Econometrica, Vol. 77, No. 2 (March, 2009), 363–402

FINANCIAL INNOVATION AND THE TRANSACTIONS DEMAND FOR CASH

BY FERNANDO ALVAREZ AND FRANCESCO LIPPI[1]

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, which qualitatively reproduces several empirical patterns. We estimate the structural parameters using micro data and show that quantitatively the model captures important economic patterns. 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, and the benefit of ATM ownership.

KEYWORDS: Money demand, technological progress, inventory models.

1. INTRODUCTION

THERE IS A LARGE LITERATURE arguing that financial innovation is important for understanding money demand, yet this literature seldom integrates the empirical analysis with a model 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 is standard in the inventory theory we assume that nonnegative cash holdings are needed to pay for cash purchases c. We extend the Baumol (1952) and Tobin (1956) model to a dynamic environment that allows for the opportunity to withdraw cash at random times at no cost. Withdrawals at any other time 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 a bank desk at a time with a low opportunity cost. Another interpretation is that

[1]We thank the co-editor and three anonymous referees for constructive criticisms on previous versions of the paper. We also thank Alessandro Secchi for his guidance in the construction and analysis of the data base. We benefited from the comments of Manuel Arellano, V. V. Chari, Luigi Guiso, Bob Lucas, Greg Mankiw, Fabiano Schivardi, Rob Shimer, Pedro Teles, Randy Wright, and seminar participants at the University of Chicago, University of Sassari, Harvard University, MIT, Wharton School, Northwestern, FRB of Chicago, FRB of Minneapolis, Bank of Portugal, European Central Bank, Bank of Italy, CEMFI, EIEF, University of Cagliari, University of Salerno, Austrian National Bank, Tilburg University, and Erasmus University Rotterdam. Alvarez thanks the Templeton Foundation for support and the EIEF for their hospitality.

 DOI: 10.3982/ECTA7451

p measures the probability that an ATM 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.

It is useful to split the agent's decision on financing of her purchases into three parts: (i) choose a technology, which is indexed by p and b (e.g., adopt the ATM card); (ii) decide the amount of expenditure to be made in cash (c), as opposed to credit or debit; (iii) decide the optimal inventory policy to minimize the cost of cash management for a *given* technology (p, b) and level of c. This paper focuses on (iii). We stress that given p, b, and c, the inventory problem in (iii) is well defined even if the optimal choice of c in (ii) depends on p and b. In Section 5 we show that the presence of a systematic relationship between c and (p, b) does not bias the estimates of the technological parameters.

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 on 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. This ratio ranges between 0 and 1, and is increasing in p. Using household data for Italy and the United States, we document that $\underline{M}/M$ is about 0.4, instead of being 0 as predicted by the BT model. Another property of our model is that a higher p increases the number of withdrawals n and decreases the average withdrawal size W, with W/M ranging between 2 and 0. Using data from Italian households, we measure values of n and W/M that are inconsistent with those predicted by the BT model, but can be rationalized by our 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 and 4 present the theory. Section 3 analyzes the effect of financial innovation 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 money demand, that is, its interest and expenditure elasticities. Section 4 introduces our benchmark inventory model, a stochastic dynamic version of the one in Section 3. In this model agents have random meetings with a financial intermediary in which they can withdraw money at no cost. 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 of 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. The structural estimation of the model parameters will allow us to shed light on this issue. We conclude the section with an analysis of the welfare implications.

Sections 5, 6, and 7 contain the empirical analysis. In Section 5 we estimate the model using the panel data for Italian households. We discuss identification and show that the two parameters p and b are overidentified because we observe four dimensions of household behavior: M, W, $\underline{M}$, and n. The estimates reproduce the sizable precautionary holdings observed in the data. The patterns of the estimates are reasonable: for instance, the parameters for the households with an ATM card indicate their access to a better technology (higher p and lower b). Section 6 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 estimates indicate that technology is better in locations with a higher density of ATM terminals and bank branches, and that it has improved through time. Even though our model can generate interest rate elasticities between 0 and $1/2$, and expenditure elasticities between $1/2$ and 1, the values implied by the estimates are close to $1/2$ for both—the values of the BT model. We discuss how to reconcile this finding with the smaller estimates of the interest rate elasticity that are common in the literature.[2] In Section 7 we use the estimates to quantify the welfare cost of inflation—relating it to the one in Lucas (2000)—and to measure the benefits of ATM card ownership. In spite of the finding that the interest elasticity is close to the one in BT over the sample, our estimate of the welfare cost is about half of the cost in BT. This happens because the interest elasticity is constant in BT, while it converges to zero in our model as the interest rate becomes nil.

2. CASH HOLDINGS PATTERNS OF ITALIAN HOUSEHOLDS

Table I presents some statistics on the cash holdings patterns of Italian households based on the *Survey of Household Income and Wealth.*[3] All households have checking accounts that pay interest at rates documented below. We report statistics separately for households with and without an ATM card. The

[2]We remark that our interest rate elasticity, as in the BT model, refers to the ratio of money stock to cash consumption. Of course if cash consumption relative to total consumption is a function of interest rates, as in the Lucas and Stokey (1987) cash–credit model, the elasticity of money to total consumption will be even higher. A similar argument applies to the expenditure elasticity. The distinction is important for comparing our results with estimates in the literature, that typically use money/total consumption. See for instance Lucas (2000), who used aggregate US data, or Attanasio, Guiso, and Jappelli (2002), who used the same household data used here.

[3]A periodic survey of the Bank of Italy that collects data on several social and economic characteristics. The cash management information used below is available only since 1993.

TABLE I

HOUSEHOLDS' CURRENCY MANAGEMENT[a]

Variable	1993	1995	1998	2000	2002	2004
Expenditure share paid w/ currency[b]						
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[c]						
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[d]						
w/o ATM	430	490	440	440	410	410
w. ATM	370	410	370	340	330	350
Currency at withdrawals[e] $\underline{M}/M$						
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[f] 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
No. of withdrawals						
n (per year)[g]						
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)[g]						
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
No. of observations w ATM card[h]	2322	2781	2998	3562	3729	3866
No. of observations w/o ATM card[h]	3421	3020	2103	2276	2275	2190

[a]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 in each year.

[b]Ratio of expenditures paid with cash to total expenditures (durables, nondurables, and services).

[c]Average currency during the year divided by daily expenditures paid with cash.

[d]The average number of adults per household is 2.3. In 2004, 1 euro in Italy was equivalent to \$1.25 in the United States, PPP adjusted (Source: World Bank International Comparison Program (ICP) tables).

[e]Average currency at the time of withdrawal as a ratio to average currency.

[f]Average withdrawal during the year as a ratio to average currency.

[g]The entries with $n = 0$ are coded as missing values.

[h]Number of households with bank account for whom the currency and the cash consumption data were available in each survey. Data on withdrawals were supplied by a smaller number of respondents (Source: Bank of Italy *Survey of Household Income and Wealth*).

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), which the table displays as a fraction of total consumption. The fraction is smaller for households with an ATM card and displays a downward trend for both type of households. These percentages are comparable to those for the United States between 1984 and 1995.[4] 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 with currency. We notice that relative to c, Italian households held about twice as much cash as U.S. households between 1984 and 1995.[5]

Table I reports three statistics that are useful to assess the empirical performance of deterministic inventory models, such as the classic one by Baumol and Tobin. 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-theoretic models, its sample mean in the data is about 0.4. A comparable statistic for U.S. households is about 0.3 in 1984, 1986, and 1995 (see Table 1 in Porter and Judson (1996)). The second statistic is the ratio between the withdrawal amount (W) and average currency holdings.[6] 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; for those without an ATM it is slightly below 2. Inspection of the raw data shows that there is substantial variation across provinces: indeed, the median across households is about 1.0 for households with and without an ATM. The third statistic is the normalized number of withdrawals, $\frac{n}{c/(2M)}$. The normalization is chosen so that this statistic is equal to 1 in BT.[7] As the table shows, the sample mean of this statistic is well above 1, especially for households with an ATM card.

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 consider the effect of measurement error systemat-

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

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

[6] The withdrawal amount is computed as the weighted average of ATM and bank desk withdrawals. Since in Italy there is no cash back at withdrawals, this measures the withdrawal amount quite accurately. See Appendix A in Alvarez and Lippi (2009) for more documentation.

[7] In the BT model, the accounting identity $nW = c$ holds and since withdrawals only happen when cash balances reach zero, then $M = W/2$.

TABLE II

FINANCIAL INNOVATION AND THE OPPORTUNITY COST OF CASH[a]

Variable	1993	1995	1998	2000	2002	2004
Bank branches[b]	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[b]	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[c]	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 theft[d]	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

[a]Mean (standard deviation in parentheses) across provinces.

[b]Per thousand residents (Source: supervisory reports to the Bank of Italy and the Italian Central Credit Register).

[c]Net nominal interest rates in percent. 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).

[d]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 cash stolen for different types of crimes using survey data on victimization rates (Source: Istituto nazionale di statistica (Istat) and authors' computations; see Appendix B in Alvarez and Lippi (2009) for details).

ically, without altering the conclusion about the drawbacks of deterministic inventory-theoretic models.

For each year, Table II 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 theft. The diffusion of bank branches and ATM terminals varies significantly across provinces and is increasing through time. Differences in the nominal interest rate across time are due mainly to the disinflation. The variation of nominal interest rates across provinces mostly reflects the segmentation of banking markets. The large differences in the probability of cash theft 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 (2009) reported 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 1/2.[8] Tables I and II show 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

[8]They estimated that the elasticity of cash holdings with respect to the interest rate is about 0 for agents who hold an ATM card and −0.2 for agents who do not.

