewise for a deposit.

There are two sources of randomness in the environment, described by independent Poisson processes with intensities p_1 and p_2 . The first process describes the arrivals of “free adjustment opportunities” (see the Introduction for examples). The second Poisson process describes the arrivals of times where the agent loses (or is stolen) her cash balances. We assume that a fixed cost b is paid for each adjustment, unless it happens exactly at the time of a free adjustment opportunity.

We can write the problem of the agent as:

$$C(m) = \min_{\{m(t), \tau_j\}} E_0 \left\{ \sum_{j=0}^{\infty} e^{-r \tau_j} [I_{\tau_j} b + (m(\tau_j^+) - m(\tau_j^-))] \right\} \quad (4)$$

subject to (3) and $m(t) \geq 0$, where τ_j denote the stopping times at which an

adjustment (jump) of m takes place, and $m(0) = m$ is given. The indicator I_{τ_j} is zero – so the cost is not paid – if the adjustment takes place at a time of a free adjustment opportunity, otherwise is equal to one. The expectation is taken with respect to the two Poisson processes. The parameters that define this problem are r, π, p_1, p_2, b and c .

4.2 Bellman equations and optimal policy

We turn to the characterization of the Bellman equations and of its associated optimal policy. We will guess, and later verify, that the optimal policy is described by two thresholds for m : $0 < m^* < m^{**}$. The threshold m^* is the value of cash that agents choose to have after a contact with a financial intermediary: we refer to it as the optimal cash replenishment level. The threshold m^{**} is a value of cash beyond which agents will pay the cost b , contact the intermediary, and make a deposit so as to leave her cash balances at m^* . Assuming that the optimal policy is of this type and that for $m \in (0, m^{**})$ the value function C is differentiable, it must satisfy:

$$\begin{aligned} rC(m) = & C'(m)(-c - \pi m) + p_1 \min_{\hat{m} \geq 0} [\hat{m} - m + C(\hat{m}) - C(m)] + \\ & + p_2 \min_{\hat{m} \geq 0} [b + \hat{m} + C(\hat{m}) - C(m)] \quad . \end{aligned} \quad (5)$$

If the agent chooses not to contact the intermediary then, as standard, the Bellman equation states that the return on the value function $rC(m)$ must equal the flow cost plus the expected change per unit of time. The first term of the summation gives the change in the value function per unit of time, conditional on no arrival of either free adjustment or of a loss of cash (theft). This change is given by the change in the state m , times the derivative of the value function $C'(m)$. The second term gives the expected change conditional on the arrival of free adjustment opportunity: an adjustment $\hat{m} - m$ is incurred instantly with its associated “capital gain” $C(\hat{m}) - C(m)$. Likewise, the third term gives the change in the value function conditional on the money stock m being stolen. In this case the cost b must be paid and the adjustment equals \hat{m} , since m is “lost”. Regardless of how the agent ends up matched with a financial intermediary, upon the match she chooses the optimal level of real balances, which we denote by m^* , which solves

$$m^* = \arg \min_{\hat{m} \geq 0} \hat{m} + C(\hat{m}) \quad . \quad (6)$$

Note that the optimal replenishment level m^* is constant. There are two boundary conditions for this problem. First, if money balances reach zero ($m = 0$) the agent must withdraw, otherwise she will violate the non-negativity constraint in the next instant. Second, for values of $m \geq m^{**}$ we conjecture that the agent chooses to pay b and deposit the extra amount, $m - m^*$. Combining these boundary conditions with (5) we have:

$$C(m) = \begin{cases} b + m^* + C(m^*) & \text{if } m = 0 \\ \frac{-C'(m)(c + \pi m) + (p_1 + p_2)[m^* + C(m^*)] + p_2 b - p_1 m}{r + p_1 + p_2} & \text{if } m \in (0, m^{**}) \\ b + m^* - m + C(m^*) & \text{if } m \geq m^{**} \end{cases} \quad (7)$$

For the assumed configuration to be optimal it must be the case that the agent prefers not to pay the cost b and adjust money balances in the relevant range:

$$m + C(m) < b + m^* + C(m^*) \quad \text{all } m \in (0, m^{**}) . \quad (8)$$

Summarizing, we say that $0 < m^* < m^{**}$, $C(\cdot)$ solve the Bellman equation for the total cost problem (4) if they satisfy (6)-(7)-(8).

We find it convenient to reformulate this problem so that it is closer to the standard inventory theoretical models. We define a related problem where the agent minimizes the shadow cost

$$V(m) = \min_{\{m(t), \tau_j\}} E_0 \left\{ \sum_{j=0}^{\infty} e^{-r\tau_j} \left[I_{\tau_j} b + \int_0^{\tau_{j+1} - \tau_j} e^{-rt} R m(t + \tau_j) dt \right] \right\} \quad (9)$$

subject to (3), $m(t) \geq 0$, where τ_j denote the stopping times at which an adjustment (jump) of m takes place, and $m(0) = m$ is given. The indicator I_{τ_j} equals zero if the adjustment takes place at the time of a free adjustment, otherwise is equal to one. In this formulation R is the opportunity cost of holding cash. In this problem there is *only one* Poisson process with intensity p describing the arrival of a free opportunity to adjust. The parameters of this problem are r, R, π, p, b and c .⁶

⁶The shadow cost formulation is the standard one used in the literature for inventory theoretical models, as in the models of Baumol-Tobin, Miller and Orr (1966), Constantinides (1976), among others. In these papers the problem aims to minimize the steady state cost implied by a stationary inventory policy. This differs from our formulation, where the agent minimizes the expected discounted cost in (9). In this regard our analysis follows the one of Constantinides and Richards (1978). For a related model, Frenkel and Jovanovic (1980) compare the resulting money demand arising from minimizing the steady state vs. the expected discounted cost.

The derivation of the Bellman equation for an agent unmatched with a financial intermediary and holding a real value of cash m follows by the same logic used to derive equation (5). The only decision that the agent must make is whether to remain unmatched, or to pay the fixed cost b and be matched with a financial intermediary. Denoting by $V'(m)$ the derivative of $V(m)$ with respect to m , the Bellman equation satisfies

$$rV(m) = Rm + p \min_{\hat{m} \geq 0} (V(\hat{m}) - V(m)) + V'(m)(-c - m\pi) \quad . \quad (10)$$

Regardless of how the agent ends up matched with a financial intermediary, she chooses the optimal adjustment and sets $m = m^*$, or

$$V^* \equiv V(m^*) = \min_{\hat{m} \geq 0} V(\hat{m}) \quad . \quad (11)$$

As in problem (4) we will guess that the optimal policy is described by two threshold values satisfying $0 < m^* < m^{**}$. This requires two boundary conditions. At $m = 0$ the agent must pay the cost b and withdraw, and for $m \geq m^{**}$ the agent chooses to pay the cost b and deposit the cash in excess of m^* . Combining these boundary conditions with (10) we have:

$$V(m) = \begin{cases} V^* + b & \text{if } m = 0 \\ \frac{Rm + pV^* - V'(m)(c + m\pi)}{r + p} & \text{if } m \in (0, m^{**}) \\ V^* + b & \text{if } m \geq m^{**} \end{cases} \quad (12)$$

To ensure that it is optimal not to pay the cost and contact the intermediary in the relevant range we require:

$$V(m) < V^* + b \quad \text{for } m \in (0, m^{**}) \quad . \quad (13)$$

Summarizing, we say that $0 < m^* < m^{**}$, V^* , $V(\cdot)$ solve the Bellman equation for the shadow cost problem (9) if they satisfy (11)-(12)-(13). We are now ready to show that, first, (4) and (9) are equivalent and, second, the existence and characterization of the solution.

Proposition 1. *Assume that the opportunity cost is given by $R = r + \pi + p_2$, and that the contact rate with the financial intermediary is $p = p_1 + p_2$. Assume that the*

functions $C(\cdot), V(\cdot)$ satisfy

$$C(m) = V(m) - m + c/r + p_2 b/r \quad (14)$$

for all $m \geq 0$. Then, $m^*, m^{**}, C(\cdot)$ solve the Bellman equation for the total cost problem (4) if and only if $m^*, m^{**}, V^*, V(\cdot)$ solve the Bellman equation for the shadow cost problem (9).

Proof. See Appendix A.

We briefly comment on the relation between the total and shadow cost problems. Notice that they are described by the same number of parameters. They have r, π, c, b in common, the total cost problem uses p_1 and p_2 , while the shadow cost problem uses R and p . That $R = r + \pi + p_2$ is quite intuitive: the shadow cost of holding money is given by the real opportunity cost of investing, r , plus the fact that cash holdings lose real value continually at a rate π and they are lost entirely with probability p_2 per unit of time. Likewise that $p = p_1 + p_2$ is clear too: since the effect of either shock is to force an adjustment on cash. The relation between C and V in (14) is quite intuitive. First the constant c/r is required, since even if withdrawals were free (say $b = 0$) consumption expenditures must be financed. Second, the constant $p_2 b/r$ is the present value of all the withdrawals cost that is paid after cash is “lost”. This adjustment is required because in the shadow cost problem there is no “theft”. Third, the term m has to be subtracted from V since this amount has already been extracted from the interest bearing account.

From now on, we use the shadow cost formulation, since it is closer to the standard inventory decision problem. On the theoretical side, having the effect of “theft” as part of the opportunity cost allows us to parameterize R as being, at least conceptually, independent of r and π . On the quantitative side we think that, at least for low nominal interest rates, the presence of other opportunity costs may be important.

4.3 Characterization of the optimal return point m^*

The next proposition gives one non-linear equation whose unique solution determines the cash replenishment value m^* as a function of the model parameters: R, π, r, p, c and b .

Proposition 2. *Assume that $r + \pi + p > 0$. The optimal return point $\frac{m^*}{c}$ has three arguments: $\beta, r + p, \pi$, where $\beta \equiv \frac{b}{cR}$. The return point m^* is given by the unique*

positive solution to

$$\left(\frac{m^*}{c}\pi + 1\right)^{1+\frac{r+p}{\pi}} = \frac{m^*}{c}(r+p+\pi) + 1 + (r+p)(r+p+\pi)\frac{b}{cR} . \quad (15)$$

Proof. See Appendix A.

Note that, keeping r and π fixed, the solution for m^*/c is a function of $b/(cR)$, as it is in the steady state money demand of Section 3. This immediately implies that m^* is homogenous of degree one in (c, b) . The next proposition gives a closed form solution for the function $V(\cdot)$, and the scalar V^* in terms of m^* .

Proposition 3. Assume that $r + \pi + p > 0$. Let m^* be the solution of (15).

(i) The value for the agents not matched with a financial institution, for $m \in (0, m^{**})$, is given by the convex function:

$$V(m) = \left[\frac{pV^* - Rc/(r+p+\pi)}{r+p} \right] + \left[\frac{R}{r+p+\pi} \right] m + \left(\frac{c}{r+p} \right)^2 A \left[1 + \pi \frac{m}{c} \right]^{-\frac{r+p}{\pi}} \quad (16)$$

where $A = \frac{r+p}{c^2} \left(R m^* + (r+p)b + \frac{Rc}{r+p+\pi} \right) > 0$.

For $m = 0$ or $m \geq m^{**}$: $V(m) = V^* + b$.

(ii) The value for the agents matched with a financial institution, V^* , is

$$V^* = \frac{R}{r} m^* . \quad (17)$$

Proof. See Appendix A.

The close relationship between the value function at zero cash and the optimal return point $V(0) = (R/r)m^* + b$ derived in this proposition will be useful to measure the gains of different financial arrangements. The next proposition uses the characterization of the solution for m^* to conduct some comparative statics.

Proposition 4. The optimal return point m^* has the following properties:

(i) $\frac{m^*}{c}$ is increasing in $\frac{b}{cR}$, $\frac{m^*}{c} = 0$ as $\frac{b}{cR} = 0$ and $\frac{m^*}{c} \rightarrow \infty$ as $\frac{b}{cR} \rightarrow \infty$.

