To Forgive or Not to Forgive: An Analysis of U.S. Consumer Bankruptcy Choices

Wenli Li

Current U.S. bankruptcy law has two separate bankruptcy procedures, known as Chapter 7 and Chapter 13. When a debtor files for bankruptcy under Chapter 7, he or she must give up all assets not legally sheltered from creditor seizure in exchange for a discharge of almost all preexisting debts. Under Chapter 13, a debtor may keep all property in exchange for a promise to pay all or some specified part of his or her debts under a payment plan approved by the court.[1] Between 1980 and 1999, the total number of U.S. personal bankruptcy filings rose from 331,257 to nearly 1.4 million per year, and the rate of consumer bankruptcies per 100,000 adults increased from 201 to 650. Most bankruptcy filings during that period (about 70 percent) were under Chapter 7 as opposed to Chapter 13, which accounts for much of the increase in the total rates. As a result, net losses to creditors grew twice as fast as consumer installment credit during those years; today, those losses are counted in the tens of billions of dollars.

The continued climb in consumer bankruptcy rates and the resulting losses to creditors have generated considerable debate and led to a number of bankruptcy reform proposals. Although there is as yet no consensus concerning the driving force behind the drastic upward trend in U.S. personal bankruptcy

This article was written while the author was with the Federal Reserve Bank of Richmond; she is now with the Board of Governors of the Federal Reserve System. She would like to thank Kartik Athreya, Margarida Duarte, Jeff Lacker, Pierre-Daniel Sarte, and John Weinberg for their helpful comments. The views expressed in the article are those of the author and do not necessarily reflect those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

[1] Section 1 provides more detailed information on the basic law.

filings,[2] researchers, practitioners, and politicians are in agreement that the current U.S. consumer bankruptcy system needs serious reform. Many proposals for reform focus on bankruptcy choices, particularly whether Chapter 13 should be encouraged over Chapter 7. The National Bankruptcy Review Commission, for example, recently recommended reforms that would eliminate reaffirmations of debt altogether. Under this recommendation, all debts would be forgiven and there would be no Chapter 13 bankruptcy. On the other hand, the Gekas bill, which has passed the House of Representatives, H.R. 3150, views Chapter 13 as an alternative to be encouraged over Chapter 7. The bill will force bankrupt debtors whose income is above the median to use all of their post-bankruptcy earnings above a predetermined level to repay debt. The credit industry supports the Gekas bill.

Changes in bankruptcy provisions will affect which chapter bankrupts choose, should they elect bankruptcy. In addition, bankruptcy provisions will affect how attractive bankruptcy is in comparison to the option of repaying debts in full as scheduled. Thus, changes in bankruptcy provisions will have an effect on the likelihood of bankruptcy. The odds of a borrower declaring bankruptcy will affect the riskiness of the loan, and thus will affect risk premia charged by competitive lenders and, in turn, borrowers' loan demand and the rate of return to savers in the economy. The likelihood that future income will be garnished to repay a loan will affect borrowers' choices regarding income—labor "effort," for example—but this can be seen as a stand-in for a broader array of incentive effects. To fully analyze the implications of a change in bankruptcy provisions, one must take into account these incentive effects as well as the general equilibrium effects. I present a simple tractable framework for capturing these effects. It may not be detailed enough to be calibrated to current statistical observations, but it does provide reliable qualitative answers to the key questions: What are the characteristics of those that are affected by proposed changes in the code? and, What are the efficiency implications of proposed changes?