ATM cards) is negatively correlated with the opportunity cost R in the cross-section, the time-series, and the pooled time-series and cross-section data. Yet the largest estimate for the interest rate elasticity is smaller than 0.25 and in most cases about 0.05 (in absolute value). 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.

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. Let R be the opportunity cost of cash (e.g., the nominal interest rate on a deposit account). 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. Technology is thus represented by the parameters b and p.

For concreteness, 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

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

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

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

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 of 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, 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 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 a model that 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 financing a constant nominal expenditure, or alternatively assuming zero inflation. Most importantly, it assumes that the agent has a Poisson arrival of free opportunities to withdraw cash at a rate p. Relative to the deterministic model, the randomness 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 2, as observed in Table I. 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 $\hat{b} = p^2 b/c$.

We assume that the agent is subject to a cash-in-advance constraint and minimizes the cost of financing a given constant flow of cash consumption, denoted by c. Let $m \geq 0$ denote the agent nonnegative real cash balances that decrease due to consumption and inflation:

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

for almost all $t \geq 0$. The agent can withdraw or deposit at any time from an account that yields net 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, that is, $m(t^+) - m(t^-) > 0$, and likewise for a deposit.

There are two sources or randomness in the environment, described by independent Poisson processes with intensities p_1 and p_2. The first source describes the arrivals of "free adjustment opportunities" (see the Introduction for examples); the second describes the arrival times where the agent's cash balances are stolen. We assume that a fixed cost b is paid for each adjustment, unless it happens at the time of a free adjustment opportunity. We can write the problem of the agent as

$$G(m) = \min_{\{m(t),\tau_j\}} \mathbb{E}_0 \left\{ \sum_{j=0}^{\infty} e^{-r\tau_j} \left[I_{\tau_j} b + (m(\tau_j^+) - m(\tau_j^-)) \right] \right\} \tag{4}$$

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 0—so the cost is not paid—if the adjustment occurs upon a free opportunity; otherwise it is 1. The expectation is taken with respect to the two Poisson processes. The parameters of this problem are r, π, p_1, p_2, b, and c.

4.1. *Bellman Equations and Optimal Policy*

We turn to the characterization of the Bellman equation 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 the agent chooses after a contact with a financial intermediary: we refer to it as the optimal cash replenishment level. The threshold m^{**} is the value of cash beyond which the agent pays the cost b, contacts the intermediary, and makes 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 G is differentiable, it must satisfy

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

The first term gives the change in the value function per unit of time, conditional on no arrival of either free adjustment or of a cash theft. 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" $G(\hat{m}) - G(m)$. Likewise, the third term gives the change in the value function conditional on a cash theft. In this case, the cost b must be paid and the cash adjustment equals $\hat{m}$. Upon being matched with a financial intermediary the agent replenishes her balances to $m = m^*$, which solves

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

This problem has two boundary conditions. First, if $m = 0$, the agent withdraws to prevent violation of the nonnegative cash constraint in the next instant. Second, for $m \geq m^{**}$, the agent pays b and deposits cash in excess of m^*. Combining these boundary conditions with (5), we have

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

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:

$$
(8)\qquad m + G(m) < b + m^* + G(m^*) \quad \text{for} \quad m \in (0, m^{**}).
$$

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

We define a related problem that it is closer to the standard inventory-theoretic problem where the agent minimizes the shadow cost

$$
(9)\qquad V(m) = \min_{\{m(t), \tau_j\}} \mathbb{E}_0 \left\{ \sum_{j=0}^{\infty} e^{-r\tau_j} \left[I_{\tau_j} b + \int_0^{\tau_{j+1} - \tau_j} e^{-rt} R m(t + \tau_j)\, dt \right] \right\}
$$

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.[9] The indicator I_{τ_j} equals 0 if the adjustment takes place at the time of a free adjustment; otherwise it is 1. In this formulation, R is the opportunity cost of holding cash and there is *only one* Poisson process with intensity p describing the arrival of a free opportunity to adjust. The parameters of the problem are r, R, π, p, b, and c.

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

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

[9]The shadow cost formulation is standard in the literature on inventory-theoretic models, as in, for example, Baumol (1952), Tobin (1956), Miller and Orr (1966), and Constantinides (1976). In these papers the problem is to minimize the steady state cost of 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 that of Constantinides and Richard (1978). In a related model, Frenkel and Jovanovic (1980) compared the resulting money demand arising from minimizing the steady state vs. the expected discounted cost.

Upon being matched with a financial intermediary, the agent chooses the optimal adjustment setting $m = m^*$ or

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

As in problem (4), we conjecture 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; for $m \geq m^{**}$, the agent pays the cost b and deposits cash in excess of m^*.[10] Combining these boundary conditions with (10), we have

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

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

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

Summarizing, we say that m^*, m^{**}, and $V(\cdot)$ solve the Bellman equation for the shadow cost problem (9) if they satisfy (11), (12), and (13). We are now ready to show that (4) and (9) are equivalent and to characterize 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* $G(\cdot)$ *and* $V(\cdot)$ *satisfy*

$$G(m) = V(m) - m + c/r + p_2 b/r \tag{14}$$

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

See the Appendix for the proof.

Notice that the total and the shadow cost problems are described by the same number of parameters. They have r, π, c, and 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 simple: the shadow cost of holding money is given by the real opportunity cost of investing, r, plus the fact that cash holdings loose real value continually at a rate π and are lost entirely with probability p_2 per

[10] Since withdrawals are the agent's only source of cash in this economy, in the invariant distribution money holdings are distributed on the interval $(0, m^*)$ and m^{**} is never reached.

unit of time. That $p = p_1 + p_2$ is clear too, since the effect of either shock is to force an adjustment on cash. 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. Quantitatively, we think that, at least for low nominal interest rates, the presence of other opportunity costs may be important. The relation between G and V in (14) is intuitive. First, the constant c/r is required, since even if all withdrawals were free, consumption expenditures must be financed. Second, the constant $p_2 b/r$ is the present value of all the withdrawal costs paid upon a cash theft. 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 withdrawn from the interest bearing account.

From now on we use the shadow cost formulation since it is closer to the standard inventory decision problem. The predictions of the two models concerning cash holdings statistics are going to be identical for M and n, and display a small difference for W and $\underline{M}$, which is discussed later.

The next proposition gives one nonlinear 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* m^*/c *has three arguments*: β, $r + p$, *and* π, *where* $\beta \equiv b/(cR)$. *The return point* m^* *is given by the unique positive solution to*

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

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

(i) m^*/c *is increasing in* $b/(cR)$, $m^*/c = 0$ *as* $b/(cR) = 0$, *and* $m^*/c \to \infty$ *as* $b/(cR) \to \infty$.

(ii) *For small* $b/(cR)$, $m^*/c = \sqrt{2b/(cR)} + o(\sqrt{b/(cR)})$, *where* $o(z)/z \to 0$ *as* $z \to 0$.

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

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

(iv) *The elasticity of* m^* *with respect to* R *evaluated at zero inflation satisfies* $0 \le -(R/m^*)(dm^*/dR)|_{\pi=0} \le \frac{1}{2}$, *is decreasing in* p, *and satisfies*

$$-\frac{R}{m^*}\frac{\partial m^*}{\partial R}\bigg|_{\pi=0} \to \frac{1}{2} \quad as \quad \frac{\hat{b}}{R} \to 0$$

and

$$-\frac{R}{m^*}\frac{\partial m^*}{\partial R}\bigg|_{\pi=0} \to 0 \quad as \quad \frac{\hat{b}}{R} \to \infty,$$

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

See the Appendix for the proof.

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. Hence m^* is homogenous of degree 1 in (c, b). Result (ii) shows that when $b/(cR)$ is small, the resulting money demand is well approximated by the BT model. Part (iv) shows that the absolute value of the interest elasticity ranges between 0 and 1/2, and that it is decreasing in p (at low inflation). 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 6. Since m^* is a function of $b/(cR)$, then the elasticity of m^* with respect to b/c equals that with respect to R with an opposite sign.

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*

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

where $A = (r+p)/c^2(Rm^* + (r+p)b + Rc/(r+p+\pi)) > 0$. *For* $m = 0$ *or* $m \geq m^{**}$, $V(m) = V^* + b$.

(ii) *The value for the agents matched with a financial institution is* $V^* = (R/r)m^*$.

See the Appendix for the proof.

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.

4.2. *Cash Holdings Patterns of the Model*

This section derives the model predictions concerning observable cash management statistics produced by the model under the invariant distribution of real cash holdings when a policy (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 to m^* when a free withdrawal occurs or when $m = 0$.

Our first result is to compute the expected number of withdrawals per unit of time, denoted by n. By the fundamental renewal theory, n equals the reciprocal of the expected time between withdrawals, which gives[11]

$$n\left(\frac{m^*}{c}, \pi, p\right) = \frac{p}{1 - \left(1 + \pi \frac{m^*}{c}\right)^{-p/\pi}}. \tag{17}$$

As can be seen from expression (17), the ratio $n/p \geq 1$, since in addition to the p free withdrawals, it includes the costly withdrawals that agents make 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, and c and m^*/c.