(ii) For small $\frac{b}{cR}$, we can approximate $\frac{m^*}{c}$ by the solution in BT model, or

$$\frac{m^*}{c} = \sqrt{2 \frac{b}{cR}} + o\left(\sqrt{\frac{b}{cR}}\right)$$

where $o(z)/z \rightarrow 0$ as $z \rightarrow 0$.

(iii) Assuming that the Fisher equation holds, in that $\pi = R - r$, the elasticity of

m^* with respect to p evaluated at zero inflation satisfies

$$0 \leq -\frac{p}{m^*} \frac{dm^*}{dp} \Big|_{\pi=0} \leq \frac{p}{p+r} .$$

(iv) The elasticity of m^* with respect to R evaluated at zero inflation satisfies

$$0 \leq -\frac{R}{m^*} \frac{dm^*}{dR} \Big|_{\pi=0} \leq \frac{1}{2} .$$

The elasticity is decreasing in p and satisfies:

$$-\frac{R}{m^*} \frac{\partial m^*}{\partial R} \Big|_{\pi=0} \rightarrow 1/2 \text{ as } \frac{\hat{b}}{R} \rightarrow 0 \text{ and } -\frac{R}{m^*} \frac{\partial m^*}{\partial R} \Big|_{\pi=0} \rightarrow 0 \text{ as } \frac{\hat{b}}{R} \rightarrow \infty$$

where $\hat{b} \equiv (p+r)^2 b/c$.

Proof. See Appendix A.

The proposition shows that when $b/(cR)$ is small the resulting money demand is well approximated by the one for the BT model. Part (iv) shows that the absolute value of the interest elasticity (when inflation is zero) ranges between zero and $1/2$, and that it is decreasing in p . In the limits we use \hat{b} to write a comparative static result for the interest elasticity of m^* with respect to p . Indeed, for $r = 0$, we have already given an economic interpretation to \hat{b} in Section 3, to which we will return in Proposition 8. Since in Proposition 2 we show that m^* is a function of $b/(cR)$, then the elasticity of m^* with respect to b/c equals the one with respect to R with an opposite sign.

4.4 Number of withdrawals and cash holdings distribution

This section derives the invariant distribution of real cash holdings when the policy characterized by the parameters (m^*, p, c) is followed and the inflation rate is π . Throughout the section m^* is treated as a parameter, so that the policy is to replenish cash holdings up to the return value m^* , either when a match with a financial intermediary occurs, which happens at a rate p per unit of time, or when the agent runs out of money (i.e. real balances hit zero). Our first result is to compute the expected number of withdrawals per unit of time, denoted by n . This includes both the withdrawals that occur upon an exogenous contact with the financial intermediary and the ones initiated by the agent when her cash balances reach zero. By the fundamental theorem of Renewal Theory n equals the reciprocal of the expected time between withdrawals, which after straightforward calculations gives

Proposition 5. *The expected number of cash withdrawals per unit of time, n , is*

$$n\left(\frac{m^*}{c}, \pi, p\right) = \frac{p}{1 - \left(1 + \pi \frac{m^*}{c}\right)^{-\frac{p}{\pi}}} . \quad (18)$$

Proof. See Appendix A.

As can be seen from expression (18) the ratio $n/p \geq 1$, since in addition to the p free withdrawals it includes the costly withdrawals that agents do when they exhaust their cash. Note how this formula yields exactly the expression in the BT model when $p = \pi = 0$. The next proposition derives the density of the invariant distribution of real cash balances as a function of p, π, c and m^*/c .

Proposition 6. (i) *The density for the real balances m is:*

$$h(m) = \left(\frac{p}{c}\right) \frac{\left[1 + \pi \frac{m}{c}\right]^{\frac{p}{\pi} - 1}}{\left[1 + \pi \frac{m^*}{c}\right]^{\frac{p}{\pi}} - 1} . \quad (19)$$

(ii) *Let $H(m, m_1^*)$ be the CDF of m for a given m^* . Let $m_1^* < m_2^*$, then $H(m, m_2^*) \leq H(m, m_1^*)$, i.e. $H(\cdot, m_2^*)$ first order stochastically dominates $H(\cdot, m_1^*)$.*

Proof. See Appendix A.

The density of m solves the following ODE (see the proof of Proposition 6)

$$\frac{\partial h(m)}{\partial m} = \frac{(p - \pi)}{(\pi m + c)} h(m) \quad (20)$$

for any $m \in (0, m^*)$. There are two forces determining the shape of this density. One is that agents meet a financial intermediary at a rate p , where they replenish their cash balances. The other is that inflation eats away the real value of their nominal balances. Notice that if $p = \pi$ these two effects cancel and the density is constant. If $p < \pi$ the density is downward sloping, with more agents at low values of real balances due to the greater pull of the inflation effect. If $p > \pi$, the density is upward sloping due the greater effect of the replenishing of cash balances. This uses that $\pi m + c > 0$ in the support of h because $\pi m^* + c > 0$ (see equation 38).

We define the average money demand as $M = \int_0^{m^*} m h(m) dm$. Using the expression for $h(m)$, integration gives

$$\frac{M}{c} \left(\frac{m^*}{c}, \pi, p\right) = \frac{\left(1 + \pi \frac{m^*}{c}\right)^{\frac{p}{\pi}} \left[\frac{m^*}{c} - \frac{(1 + \pi \frac{m^*}{c})}{p + \pi}\right] + \frac{1}{p + \pi}}{\left[1 + \pi \frac{m^*}{c}\right]^{\frac{p}{\pi}} - 1} . \quad (21)$$

Next we analyze how M depends on m^* and p . The function $\frac{M}{c}(\cdot, \pi, p)$ is increasing in m^* , which follows immediately from part (ii) of Proposition 6: with a higher target replenishment level the agents end up holding more money on average. The next proposition shows that for a fixed m^* , M is increasing in p :

Proposition 7. *The ratio $\frac{M}{m^*}$ is increasing in p with:*

$$\frac{M}{m^*}(\pi, p) = \frac{1}{2} \text{ for } p = \pi \quad \text{and} \quad \frac{M}{m^*}(\pi, p) \rightarrow 1 \text{ as } p \rightarrow \infty.$$

Proof. See Appendix A.

It is useful to compare this result with the corresponding one for the BT case, which is obtained when $\pi = p = 0$. In this case agents withdraw m^* hence $M/m^* = 1/2$. The other limit corresponds to the case where withdrawals happen so often that at all times the average amount of money coincides with the amount just after a withdrawal.

The average withdrawal, W , is

$$W = m^* \left[1 - \frac{p}{n} \right] + \left[\frac{p}{n} \right] \int_0^{m^*} (m^* - m) h(m) dm. \quad (22)$$

To understand the expression for W notice that $(n-p)$ is the number of withdrawals in a unit of time that occur because the zero balance is reached, so if we divide it by the total number of withdrawals per unit of time (n) we obtain the fraction of withdrawals that occur at a zero balance. Each of these withdrawals is of size m^* . The complementary fraction gives the withdrawals that occur due to a chance meeting with the intermediary. A withdrawal of size $m^* - m$ happens with frequency $h(m)$. Inspection of (22) shows that W/c is a function of three arguments: $m^*/c, \pi, p$.

Combining the previous results we can see that for $p \geq \pi$, the ratio of withdrawals to average cash holdings is less than two. To see this, using the definition of W we can write

$$\frac{W}{M} = \frac{m^*}{M} - \frac{p}{n}. \quad (23)$$

Since $M/m^* \geq 1/2$, then it follows that $W/M \leq 2$. Indeed notice that for p large enough this ratio can be smaller than one. We mention this property because for the Baumol - Tobin model the ratio W/M is exactly two, while in the data of Table 1 for households with an ATM card the average ratio is below 1.5 and its median value is 1. The intuition for this result in our model is clear: agents take advantage of the free random withdrawals regardless of their cash balances, hence

the withdrawals are distributed on $[0, m^*]$, as opposed to be concentrated on m^* , as in the BT model.

We let \underline{M} be the average amount of money that an agent has at the time of withdrawal. A fraction $[1 - p/n]$ of the withdrawals happens when $m = 0$. For the remaining fraction, p/n , an agent has money holdings at the time of the withdrawal distributed with density h , so that: $\underline{M} = 0 [1 - \frac{p}{n}] + [\frac{p}{n}] \int_0^{m^*} m h(m) dm$. Inspection of this expression shows that \underline{M}/c is a function of three arguments: $m^*/c, \pi, p$. Simple algebra shows that $\underline{M} = m^* - W$ or, inserting the definition of \underline{M} into the expression for M :

$$\underline{M} = \frac{p}{n} M . \quad (24)$$

The ratio \underline{M}/M is a measure of the precautionary demand for cash: it is zero only when $p = 0$, it goes to 1 as $p \rightarrow \infty$ and, at least for $\pi = 0$, it is increasing in p . This is because as p increases the agent has more opportunities for a free withdrawal, which directly increases \underline{M}/M (see equations 18 and 24), and from part (iii) in Proposition 4 the induced effect of p on m^* cannot outweigh the direct effect.

Other researchers noticing that currency holdings are positive at the time of withdrawals account for this feature by adding a constant \underline{M}/M to the sawtooth path of a deterministic inventory model, which implies that the average cash balance is $M_1 = \underline{M} + 0.5 c/n$ or $M_2 = \underline{M} + 0.5 W$. See e.g. equations 1 and 2 in Attanasio, Guiso and Jappelli (2002) and Table 1 in Porter and Judson (1996). Instead, when we model the determinants of the precautionary holdings \underline{M}/M in a random setup, we find that $W/2 < M < \underline{M} + W/2$. The leftmost inequality is a consequence of Proposition 7 and equation (23), the other can be easily derived using the form of the optimal decision rules and the law of motion of cash flows (see the working paper version for details). The discussion above shows that the expressions for the demand for cash proposed in the literature to deal with the precautionary motive are upward biased. Using the data of Table 1 shows that both expressions M_1 and M_2 overestimate the average amount of cash held by Italian households by a large margin.⁷

4.5 Comparative statics on M , \underline{M} , W and welfare

We begin with a comparative statics exercise on M , \underline{M} and W in terms of the primitive parameters b/c , p , and R . To do this we combine the results of Section

⁷The expression for M_1 overestimates the average cash by 20% and 140% for household with and without ATMs, respectively; the one for M_2 by 7% and 40%, respectively.

4.3, where we analyzed how the optimal decision rule m^*/c depends on p , b/c and R , with the results of Section 4.4 where we analyze how M , \underline{M} , and W change as a function of m^*/c and p . The next proposition defines a one dimensional index $\hat{b} \equiv (b/c)p^2$ that characterizes the shape of the money demand and the strength of the precautionary motive focusing on $\pi = r = 0$. When $r \rightarrow 0$ our problem is equivalent to minimizing the steady state cost. The choice of $\pi = r = 0$ simplifies the comparison of the analytical results with the ones for the original BT model and with the ones of Section 3.

Proposition 8. *Let $\pi = 0$ and $r \rightarrow 0$, the ratios: W/M , \underline{M}/M and $(M/c)p$ are determined by three strictly monotone functions of \hat{b}/R that satisfy:*

$$\begin{aligned} \text{As } \frac{\hat{b}}{R} \rightarrow 0 : \quad & \frac{W}{M} \rightarrow 2, \quad \frac{\underline{M}}{M} \rightarrow 0, \quad \frac{\partial \log \frac{Mp}{c}}{\partial \log \frac{\hat{b}}{R}} \rightarrow \frac{1}{2}. \\ \text{As } \frac{\hat{b}}{R} \rightarrow \infty : \quad & \frac{W}{M} \rightarrow 0, \quad \frac{\underline{M}}{M} \rightarrow 1, \quad \frac{\partial \log \frac{Mp}{c}}{\partial \log \frac{\hat{b}}{R}} \rightarrow 0. \end{aligned}$$

Proof. See Appendix A.