The basic economic argument for having a personal bankruptcy procedure is that it helps risk-averse borrowers by providing them with insurance against the possibility that their income or wealth might fall at the time when they have to repay their loans. Thus, borrowers and creditors share the risk of a fall in borrowers' income or wealth. The implicit assumption here is that consumers cannot fully insure against their idiosyncratic income risk. This market incompleteness can arise because of informational problems. For instance, if individual incomes are private information and hence unverifiable, then it may be impossible to provide any insurance against the risks that

[2] Two competing explanations are declining bankruptcy stigma, which led to increasing abuse of the system by borrowers, and the increasingly reckless practices of major consumer creditors, especially credit card issuers.

individuals face. If there are adverse selection problems (different groups of individuals have different risk characteristics; these are private information, and therefore insurance companies protect themselves by penalizing entire groups rather than single persons), then some groups may be prevented from buying as much insurance as they would like (Aiyagari 1997).[3] Accordingly, in the economic environment studied here, individuals face fluctuating income streams and can save through a riskless saving instrument and can borrow as well. They borrow in order to smooth their consumption intertemporally.[4]

My analysis suggests that given labor income and outstanding debt, an individual will file for bankruptcy if his or her assets fall below a certain threshold. This threshold varies negatively with income. Among those who file for bankruptcy, individuals with higher assets and lower income tend to choose Chapter 13, while those with lower assets and higher income tend to choose Chapter 7. These findings, then, confirm the general view on consumers' bankruptcy choices. My discussion also indicates that ex ante, individuals may hold assets and debts simultaneously. In other words, bankruptcy provisions provide an incentive for people to borrow in order to save. Furthermore, Chapter 13 decreases a person's labor effort. To the extent that this labor effort is not directly observable, it can be measured by a person's labor productivity.

In terms of policy experiments, I focus on three policy instruments: asset exemption levels under Chapter 7, the percentage of labor income that can be garnished under Chapter 13, and a mandated income rule for Chapter 13. My analysis shows that there is an efficiency tradeoff among the policies and that their distributive effects differ greatly. In particular, an increase in the asset exemption level under Chapter 7 benefits people with medium assets and medium labor income; ex ante, they will save less and borrow more.

---

[3] Brunstad (2000) discusses a list of issues, which he refers to as "problems of economic futility," that require a unique legal system of bankruptcy law.

[4] This article's modeling strategy is most closely related to those of Athreya (2000) and Lehnert and Maki (1999). Both Athreya and Lehnert and Maki, however, study only bankruptcy filings under Chapter 7. Wang and White (2000) and Adler, Polak, and Schwartz (2000) research issues similar to those I investigate. Although Wang and White make different assumptions about consumer behavior (they assume that some strategic households can hide part of their wealth so as to get maximum financial benefit under the bankruptcy system) and focus on different aspects of policy analysis (optimal personal bankruptcy procedures), my article is consistent with theirs in that both studies confirm the general view that households with relatively more wealth file for Chapter 13. Adler, Polak, and Schwartz treat loan borrowing as exogenous and investigate how private contractual arrangement affects consumers' bankruptcy choices within a principal/agent framework. Empirically, Domowitz and Sartain (1999) estimate qualitative choice models of consumers' decisions to file for bankruptcy and their choice of bankruptcy chapter. They find that medical and credit card debt are the strongest contributors to bankruptcy. Higher marriage rates, employment rates, and income all encourage the choice of Chapter 13 rehabilitation over Chapter 7. Additionally, higher levels of equity relative to debt push debtors into Chapter 13. Nelson (1999) also studies consumer bankruptcy and chapter choice using state panel data. He finds that both homestead exemption laws and garnishment laws are statistically significant for bankruptcy choices. Sullivan, Warren, and Westbrook (1989) do case studies and discuss in great detail the characteristics of a sample of personal bankruptcy filings.

The general equilibrium price effects may dampen these results. A reduction in the percentage of labor income that can be collected by creditors under Chapter 13 benefits individuals with high assets and good labor income; ex ante, people exert less labor effort. Finally, the implementation of the labor income threshold for Chapter 13 will affect negatively those with few assets and medium labor income; ex ante, people exert less labor effort.

1. THE U.S. CONSUMER BANKRUPTCY SYSTEM

An Overview of the Basic Law

The key aspect of the current U.S. bankruptcy law is that there are two separate personal bankruptcy procedures, known as Chapter 7 and Chapter 13, and debtors are allowed to choose between them. Both procedures discharge many types of debts, causing losses for creditors.