PROPOSITION 4: (i) *The density for the real balances m is*

$$h(m) = \left(\frac{p}{c}\right) \frac{\left[1 + \pi \frac{m}{c}\right]^{p/\pi - 1}}{\left[1 + \pi \frac{m^*}{c}\right]^{p/\pi} - 1}. \tag{18}$$

(ii) *Let $H(m, m_1^*)$ be the cumulative distribution function (CDF) of m for a given m^*. Let $m_1^* < m_2^*$. Then $H(m, m_2^*) \leq H(m, m_1^*)$, that is, $H(\cdot, m_2^*)$ first order stochastically dominates $H(\cdot, m_1^*)$.*

For the proof, see the Appendix.

The density of m solves the ordinary differential equation (ODE) (see the proof of Proposition 4)

$$\frac{\partial h(m)}{\partial m} = \frac{(p - \pi)}{(\pi m + c)} h(m) \tag{19}$$

for any $m \in (0, m^*)$. There are two forces that determine the shape of this density. One is that agents replenish their balances to m^* at a rate p. The other is that inflation erodes the real value of their nominal balances. For $p = \pi = 0$, the two effects cancel and the distribution is uniform, as in BT.

[11]The time between withdrawals is distributed as a truncated exponential with parameter p. It is exponential because free withdrawals arrive at a rate p. Since agents must withdraw when $m = 0$, the distribution is truncated at $\bar{t} = (1/\pi)\log(1 + m^*/c\pi)$, which is the time to deplete cash balances from m^* to 0 conditional on not having a free withdrawal. Simple algebra gives the equation in the text.

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

$$\frac{M}{c}\left(\frac{m^*}{c}, \pi, p\right) = \frac{\left(1+\pi\frac{m^*}{c}\right)^{p/\pi}\left[\frac{m^*}{c} - \frac{(1+\pi m^*/c)}{p+\pi}\right] + \frac{1}{p+\pi}}{\left[1+\pi\frac{m^*}{c}\right]^{p/\pi} - 1}. \tag{20}$$

The function $\frac{M}{c}(\cdot, \pi, p)$ is increasing in m^*, which follows immediately from part (ii) of Proposition 4. Next we show that for a fixed m^*, M is increasing in p:

PROPOSITION 5: *The ratio M/m^* is increasing in p with*

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

For the proof, see the Appendix.

Note that $p = \pi = 0$ gives BT, that is, $M/m^* = 1/2$. The other limit corresponds to the case where the agent is continuously replenishing her balances. The average withdrawal $W = m^*[1 - \frac{p}{n}] + [\frac{p}{n}]\int_0^{m^*}(m^* - m)h(m)\,dm$. The expression $[1 - \frac{p}{n}]$ is the fraction of withdrawals, of size m^*, that occur when $m = 0$. The complementary fraction gives the withdrawals that occur upon a chance meeting with the intermediary, which are of size $m^* - m$ and happen with frequency $h(m)$. Combining the previous results, we can see that for $p \geq \pi$, the ratio of withdrawals to average cash holdings is less than 2. To see this, using the definition of W, we can write

$$\frac{W}{M} = \frac{m^*}{M} - \frac{p}{n}. \tag{21}$$

Since $M/m^* \geq 1/2$, then $W/M \leq 2$. Notice that for p large enough this ratio can be smaller than 1. In the BT model $W/M = 2$, while in the data of Table I the average ratio is below 1.5 and its median value is 1 for households with an ATM card. The intuition for this result is clear: the agent takes advantage of the free random withdrawals regardless of her cash balances, hence the withdrawals are distributed on $[0, m^*]$, as opposed to being concentrated on m^*, as in the BT model.

Let $\underline{M}$ be the average amount of money at the time of withdrawal. The derivation, analogous to that for W, gives $\underline{M} = 0 \cdot [1 - \frac{p}{n}] + [\frac{p}{n}]\int_0^{m^*} mh(m)\,dm$ or

$$\underline{M} = \frac{p}{n}M, \tag{22}$$

which implies that $0 < \underline{M}/M < 1$, $\underline{M}/M \to 0$, as $p \to 0$ and $\underline{M}/M \to 1$ as $p \to \infty$.

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 saw-toothed path of a deterministic inventory model, which implies that the average cash balance is $M_1 = \underline{M} + 0.5c/n$ or $M_2 = \underline{M} + 0.5W$; see, for example, equations (1) and (2) in Attanasio, Guiso, and Jappelli (2002) and Table 1 in Porter and Judson (1996). Instead, in our model $W/2 < M < \underline{M} + W/2$. The leftmost inequality is a consequence of Proposition 5 and equation (21); the other can be derived using the form of the optimal decision rules and the law of motion of cash flows (see Appendix C in Alvarez and Lippi (2009)). Hence in our model M_1 and M_2 are upward biased estimates of M. Indeed the data of Table I show they overestimate M by a large margin.[12]

We finish with two comments on W/M and $\underline{M}/M$. First, these ratios, as is the case with M/c and n, are functions of three arguments: m^*/c, p, and π. This property is useful in the estimation, where we use the normalization m^*/c. Second, as mentioned in the comment to Proposition 1, the statistics for W/M and $\underline{M}/M$ produced by the total cost problem differ from those for the shadow cost problem displayed in (21) and (22). The expression for the total cost problem are given by $W/M = m^*/M - p/n + p_2/n$ and $\underline{M}/M = p/n - p_2/n$. The correction term p_2/n is due to the effect of cash theft. Quantitatively, the effect of p_2/n on W/M and $\underline{M}$ is negligible compared to p, and hence we ignore this term in the expressions for W and $\underline{M}$ below. Note that, instead, p_2 is quantitatively important in the computation of the opportunity cost of cash, $R = r + \pi - p_2$, at low inflation rates.

4.3. *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 Proposition 2 that describe how m^*/c depends on p, b/c, and R, and the results of Section 4.2 that 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 $r \to 0$, our problem is equivalent to minimizing the steady state cost. The choice of $\pi = r = 0$ simplifies comparison of the analytical results with those of the original BT model and with those of Section 3.

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

PROPOSITION 6: *Let* $\pi = 0$ *and* $r \to 0$. *The ratios* W/M, $\underline{M}/M$, *and* $(M/c)p$ *are determined by three strictly monotone functions of* $\hat{b}/R$ *that satisfy*

$$\text{as} \quad \frac{\hat{b}}{R} \to 0: \quad \frac{W}{M} \to 2, \quad \frac{\underline{M}}{M} \to 0, \quad \frac{\partial \log \frac{Mp}{c}}{\partial \log \frac{\hat{b}}{R}} \to \frac{1}{2};$$

$$\text{as} \quad \frac{\hat{b}}{R} \to \infty: \quad \frac{W}{M} \to 0, \quad \frac{\underline{M}}{M} \to 1, \quad \frac{\partial \log \frac{Mp}{c}}{\partial \log \frac{\hat{b}}{R}} \to 0.$$

See the Appendix for the proof.

The limit where $\hat{b}/R \to 0$ corresponds to the BT model, where $W/M = 2$, $\underline{M}/M = 0$, and $\frac{\partial \log(M/c)}{\partial \log R} = -1/2$ for all b/c and R. 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

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

Direct computation gives that

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

The effects of p, b/c, and R on M/c are smooth versions of those described in the model with p deterministic free withdrawals in Section 3; the effects on W and $\underline{M}$ differ due to the precautionary demand generated by the random withdrawal cost.

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 V. 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 of 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

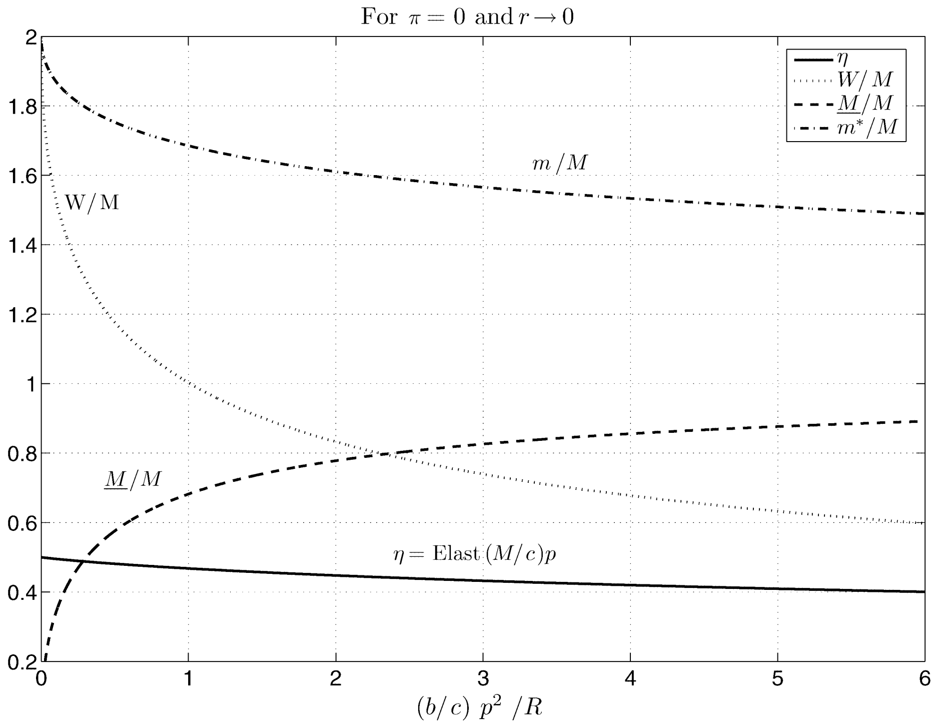

FIGURE 1.—W/M, $\underline{M}/M$, m^*/M, and $\eta =$ elasticity of $(M/c)p$.