The elasticity of $(M/c)p$ with respect to \hat{b}/R determines the effect of the technological parameters b/c and p on the level of money demand, as well as on the interest rate elasticity of M/c with respect to R since

$$\eta(\hat{b}/R) \equiv \frac{\partial \log(M/c)p}{\partial \log(\hat{b}/R)} = -\frac{\partial \log(M/c)}{\partial \log R}. \quad (25)$$

Direct computation gives that

$$\frac{\partial \log(M/c)}{\partial \log p} = -1 + 2\eta(\hat{b}/R) \leq 0 \quad \text{and} \quad 0 \leq \frac{\partial \log(M/c)}{\partial \log(b/c)} = \eta(\hat{b}/R). \quad (26)$$

The previous sections showed that p has two opposing effects on M/c : for a given m^*/c , the value of M/c increases with p , but the optimal choice of m^*/c decreases with p . Proposition 8 and equation (26) show that the net effect is always negative. For low values of \hat{b}/R , where $\eta \approx 1/2$, the elasticity of M/c with respect to p is close to zero and the one with respect to b/c is close to $1/2$, which is the BT case. For large values of \hat{b}/R , the elasticity of M/c with respect to p goes to -1 , and the one with respect to b/c goes to zero. Likewise, equation (26) implies that $\partial \log M / \partial \log c = 1 - \eta$ and hence that the expenditure elasticity of the money demand ranges between $1/2$ (the BT value) and 1 as \hat{b}/R becomes large.

In the original BT model $W/M = 2$, $\underline{M}/M = 0$ and $\frac{\partial \log(M/c)}{\partial \log R} = -1/2$ for all b/c and R . These are the values that correspond to our model as $\hat{b}/R \rightarrow 0$. This limit includes the standard case where $p \rightarrow 0$, but it also includes the case where b/c is much smaller than p^2/R . As \hat{b}/R grows, our model predicts smaller interest rate elasticity than the BT model, and in the limit, as $\hat{b}/R \rightarrow \infty$, that the elasticity goes to zero. This result is a smooth version of the one for the model with p deterministic free withdrawal opportunities of Section 3. In that model the elasticity $\partial \log(Mp/c)/\partial \log(\hat{b}/R)$ is a step function that takes two values, $1/2$ for low values of \hat{b}/R , and zero otherwise. The smoothness is a natural consequence of the randomness on the free withdrawal opportunities. One key difference is that the deterministic model of Section 3 has no precautionary motive for money demand, hence $W/M = 2$ and $\underline{M}/M = 0$. Instead, as Proposition 8 shows, in the model with random free withdrawal opportunities, the strength of the precautionary motive, as measured by W/M and \underline{M}/M , is a function of \hat{b}/R .


Figure 1 plots W/M , \underline{M}/M and η as functions of \hat{b}/R . This figure completely characterizes the shape of the money demand and the strength of the precautionary motive since the functions plotted in it depend only on \hat{b}/R . The range of the \hat{b}/R values used in this figure is chosen to span the variation of the estimates presented in Table 5. While this figure is based on results for $\pi = r = 0$, the figure obtained using the values of π and r that correspond to the averages for Italy during 1993-2004 is quantitatively indistinguishable.

We conclude this section with a result on the welfare cost of inflation and the effect technological change. Let (R, κ) be the vector of parameters that index the value function $V(m; R, \kappa)$ and the invariant distribution $h(m; R, \kappa)$, where $\kappa = (\pi, r, b, p, c)$. We define the average flow cost of cash purchases borne by households $v(R, \kappa) \equiv \int_0^{m^*} rV(m; R, \kappa)h(m; R, \kappa)dm$. We measure the benefit of lower inflation for households, say as captured by a lower R and π , or of a better technology, say as captured by a lower b/c or a higher p , by comparing $v(\cdot)$ for the corresponding values of (R, κ) . A related concept is $\ell(R, \kappa)$, the expected withdrawal cost borne by households that follow the optimal rule

$$\ell(R, \kappa) = [n(m^*(R, \kappa), p, \pi) - p] \cdot b \quad (27)$$

where n is given in (18) and the expected number of free withdrawals, p , are subtracted. The value of $\ell(R, \kappa)$ measures the resources wasted trying to economize on cash balances, i.e. the deadweight loss for the society corresponding to R . While ℓ

Figure 1: W/M , \underline{M}/M , m^*/M and $\eta = \text{elasticity of } (M/c)p$
For $\pi = 0$ and $r \rightarrow 0$



is the relevant measure of the cost for the society, we find useful to define v separately to measure the consumers' benefit of using ATM cards. The next proposition characterizes $\ell(R, \kappa)$ and $v(R, \kappa)$ as $r \rightarrow 0$. This limit is useful for comparison with the BT model and it also turns out to be an excellent approximation for the values of r that we use in our estimation.

Proposition 9. *Let $r \rightarrow 0$: (i) $v(R, \kappa) = R m^*(R, \kappa)$; (ii) $v(R, \kappa) = \int_0^R M(\tilde{R}, \kappa) d\tilde{R}$, and (iii) $\ell(R, \kappa) = v(R, \kappa) - R M(R, \kappa)$.*

Proof. See Appendix A.

This proposition allows us to estimate the effect of inflation or technology on agents' welfare using data on W and \underline{M} , since $W + \underline{M} = m^*$. In the BT model $\ell = RM = \sqrt{Rbc}/2$ since $m^* = W = 2M$. In our model $m^*/M = W/M + \underline{M}/M < 2$, as can be seen in Figure 1, thus using RM as an estimate of $R(m^* - M)$ produces an overestimate of the cost of inflation ℓ . For instance, for $\hat{b}/R = 1.8$, the BT welfare cost measure overestimates the cost of inflation by about 60%, since $m^*/M \cong 1.6$.

Clearly the loss for society is smaller than the cost for households; using (i)-(iii) and Figure 1 the two can be easily compared. As \hat{b}/R ranges from zero to ∞ , the

ratio of the costs ℓ/v decreases from 1/2, the BT value, to zero. Not surprisingly (ii)-(iii) implies that the loss for society coincides with the consumer surplus that can be gained by reducing R to zero, i.e. $\ell(R) = \int_0^R M(\tilde{R})d\tilde{R} - RM(R)$. This extends the result of Lucas (2000), derived from a money-in-the-utility-function model, to an explicit inventory-theoretic model. Measuring the welfare cost of inflation using the consumer surplus requires the estimation of the money demand for different interest rates, while the approach using (i) and (iii) can be done using information on M , W and \underline{M} . Section 8 presents an application of these results and a comparison with the ones by Lucas (2000).⁸

5 A model with costly random withdrawals

The dynamic model discussed above has the unrealistic feature that agents withdraw every time a match with a financial intermediary occurs, so that many of the withdrawals have a very small size. This section extends the model to the case where the withdrawals done upon the random contacts are subject to a fixed cost f , assuming $0 < f < b$. The model produces a more realistic distribution of withdrawals, by limiting the minimum withdrawal size.

We skip the formulation of the total cost problem, that is exactly parallel to the one for the case of $f = 0$. With notation analogous to the one used above, the Bellman equation when the agent is not matched with an intermediary is

$$rV(m) = Rm + p \min \{V^* + f - V(m), 0\} + V'(m)(-c - m\pi) \quad (28)$$

where $V^* \equiv \min_{\hat{m}} V(\hat{m})$ and $\min \{V^* + f - V(m), 0\}$ takes into account that it may not be optimal to withdraw for all contacts. Indeed, whether the agent chooses to do so depend on the level of cash balances.

It can be shown that $V(\cdot)$ is strictly decreasing at $m = 0$ with a unique value of m^* such that $V'(m^*) = 0$. Then there will be two thresholds, \underline{m} and \bar{m} , that satisfy $V^* + f = V(\underline{m}) = V(\bar{m})$. Thus solving the Bellman equation is equivalent to finding 5 numbers m^* , m^{**} , \underline{m} , \bar{m} , V^* and a function $V(\cdot)$ such that: $V^* = V(m^*)$,

⁸In (ii)-(iii) we measure welfare and consumer surplus with respect to variations in R , keeping π fixed. The effect on M and v of changes in π for a constant R are quantitatively small.

$$V'(m^*) = 0,$$

$$V(m) = \begin{cases} \frac{Rm + p(V^* + f) - V'(m)(c + m\pi)}{r + p} & \text{if } m \in (0, \underline{m}) \\ \frac{Rm - V'(m)(c + m\pi)}{r + p} & \text{if } m \in (\underline{m}, \bar{m}) \\ \frac{Rm + p(V^* + f) - V'(m)(c + m\pi)}{r + p} & \text{if } m \in (\bar{m}, m^{**}) \end{cases}$$

and the boundary conditions: $V(0) = V^* + b$, $V(m) = V^* + b$ for $m > m^{**}$.

The optimal policy in this model is to pay the fixed cost f and withdraw cash if the contact with the intermediary occurs when cash balances are in $(0, \underline{m})$ range, or to deposit if cash balances are larger than \bar{m} . The withdrawal (or the deposit) is such that the post transfer cash balances are equal to m^* . If the agent contacts an intermediary when her cash balances are in (\underline{m}, \bar{m}) then no action is taken. If the agent cash balances get to zero, then the fixed cost b is paid and the cash balances are set to m^* . Notice that $m^* \in (\underline{m}, \bar{m})$. Hence in this model withdrawals have a minimum size given by $m^* - \underline{m}$.

Following steps that are analogous to the ones for the model where $f = 0$, one can characterize the value function, obtain expressions for m^* , \underline{m} , the invariant distribution $h(m)$ and the statistics: n , W , M , \underline{M} . It can be shown that the minimum withdrawal size is determined by the fixed cost relative to the interest cost, i.e. f/R , and it is independent of p . For instance, when $\pi = 0$ the range of inaction satisfies $(m^* - \underline{m})/c = \sqrt{(2f)/(Rc)} + o(\sqrt{(2f)/(Rc)})$. Hence the minimum withdrawal does not depend on p and b , and is analogous to the withdrawal of the BT model with a fixed cost f and an interest rate R .⁹

6 Estimation of the model

This section estimates the parameters (p, b) of the model presented in Section 4 using the household data set described in Section 2. As we explain below, this data is not rich enough to estimate f precisely, so we concentrate on the version of the model with $f = 0$. Our estimation procedure selects parameter values for (p, b) to produce values for $(M/c, W/M, n, \underline{M}/M)$ that are closest to the corresponding quantities in the data, for each year, geographic-location and household type. In this section we also discuss the nature of the measurement error, and the identification of the

⁹The working paper version contains statements and proofs of the results in this paragraph.

parameters. We finish the section by assessing the goodness of fit of model in a variety of ways.

For estimation we aggregate the household level data for each year, geographical location, and household type. In the baseline case the geographic location is a province. The household type is defined by grouping households in each year and province according to the level of cash consumption and whether they own an ATM card or not. In the baseline case we use three cash consumption groups containing an equal number of households. This yields about 3,600 cells, the product of 103 provinces, 6 years, 2 ATM ownership status, and 3 cash consumption levels.

In the following discussion we fix a particular combination of year-province-type. We let i index the household in that province-year-type combination. For all households in that cell we assume that b_i/c_i and p_i are identical. Given the homogeneity of the optimal decision rules, these assumptions allow us to aggregate the decisions of different households in a given province-year-type.