Debtors who file under Chapter 7 of the U.S. Bankruptcy Code are not obliged to use future income to repay their debts and are only obliged to use wealth to repay debt to the extent that their wealth exceeds predetermined exemption levels. In other words, under Chapter 7 the bankruptcy court discharges all eligible debts so that the debtor enjoys a "fresh start." Chapter 7 asset exemptions fall into two categories: homestead (applied to equity in a home used as a primary residence) and non-homestead. Although bankruptcy is a matter of federal law and the rules are uniform across the United States, asset exemption levels under Chapter 7 are set by the state in which the debtor lives and vary widely. For instance, 7 states (Arizona, Florida, Kansas, Minnesota, Oklahoma, South Dakota, and Texas) have unlimited homestead exemptions, while 20 others (including Alabama, California, and Georgia) have homestead exemptions of $7,500 or less for individual debtors (see White [1998], Table 1, for Chapter 7 exemptions for all states and the District of Columbia). Non-homestead exemptions are less clear cut. State laws frequently allow households to exempt 100 percent of the value of a specific type of asset. For example, many states exempt 100 percent of the value of clothing for personal use.

Alternatively, debtors can file for bankruptcy under Chapter 13, which offers virtually the opposite option. In Chapter 13 the law allows debtors to keep all property, exempt and nonexempt, in exchange for a promise to pay all or some specified part of their debts under a three- to five-year payment plan approved by the court. The remainder of the debt will be discharged. There are several restrictions attached to the repayment plan. First, Chapter 13 sets minimum amounts that must be paid under a plan. For secured creditors, the minimum payment must at least match the value of the collateral, plus interest, but may be made over a time longer than the requirements of the statute. Certain unsecured creditors with special priorities (such as tax authorities, former spouses, and children with support orders) receive payment in full.

Figure 1 Annual Bankruptcy Filings as a Percent of Total Adult Population—United States

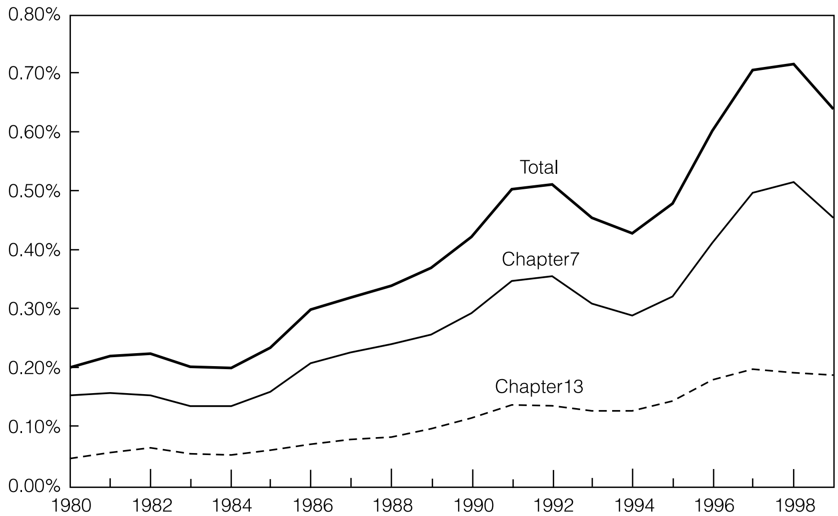

The minimum amount that must be paid to the rest of the unsecured creditors equals the disposable income remaining after necessary living expenses are paid and the secured and priority debt payments are made. Second, debtors must propose to pay at least as much as the creditors would have received under Chapter 7. Third, the statute requires that debtors propose plans in good faith, which is generally interpreted to mean that debtors must make some repayment even if there would have been none under Chapter 7; the amount required varies tremendously from district to district.