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

$$\ell(R, \kappa) = \big[n(m^*(R, \kappa), p, \pi) - p\big] \cdot b, \tag{23}$$

where n is given in (17) 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, that is, the deadweight loss for the society corresponding to R. While ℓ is the relevant measure of the cost for the society, we find it 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 $r \to 0$. 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 used in our estimation.

PROPOSITION 7: *Let* $r \to 0$. *Then* (i) $v(R,\kappa) = Rm^*(R,\kappa)$; (ii) $v(R,\kappa) = \int_0^R M(\tilde{R},\kappa)\,d\tilde{R}$; *and* (iii) $\ell(R,\kappa) = v(R,\kappa) - RM(R,\kappa)$.

For the proof, see the Appendix.

The proposition implies that the loss for society coincides with the consumer surplus that can be gained by reducing R to 0, that is, $\ell(R) = \int_0^R M(\tilde{R})\,d\tilde{R} - RM(R)$.[13] This extends the result of Lucas (2000), derived from a money-in-the-utility-function model, to an explicit inventory-theoretic model. Note that measuring the welfare cost of inflation using consumer surplus requires estimation of the money demand for different interest rates, while the approach using (i) and (iii) can be implemented once M, W, and $\underline{M}$ are known, since $W + \underline{M} = m^*$.

In the BT model $\ell = RM$, since $m^* = W = 2M$. Relative to BT, the welfare cost in our model is $\ell/(RM) = m^*/M - 1$, a value that ranges between 1 and 0, as can be seen in Figure 1. Hence the cost of inflation in our model is smaller than in BT. The difference is due to the behavior of the interest elasticity: while it is constant and equal to 1/2 in BT, the elasticity is between 1/2 and 0 in our model, and is smaller at lower interest rates (recall Proposition 6). Section 7 presents a quantitative application of this result. Finally, note that 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 0 to ∞, the ratio of the costs ℓ/v decreases from 1/2 (the BT value) to 0.

5. ESTIMATION OF THE MODEL

We estimate the parameters $(p, b/c)$ using the data described in Section 2 under two alternative sets of assumptions. Our baseline assumptions are that all households in the same cell (to be defined below) have the same parameters $(p, b/c)$. Alternatively, in Section 5.3 we assume that $(p, b/c)$ comprises a simple parametric function of individual household characteristics, that is, an instance of "observed heterogeneity." In both cases we take the opportunity cost R as observable (see Table II), and assume that households' values of $(M/c, n, W/M, \underline{M}/M)$ are observed with classical normally distributed measurement error (in logs). Appendices F and G in Alvarez and Lippi (2009) explore alternative estimation setups, including one with unobserved heterogeneity, all of which lead to similar results.

The assumption of classical measurement error is often used when estimating models based on household survey data. We find that the pattern of violations of a simple accounting identity, $c = nW - \pi M$, is consistent with large classical measurement error. In particular, a histogram of the deviations of

[13] In (ii) and (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.

this identity (in log points) is centered around zero, symmetric, and roughly bell shaped (see Appendix D in Alvarez and Lippi (2009)).

We stressed in the Introduction that ignoring the endogeneity of c with respect to (p, b) does not impair our estimation strategy. Suppose, for instance, that the cash expenditure $c(p, b)$ depends on p and b (this would be the solution of problem (ii) in the Introduction). At this point, the agent solves the inventory problem, that is, chooses the number of withdrawals $n(c, p, b), W(c, p, b), \ldots$ to finance $c(p, b)$, where the notation emphasizes that $n(\cdot)$ and $W(\cdot)$ depend on (c, p, b). Since we have data on c, we can invert the decisions on $n(c, p, b)$, $W(c, p, b), \ldots$ to estimate p and b without the need to know the mapping: $c(p, b)$. Instead, if c was not observed, the mapping $c(p, b)$ would be needed to estimate p and b by inverting $n(c(p, b), p, b)$, $W(c(p, b), p, b), \ldots.$

5.1. *Identification of p and $b/(cR)$*

Our identification strategy uses the fact that the model is described by two parameters $(p, b/(cR))$ and that, for each observation, the model has implications for four variables $(M/c, n, W/M, \underline{M}/M)$ as shown below. Under the hypothesis that the model is the data generating process, the parameters $(p, b/(cR))$ can be estimated independently for each observation, regardless of the distribution of $(M/c, n, W/M, \underline{M}/M, R)$ across observations. An advantage of this strategy is that the estimates of $(p, b/(cR))$ would be unbiased (or unaffected by selection bias) even if agents were assigned to ATM and non-ATM holder classes in a systematic way.

For simplicity this subsection assumes that $\pi = 0$ and ignores measurement error. Both assumptions are relaxed in the estimation. We consider three cases, each of which exactly identifies p and $b/(cR)$ using a different pair of variables. In the first case we show how M/c and n exactly identify p and $b/(cR)$. For the BT model, that is, for $p = 0$, we have $W = m^*$, $c = m^* n$, and $M = m^*/2$, which implies $2M/c = 1/n$. 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 model in Section 4. From (17) and (20), for a given value of p we have

$$\frac{M}{c} = \frac{1}{p}\left[n\frac{m^*}{c} - 1\right] \quad \text{and} \quad n = \frac{p}{1 - \exp(-pm^*/c)} \tag{24}$$

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

$$\frac{M}{c} = \xi(n, p) = \frac{1}{p}\left[-\frac{n}{p}\log\left(1 - \frac{p}{n}\right) - 1\right]. \tag{25}$$

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. The function ξ is defined only for $n \geq p$.

Analysis of ξ shows that p is identified. Specifically, consider a pair of observations on M/c and n: if $M/c \geq 1/(2n) = \xi(n, 0)$, then there is a unique value of p that solves $M/c = \xi(n, p)$; if $M/c < 1/(2n)$, then there is no solution.[14] Thus, for any $n \geq /(2M)$, our model can rationalize values of $M/c > \xi(n, 0)$, where $\xi(n, 0)$ is the value of M/c in the BT model. In fact, fixing M/c, a higher value of n implies a higher value of p.

The identification of $\beta \equiv b/(cR)$ uses the first order condition for m^*. In particular, given the values of p and the corresponding pair $(M/c, n)$, we use (24) to solve for m^*/c. Finally, using the equation for m^* given in Proposition 2 gives

$$\beta \equiv \frac{b}{cR} = \frac{\exp[(r+p)m^*/c] - [1 + (r+p)(m^*/c)]}{(r+p)^2}. \tag{26}$$

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, because when trips are expensive relative to the opportunity cost of cash (high β), agents visit the intermediary less often. Hence, data on M/c and n identify p and β. An estimate of b/c can then be retrieved using data on R.

The second case shows that W/M and n exactly identify p and $b/(cR)$. Consider an agent who follows an arbitrary policy of replenishing her cash to m^* either as $m = 0$ or when a free withdrawal occurs. Using the cash flow identity $nW = c$ and (25) yields

$$\frac{W}{M} = \delta(n, p) \equiv \left[\frac{1}{p/n} + \frac{1}{\log(1 - p/n)}\right]^{-1} - \frac{p}{n} \tag{27}$$

for $n \geq p$ and $p \geq 0$. Notice that the ratio W/M is a function only of the ratio p/n. As in the previous case, given a pair of observations on $0 < W/M \leq 2$ and $n > 0$, we can use $\delta(n, p)$ to solve for the unique corresponding value of p.[15] The interpretation of this is clear: for $p = 0$ we have $W/M = 2$, the highest value that can be achieved by W/M. A smaller W/M observed for a given n implies a larger value of p. Indeed, as n converges to 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 first case, the identification of $\beta \equiv b/(cR)$ uses the first order condition for m^*. In particular, we can find the value of m^*/c using $W/M = (m^*/c)/(M/c) - p/n$ (equation (21)). With the values of

[14]Since for any $n > 0$ the function ξ satisfies $\xi(n, 0) = 1/(2n)$, $\partial\xi(n, p)/\partial p > 0$ for all $p > 0$, and $\lim_{p \uparrow n} \xi(n, p) = \infty$.

[15]This follows since for all $n > 0$, $\delta(n, 0) = 2$, $\partial\delta(n, p)/\partial p < 0$, and $\lim_{p \uparrow n} \delta(n, p) = 0$.

$(m^*/c, p)$ we can find the unique value of $\beta = (b/c)/R$ that rationalizes this choice, using (26). Thus, data on W/M and n identify β and p.

The third case shows that observations on $\underline{M}/M$ and n exactly identify p and $b/(cR)$. Equation (22) gives $p = n(\underline{M}/M)$. If $\underline{M}/M < 1$, then p is immediately identified; otherwise, there is no solution with $n \leq p$. As in the previous cases the identification of β uses the first order condition for m^*. For a fixed p, different combinations 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 on the number of withdrawals and keep larger cash balances (see Figure 2 in Alvarez and Lippi (2007) for a graphical analysis of the identification problem).

We have discussed how data on each 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 to estimate (p, β) should produce identical estimates. In other words, the model is overidentified. In the next section, we will use this idea to report how well the model fits the data or, more formally, to test for the overidentifying restrictions. 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, 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.

5.2. *Baseline Case: Cell Level Estimation*

In the baseline estimation we define a cell as a particular combination of year–province–household type, where the latter is defined by the cash-expenditure group (lowest, middle, and highest third of households ranked by cash expenditure) and ATM ownership. This yields about 3700 cells, the product of the 103 provinces of Italy $\times$ 6 time periods (spanning 1993–2004) $\times$ 2 ATM ownership statuses (whether a household has an ATM card or not) $\times$ 3 cash expenditure group. For each year we observe the inflation rate π, and for each year–province–ATM ownership type we observe the opportunity cost R. Let i index the households in a cell. 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, this implies that all household i have the same values of M/c, W/M, n, and $\underline{M}/M$.