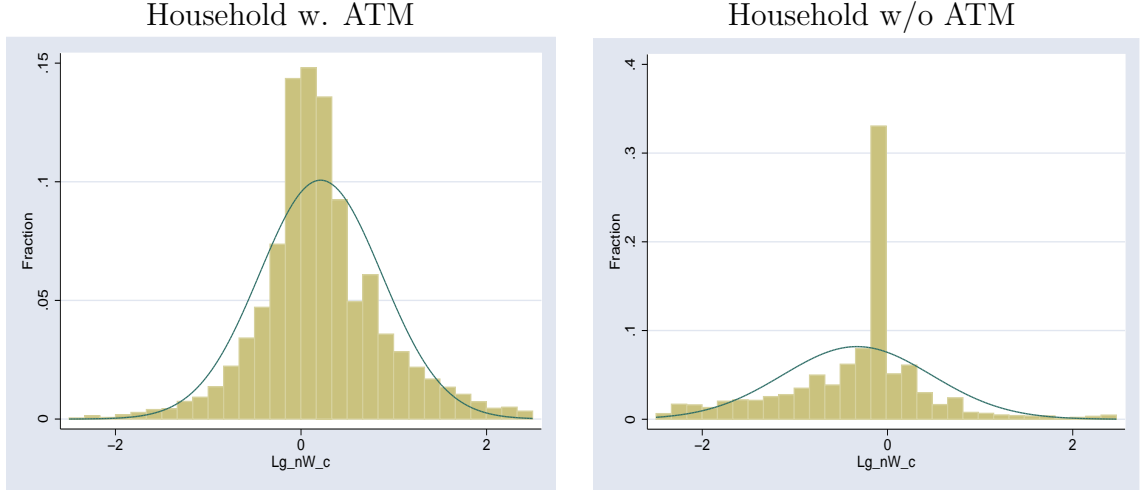
We assume that the variables M/c , W/M , n and \underline{M}/M , which we index as $j = 1, 2, 3$ and 4, are measured with a multiplicative error (additive in logs). Let z_i^j be the (log of the) i -th observation on variable j , and $\zeta^j(\theta)$ the (log of the) model prediction of the j variable for the parameter vector $\theta \equiv (p, b/c)$. The number N_j is the sample size of the variable j (the data set has different number of observations for different variables j). The idea behind this formulation is that the variable z_i^j is observed with a measurement error ε_i^j which has zero expected value and variance σ_j^2 so that $z_i^j = \zeta^j(\theta) + \varepsilon_i^j$ where the errors ε_i^j are assumed to be independent across households i and across variables j .

An illustration of the extent of the measurement error can be derived by assuming that the data satisfy the identity for the cash flows $c = nW - \pi M$ which holds in a large class of models (see the working paper version). Figure 2 reports a histogram of the logarithm of $n(W/c) - \pi(M/c)$ for each type of household. In the absence of measurement error, all the mass should be located at zero. It is clear that the data deviate from this value for many households.¹⁰

We estimate the vector of parameters θ for each province-year-type by minimizing

¹⁰Besides measurement error in reporting, which is important in this type of survey, there is also the issue of whether households have an alternative source of cash. An example of such as source occurs if households are paid in cash. This will imply that they do require fewer withdrawals to finance the same flow of consumption or, alternatively, that they effectively have more trips per periods.

Figure 2: Measurement error: deviation from the cash flow identity



the objective function

$$F(\theta; z) \equiv \sum_{j=1}^4 \left(\frac{N_j}{\sigma_j^2} \right) \left(\frac{1}{N_j} \sum_{i=1}^{N_j} z_i^j - \zeta^j(\theta) \right)^2 \quad (29)$$

where σ_j^2 is the variance of the measurement error for the variable j . Minimizing F yields the maximum likelihood estimator provided the ε_i^j are independent across j for each i . The average number of observations (N_j) available for each variable is similar for households with and without ATM cards. There are more observations on M/c than for each of the other three variable, and its average weight (N_1/σ_1^2) is about 1.5 times larger than each of the other three weights (see the working paper version for further documentation).

6.1 Estimation and Identification

In this section we discuss the features of the data that identify our parameters. We argue that with our data set we can identify $(p, \frac{b}{cR})$ and test the model with $f = 0$. As a first step we study how to select the parameters to match M/c and n only, as opposed to $(M/c, n, W/M, \underline{M}/M)$. To simplify the exposition here, assume that inflation is zero, so that $\pi = 0$. For the BT model, i.e. for $p = 0$, we have $W = m^*$, $c = m^* n$ and $M = m^*/2$ which implies $2 M/c = 1/n$. Hence, if

the data were generated by the BT model, M/c and n would have to satisfy this relation. Now consider the average cash balances generated by a policy like the one of the model of Section 4 with zero inflation. From (18) and (21), for a given value of p and setting $\pi = 0$, we have:

$$\frac{M}{c} = \frac{1}{p} [n m^*/c - 1] \quad \text{and} \quad n = \frac{p}{1 - \exp(-pm^*/c)} \quad (30)$$

or, solving for M/c as a function of n :

$$\frac{M}{c} = \xi(n, p) = \frac{1}{p} \left[-\frac{n}{p} \log \left(1 - \frac{p}{n} \right) - 1 \right]. \quad (31)$$

For a given p , the pairs $M/c = \xi(n, p)$ and n are consistent with a cash management policy of replenishing balances to some value m^* either when the zero balance is reached or when a chance meeting with an intermediary occurs. Notice first that setting $p = 0$ in this equation we obtain BT, i.e. $\xi(n, 0) = (1/2)/n$. Second, notice that this function is defined only for $n \geq p$. Furthermore, note that for $p > 0$: $\frac{\partial \xi}{\partial n} \leq 0$, $\frac{\partial^2 \xi}{\partial n^2} > 0$, and $\frac{\partial \xi}{\partial p} > 0$. Consider plotting the target value of the data on the $(n, M/c)$ plane. For a given M/c , there is a minimum n that the model can generate, namely the value $(1/2)/(M/c)$. Given that $\partial \xi / \partial p > 0$, any value of n smaller than the one implied by the BT model cannot be made consistent with our model, regardless of the values for the rest of the parameters. By the same reason, any value of n higher than $(1/2)/(M/c)$ can be accommodated by an appropriate choice of p . This is quite intuitive: relative to the BT model, our model can generate a larger number of withdrawals for the same M/c if the agent meets an intermediary often enough, i.e. if p is large enough. On the other hand there is a minimum number of expected chance meetings, namely $p = 0$.

The previous discussion showed that p is identified. Specifically, fix a province-year-type of household combination, with its corresponding values for M/c and n . Then, solving $M/c = \xi(n, p)$ for p gives an estimate of p . Taking this value of p , and those of M/c and n for this province-year-type combination, we use (30) to solve for m^*/c . Finally, we find the value of $\beta \equiv b/(cR)$ consistent with this replenishment target by solving the equation for m^* given in Proposition 2,

$$\beta \equiv \frac{b}{cR} = \frac{\exp[(r+p)m^*/c] - [1 + (r+p)(m^*/c)]}{(r+p)^2}. \quad (32)$$

To understand this expression consider two pairs $(M/c, n)$, both on the locus defined

by $\xi(\cdot, p)$ for a given value of p . The pair with higher M/c and lower n corresponds to a higher value of β . This is quite simple: agents will economize on trips to the financial intermediary if β is high, i.e. if these trips are expensive relative to the opportunity cost of cash. Hence, data on M/c and n identify p and β . Using data on R for this province-year, we can estimate b/c .

Figure 3 plots the function $\xi(\cdot, p)$ for several values of p , as well as the average value of M/c and n for all households of a given type (i.e. with and without ATM cards) for each province-year in our data (to make the graph easier to read we do not plot different consumption cells for a given province-year-ATM ownership). Notice that 46 percent of province-year pairs for households without an ATM card are below the $\xi(\cdot, 0)$ line, so no parameters in our model can rationalize those choices. The corresponding value for those with an ATM card is only 3.5 percent of the pairs. The values of p required to rationalize the average choice for most province-year pairs for those households without ATM cards are in the range $p = 0$ to $p = 20$. The corresponding range for those with ATM cards is between $p = 5$ and $p = 60$. Inspecting this figure we can also see that the observations for households with ATM cards are to the south-east of those for households without ATM cards. Equivalently, we can see that for the same value of p , the observations that correspond to households with ATM tend to have lower values of β .

Now we turn to the analysis of the ratio of the average withdrawal to the average cash balances, W/M . As in the previous case, consider an agent that follows an arbitrary policy of replenishing her cash to a return level m^* , either as her cash balances gets to zero, or at the time of chance meeting with the intermediary. Again, to simplify consider the case of $\pi = 0$. Using the cash flow identity $nW = c$ and (31) yields


$$\frac{W}{M} = \delta(n, p) \equiv \left[\frac{1}{p/n} + \frac{1}{\log(1 - p/n)} \right]^{-1} - \frac{p}{n} \quad (33)$$

for $n \geq p$, and $p \geq 0$. Some algebra shows that:

$$\delta(n, 0) = 2, \quad \delta(n, n) = 0, \quad \frac{\partial \delta(n; p)}{\partial p} < 0, \quad \frac{\partial \delta(n; p)}{\partial n} > 0.$$

Notice that the ratio W/M is a function only of the ratio p/n . The interpretation of this is clear: for $p = 0$ we have $W/M = 2$, as in the BT model. This is the highest value that can be achieved of the ratio W/M . As p increases for a fixed n , the replenishing level of cash m^*/c must be smaller, and hence the average withdrawal becomes smaller relative the average cash holdings M/c . Indeed, as n converges to

Figure 3: Theory vs. data (province-year mean): $M/c, n$



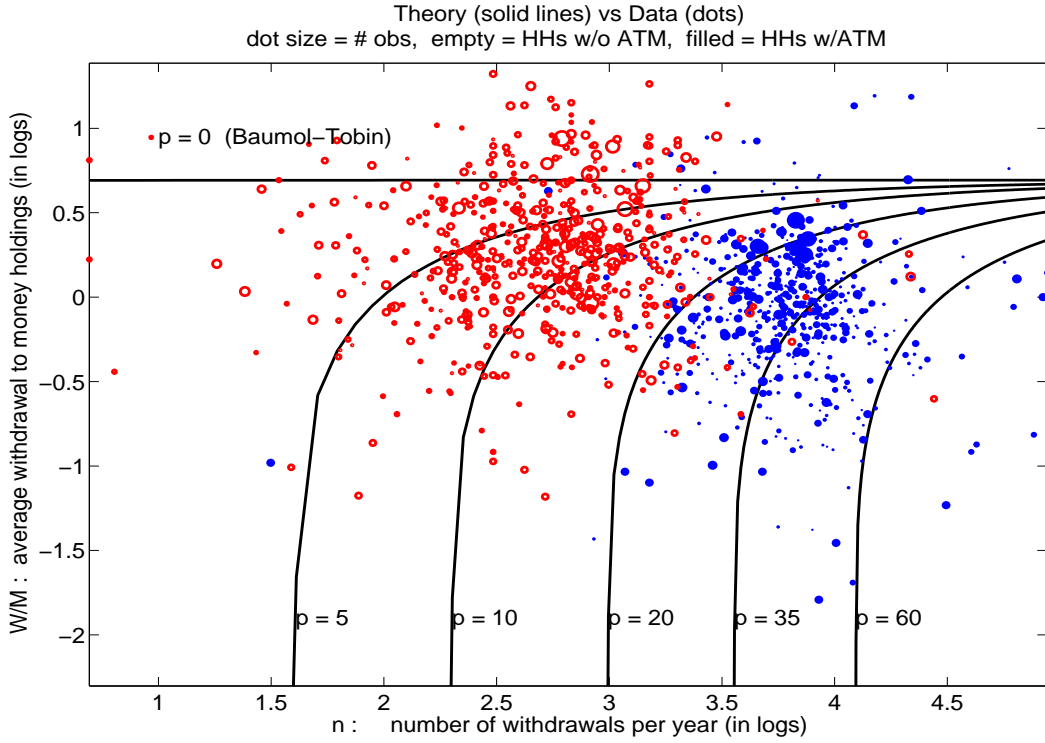
p – a case where almost all the withdrawals are due to chance meetings with the intermediary–, then W/M goes to zero.

As in the previous case, given a pair of observations on W/M and n , we can use δ to solve for the corresponding p . Then, using the values of $(W/M, p, n)$ we can find a value of $(b/c)/R$ to rationalize the choice of W/M . To see how, notice that given W/M , M/c , and p/n , we can find the value of m^*/c using $\frac{W}{M} = \frac{m^*/c}{M/c} - \frac{p}{n}$ (equation 23). With the values of $(m^*/c, p)$ we can find the unique value of $\beta = (b/c)/R$ that rationalizes this choice, using (32). Thus, data on W/M and n identifies p .