Debtors whose debts are discharged under either Chapter 7 or Chapter 13 are released from prebankruptcy debts, but they are ineligible for Chapter 7 for six years. Only debtors who pay in full under Chapter 13 remain eligible for a Chapter 7 bankruptcy discharge during the following six years. Creditors, especially unsecured creditors, generally favor a Chapter 13 arrangement since they may actually receive some payment under Chapter 13; under Chapter 7, they will be repaid only if debtors' assets exceed the exemption level. Nevertheless, most credit agencies identify all Chapter 13 debtors, regardless of their payment success, as having taken bankruptcy.

Recent Trends

Figure 1 depicts annual bankruptcy filings as a percent of the total adult population from 1980 to 1999 in the United States. Bankruptcy rates per 1,000 adults remained relatively stable from 1980 to 1984 but later exploded, moving from 1.2 in 1984 to 7.2 in 1998. The increase in bankruptcy rates lessened slightly in 1999. At current levels, 2 percent of American adults file for personal bankruptcy every three years. What was formerly a rarity is now almost commonplace. Also notable in Figure 1 is that about 70 percent of consumer bankruptcies are filed under Chapter 7. Moreover, Chapter 7 filings have increased at a faster rate than Chapter 13 filings. On the basis of this observation and the notion that Chapter 7 may be more harmful to the society since debtors do not need to repay any of their debts, Congress and the credit industry have formed the view that Chapter 13 should be greatly encouraged (or even mandated) so as to prevent the increasing abuse of the current system.[5]

2. A THEORETICAL MODEL

The Economic Environment

I will now consider a two-period economic model. The model has several important features. First, agents[6] face an uninsurable idiosyncratic shock to period 2 wages, and they smooth their consumption over the two periods through borrowing and saving. Specifically, agents can save through a riskless saving instrument and can borrow as well. As mentioned earlier, the introduction of this feature is to provide a potential role for bankruptcy laws. Second, agents need to exert effort in order to receive positive labor income for period 2. More importantly, an individual's choice of effort level is unobservable. It is therefore impossible to make contracts contingent on effort levels. I assume that the income an agent receives in period 2 is a product of his or her effort level and the idiosyncratic shock. This product can also be viewed as total output produced by the agent. In that sense, the labor effort of each agent affects the total output of the economy. For the remaining analysis, I will first treat a representative agent in isolation, taking borrowing and lending rates as given. I will then embed that model in a general equilibrium setting in order to endogenize the borrowing and lending rates.

In period 1, an agent enters the economy endowed with some assets a_1 and a signal θ_1, which indicates the quality of the agent's period 2 labor income. An agent draws utility from consumption in both periods, disutility from working in period 2, and a utility penalty from filing for bankruptcy in

[5] Some researchers, however, dispute this view. White (1998), for example, argues that a much higher fraction of U.S. households would benefit financially from bankruptcy than actually file under current bankruptcy provisions.

[6] I use the words "agent," "person," and "individual" interchangeably throughout this article. I assume that there is a continuum of agents in the economy.

period 2.[7] There exists a single creditor in the economy, and agents borrow and deposit with the creditor. Let r^b denote the rate at which the agent borrows, and let r^d denote the deposit rate.

In period 1, an agent chooses consumption debt and/or asset holdings for period 1 and effort for period 2. I denote these decision rules by c_1, d_2/a_2, and l_2 respectively.[8] The agent's labor income shock described by θ_2 is revealed at the beginning of period 2. The probability distribution of θ_2 conditional on the signal θ_1 increases in θ_1 in the sense of first order stochastic dominance, i.e.,

$$F(\theta_2|\theta_1') \leq F(\theta_2|\theta_1) \text{ for } \theta_1' > \theta_1,$$

$F(\theta_2|\theta_1)$ is the cumulative distribution function for period 2's labor income shock θ_2 given period 1's signal θ_1.