Let $j = 1, 2, 3, 4$ index the variables M/c, W/M, n, and $\underline{M}/M$, let z_i^j be the (log of the) ith household observation on variable j, and let $\zeta^j(\theta)$ be the (log of the) model prediction of the j variable for the parameter vector $\theta \equiv (p, b/c)$. The variable z_i^j is observed with a zero-mean measurement error ε_i^j with variance σ_j^2, so that $z_i^j = \zeta^j(\theta) + \varepsilon_i^j$. It is assumed that the parameter σ_j^2 is common across cells (we allow one set of variances for households with ATM cards and one for those without).

The estimation proceeds in two steps. We first estimate σ_j^2 by regressing each of the four observables, measured at the individual household level, on a vector of cell dummies. The variance of the regression residual is our estimate of σ_j^2. We treat σ_j^2 as known parameters because there are about 20,000 degrees of freedom for each estimate. Since the errors ε_i^j are assumed to be independent across households i and variables j, in the second step we estimate the vector of parameters θ for each cell separately, by minimizing the likelihood criterion

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

where σ_j^2 is the measurement error variance estimated above and N_j is the sample size of the variable j.[16] Minimizing F (for each cell) yields the maximum likelihood estimator provided the ε_i^j are independent across j for each i.

Table III reports some summary statistics for the baseline cell (province–year–type combination) estimates. The first two panels in the table report the mean, median, and 95th and 5th percentiles of the estimated values for p and b/c across all cells. 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 theft described in Table II. Inflation in each year is measured by the Italian consumer price index (CPI) (the same across provinces); the real return r is fixed at 2% per year.

The parameter p gives the average number of free withdrawal opportunities per year. The parameter $b/c \cdot 100$ is the cost of a withdrawal in percentage of daily cash expenditure. We also report the mean value of the t statistics for these parameters. The asymptotic standard errors are computed by solving for the information matrix. The estimates reported in the first two columns of the table concern households who possesses an ATM card, shown separately for those in the lowest and highest cash-expenditure group. The corresponding statistics for households without ATM cards appear in the third and fourth columns. The difference between the 95th and the 5th percentile indicates that there is a significant amount of heterogeneity across cells. The relatively low values for the mean t-statistics reflects that the number of households used in each cell is small. Indeed, in Appendix F in Alvarez and Lippi (2009) 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 as much, the average

[16]The average number of observations (N_j) available for each variable varies. It is similar for households with and without ATM cards. There are more observations on M/c than for each of the other variables, and its average weight (N_1/σ_1^2) is about 1.5 times larger than each of the other three weights (see Appendix E in Alvarez and Lippi (2009) for further documentation). The number of household–year–type combinations used to construct all the cells is approximately 40,000.

TABLE III

SUMMARY OF $(p, b/c)$ ESTIMATES ACROSS PROVINCE–YEAR–TYPE CELLS

	Cash Expenditure[a]			
	Household w/o ATM		Household w. ATM	
	Low	High	Low	High
Parameter p (avg. no. of opportunities per year)				
Mean[b]	6.8	8.7	20	25
Median[b]	5.6	6.2	17	20
95th percentile[b]	17	25	49	61
5th percentile[b]	1.1	0.8	3	4
Mean t-statistics[b]	2.5	2.2	2.7	3.5
Parameter b/c (in % of daily cash expenditure)				
Mean[b]	10.5	5.5	6.5	2.1
Median[b]	7.3	3.6	3.5	1.1
95th percentile[b]	30	17	24	7
5th percentile[b]	1.5	0.4	0.6	0.3
Mean t-statistics[b]	2.8	2.5	2.4	3.3
No. of cells[c]	504	505	525	569

	Goodness of Fit: Likelihood Criterion $F(\theta; z) \sim \chi^2_{(2)}$	
	Household w/o ATM	Household w. ATM
Percentage of cells where[d]		
$F(\theta; z) \leq 4.6 = 90$th percentile of $\chi^2_{(2)}$	59%	48%
$F(\theta; z) \leq 1.4 = 50$th percentile of $\chi^2_{(2)}$	28%	22%
Average no. of households per estimate	10.7	13.5

[a] Low (high) denotes the lowest (highest) third of households ranked by cash expenditure c.

[b] Statistics computed across cells.

[c] The total number of cells, which includes the group with middle cash expenditure, is 1539 and 1654 for households without and with ATM, respectively.

[d] Only cells where all four variables $(M/c, n, W/M, \underline{M}/M)$ are available are used to computed these statistics (about 80% of all cells).

t-statistics increase roughly with the (square root) of the average number of observations per cell.[17]

Table III shows that the average value of b/c across all cells is between 2% and 10% of daily cash consumption. Fixing an ATM ownership type and comparing the average estimates for p and b/c across cash consumption cells, we

[17] Concerning 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.

see that there are small differences for p, but that b/c is substantially smaller for the those in the high cash-expenditure group. 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 6. Using information from Table I for the corresponding cash expenditure to which these percentages refer, the mean values of b for households with and without ATMs 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.[18]

Next we discuss three patterns that emerge from the estimates of $(p, b/c)$ that are consistent with the economics of the withdrawal technology, but that were not imposed in the estimation of these parameters, which we take as supportive of the economic significance of the model and its estimates. The first pattern is that households with ATM cards have higher values of p and correspondingly lower values of b/c than those without cards. This can be seen for the mean and median values in Table III, but more strikingly (not shown in the table), the estimated value of p is higher for those with ATM cards in 88% of the cells, and the value for b/c is smaller in 82% of the cells. This pattern supports the hypothesis that households with ATM cards have access to a superior withdrawal technology.

The second pattern is the positive correlation (0.69) of the estimated values of b/c between households with and without ATM across province–year–consumption cells. Likewise we find a correlation (0.30) of the estimated values of p between households with and without ATM cards. This pattern supports the hypothesis that province–year–consumption cells are characterized by different levels of efficiency on the withdrawal technology for both ATM and non-ATM card holders.

The third pattern is that b/c shows a strong negative correlation with indicators of the density of financial intermediaries (bank branches and ATMs per resident, shown in Table II) that vary across provinces and years. Likewise, the correlation of p with those indicators is positive, although it is close to zero (see Alvarez and Lippi (2007) for details). As greater financial diffusion raises the chances of a free withdrawal opportunity (p) and reduces the cost of contacting an intermediary (b/c), we find that these correlation are consistent with the economics of the model.

We find these patterns reassuring since we have estimated the model independently for each cell, that is, for ATM holders/nonholders (first pattern), for province–year–consumption combinations (second pattern), and without using information on indicators of financial diffusion (third pattern).

Finally, we report on the statistical goodness of fit of the model. The bottom panel of Table III reports some statistics on the goodness of fit of the

[18]The sources are RBR (2005) and an internal report by the Bank of Italy.

model. Let S be the number of estimation cells and consider a cell $s \leq S$ with data z_s and estimated parameter θ_s. Under the assumption of normally distributed errors, or as the number of households in the cell is large, the minimized likelihood criterion $F(\theta_s; z_s)$ is distributed as a $\chi^2_{(2)}$. The 2 degrees of freedom result from having four observable variables—that give four moments—and two parameters, that is, two overidentifying restrictions. As standard, the cell s passes the overidentifying restriction test with a 10% confidence level if $F(\theta_s; z_s) < 4.6$, the 90th percentile of $\chi^2_{(2)}$. As shown in the table, this happens for 48% and 59% of the cells for households with and without ATM cards, respectively. In this sense the statistical fit of the model is relatively good. Alternatively, under the assumption that errors are independent across cells, the vector $\{F(\theta_s; z_s)\}_{s=1}^{S}$ are is a sample of size S from a chi square with 2 degrees of freedom. Since S is a large number, the fraction of cells with $F < 4.6$ should be around 0.90 and with $F < 1.4$, the median of a $\chi^2_{(2)}$, should be around 0.50. As the corresponding values in the table are smaller, the joint statistical fit of the model is poor.

5.3. *Estimates With Observed Household Heterogeneity*

This section explores an alternative estimation strategy that incorporates observed household level heterogeneity. It is assumed that the four variables $(M/c, W/M, n, \underline{M}/M)$ are observed with classical measurement error, and that households differ in the parameters b/c and p which are given by a simple parametric function of household observables. In particular, let X_i be a k dimensional vector containing the value of households i covariates. We assume that for each household i the values of b/c and p are given by $(b/c)_i = \exp(\lambda_{b/c} \cdot X_i)$ and $p_i = \exp(\lambda_p \cdot X_i)$, where λ_p and $\lambda_{b/c}$ are the parameters to be estimated. The vector X_i contains $k = 8$ covariates: a constant, calendar year, the (log) household cash expenditure, an ATM dummy, a measure of the financial diffusion of bank branches (BB) and ATM terminals at the province level, a credit card dummy, the (log) income level per adult, and the household (HH) size.