Figure 4 plots the function $\delta(n, p)$ for several values of p , as well as the average values of n and W/M for the different province-year-household type combinations for our data set (as done above, we omit the cash expenditure split to make the figure easier to read). We note that about 3 percent of the province-year pairs for households with an ATM cards have W/M above 2, while for those without ATM card the corresponding value is 15 percent. In this case, as opposed to the experiment displayed in Figure 3, no data on the average cash expenditure flow (c) is used, thus it may be that these smallest percentages are due to larger measurement error on c . The implied values of p needed to rationalize these data are similar to

the ones found using the information of M/c and n displayed in Figure 3. Also the implied values of β that corresponds to the same p tend to be smaller for households with an ATM card since the observations are to the south-east.

Figure 4: Theory vs. data (province-year mean): $W/M, n$

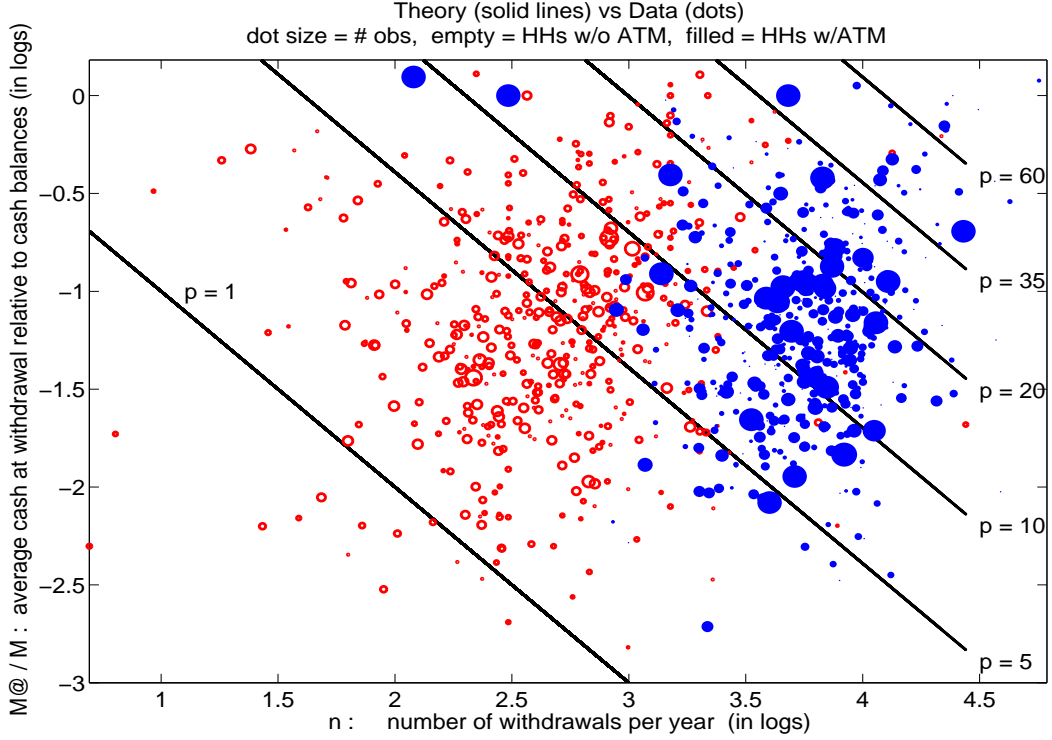


Finally we discuss the ratio between the average cash at withdrawals and the average cash: \underline{M}/M . In (24) we have derived that $p = n \cdot (\underline{M}/M)$. We use this equation as a way to estimate p . If \underline{M} is zero, then p must be zero. Hence the fact that $\underline{M}/M > 0$, documented in Table 1, is an indication that our model requires $p > 0$. We can readily use this equation to estimate p since we have data on both n and (\underline{M}/M) . According to this formula a large value of p is consistent with either a large ratio of cash at withdrawals, \underline{M}/M , or a large number of withdrawals, n . Also, for a fixed p , different combination of n and \underline{M}/M that give the same product are due to differences in $\beta = (b/c)/R$. If β is high, then agents economize in the number of withdrawals n and keep larger cash balances.

Figure 5 plots the average logarithm of \underline{M}/M and n , as well as lines corresponding different hypothetical values of p for each province-year for households with and without ATM. The fraction of province-years where $\underline{M}/M > 1$, is less than 3 percent for both types of households. The ranges of values of p needed to rationalize the

choices of households with and without ATM across the province-years is similar than the ones in the previous two figures. Also, for a given p the observations for households with ATM correspond to lower values of β (i.e. they are to the south-east of those without ATM cards).

Figure 5: Theory vs. data (province-year mean): $\underline{M}/M, n$



We have discussed how data on either of the pairs $(M/c, n)$, $(W/M, n)$ or $(\underline{M}/M, n)$ identify p and β . Of course, if the data had been generated by the model, the three ways of estimating (p, β) would produce identical estimates. In other words, the model is overidentified. We will use this idea to report how well the model fits the data or, more formally, to test for the overidentifying restrictions in the next subsection.

Considering the case of $\pi > 0$ makes the expressions more complex, but, at least qualitatively, does not change any of the properties discussed above. Moreover, quantitatively, since the inflation rate in our data set is quite low the expressions for $\pi = 0$ approximate the relevant range for $\pi > 0$ very well.

6.2 Estimation results

We estimate the $f = 0$ model for each province-year-type of household and report statistics of the estimates in Table 3. For each year we use the inflation rate corresponding to the Italian CPI for all provinces and fix the real return r to be 2% per year. The first two panels in the table report the mean, median, 95th and 5th percentile of the estimated values for p and b/c across all province-year. As explained above, our procedure estimates $\beta \equiv \frac{b}{cR}$, so to obtain b/c we compute the opportunity cost R as the sum of the nominal interest rate and the probability of cash being stolen described in Table 2. The parameter p gives the average number of free withdrawals opportunities per year. The parameter $b/c \cdot 100$ is the cost of a withdrawal in percentage of the daily cash-expenditure. We also report the mean value of the t statistics for these parameters. The standard errors are computed by solving for the information matrix.

The results reported in the first two columns of the table concern households who possess an ATM card, shown separately for those in the lowest and highest cash expenditure levels. The corresponding statistics for households without ATM card appear in the third and fourth columns. The results in this table confirm the graphical analysis of figures 3-5 discussed in the previous section: the median estimates of p are just where one would locate them by the figures. The difference between the 95th and the 5th percentiles indicates that there is a tremendous amount of heterogeneity across province-years. The relatively low values for the mean t -statistics reflect the fact that the number of households used in each estimation cell is small. Indeed, in the working paper version we consider different levels of aggregation and data selection. In all the cases considered we find very similar values for the average of the parameters p and b/c , and we find that when we do not disaggregate the data so much the average t -stats increase roughly with the (square root) of the average number of observations per cell.¹¹

Table 3 shows that the average value of b/c across all province-year-type is between 2 and 10 per cent of daily cash consumption. Fixing an ATM ownership type, and comparing the average estimates for p and b/c across cash consumption cells we see that there are small differences for p , but that b/c is substantially smaller for

¹¹Concerning aggregation, we repeat all the estimates without disaggregating by the level of cash consumption, so that N_j is three times larger. Concerning data selection, we repeat all the estimates excluding those observations where the cash holding identity is violated by more than 200% or where the share of total income received in cash by the household exceeds 50%. The goal of this data selection, that roughly halves the sample size, is to explore the robustness of the estimates to measurement error.

Table 3: Summary of $(p, b/c)$ estimates across province-year-types

<i>Cash expenditure</i> ^a :	Household w/o ATM		Household w. ATM	
	<i>Low</i>	<i>High</i>	<i>Low</i>	<i>High</i>
	<i>Parameter p</i>			
Mean	6.8	8.7	20	25
Median	5.6	6.2	17	20
95 th percentile	17	25	49	61
5 th percentile	1.1	0.8	3	4
Mean t-stat	2.5	2.2	2.7	3.5
	<i>Parameter b/c (in % of daily cash expenditure)</i>			
Mean	10.5	5.5	6.5	2.1
Median	7.3	3.6	3.5	1.1
95 th percentile	30	17	24	7
5 th percentile	1.5	0.4	0.6	0.3
Mean t-stat	2.8	2.5	2.4	3.3
# prov-year-type estimates	504	505	525	569
	<i>Goodness of fit: Objective function $F(\theta, x) \sim \chi^2$</i>			
% province-years-type where:				
- $F(\theta, x) < 4.6$ ^b	64%		57%	
- Hp. $f = 0$ is rejected ^c	2%		19%	
# prov-year-type estimates	1,539		1,654	
Avg. # of households per estimate	10.7		13.5	

Summary statistics for the estimates of $(p, b/c)$ obtained from each of the 1,854 province-year-type cells. All the lines except one (see note ^c) report statistics from the model with $f = 0$.

Notes: - ^a Low (high) denotes the lowest (highest) third of households ranked by cash expenditure. - ^b Percentage of province-year-type estimates where the overidentifying restriction test is not rejected at the 10 per cent confidence level. - ^c Percentage of estimates where the null hypothesis of $f = 0$ is rejected by a likelihood ratio test at the 5% confidence level.

the those in the highest cash consumption cell. Indeed, combining this information with the level of cash consumption that corresponds to each cell we estimate b to be uncorrelated with cash consumption levels, as documented in Section 7. Using information from Table 1 for the corresponding cash expenditure to which these percentages refer, the mean values of b for households with and without ATM are 0.8 and 1.7 euros at year 2004 prices, respectively. For comparison, the cash withdrawal charge for own-bank transactions was zero, while the average charge for other-bank transactions, which account for less than 20 % of the total, was 2.0 euros.¹²

Next we discuss four different types of evidence that indicate a successful empirical performance of the model. First, Table 3 shows that households with ATM

¹²The sources are Retail Banking Research (2005) and an internal report by the Bank of Italy.

cards have a higher mean and median value of p and correspondingly lower values of b/c . The comparison of the $(p, b/c)$ estimates across province-year-consumption cells shows that 88 percent of the estimated values of p are higher for households with ATM, and for 82 percent of the estimated values of b/c are lower. Also, there is evidence of an effect at the level of the province-year-consumption cell, since we find that the correlation between the estimated values of b/c for households with and without ATM across province-year-consumption cell is 0.69. The same statistic for p is 0.3. These patterns are consistent with the hypothesis that households with ATM cards have access to a more efficient transactions system, and that the efficiency of the transaction technology in a given province-year-consumption cell is correlated for both ATM and non-ATM adopters. We find this result reassuring since we have estimated the model for ATM holders and non-holders and for each province-year-consumption cell separately.

Second, in the third panel of Table 3 we report statistics on the goodness of fit of the model. For each province-year-type cell, under the assumption of normally distributed errors, or as an asymptotic result, the minimized objective function is distributed as a $\chi^2_{(2)}$. According to the statistic reported in the first line of this panel, in more than half of the province-years-consumption cells the minimized objective function is smaller than the critical value corresponding to a 10% probability confidence level.

Third, we examine the extent to which imposing the constraint that $f = 0$ diminishes the ability of the model to fit the data. To do so we reestimated the model letting f/c vary across province-years-households type, and compare the fit of the restricted ($f = 0$) with the unrestricted model using a likelihood ratio test. The second line of the panel reports the percentage of province-years-consumption cells where the null hypothesis of $f = 0$ is rejected at a 5% confidence level. Only for a small fraction of cases (19% for those cells that correspond to households with ATM cards, and 2% for those without cards) there is evidence of an improvement in the fit of the model by letting $f > 0$. Indeed we find (not reported here) that when we let $f > 0$ and estimate the model for each province-year-type, the average t-statistic of the parameters $(p, b/c, f/c)$ are very low, the average correlation between the estimates is extremely high, and there is an extremely high variability in the estimated parameters across province-years. We conclude that our data set does not allow us to estimate p , b/c and f/c with a reasonable degree of precision. As explained above, the reason we consider the $f > 0$ model is to eliminate the extremely small withdrawals that $f = 0$ model implies. Hence, what would be

helpful to estimate f is information on the minimum size of withdrawals, or some other feature of the withdrawal distribution.