After observing his or her labor income shock for period 2, θ_2, the agent works at the effort level decided in period 1. The agent's period 2 income, therefore, consists of labor income and any interest earned on deposits. The agent then decides whether to repay the debt, d_2. If he or she does not repay the debt, the agent will file for either Chapter 7 or Chapter 13 bankruptcy. Under Chapter 7 bankruptcy, the agent will keep his or her assets up to the maximum amount that is exempted under the bankruptcy law, which I denote by E. The agent surrenders remaining assets to the creditor, but keeps all labor income. If Chapter 13 bankruptcy is chosen, the agent may keep all assets; however, a portion ρ ($0 \leq \rho \leq 1$) of period 2 labor income will be used to pay off debts.[9] In the theoretical discussion of this paper, I treat ρ as fixed, independent of an agent's debt and earnings. I use x to denote an agent's decision; x equals 1 if the agent pays off his or her debt, x equals 7 if Chapter 7 bankruptcy is filed, and x equals 13 if Chapter 13 bankruptcy is filed.

An agent suffers a utility loss, S, from a bankruptcy filing of either type. This utility loss represents either the cost of having to borrow at a much higher rate in the future had the model been of infinite horizon (a more realistic case), or the stigma, the level of social disapproval of bankruptcy, or both. In the first case, S would be a function of the interest rate that an agent is charged if he or she has to borrow again after filing for bankruptcy and the agent's average income, income volatility, and desire to smooth consumption. The higher the desire to borrow after bankruptcy, and the higher the postbankruptcy borrowing rate, the higher the cost of filing for bankruptcy will be. Consumption for

[7] To simplify my analysis, without loss of generality, I assume away any labor decision in period 1.

[8] The labor effort here measures how hard a person works in terms of whether he or she tries hard enough to get a well-paying job. Therefore, it is not something a court can mandate.

[9] The implicit assumption is that while private agents cannot enforce a contract contingent on an agent's income, the government can, due to its special enforcement technology (by putting an individual in prison, for example).

period 2 is simply the agent's remaining income after the payment/bankruptcy decision.

To summarize, the time line of the economy is as follows:

time	period 1	period 2
information	asset a_1, period 2 labor income signal θ_1	labor income shock θ_2
decision rules	consumption c_1, asset a_2, borrowing d_2, and effort l_2	payment decision x, and consumption c_2

For the analysis that follows, I assume that an agent's utility function takes the form $Q(c_1)+E[U(c_2)-S\cdot 1(x=7 \text{ or } 13)]-V(l_2)$, where $Q'>0$, $Q''<0$, $Q'''>0$, $U'>0$, $U''<0$, $U'''>0$, $V'>0$, $V''>0$, and $1(.)$ is an indicator function that takes the value 1 if the statements inside the parentheses are true and takes the value 0 otherwise.

An Agent's Problem

To correctly present the choice problem of an agent in our economy, it is helpful to think about it in reverse chronological order.

An Agent's Period 2 Problem: The Bankruptcy Decision

In period 2, agents are described by their assets, debt positions, labor ability, and labor effort decision, namely $(a_2, d_2, \theta_2, l_2)$. They make payment and consumption decisions to maximize period 2 utility as follows,

$$\max_{x, c_2>0} U(c_2) - S\cdot 1(x=7 \text{ or } 13) - V(l_2).$$

An agent faces three choices here: file for bankruptcy under Chapter 7, file for bankruptcy under Chapter 13, or repay the debt. Let W^7, W^{13}, and W^R denote an agent's period 2 utility under the three choices respectively. The last option, of course, requires that the agent's income exceeds the debt payment, i.e., $r^d a_2+\theta_2 l_2 \geq r^b d_2$. We then have the following expressions,

$$W^7 = U(\min\{E, r^d a_2\}+\theta_2 l_2) - S, \quad (1)$$
$$W^{13} = U(r^d a_2 + (1-\rho)\theta_2 l_2) - S, \quad (2)$$
$$W^R = U(r^d a_2+\theta_2 l_2 - r^b d_2), \text{ and } r^d a_2+\theta_2 l_2 \geq r^b d_2. \quad (3)$$

When the agent's income is not enough to repay the debt, i.e., $r^d a_2+\theta_2 l_2 \leq r^b d_2$, he or she has no choice but to file for bankruptcy. I call this type of