Assuming that the measurement error is independent across households and variables, the maximum likelihood estimate of λ minimizes

$$F(\lambda; X, z) \equiv \sum_{j=1}^{4} \frac{1}{\sigma_j^2} \sum_{i=1}^{N} \left[z_i^j - \zeta^j(\theta(\lambda, X_i, R_i))\right]^2,$$

where, as above, z_i^j is the log of the jth observable for household i, $\zeta^j(\theta)$ is the model solution given the parameters θ, and N is the number of households in the sample.[19] The estimation proceeds in two steps. We first estimate σ_j^2

[19]We treat the opportunity cost R_i as known. To speed up the calculations, we estimate the model by assuming that inflation is zero, which has almost no effect on the estimates. We also

TABLE IV

HOUSEHOLD LEVEL $(p, b/c)$ ESTIMATES WITH OBSERVED HETEROGENEITY[a]

X_i Covariates	λ_p	t-Stat	$\bar{\lambda}_p$	$\lambda_{b/c}$	t-Stat	$\bar{\lambda}_{b/c}$
Constant	−87.7	(−1370)	−87.7	225	(3340)	217
Year	0.04	(0.67)	0.04	−0.11	(−1.64)	−0.11
log cash expenditure	0.04	(0.03)	−0.01	−0.96	(−0.62)	−0.97
ATM dummy	1.24	(64.70)	1.28	−0.66	(−32.7)	−0.75
log ATM and BB density	−0.15	(−1.30)	−0.16	−0.37	(−2.8)	−0.34
Credit card dummy	0.30	(2.96)	0.21	−0.01	(−0.05)	0.08
log income	0.25	(4.18)	0.30	0.26	(4.05)	0.33
log HH size	0.35	(4.05)	0.28	0.26	(2.74)	0.27

[a]Estimates for p and b/c under the assumption that $(b/c)_i = \exp(\lambda_{b/c} X_i)$ and $p_i = \exp(\lambda_p X_i)$, X_i is at the household level, and $(M/c, W/M, n, \underline{M}/M)$ is measured with error.

for each of the four variables by running a regression at the household level of each of the four variables against the household level X_i. We then minimize the likelihood criterion $F(\cdot; X, z)$, taking the estimated $\hat{\sigma}_j^2$ as given. The asymptotic standard errors of λ are computed by inverting the information matrix.

Table IV presents the estimates of λ. The first data column displays the point estimates of λ_p and the fourth data column displays the point estimates for $\lambda_{b/c}$. The numbers in parentheses next to the point estimates are the corresponding t-statistics. To compare the results with the baseline estimates of Section 5.2, the table also includes the coefficients of two regressions, labeled $\bar{\lambda}_p$ and $\bar{\lambda}_{b/c}$. The dependent variables of these regression are the baseline estimates of p and b/c, and hence they are the same for all households in a cell (i.e., combination of a year, province, ATM card ownership, and third-tile cash consumption). The right hand side variables are the cell means of the X_i covariates.

We summarize the findings of the household level observed heterogeneity estimates displayed in Table IV. First, and most importantly, the values of λ and $\bar{\lambda}$ are extremely close, which shows that the benchmark cell estimates and the household level estimates provide the same information on the variations of $(p, b/c)$ on observables.[20] The estimates of p and b/c that correspond to a household with average values of each of the X_i variables and our estimated parameters λ_p and $\lambda_{b/c}$ are, respectively, 11 and 5.2%. These values are similar to the estimates reported in Table III (in particular they are close to the median across cells). The mean estimate for p, greater than zero, supports the introduction of this dimension of the technology, as opposed to having only the BT parameter b/c. The estimates of both ATM dummies are economically

restrict the sample to households that have information for all four variables of interest. This gives us a sample of about $N = 17{,}000$ (as opposed to 40,000 in our baseline estimates).

[20]The exceptions are two values which are small and statistically insignificant.

important and statistically significant. Households with an ATM card have a value of p approximately three times larger ($\exp(1.24) \approx 3.46$) and a value of b/c about half ($\exp(-0.66) \approx 0.52$) relative to households without ATM cards. There is a small positive time trend on p and a larger negative time trend on b/c, although neither estimate is statistically significant. The value of b/c is smaller in locations with a higher density of ATMs or bank branches, with an elasticity of -0.37, which is borderline statistically significant, but this measure has a small negative effect on p. Credit card ownership has no effect on b/c and a small positive (borderline significant) effect on p. A possible interpretation for the effect on p is that households with a credit card have better access to financial intermediaries. We find a positive effect of the household size (number of adults) on both p and b/c.[21] The coefficient of cash expenditure indicates no effect on p and a negative near-unit elasticity with respect to b/c, though it is imprecisely estimated (this elasticity is very close to that estimated using cell level aggregated data). The income per adult has a positive elasticity of about 0.25 for both p and b/c. We interpret the effect of income per capita on p as reflecting better access to financial intermediaries, and with respect to b/c as measuring a higher opportunity cost of time. The combination of the effects of income per capita and cash expenditures yields the following important corollary: the value of b is estimated to be independent of the level of cash expenditure of the household, implying a cash-expenditure elasticity of money demand of approximately one-half *provided that the opportunity cost of time is the same*.

Under the assumption of independent measurement error, the value of the likelihood criterion F is asymptotically distributed as a χ^2 with $N \times 4 - 2 \times k = 54{,}260$ degrees of freedom.[22] The minimized value for F, given by $F = 62{,}804$, implies a relatively poor statistical fit of the model since the tail probability for the corresponding χ^2 of such value of F is essentially zero.

6. 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 baseline estimates of p and b/c. Table V shows that p has approximately doubled and that (b/c) has approximately halved over the sample period. In words, our point estimates indicate that the withdrawal technology

[21] This result is hard to interpret because if the withdrawal technology had increasing (decreasing) returns with respect to the household size, we would have expected the p and b/c to vary in opposite ways as the size changed.

[22] F equals half of the log-likelihood minus a constant not involving λ. We estimate k loadings $\lambda_{b/c}$ and k loadings λ_p using N households with four observations each.

TABLE V

TIME SERIES PATTERN OF ESTIMATED MODEL PARAMETERS[a]

	1993	1995	1998	2000	2002	2004	All Years
Households with ATMs							
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
Households without ATMs							
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

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

has improved through time.[23] The table also reports $\hat{b}/R \equiv (b/c)p^2/R$, which, as shown in Proposition 6 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.

By Proposition 6, 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 V 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 cards, respectively. Even for the largest values of $\hat{b}/R$ recorded in Table V, the value of η remains above 0.4. In fact, further extending the range of Figure 1 shows 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 (2009) for Italian data and Daniels and Murphy (1994) using U.S. 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

[23] Since we have only six time periods, the time trends are imprecisely estimated, as can be seen from the t-statistics corresponding to years in Table IV.

TABLE VI

A LABORATORY EXPERIMENT ON THE INTEREST ELASTICITY OF MONEY DEMAND[a]

Dependent Variable: $\log(M/c)$	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 (cells)	1654	1654	1539	1539

[a] All regressions include a constant.

interest rate elasticity is biased toward zero. To explore this hypothesis, in Table VI 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 6, $(M/c)p$ has elasticity $\eta(\hat{b}/R)$ so that we approximate it using a constant elasticity: $\log M/c = -\log p + \eta(\log(b/c) + 2\log(p)) - \eta\log(R)$. The regression coefficient for η estimated from this equation gives virtually the same value obtained from Figure 1. Since the left hand side of the equation uses the values of M/c produced by the model using the estimated p and b/c and no measurement error, the only reason why the regression R^2 does not equal 1 is that we are approximating a nonlinear function with a linear one. Yet the R^2 is pretty close to 1 because the elasticity for this range of parameters is close to constant. To estimate the size of the bias due to omission of 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, pointing to a large omitted variable bias: the correlation between $(\log(b/c) + 2\log(p))$ and $\log R$ is 0.12 and 0.17 for households with and without ATM cards, 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 VI using the actual, as opposed to the fitted, value of M/c yields very similar results (not reported here).

We now estimate the expenditure elasticity of the money demand. An advantage of our data is that we use direct measures of cash expenditures (as opposed to income or wealth).[24] By Proposition 6, the expenditure elasticity is

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

[24] Dotsey (1988) argued for the use of cash expenditure as the appropriate scale variable.

For instance, if the ratio b/c is constant across values of c, then the elasticity is 1; 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 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 cards, respectively. Using the estimates for η, we obtain that the mean expenditure elasticity is $1 + 0.48 \cdot (-0.82) = 0.61$ for households without ATMs, and 0.56 for those with.

7. COST OF INFLATION AND BENEFITS OF ATM CARD

This section uses our model to quantify the cost of inflation and the benefits of ATM card ownership. Section 4.3 shows that the loss is $\ell = R(m^* - M)$ and the household cost is $v = Rm^*$. We use the baseline estimates of $(p, \frac{b}{c})$ from Section 5.2 to compute m^* and M and the implied losses for each estimation cell. The analysis shows that the cost v is lower for households with ATM cards, reflecting their access to a better technology, and that it is lower for households with higher cash expenditures c, reflecting that our estimates of b/c are uncorrelated with c. Quantitatively, the sample mean value of ℓ across all years and households in our sample is about 15 euros or approximately 0.6 day of cash purchases per year.

To put this quantity in perspective, we relate it to the one in Lucas (2000), obtained by fitting a log-log money demand with constant interest elasticity of 1/2, which corresponds to the BT model. Our model predicts a smaller welfare cost of inflation relative to BT: $\ell/(RM) = m^*/M - 1$ (see Section 4.3). For $R = 0.05$ and $\hat{b}/R = 1.8$, which are about the mean of our baseline estimates, $\ell/(RM) = 0.6$, which shows that the welfare cost in our model is 40% smaller than in BT. As discussed after Proposition 7, the discrepancy is due to the different behavior of the interest rate elasticity in our model. As indicated by Lucas, the behavior of the elasticity at low interest rates is key to quantifying the inflation costs. Despite the fact that the interest elasticity is about 1/2 in both models at the sample mean estimates, the elasticity is constant in BT while it is decreasing and eventually zero in our model (recall Proposition 6). Another difference between these estimates is the choice of the monetary aggregate. In both models the welfare cost is proportional to the level of the money demand. While we focus on currency held by households, Lucas used the stock of M1, an aggregate much larger than ours.[25]

[25] Attanasio, Guiso, and Jappelli (2002) fitted a different model to the same data set, focusing mostly on cash balances M—as opposed to W, n, and $\underline{M}$—but endogenizing the decision to obtain an ATM card. They also found a first order difference compared to Lucas' estimates that originates from the use of a smaller monetary aggregate.