Table 4: Correlations between $(p, \frac{b}{c}, V(0))$ estimates and financial diffusion indices

Household with ATM			
	p	b/c	$V(0)$
Bank-branch per 1,000 head	0.08	-0.19	-0.18
ATM per 1,000 head	0.10	-0.27	-0.27
Household with No ATM			
	p	b/c	$V(0)$
Bank-branch per 1,000 head	0.00 ^a	-0.26	-0.20

Notes: All variables are measured in logs. The sample size is 1,654 for HH w. ATM and 1,539 for HH without ATM. P-values (not reported), computed assuming that the estimates are independent, are smaller than 1 per cent with the exception of the one denoted by a .

Fourth, in Table 4 we compute correlations of the estimates of the technological parameters p , b/c and the cost of financing cash purchases $V(0)$ with indicators that measure the density of financial intermediaries: bank branches and ATMs per resident that vary across province and years. A greater financial diffusion raises the chances of a free withdrawal opportunity (p) and reduces the cost of contacting an intermediary (b/c). Hence we expect $V(0)$ to be negatively correlated with the diffusion measure. We find that the estimates of b/c and $V(0)$ are negatively correlated with these measures, and that the estimated p are positively correlated, though the latter correlation is smaller. This finding is reassuring since the indicators of financial diffusion are not used in the estimation of $(p, b/c)$.

7 Implications for money demand

In this section we study the implications of our findings for the time patterns of technology and for the expenditure and interest elasticity of the demand for currency.

We begin by documenting the trends in the withdrawal technology, as measured by our estimates of p and b/c . Table 5 shows that p has approximately doubled, and that (b/c) has approximately halved over the sample period. In words, the withdrawal technology has improved through time. The table also reports $\hat{b}/R \equiv (b/c)p^2/R$, which as shown in Proposition 8 and illustrated in Figure 1 determines the elasticity of the money demand and the strength of the precautionary motive. In particular, the proposition implies that W/M and \underline{M}/M depend only on \hat{b}/R . The

upward trend in the estimates of \hat{b}/R , which is mostly a reflection of the downward trend in the data for W/M , implies that the interest rate elasticity of the money demand has decreased through time.

Table 5: Time series pattern of estimated model parameters

	1993	1995	1998	2000	2002	2004	All years
<i>Households with ATM</i>							
p	17	16	20	24	22	33	22
$b/c \times 100$	6.6	5.7	2.8	3.1	2.8	3.5	4.0
\hat{b}/R	1.1	1.4	1.9	5.6	3.0	5.8	3.2
<i>Households without ATM</i>							
p	6	5	8	9	8	12	8
$b/c \times 100$	13	12	6.2	4.9	4.5	5.7	7.7
\hat{b}/R	0.2	0.2	0.4	0.4	0.4	1.6	0.5
$R \times 100$	8.5	7.3	4.3	3.9	3.2	2.9	5.0

R and p are *annual* rates, c is the *daily* cash expenditure rate, and for each province-year-type $\hat{b}/R = (b/c) p^2 / (365 R)$, which has no time dimension. Entries in the table are sample means across province-type in a year.

By Proposition 8, the interest rate elasticity $\eta(\hat{b}/R)$ implied by those estimates is smaller than $1/2$, the BT value. Using the mean of \hat{b}/R reported in the last column of Table 5 to evaluate the function η in Figure 1 yields values for the elasticity equal to 0.43 and 0.48 for households with and without ATM card, respectively. Even for the largest values of \hat{b}/R recorded in Table 5, the value of η remains above 0.4. In fact, further extending the range of Figure 1 it can be shown that values of \hat{b}/R close to 100 are required to obtain an elasticity η smaller than 0.25. For such high values of \hat{b}/R , the model implies \underline{M}/M of about 0.99 and W/M below 0.3, values reflecting much stronger precautionary demand for money than those observed for most Italian households. On the other hand, studies using cross sectional household data, such as Lippi and Secchi (2007) for Italian data, and Daniels and Murphy (1994) using US data, report interest rate elasticities smaller than 0.25.

A possible explanation for the difference in the estimated elasticities is that the cross sectional regressions in the studies mentioned above fail to include adequate measures of financial innovations, and hence the estimate of the interest rate elasticity is biased towards zero. To make this clear, in Table 6 we estimate the interest elasticity of M/c by running two regressions for each household type where M/c is

the model fitted value for each province-year-consumption type. The first regression includes the log of p , b/c and R . According to Proposition 8, (M/c) p has elasticity $\eta(\hat{b}/R)$ so that we approximate it using a constant elasticity:

$$\log M/c = -\log p + \eta(\log(b/c) + 2 \log(p)) - \eta \log(R) . \quad (34)$$

As expected the coefficient of the regressions following (34) gives essentially the same values for η as those obtained above using Figure 1. To estimate the size of the bias from omitting the variables $\log p$ and $\log b/c$, the second regression includes only $\log R$. The regression coefficient for $\log R$ is an order of magnitude smaller than the value of η , reflecting a large omitted variable bias. For instance, the correlation between $(\log(b/c) + 2 \log(p))$ and $\log R$ is 0.12 and 0.17 for households with and without ATM card, respectively. Interestingly, the regression coefficients on $\log R$ estimated by omitting the log of p and b/c are similar to the values that are reported in the literature mentioned above. Replicating the regressions of Table 6 using the actual, as opposed to the fitted, value of M/c as a dependent variable yields very similar results (not reported here).

Table 6: Interest elasticity of money demand

<i>Dependent variable:</i> $\log(M/c)$	Household w. ATM	Household w/o ATM		
$\log(p)$	-0.05	-	-0.01	-
$\log(b/c)$	0.45	-	0.48	-
$\log(R)$	-0.44	-0.07	-0.48	-0.04
R^2	0.985	0.01	0.996	0.004
# observations	1,654	1,654	1,539	1,539

Notes: All regressions include a constant.

We now estimate the expenditure elasticity of the money demand. An advantage of our dataset is that we use direct measures of cash expenditures (as opposed to income or wealth).¹³ By Proposition 8, the expenditure elasticity is

$$\frac{\partial \log M}{\partial \log c} = 1 + \eta(\hat{b}/R) \frac{\partial \log b/c}{\partial \log c} . \quad (35)$$

For instance, if the ratio b/c is constant across values of c then the elasticity is one; alternatively, if b/c decreases proportionately with c the elasticity is $1 - \eta$. Using the variation of the estimated b/c across time, locations and household groups with

¹³Dotsey (1988) argues for the use of cash expenditure as the appropriate scale variable.

different values of c , we estimate the elasticity of b/c with respect to c equal to -0.82 and -1.01 for households without and with ATM card, respectively. Using the estimates for η we obtain that the mean expenditure elasticity is $1 + 0.48 \times (-0.82) = 0.61$ for households without ATM, and 0.56 for those with.

8 Cost of inflation and Benefits of ATM card

We use the estimates of $(p, \frac{b}{c})$ to quantify the deadweight loss for the society and the cost for households of financing cash purchases and to discuss the benefits of ATM card ownership. In Section 4.5 we showed that the loss is $\ell = R(m^* - M)$ and the household cost is $v = Rm^*$. In the first panel of Table 7 we display the average of ℓ and of ℓ/c for each year. In 1993 the loss is 24 euros or 0.99 days of cash purchases.

Table 7: Deadweight loss ℓ and household cost v of cash purchases

	1993	1995	1998	2000	2002	2004	mean
ℓ (2004 euros, per household)	24	23	11	11	10	10	15
ℓ/c (in days of cash purchases)	0.99	0.85	0.46	0.42	0.39	0.40	0.59
ℓ/c under 1993 technology	0.99	0.90	0.72	0.71	0.67	0.66	0.78
v (2004 euros, per household)	51	49	25	25	22	25	33
v/c : avg. group / avg. all groups	w. ATM			w/o ATM			
- high c (top third ranked by c)	0.61			1.00			
- low c (bottom third ranked by c)	1.11			1.48			

Note: ℓ and v are averages weighted by the number of household type, and measured as annual flows. The average value of v/c across all groups is 1.31 days of cash purchases.

To put this quantity in perspective we relate it to the one in Lucas (2000), obtained by fitting a log-log money demand with an interest elasticity of $1/2$, which corresponds to the BT model. Figure 5 in his paper plots the welfare cost of inflation, denoted by w and defined as our ℓ , which for an opportunity cost R of 5%, is about 1.1% of US GDP. At the same R our deadweight loss ℓ is about 14 times smaller, or 0.08% of the annual income for Italian households y ($c \cdot 365$ accounts for about half of annual Italian GDP, $\ell/y = 0.59/(2 \cdot 365) \cong 0.08\%$). There are two reasons for this difference. The first is that for a given cost R and money demand M/c , the deadweight loss in our model is smaller than in BT (see Section 4.5). For instance for $R = 0.05$ and $\hat{b}/R = 1.8$, which is about our sample average, w/ℓ is about 1.6. The second is that the welfare cost is proportional to the level of the money demand:

multiplying M/y by a constant, multiplies ℓ/y by the same constant. In particular, Lucas fits US data with a much higher value of M/y than the one we use for Italy: 0.225 versus 0.026 at $R = 0.05$. This is because while we focus on currency held by households, he uses the stock of M1, an aggregate much larger than ours (including cash holdings of non-residents and firms).¹⁴

Table 7 also shows that by the end of the sample the welfare loss is about 40% smaller than its initial value. The reduction is explained by decreases in the opportunity cost R and by advances in the withdrawal technology, i.e. decreases in b/c and increases in p . To account for the contribution of these two determinants on the reduction of the deadweight loss we compute a counterfactual. For each province-type of household we freeze the values of p and b/c at those estimated for 1993, and compute ℓ/c for the opportunity cost R and inflation rates π corresponding to the subsequent years. We interpret the difference between the value of ℓ/c in 1993 and the value corresponding to subsequent years as the increase in welfare due to the Italian disinflation. We find that the contributions of the disinflation and of technological change to the reduction in the welfare loss are of similar magnitude (see the working paper version for details).

The bottom panel of Table 7 examines the cross section variation in the cost v/c . Comparing the values across columns shows that the cost is lower for households with ATM cards, reflecting their access to a better technology. Comparing the values across rows shows that the cost is lower for households with higher consumption purchases c , reflecting that our estimates of b/c are uncorrelated with the c .

We use v/c to quantify the benefits associated to the ownership of the ATM card. Under the maintained assumption that b is proportional to consumption *within* each year-province-consumption group type, the value of the benefit for an agent without ATM card, keeping cash purchases constant, is defined as: $v_0 - v_1 \frac{c_0}{c_1} = R(m_0^* - m_1^* \frac{c_0}{c_1})$, where the 1/0 subscript indicates ownership (lack of) ATM card. The benefit is thus computed assuming that the only characteristic that changes when comparing costs is ATM ownership (i.e. c is kept constant).¹⁵ Table 8 shows that the mean benefit of ATM card ownership ranges between 15 and 30 euros per year in the early sample and that it is smaller, between 4 and 13 euros, in 2004. The population weighted average of the benefits across all years and types is 17 euros (not reported in the table). The downward trend in the benefits is due to both the disinflation and the

¹⁴Hence the 14-fold difference in ℓ/y is given by the product of the factor 1.6 (the welfare cost ratio for a given level of money demand), and the factor 8.6 (the ratio of money demand levels).

¹⁵It is however noteworthy that within a consumption class (e.g. bottom or top third) the difference in c between households with and without ATM is very small (on average 4 per cent).

improvements in the technology, as discussed above. Table 8 also shows that the benefit is higher for household in the top third of the distribution of cash expenditure. This mainly reflects the different level of c of this group, since the benefit per unit of c is roughly independent of its level. The bottom panel of Table 8 shows that the benefit associated to ATM ownership is estimated to be positive for over 91 % of the province-year-type estimates. Two statistical tests are presented: the null hypothesis that the gain is positive cannot be rejected (at the 10 % confidence level) in 99.5 % of our estimates. Conversely, we are able to reject the null hypothesis that the benefit is negative in about 64% of the cases. Since our estimates of the parameters for households with and without ATM are done independently, we think that the finding that the estimated benefit is positive for most province-years provides additional support for the model.