Table 1 Parameter Values for Example 1

Parameter	Values
bankruptcy cost (S)	1.50
Chapter 7 asset exemption level (E)	4.00
deposit rate (r^d)	1.00
borrowing rate (r^b)	1.20
debt holding (d_2)	7.00
period 2 labor decision (l_2)	1.00
portion of income garnished under Chapter 13 (ρ)	0.45

bankruptcy filing "involuntary bankruptcy." The bankruptcy choice between the two chapters will depend only on the agent's period 2 consumption since the penalty S applies in either case. When $\min\{E, r^d a_2\} + \theta_2 l_2 \geq r^d a_2 + (1 - \rho)\theta_2 l_2$, the agent will file for bankruptcy under Chapter 7, but otherwise will file under Chapter 13. More specifically, when the agent's assets are below the exemption level, $r^d a_2 \leq E$, he or she will always file for bankruptcy under Chapter 7; when the agent's assets are above the exemption level, $r^d a_2 \geq E$, he or she will file for bankruptcy under Chapter 7 only if $a_2 \leq \frac{1}{r^d}(E + \rho\theta_2 l_2)$.

The case where the agent has enough income to repay the debt is more involved. An agent will file for Chapter 7 if his or her assets are below the threshold discussed above and the consumption benefit of filing for bankruptcy is at least S, i.e., $U(\min\{E, r^d a_2\} + \theta_2 l_2) - U(r^d a_2 + \theta_2 l_2 - r^b d_2) \geq S$. Let Δ denote the utility difference $U(\min\{E, r^d a_2\} + \theta_2 l_2) - U(r^d a_2 + \theta_2 l_2 - r^b d_2) - S$. That is, Δ is the net benefit of filing under Chapter 7 compared to paying off one's debt.

$$\frac{\partial \Delta}{\partial a_2} = \begin{cases} r^d U'(r^d a_2 + \theta_2 l_2) - r^d U'(r^d a_2 + \theta_2 l_2 - r^b d_2), & \text{if } r^d a_2 \leq E; \\ -r^d U'(r^d a_2 + \theta_2 l_2 - r^b d_2), & \text{otherwise.} \end{cases} \tag{4}$$

Given our assumption that U is concave, it is obvious that $\frac{\partial \Delta}{\partial a_2} \leq 0$. Similarly, we can show that the utility difference Δ decreases in the labor income shock θ_2 and increases in the agent's debt holding d_2. It follows that the higher an agent's assets or labor income, and the lower the agent's debt holding, the likelier he or she is to repay the debt.

When the agent's assets are above the threshold described earlier and the cost of filing for Chapter 13 is smaller than the bankruptcy cost S, $U(r^d a_2 + (1 - \rho)\theta_2 l_2) - U(r^d a_2 + \theta_2 l_2 - r^b d_2) \geq S$, he or she will file for Chapter 13. If we let $\tilde{\Delta}$ denote the utility difference $U(r^d a_2 + (1 - \rho)\theta_2 l_2) - U(r^d a_2 + \theta_2 l_2 - r^b d_2) - S$, analysis similar to that above shows that this utility difference exhibits the same properties as the utility difference between filing for Chapter 7 and repaying the debt. The agent is more likely to repay the debt when assets and labor income shock are higher and the debt holding is lower. When the agent has the income to repay the debt, filing is often referred to as "voluntary bankruptcy."

Result 1 *Given the period 2 labor income shock and debt holdings, as assets increase an agent will file for bankruptcy first under Chapter 7, then under Chapter 13, and will repay the debt only if assets exceed some threshold. For a given level of assets and as labor income increases, an agent will file for bankruptcy first under Chapter 13, then under Chapter 7, and will repay the debt only if the labor shock exceeds some threshold.*