TABLE VII

DEADWEIGHT LOSS ℓ AND HOUSEHOLD COST v OF CASH PURCHASES[a]

Variable[b]	1993	1995	1998	2000	2002	2004	Mean
ℓ	24	23	11	11	10	10	15
v	51	49	25	25	22	25	33

[a] ℓ and v are weighted sample averages, measured as annual flows.
[b] Per household in 2004 euros.

Table VII shows that the welfare loss in 2004 is about 40% smaller than in 1993. The reduction is due to the decrease in R and advances in the withdrawal technology (decreases in b/c and increases in p).[26]

We use v to quantify the benefits associated with ownership of an 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 an ATM card, keeping cash purchases constant, is defined as $v_0 - v_1 c_0/c_1 = R(m_0^* - m_1^* c_0/c_1)$, where the subscript 1 (0) indicates ownership of (lack of) an ATM card. Our computations show that the mean benefit of ATM card ownership, computed as the weighted average of the benefits across all years and households, is 17 euros. The benefit associated with ATM card ownership is estimated to be positive for over 91% of the province–year–type estimates. The null hypothesis that the gain is positive cannot be rejected (at the 10% confidence level) in 99.5% of our estimates. Since our estimates of the parameters for households with and without ATMs are done independently, we think that the finding that the estimated benefit is positive for most province–years provides additional support for the model.

There are two important limitations of this counterfactual exercise. First, the estimated benefit assumes that households without ATM cards differ from those with a card *only* in terms of the withdrawal technology that is available to them $(p, b/c)$. The second is that ATM cards provide other benefits, such as access to electronic retail transactions. In future work we plan to study the household card adoption choice, which will be informative on the size of the estimates' bias. Yet, we find it interesting that our estimated benefit of ATM cards is 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 RBR (2005)).

[26] A counterfactual exercise suggests that the contribution of the disinflation and of technological change to the reduction in the welfare loss is of similar magnitude; see Section 8 in Alvarez and Lippi (2007).

8. CONCLUSIONS

This paper proposes a simple, tightly parameterized, extension of the classic Baumol–Tobin model to capture empirical regularities of households' cash management. We now discuss some extensions of the model that we plan to develop fully in the future.

Our model has some unrealistic features: all random withdrawals are free and all cash expenditures are deterministic. Two variations of our model that address these issues are sketched below. The first one introduces an additional parameter f, which denotes a fixed cost for withdrawals upon random contacts with the financial intermediary (see Appendix F in Alvarez and Lippi (2009)). The motivation for this is that when random withdrawals are free, the model has the unrealistic feature that agents withdraw every time they match with an intermediary, making several withdrawals of extremely small size. Instead, the model with $0 < f < b$ has a strictly positive minimum withdrawal size. In Appendix I in Alvarez and Lippi (2009) we use a likelihood ratio test to compare the fit of the $f > 0$ model with our benchmark $f = 0$ model. It is shown that the fit does not improve much. Additionally, we show that the parameter f is nearly not identified. To understand the intuition behind this result, notice that the BT model is obtained for $p = 0$, $f = 0$, and $b > 0$ or, equivalently, for $f = b > 0$, and $p > 0$. More data, such as information on the minimum withdrawal size, would be needed to estimate $f > 0$. We left this exploration for future work.

The second variation explores the consequences of assuming that the cash expenditure has a random component. One interesting result of this model is that it may produce $W/M \geq 2$ or, equivalently, $M < \mathbb{E}(c)/2n$, where $\mathbb{E}(c)$ stands for expected cash consumption per unit of time. These inequalities are indeed observed for a small number of households, especially those without ATM cards (see Table I). However, this model is less tractable than our benchmark model, and it is inconsistent with the large number of withdrawals and the values of W/M that characterize the behavior of most households in the sample. Although we solved for the dynamic programming problem for both variations, as well as for the implied statistics for cash balances and withdrawals, we do not develop them further here to keep the discussion simple. Moreover, as briefly discussed, while the models incorporate some realistic features of cash management, they deliver only a modest improvement on the fit of the statistics that we focussed on.

Our model, like the one by BT, takes as given the household cash expenditure. We think that this framework should work well as an input for a cash–credit model, and view this as an important extension for future work. New household level data sets with information on cash management, similar to the one we have used, as well as detailed diary information on how different purchases were paid (cash, debit, credit card, check, etc.) will allow careful quantitative work in this area.[27]

APPENDIX: PROOFS FOR THE MODEL WITH FREE WITHDRAWALS

PROOF OF PROPOSITION 1: Given two functions G and V that satisfy (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,

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

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

$$\begin{aligned}&[r + p_1 + p_2]V(m)\\&\quad = V'(m)(-c - \pi m) + [p_1 + p_2]V(m^*) + [r + p_2 + \pi]m.\end{aligned}$$

Using $R = r + \pi + p_2$ and $p = p_1 + p_2$, we obtain the desired result, that is, (12). The proof that if V solves the Bellman equation for $m \in (0, m^{**})$, then so does G defined as in (14) follows from analogous steps. *Q.E.D.*

LEMMA 1: *Let V^* be an arbitrary nonnegative value.*

(a) *For $m \in (0, m^{**})$ the ODE in (10) is solved by (16) for some constant $A > 0$.*

(b) *Imposing that (16) satisfies $V(0) = V^* + b$ gives*

$$A = \frac{V^*(r+p)r + Rc\Big/\left(1 + \dfrac{\pi}{r+p}\right) + (r+p)^2 b}{c^2} > 0. \tag{28}$$

(c) *The expressions in (16) and (28) imply that $V(\cdot)$ is a convex function of m.*

(d) *Let A be the constant that indexes the expression for $V(\cdot)$ in (16). The value m^* that solves $V'(m^*) = 0$ is*

$$m^* = \frac{c}{\pi}\left(\left[\frac{R}{Ac}\Big/\left(1 + \frac{\pi}{r+p}\right)\right]^{-\pi/(r+p+\pi)} - 1\right). \tag{29}$$

[27]One such data set was developed by the Austrian National Bank and was used, for instance, by Stix (2004) and Mooslechner, Stix, and Wagner (2006).

(e) *The value of V^* is*

$$V^* = \frac{R}{r} m^*. \tag{30}$$

PROOF: (a) Follows by differentiation. (b) Follows using simple algebra. (c) Direct differentiation gives $V''(m) > 0$. (d) Follows using simple algebra. (e) Replacing $V'(m^*) = 0$ and $V(m^*) = V^*$ in (10) evaluated at $m = m^*$ yields $rV^* = Rm^*$. *Q.E.D.*

PROOF OF PROPOSITION 2: Lemma 1 yields a system of three equations, (28), (29), and (30), in the three unknowns A, m^*, and V^*. Replacing equation (30) into (28) yields one equation for A. Rearranging equation (29), we obtain another equation for A. Equating these expressions for A, collecting terms, and rearranging yields equation (15). Let $f(m^*)$ and $g(m^*)$ be the left and the right hand sides of equation (15), respectively. We know that $f(0) < g(0)$ for $b > 0$, $g'(0) = f'(0) > 0$, $g''(m^*) = 0$, and $f''(m^*) > 0$ for all $m^* > 0$. Thus there exists a unique value of m^* that solves (15).

(i) 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^*.

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

$$\frac{b}{cR} = \left(\frac{m^*}{c}\right)^2 \left[\frac{1}{2} + \sum_{j=1}^{\infty} \frac{1}{(2+j)!} \left[\prod_{s=1}^{j} (r + p - s\pi)\right] \left(\frac{m^*}{c}\right)^j\right], \tag{31}$$

which follows by expanding $(\frac{m}{c}\pi + 1)^{1+(r+p)/\pi}$ around $m = 0$. We notice that $m^*/c = \sqrt{2b/(cR)} + o(\sqrt{b/c})$ is equivalent to $(m^*/c)^2 = 2b/(cR) + [o(\sqrt{b/c})]^2 + 2\sqrt{2b/(cR)}\,o(\sqrt{b/c})$. Inserting this expression into (31), dividing both sides by $b/(cR)$, and taking the limit as $b/(cR) \to 0$ verifies our approximation.

(iii) For $\pi = R - r = 0$, using (31) 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 right hand side 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 brackets, 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

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

Thus

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

(iv) For $\pi \to 0$, equation (15) yields $\exp(m^*/c(r+p)) = 1 + m^*/c(r+p) + (r+p)^2 b/(cR)$. 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 with respect to (w.r.t.) $\log R$,

$$z'\left[1 + \left(\frac{1}{2}\right)\frac{\varphi'(\exp(z))\exp(z)}{\big(1+\varphi(\exp(z))\big)}\right] = -\frac{1}{2}$$

or

$$\eta_{x,R} \equiv -\frac{R}{x}\frac{dx}{dR} = \frac{(1/2)}{1 + (1/2)\dfrac{\varphi'(x)x}{1+\varphi(x)}}.$$

Direct computation gives

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

where

$$\kappa_j(x) = \frac{\dfrac{2}{(j+2)!}[x]^j}{1 + \displaystyle\sum_{s=1}^{\infty} \frac{2}{(s+2)!}[x]^s} \quad \text{for} \quad j \geq 1$$

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 $\kappa_{j+1}(x)/\kappa_j(x) = x/(j+3)$ 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 $x \to 0$, the distribution κ puts all the mass in $j = 0$ and hence $\eta_{x,R} \to 1/2$. As $x \to \infty$, the distribution κ concentrates all the mass in arbitrarily large values of j, hence $E^x[j] \to \infty$ and $\eta_{x,R} \to 0$. *Q.E.D.*

PROOF OF PROPOSITION 3: (i) The function $V(\cdot)$ and the expression for A are derived in parts (a) and (b) of Lemma 1.