Table 8: Annual benefit of ATM ownership (in euros at 2004 prices)

	1993	1995	1998	2000	2002	2004
<i>Top third of households ranked by c</i>						
Mean across province years	29	35	17	15	13	13
<i>Bottom third of households ranked by c</i>						
Mean across province years	17	14	6.6	5.5	3.6	4.4
Point estimate benefit > 0 91% of cells	H_o : benefit > 0 rejected		0.5% of cells		H_o : benefit < 0 rejected 64% of cells	

Note: Both hypothesis are rejected at the 10% confidence level. There are about 1,500 province-year-consumption group cells.

Two caveats are noteworthy about the above counterfactual exercise. First, the estimated benefit assumes that within a given province-year-consumption group households without ATM card differ from those with a card *only* in terms of the withdrawal technology that is available to them ($p, b/c$). In future work we plan to study the household choice of whether or not to have an ATM card, which will be informative on the size of the estimates' bias. The second caveat is that ATM cards provide other benefits, such as access to banking information and electronic funds transfers for retail transactions (EFTPOS payments), where the latter is particularly important in Italy. In spite of these caveats, our estimates of the annual benefit of ATM card ownership are close to annual cardholder fees for debit cards, which vary from 10 to 18 euros for most Italian banks over 2001-2005 (see page 35 and Figure 3.8.2 in Retail Banking Research Ltd., 2005).

References

- [1] Attanasio, O., L. Guiso and T. Jappelli, 2002. “The Demand for Money, Financial Innovation and the Welfare Cost of Inflation: An Analysis with Household Data”, **Journal of Political Economy**, Vol. 110, No. 2, pp. 318-351.
- [2] Cooley, Thomas and Gary Hansen, 1991. “The Welfare Costs of Moderate Inflation”, **Journal of Money, Credit and Banking**, Vol. 23, No. 3, pp. 483-503.
- [3] Constantinides, George M. and Scott F. Richard, 1978. “Existence of Optimal Simple Policies for Discounted-Cost Inventory and Cash Management in Continuous Time”, **Operations Research**, Vol. 26, No. 4, pp. 620-636.
- [4] Constantinides, George M., 1976. “Stochastic Cash Management with Fixed and Proportional Transaction Costs”, **Management Science**, Vol. 22 , pp. 1320-1331.
- [5] Dotsey, Michael , 1988. “The demand for currency in the United States”, **Journal of Money, Credit and Banking** , Vol. 20 , No. 1, pp. 22-40.
- [6] European Central Bank, 2006. *Blue Book. Payment and securities settlement systems in the european union and in the acceding countries*. December.
- [7] Frenkel, Jacob A. and Boyan Jovanovic (1980). “On transactions and precautionary demand for money ” **The Quarterly Journal of Economics**, Vol. 95, No. 1, pp. 25-43.
- [8] Humphrey, David B. , 2004. “Replacement of cash by cards in US consumer payments ”, **Journal of Economics and Business**, Vol. 56, Issue 3, pp. 211-225.
- [9] Lippi, F. and A. Secchi, 2007. “Technological change and the demand for currency: An analysis with household data”, CEPR discussion paper No. 6023.
- [10] Lucas, Robert E. Jr, 2000. “Inflation and Welfare”, **Econometrica**, Vol. 68(2), 247—74.
- [11] Miller, Merton and Daniel Orr, 1966. “A model of the demand for money by firms”, **Quarterly Journal of Economics**, Vol. 80, pp. 413-35.
- [12] Porter Richard D. and Ruth A. Judson, 1996. “The Location of US currency: How Much Is Abroad? ”, **Federal Reserve Bulletin**, October 1996, pp. 883–903.
- [13] Retail Banking Research Ltd., 2005. “Study of the Impact of Regulation 2560/2001 on Bank Charges for National Payments”, Prepared for the European Commission, London, September 2005.

Appendix

A Proofs for the model with free withdrawals

Proof of Proposition 1. Given two functions C , V satisfying (14) it is immediate to verify that the boundary conditions of the two systems at $m = 0$ and $m \geq m^{**}$ are equivalent. Also, it is immediate to show that for two such functions

$$m^* = \arg \min_{\hat{m} \geq 0} V(\hat{m}) = \arg \min_{\hat{m} \geq 0} \hat{m} + C(\hat{m}).$$

It only remains to be shown that the Bellman equations are equivalent for $m \in (0, m^{**})$. Using (14) we compute $C'(m) = V'(m) - 1$. Assume that $C(\cdot)$ solves the Bellman equation (7) in this range, inserting (14) and its derivative into (7) gives

$$[r + p_1 + p_2]V(m) = V'(m)(-c - \pi m) + [p_1 + p_2]V(m^*) + [r + p_2 + \pi]m.$$

Using $R = r + \pi + p_2$ and $p = p_1 + p_2$ we obtain the desired result, i.e. (12). The proof that if V solves the Bellman equation for $m \in (0, m^{**})$ so does C defined as in (14) follows from analogue steps. \square

Proof of Proposition 2. To solve for V^* , m^* , m^{**} and $V(\cdot)$ satisfying (11) and (12) we proceed as follows. Lemma 1 solves for $V(A, V^*)$, Lemma 2 gives $A(V^*)$. Lemma 3 shows that $V(\cdot)$ is convex for any $V^* > 0$. Lemma 4 solves for m^* using that, since V is convex, m^* must satisfy $V'(m^*) = 0$. Finally, Lemma 5 gives $V^* = V(m^*)$.

Lemmas 2, 4 and 5 yield a system of 3 equations in the 3 unknowns A, m^*, V^* :

$$A = \frac{V^*(r+p)r + Rc / \left(1 + \frac{\pi}{r+p}\right) + (r+p)^2 b}{c^2} > 0 \quad (36)$$

$$V^* = \frac{R}{r} m^* \quad (37)$$

$$m^* = \frac{c}{\pi} \left(\left[\frac{R}{Ac} / \left(1 + \frac{\pi}{r+p}\right) \right]^{-\frac{\pi}{r+p+\pi}} - 1 \right) \quad (38)$$

Replacing equation (37) into (36) yields one equation for A . Rearranging equation (38) we obtain another equation for A . Equating these expressions for A , collecting terms and rearranging yields equation (15) in the main text (which determines m^*).

To see that equation (15) has a unique non negative solution rewrite the equation as $f\left(\frac{m^*}{c}\right) = g\left(\frac{m^*}{c}\right)$ where the function f denotes the left hand side and g the right hand side. For $r + p + \pi > 0$, straightforward analysis shows that the solution exists and is unique. \square

Lemma 1. *Let V^* be an arbitrary value. The differential equation in (10) for $m \in (0, m^{**})$ is solved by the expression given in (16).*

Proof of Lemma 1. Follows by differentiation. \square

Lemma 2. Let V^* be an arbitrary non negative value. Let A be the constant that solves the ODE in Lemma 1. Imposing that this solution satisfies $V(0) = V^* + b$ the constant A is given by the expression in (36).

Proof of lemma 2. It follows using the expression in (16) to evaluate $V(0)$. \square

Lemma 3. Let V^* be an arbitrary value. The solution of V given in Lemma 1, with the value of A given in Lemma 2 is a convex function of m .

Proof of Lemma 3. Direct differentiation of V gives

$$V''(m) = \left(\frac{\pi}{r+p} \right) \left(1 + \frac{r+p}{\pi} \right) A \left[1 + \pi \frac{m}{c} \right]^{-\frac{r+p}{\pi}-2} > 0$$

since, as shown in Lemma 2, $A > 0$. \square

Lemma 4. Let A be an arbitrary value for the constant that indexes the solution of the ODE for V in Lemma 1, given by (16). The value m^* that solves $V'(m^*) = 0$ is given by the expression in (38).

Proof of Lemma 4. Follows using simple algebra. \square

Lemma 5. The value of V^* is $V^* = \frac{R}{r}m^*$

Proof of Lemma 5. Recall that at $m = m^*$ we have $V'(m^*) = 0$ and $V(m^*) = V^*$. Replacing these values in the Bellman equation (10) evaluated at $m = m^*$ yields $rV^* = Rm^*$. \square

Proof of Proposition 3. (i) The function $V(\cdot)$ is derived in Lemma 1, the expression for A in Lemma 2. (ii) The solution for V^* comes from Lemma 5. \square

Proof of Proposition 4. Proof of (i). Let $f(\cdot)$ and $g(\cdot)$ be the left hand side and the right hand side of equation (15) as a function of m^* . We know that $f(0) < g(0)$ for $b > 0$, $g'(0) = f'(0) > 0$, and $g''(m^*) = 0$, and $f''(m^*) > 0$ for all $m^* > 0$. Thus there exists a unique value of m^* that solves (15). Let $u(m^*) \equiv f(m^*) - g(m^*) + b/(cR)(r+p)(r+\pi+p)$. Notice that $u(m^*)$ is strictly increasing, convex, goes from $[0, \infty)$ and does not depend on $b/(cR)$. Simple analysis of $u(m^*)$ establishes the desired properties of m^* .

Proof of (ii). For this result we use that $f\left(\frac{m^*}{c}\right) = g\left(\frac{m^*}{c}\right)$ is equivalent to

$$\frac{b}{cR} = \left(\frac{m^*}{c} \right)^2 \left[\frac{1}{2} + \sum_{j=1}^{\infty} \frac{1}{(2+j)!} [\Pi_{s=1}^j (r+p-s\pi)] \left(\frac{m^*}{c} \right)^j \right] \quad (39)$$

which follows by expanding $\left(\frac{m}{c}\pi + 1\right)^{1+\frac{r+p}{\pi}}$ around $m = 0$. We notice that $m^*/c = \sqrt{\frac{2b}{cR}} + o\left(\sqrt{b/c}\right)$ is equivalent to $(m^*/c)^2 = \frac{2b}{cR} + \left[o\left(\sqrt{b/c}\right)\right]^2 + 2\sqrt{\frac{2b}{cR}} o\left(\sqrt{b/c}\right)$. Inserting this expression into (39), dividing both sides by $b/(cR)$ and taking the limit as $b/(cR) \rightarrow 0$ verifies our approximation.

Proof of (iii). For $\pi = R - r = 0$, using (39) we have

$$\frac{b}{cr} = \left(\frac{m^*}{c}\right)^2 \left[\frac{1}{2} + \sum_{j=1}^{\infty} \frac{1}{(j+2)!} (r+p)^j \left(\frac{m^*}{c}\right)^j \right]$$

To see that m^* is decreasing in p notice that the RHS is increasing in p and m . That $m^*(p+r)$ is increasing in p follows by noting that since $(m^*)^2$ decreases as p increases, then the term in square bracket, which is a function of $(r+p)m^*$, must increase. This implies that the elasticity of m^* with respect to p is smaller than $p/(p+r)$ since

$$0 < \frac{\partial}{\partial p} (m^* (p+r)) = m^* + (p+r) \frac{\partial m^*}{\partial p} = m^* \left[1 + \frac{(p+r)}{p} \frac{p}{m^*} \frac{\partial m^*}{\partial p} \right] \text{ thus}$$

$$\frac{(p+r)}{p} \frac{p}{m^*} \frac{\partial m^*}{\partial p} \geq -1 \text{ or } 0 \leq -\frac{p}{m^*} \frac{\partial m^*}{\partial p} \leq \frac{p}{p+r}.$$

Proof of (iv). For $\pi \rightarrow 0$, equation (15) yields: $\exp\left(\frac{m^*}{c}(r+p)\right) = 1 + \frac{m^*}{c}(r+p) + (r+p)^2 \frac{b}{cR}$. Replacing $\hat{b} \equiv (p+r)^2 b/c$ and $x \equiv m^*(r+p)/c$ into this expression, expanding the exponential, collecting terms and rearranging yields:

$x^2 \left[1 + \sum_{j=1}^{\infty} \frac{2}{(j+2)!} (x)^j \right] = 2 \frac{\hat{b}}{R}$. We now analyze the elasticity of x with respect to R . Letting $\varphi(x) \equiv \sum_{j=1}^{\infty} \frac{2}{(j+2)!} [x]^j$, we can write that x solves $x^2 [1 + \varphi(x)] = 2\hat{b}/R$. Taking logs and defining $z \equiv \log(x)$ we get: $z + (1/2) \log(1 + \varphi(\exp(z))) = (1/2) \log(2\hat{b}) - (1/2) \log R$. Differentiating z w.r.t. $\log R$:

$$z' \left[1 + (1/2) \frac{\varphi'(\exp(z)) \exp(z)}{(1 + \varphi(\exp(z)))} \right] = -1/2 \text{ or } \eta_{x,R} \equiv -\frac{R}{x} \frac{dx}{dR} = \frac{(1/2)}{1 + (1/2) \frac{\varphi'(x)x}{1 + \varphi(x)}}.$$

Direct computation gives:

$$\frac{\varphi'(x)x}{1 + \varphi(x)} = \frac{\sum_{j=1}^{\infty} j \frac{2}{(j+2)!} [x]^j}{1 + \sum_{j=1}^{\infty} \frac{2}{(j+2)!} [x]^j} = \sum_{j=0}^{\infty} j \kappa_j(x) \text{ where}$$

$$\kappa_j(x) = \frac{\frac{2}{(j+2)!} [x]^j}{1 + \sum_{s=1}^{\infty} \frac{2}{(s+2)!} [x]^s} \text{ for } j \geq 1, \text{ and } \kappa_0(x) = \frac{1}{1 + \sum_{s=1}^{\infty} \frac{2}{(s+2)!} [x]^s}.$$

so that κ_j has the interpretation of a probability. For larger x the distribution κ is stochastically larger since: $\frac{\kappa_{j+1}(x)}{\kappa_j(x)} = \frac{x}{(j+3)}$, for all $j \geq 1$ and x . Then we can write

$\frac{\varphi'(x)x}{1+\varphi(x)} = E^x [j]$, where the right hand side is the expected value of j for each x .