The intuition behind Result 1 is straightforward. An agent will lose either wealth or income after filing for bankruptcy. Agents who have sufficient wealth or income or both will, therefore, have no incentive to file for bankruptcy. Chapter 7 exempts all labor income but only part of assets. As a result, agents with good income but few assets will benefit from Chapter 7. Chapter 13 protects agents' assets at the cost of their labor income. Consequently, Chapter 13 benefits agents with large assets but low labor income. The following numerical example illustrates Result 1.[10]

Example 1 *Utility is logarithmic, i.e.,* $U(c_2) = \log(c_2)$. *Parameters are chosen as in Table 1.*

Figure 2 depicts an agent's bankruptcy decision in relation to his or her period 2 assets and labor income shock. Line A describes asset-labor income shock pairs for which the agent is indifferent between filing for Chapter 7 or Chapter 13, $r^d a_2 - \rho\theta_2 l_2 - E = 0$. The agent will file for Chapter 7 if this expression is negative and will file for Chapter 13 otherwise. Line B consists of assets and labor income shocks that are just enough to repay the debt, $r^d a_2 + \theta_2 l_2 - r^b d_2 = 0$. Agents with more income than debt lie above line B. Line C is the indifference curve of the agent between filing for bankruptcy and repaying the debt, i.e., $\max(U(\min[E, r^d a_2] + \theta_2 l_2), U(r^d a_2 + (1-\rho)\theta_2 l_2)) - U(r^d a_2 + \theta_2 l_2 - r^b d_2) - S = 0$. Agents preferring to repay the debt lie above this indifference curve.

As seen in Figure 2, an agent will repay the debt if both assets and the labor income shock exceed the threshold set by line C. Underneath the repayment line, the fewer assets agents have, the higher the labor income shock they receive, and the more likely they are to file for Chapter 7. Conversely, the more assets they have, and the lower their period 2 labor income, the more likely they are to file for Chapter 13. More specifically, as illustrated in Figure 2, for the given value of a_2, if $\theta_2 < \overline{\theta}$, the agent files for bankruptcy under Chapter 13, if $\overline{\theta} \leq \theta_2 < \overline{\overline{\theta}}$, the agent files for bankruptcy under Chapter 7,

[10] Note that the examples presented in this article consist of magnitudes that are by no means calibrated. They are included to illustrate the discussion.

Figure 2 Agent's Period 2 Bankruptcy Decision—An Example

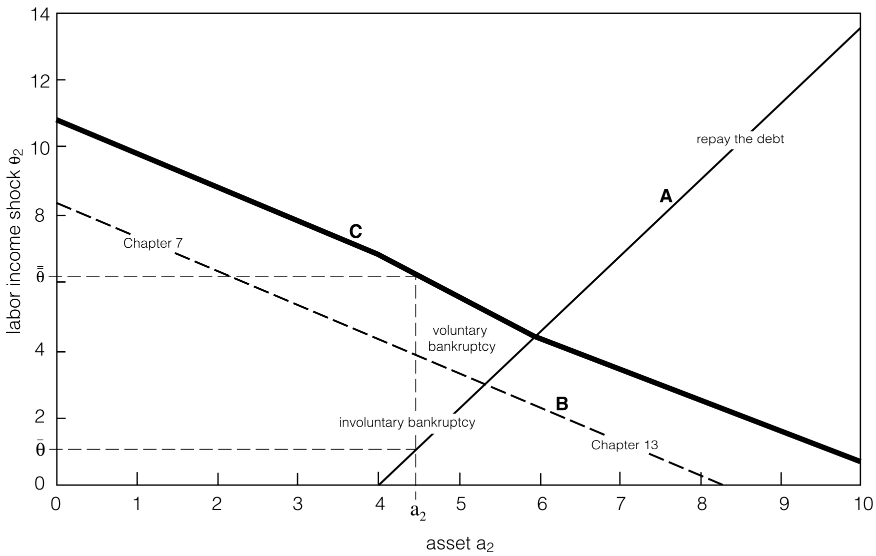

and if $\theta_2 \geq \overline{\overline{\theta}}$, the agent repays the debt. An increase in debt d_2 will move line B and line C upward, reducing the region of debt repayment.