(ii) V^* is given in part (e) of Lemma 1. *Q.E.D.*

PROOF OF PROPOSITION 4: (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. Let

$$h(m, t; \Delta) = \frac{\psi(m, t; \Delta)}{\Delta(m\pi + c)} \tag{32}$$

so that $\lim h(m, t; \Delta)$ as $\Delta \to 0$ 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

$$\psi(m, t + \Delta; \Delta) = \psi(m + \Delta(m\pi + c), t; \Delta)(1 - \Delta p). \tag{33}$$

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 (32) into (33), 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 $\Delta \to 0$, we obtain (19). The solution of this ODE is $h(m) = 1/m^*$ if $p = \pi$ and $h(m) = A[1 + \pi\frac{m}{c}]^{(p-\pi)/\pi}$ for some constant A if $p \neq \pi$. The constant A is chosen so that the density integrates to 1, so that $A = 1/\{(\frac{c}{p})([1 + \frac{\pi}{c}m^*]^{p/\pi} - 1)\}$.

(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

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

In the interval $[m_1^*, m_2^*)$ we have $H(m; m_1^*) = 1 > H(m; m_2^*)$. *Q.E.D.*

PROOF OF PROPOSITION 5: 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 (18) 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 1, 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 $\frac{\partial h(\hat{m}; p')}{\partial m} > \frac{\partial h(\hat{m}, p)}{\partial m}$. To summarize, $h(m; p) > h(m; p')$ for $0 \le m < \hat{m}$, $h(\hat{m}; p) = h(\hat{m}; p')$, and $h(m; p) < h(m; p')$ for $\hat{m} < m \le m^*$. 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

$$\begin{aligned}\frac{1}{h(m^*)} &= \frac{c}{p} \frac{\left[1 + \pi\frac{m^*}{c}\right]^{p/\pi} - 1}{\left[1 + \pi\frac{m^*}{c}\right]^{p/\pi - 1}} \\ &= \frac{c}{p}\left[1 + \pi\frac{m^*}{c}\right]\left(1 - \frac{1}{\left[1 + \pi\frac{m^*}{c}\right]^{p/\pi}}\right),\end{aligned}$$

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

PROOF OF PROPOSITION 6: 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 $x \to \infty$ as $\hat{b}/R \to \infty$ and $x \to 0$ as $\hat{b}/R \to 0$.

To see how the ratio Mp/c depends on x, notice that from (24) we have that $Mp/c = \phi(xp/(p + r))$, where $\phi(z) \equiv z/(1 - e^{-z}) - 1$. Thus $\lim_{r\to 0} Mp/c =$

$\phi(x)$. To see why the ratios W/M and $\underline{M}/M$ are functions only of x, note from (24) that $p/n = 1 - \exp(-pm^*/c) = 1 - \exp(-xp/(p+r))$ and hence as $r \to 0$, we can write $p/n = \omega(x) = \underline{M}/M$, where the last equality follows from (22) and ω is the function $\omega(x) \equiv 1 - \exp(-x)$. Using (27) 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, and α is straightforward to check. The limits for $\underline{M}/M$ and W/M as $x \to 0$ or as $x \to \infty$ follow from a tedious but straightforward calculation.

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

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

that is, 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 $\eta_{\phi,x} = (x(1-e^{-x}-xe^{-x}))/((x-1+e^{-x})(1-e^{-x}))$, where $\lim_{x\to\infty}\eta_{\phi,x} = 1$. A second order expansion of each of the exponential functions shows that $\lim_{x\to 0}\eta_{\phi,x} = 1$. Direct computations using $x = \gamma(\hat{b}/R)$ yield $\eta_{x,\hat{b}/R} = (e^x - x - 1)/(x(e^x - 1))$. It is immediate that $\lim_{x\to\infty}\eta_{x,\hat{b}/R} = 0$ and $\lim_{x\to 0}\eta_{x,\hat{b}/R} = 1/2$. *Q.E.D.*

PROOF OF PROPOSITION 7: (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

$$\begin{aligned}&\frac{\partial m^*}{\partial R}\left[\left(1+\pi\frac{m^*}{c}\right)^{(r+p)/\pi} - 1\right]\frac{(r+p+\pi)}{c}\\&\quad = -\frac{b}{cR^2}(r+p)(r+p+\pi).\end{aligned}$$

Using (15) again to replace $\frac{b}{cR}(r+p)(r+p+\pi)$, inserting the resulting expression into $m^*(R) + R\,\partial m^*(R)/\partial R$, letting $r \to 0$, and rearranging yields the expression for M obtained in (20).

(iii) Using (i) in (iii) yields $R(m^* - M) = (n-p)b$. Replacing M and n using equations for the expected values (17) and (20) for an arbitrary m^* yields an equation identical to the one characterizing the *optimal* value of m^*, (15), evaluated at $r = 0$. *Q.E.D.*

REFERENCES

ALVAREZ, F. E., AND F. LIPPI (2007): "Financial Innovation and the Transactions Demand for Cash," Working Paper 13416, National Bureau of Economic Research. [384,387,394]

——— (2009): "Supplement to 'Financial Innovation and the Transactions Demand for Cash'," *Econometrica Supplemental Material*, 77, http://econometricsociety.org/ecta/Supmat/7451_data.pdf; http://econometricsociety.org/ecta/Supmat/7451_data and programs.zip. [367,368,378,381,382,385,395]

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*, 110, 318–351. [365,378,393]

BAUMOL, W. J. (1952): "The Transactions Demand for Cash: An Inventory Theoretic Model," *Quarterly Journal of Economics*, 66, 545–556. [363,372]

CONSTANTINIDES, G. M. (1976): "Stochastic Cash Management With Fixed and Proportional Transaction Costs," *Management Science*, 22, 1320–1331. [372]

CONSTANTINIDES, G. M., AND S. F. RICHARD (1978): "Existence of Optimal Simple Policies for Discounted-Cost Inventory and Cash Management in Continuous Time," *Operations Research*, 26, 620–636. [372]

COOLEY, T., AND G. HANSEN (1991): "The Welfare Costs of Moderate Inflations," *Journal of Money, Credit and Banking*, 23, 483–503. [367]

DANIELS, K. N., AND N. B. MURPHY (1994): "The Impact of Technological Change on the Currency Behavior of Households: An Empirical Cross Section Study," *Journal of Money, Credit and Banking*, 26, 867–874. [391]

DOTSEY, M. (1988): "The Demand for Currency in the United States," *Journal of Money, Credit and Banking*, 20, 22–40. [392]

FRENKEL, J. A., AND B. JOVANOVIC (1980): "On Transactions and Precautionary Demand for Money," *The Quarterly Journal of Economics*, 95, 25–43. [372]

HUMPHREY, D. B. (2004): "Replacement of Cash by Cards in US Consumer Payments," *Journal of Economics and Business*, 56, 211–225. [367]

LIPPI, F., AND A. SECCHI (2009): "Technological Change and the Households' Demand for Currency," *Journal of Monetary Economics* (forthcoming). [368,391]

LUCAS, R. E. J. (2000): "Inflation and Welfare," *Econometrica*, 68, 247–274. [365,381,393]

LUCAS, R. E. J., AND N. L. STOKEY (1987): "Money and Interest in a Cash-in-Advance Economy," *Econometrica*, 55, 491–513. [365]

MILLER, M., AND D. ORR (1966): "A Model of the Demand for Money by Firms," *Quarterly Journal of Economics*, 80, 413–435. [372]

MOOSLECHNER, P., H. STIX, AND K. WAGNER (2006): "How Are Payments Made in Austria? Results of a Survey on the Structure of Austrian Households Use of Payment Means in the Context of Monetary Policy Analysis," *Monetary Policy and the Economy*, 2, 111–134. [396]

PORTER, R. D., AND R. A. JUDSON (1996): "The Location of U.S. Currency: How Much Is Abroad?" *Federal Reserve Bulletin*, 82, 883–903. [367,378]

RETAIL BANKING RESEARCH (RBR) (2005): "Study of the Impact of Regulation 2560/2001 on Bank Charges for National Payments," Report RBR, prepared for the European Commission, London. [387,394]

STIX, H. (2004): "How Do Debit Cards Affect Cash Demand? Survey Data Evidence," *Empirica*, 31, 93–115. [396]

TOBIN, J. (1956): "The Interest Elasticity of Transactions Demand for Money," *Review of Economics and Statistics*, 38, 241–247. [363,372]

Dept. of Economics, University of Chicago, 1126 East 59th Street, Chicago, IL 60637, U.S.A. and NBER; f-alvarez1@uchicago.edu

and

Dept. of Economics, University of Sassari, Via di Torre Tonda, 34, 07100 Sassari, Italy and Einaudi Institute for Economics and Finance, Rome, Italy and CEPR; f.lippi@uniss.it.

Manuscript received September, 2007; final revision received October, 2008.