Hence, for higher x we have that $E^x [j]$ increases and thus the elasticity $\eta_{x,R}$ decreases. As $x \rightarrow 0$ the distribution κ puts all the mass in $j = 0$ and hence $\eta_{x,R} \rightarrow 1/2$. As $x \rightarrow \infty$ the distribution κ concentrates all the mass in arbitrarily large values of j , hence $E^x [j] \rightarrow \infty$ and $\eta_{x,R} \rightarrow 0$. \square

Proof of Proposition 5. By the fundamental theorem of Renewal Theory n equals the reciprocal of the expected time between withdrawals, which is distributed as an exponential with parameter p and truncated at time \bar{t} . It is exponential because agents have an arrival rate p of free withdrawals. It is truncated at \bar{t} because agents must withdraw when balances hit the zero bound, where $\bar{t} = (1/\pi)\log\left(1 + \frac{m^*}{c}\pi\right)$, the time that it takes to deplete a cash balance from m^* to zero conditional on not having a free withdrawal opportunity. Simple algebra gives that the expected time between withdrawals is equal to: $(1 - e^{-p\bar{t}})/p$. \square

Proof of Proposition 6 . (i) Let $H(m, t)$ be the CDF for m at time t . Define $\psi(m, t; \Delta) \equiv H(m, t) - H(m - \Delta(m\pi + c), t)$. Thus $\psi(m, t; \Delta)$ is the fraction of agents with money in the interval $[m, m - \Delta(m\pi + c)]$ at time t , and let h :

$$h(m, t; \Delta) = \frac{\psi(m, t; \Delta)}{\Delta(m\pi + c)} \quad (40)$$

so that $\lim_{\Delta \rightarrow 0} h(m, t; \Delta)$ as $\Delta \rightarrow 0$ is the density of H evaluated at m at time t . In the discrete time version of the model with period of length Δ the law of motion of cash implies:

$$\psi(m, t + \Delta; \Delta) = \psi(m + \Delta(m\pi + c), t; \Delta) (1 - \Delta p) \quad (41)$$

Assuming that we are in the stationary distribution $h(m, t; \Delta)$ does not depend on t , so we write $h(m; \Delta)$. Inserting equation (40) in (41), substituting $h(m; \Delta) + \frac{\partial h}{\partial m}(m; \Delta) [\Delta(m\pi + c)] + o(\Delta)$ for $h(m + \Delta(m\pi + c); \Delta)$ canceling terms, dividing by Δ and taking the limit as $\Delta \rightarrow 0$, we obtain (20). The solution of this ODE is $h(m) = 1/m^*$ if $p = \pi$ and $h(m) = A \left[1 + \pi \frac{m}{c}\right]^{\frac{p-\pi}{\pi}}$ for some constant A if $p \neq \pi$. The constant A is chosen so that the density integrates to 1, so that $A = 1 / \left\{ \left(\frac{c}{p}\right) \left(\left[1 + \frac{\pi}{c} m^*\right]^{\frac{p}{\pi}} - 1 \right) \right\}$.

(ii) We now show that the distribution of m that corresponds to a higher value of m^* is stochastically higher. Consider the CDF $H(m; m^*)$ and let $m_1^* < m_2^*$ be two values for the optimal return point. We argue that $H(m; m_1^*) > H(m; m_2^*)$ for all $m \in [0, m_2^*)$. This follows because in $m \in [0, m_1^*]$ the densities satisfy

$$\frac{h(m; m_2^*)}{h(m; m_1^*)} = \left(\left[1 + \pi \frac{m_1^*}{c}\right]^{\frac{p}{\pi}} - 1 \right) / \left(\left[1 + \pi \frac{m_2^*}{c}\right]^{\frac{p}{\pi}} - 1 \right) < 1$$

In the interval $[m_1^*, m_2^*)$ we have: $H(m; m_1^*) = 1 > H(m; m_2^*)$. \square

Proof of Proposition 7. We first show that if $p' > p$, then the distribution associated with p' stochastically dominates the one associated with p . For this we use four properties. First, equation (19) evaluated at $m = 0$ shows that $h(0; p)$ is decreasing in p . Second, since $h(\cdot; p)$ and $h(\cdot; p')$ are continuous densities, they integrate to one, and hence there must be some value \tilde{m} such that $h(\tilde{m}; p') > h(\tilde{m}; p)$. Third, by the intermediate value theorem, there must be at least one $\hat{m} \in (0, m^*)$ at which $h(\hat{m}; p) = h(\hat{m}; p')$. Fourth, note that there is at most one such value $\hat{m} \in (0, m^*)$. To see why, recall that h solves $\frac{\partial h(m)}{\partial m} = \frac{(p-\pi)}{(\pi m+c)} h(m)$ so that if $h(\hat{m}, p) = h(\hat{m}, p')$ then $\frac{\partial h(\hat{m}; p')}{\partial m} > \frac{\partial h(\hat{m}; p)}{\partial m}$. Summarizing: $h(m; p) > h(m; p')$ for $0 \leq m < \hat{m}$, $h(\hat{m}; p) = h(\hat{m}; p')$, and $h(m; p) < h(m; p')$ for $\hat{m} < m \leq m^*$. This establishes that $H(\cdot; p')$ is stochastically higher than $H(\cdot; p)$. Clearly this implies that M/m^* is increasing in p .

Finally, we obtain the expressions for the two limiting cases. Direct computation yields $h(m) = 1/m^*$ for $p = \pi$, hence $M/m^* = 1/2$. For the other case, note that

$$\frac{1}{h(m^*)} = \frac{c}{p} \frac{\left[1 + \pi \frac{m^*}{c}\right]^{\frac{p}{\pi}} - 1}{\left[1 + \pi \frac{m^*}{c}\right]^{\frac{p}{\pi}-1}} = \frac{c}{p} \left[1 + \pi \frac{m^*}{c}\right] \left(1 - \frac{1}{\left[1 + \pi \frac{m^*}{c}\right]^{\frac{p}{\pi}}}\right)$$

hence $h(m^*) \rightarrow \infty$ for $p \rightarrow \infty$. Since h is continuous in m , for large p the distribution of m is concentrated around m^* . This implies that $M/m^* \rightarrow 1$ as $p \rightarrow \infty$. \square

Proof of Proposition 8.

Let $x \equiv m^*(r+p)/c$. Equation (15) for $\pi = 0$ and $r = 0$, shows that the value of x solves: $e^x = 1 + x + \hat{b}/R$. This defines the increasing function $x = \gamma(\hat{b}/R)$. Note that $x \rightarrow \infty$ as $\hat{b}/R \rightarrow \infty$ and $x \rightarrow 0$ as $\hat{b}/R \rightarrow 0$.

To see how the ratio Mp/c depends on x notice that from (30) we have that $Mp/c = \phi(x p/(p+r))$ where $\phi(z) \equiv z/(1 - e^{-z}) - 1$. Thus $\lim_{r \rightarrow 0} Mp/c = \phi(x)$. To see why the ratios W/M and \underline{M}/M are functions only of x , note from (30) that $\frac{p}{n} = 1 - \exp(-pm^*/c) = 1 - \exp(-x p/(p+r))$ and hence as $r \rightarrow 0$ we can write $p/n = \omega(x) = \underline{M}/M$ where the last equality follows from (24) and ω is the function: $\omega(x) \equiv 1 - \exp(-x)$. Using (33) we have $W/M = \alpha(\omega)$ where $\alpha(\omega) \equiv [1/\omega + 1/\log(1-\omega)]^{-1} - \omega$. The monotonicity of the functions ϕ, ω, α is straightforward to check. The limits for \underline{M}/M and W/M as $x \rightarrow 0$ or as $x \rightarrow \infty$ follow from a tedious but straightforward calculation.

Finally, the elasticity of the aggregate money demand with respect to \hat{b}/R is:

$$\frac{R}{M/c} \frac{\partial M/c}{\partial R} = \frac{(1/p) \phi'(x)}{M/c} R \frac{\partial x}{\partial R} = x \frac{\phi'(x)}{\phi(x)} \frac{R}{x} \frac{\partial x}{\partial R} = \eta_{\phi, x} \cdot \eta_{x, \hat{b}/R}$$

i.e. is the product of the elasticity of ϕ w.r.t. x , denoted by $\eta_{\phi, x}$, and the elasticity of x w.r.t. \hat{b}/R , denoted by $\eta_{x, \hat{b}/R}$. The definition of $\phi(x)$ gives: $\eta_{\phi, x} = \frac{x(1-e^{-x}-xe^{-x})}{(x-1+e^{-x})(1-e^{-x})}$ where $\lim_{x \rightarrow \infty} \eta_{\phi, x} = 1$. A second order expansion of each of the exponential functions shows that $\lim_{x \rightarrow 0} \eta_{\phi, x} = 1$. Direct computations using $x = \gamma(\hat{b}/R)$ yields $\eta_{x, \hat{b}/R} = \frac{e^x - x - 1}{x(e^x - 1)}$. It is immediate that $\lim_{x \rightarrow \infty} \eta_{x, \hat{b}/R} = 0$ and $\lim_{x \rightarrow 0} \eta_{x, \hat{b}/R} = 1/2$. \square

Proof of Proposition 9.

(i) By Proposition 3, $rV(m^*) = Rm^*$, $V(\cdot)$ is decreasing in m , and $V(0) = V(m^*) + b$. The result then follows since m^* is continuous at $r = 0$. (ii) Since $v(0) = 0$ it suffices to show that $\frac{\partial v(R)}{\partial R} = \frac{\partial Rm^*(R)}{\partial R} = M(R)$ or equivalently that $m^*(R) + R\frac{\partial m^*(R)}{\partial R} = M(R)$. From (15) we have that: $\frac{\partial m^*}{\partial R} \left[\left(1 + \pi \frac{m^*}{c}\right)^{(r+p)/\pi} - 1 \right] \frac{(r+p+\pi)}{c} = -\frac{b}{cR^2}(r+p)(r+p+\pi)$. Using (15) again to replace $\frac{b}{cR}(r+p)(r+p+\pi)$, inserting the resulting expression into $m^*(R) + R\frac{\partial m^*(R)}{\partial R}$, letting $r \rightarrow 0$ and rearranging yields the expression for M obtained in (21). (iii) Using (i) in (iii) yields $R(m^* - M) = (n-p)b$. Replacing M and n using equations for the expected values (18) and (21) for an arbitrary m^* yields an equation identical to the one characterizing the *optimal* value of m^* , (15), evaluated at $r = 0$. \square