An Agent's Period 1 Problem: Portfolio and Labor Effort Decisions

Agents in period 1 make portfolio, consumption, and labor decisions to maximize their lifetime utility as follows,

$$\max_{a_2, d_2, c_1, l_2} Q(c_1) + E_{\theta_2|\theta_1} \max(W^7, W^{13}, W^R) - V(l_2)$$

s.t.

$$c_1 + a_2 = a_1 + d_2, \quad (5)$$

$$c_1, a_2, d_2, l_2 \geq 0. \quad (6)$$

I assume that the labor income shock θ takes value in $[\theta_{\min}, \theta_{\max}]$. Let $\overline{\theta}$ and $\overline{\overline{\theta}}$ denote the two thresholds so that given a_2, d_2, and l_2, when $\theta_2 \leq \overline{\theta}$, an agent will file for bankruptcy under Chapter 13; when $\overline{\theta} \leq \theta_2 \leq \overline{\overline{\theta}}$, an agent will file for bankruptcy under Chapter 7; when $\theta_2 \geq \overline{\overline{\theta}}$, an agent will repay the debt (Result 1). The agent's period 1 utility can then be rewritten as $Q(a_1 + d_2 - a_2) + (\int_{\theta_{\min}}^{\overline{\theta}} W^{13} + \int_{\overline{\theta}}^{\overline{\overline{\theta}}} W^7 + \int_{\overline{\overline{\theta}}}^{\theta_{\max}} W^R) dF(\theta_2|\theta_1) - V(l_2)$.

This utility function is no longer concave; it consists of kinks where agents are indifferent between two of the three choices. In the analysis that follows, I focus on cases in which the equilibrium solution does not fall on any of these kinks, in which case Euler equations are both necessary and sufficient. The qualitative results thus obtained can be generalized to other cases where equilibrium occurs at a kink. I omit those analyses in order to save space.

Assuming interior solutions (all the choice variables take positive values), Euler equations for this maximization problem are

$$\begin{aligned} Q'(a_1 + d_2 - a_2) \\ &= r^d \int_{\theta_{\min}}^{\overline{\theta}} U'(r^d a_2 + (1-\rho)\theta_2 l_2) dF(\theta_2|\theta_1) + r^d \int_{\overline{\theta}}^{\overline{\overline{\theta}}} \\ &\quad U'(r^d a_2 + \theta_2 l_2) dF(\theta_2|\theta_1) 1(r^d a_2 \leq E) \\ &\quad + r^d \int_{\overline{\overline{\theta}}}^{\theta_{\max}} U'(r^d a_2 + \theta_2 l_2 - r^b d_2) dF(\theta_2|\theta_1), \quad (7) \end{aligned}$$

$$Q'(a_1 + d_2 - a_2) = r^b \int_{\overline{\overline{\theta}}}^{\theta_{\max}} U'(r^d a_2 + \theta_2 l_2 - r^b d_2) dF(\theta_2|\theta_1), \quad (8)$$

$$\begin{aligned} V'(l_2) &= \int_{\theta_{\min}}^{\overline{\theta}} (1-\rho)\theta_2 U'(r^d a_2 + (1-\rho)\theta_2 l_2) dF(\theta_2|\theta_1) \\ &\quad + \int_{\overline{\theta}}^{\overline{\overline{\theta}}} \theta_2 U'(\min\{E, r^d a_2\} + \theta_2 l_2) dF(\theta_2|\theta_1) \\ &\quad + \int_{\overline{\overline{\theta}}}^{\theta_{\max}} \theta_2 U'(r^d a_2 + \theta_2 l_2 - r^b d_2) dF(\theta_2|\theta_1). \quad (9) \end{aligned}$$

When there are no bankruptcy provisions, Euler equations (7) and (8) for a_2 and d_2 respectively become

$$Q'(a_1 + d_2 - a_2) = r^d \int_{\theta_{\min}}^{\theta_{\max}} U'(r^d a_2 + \theta_2 l_2 - r^b d_2) dF(\theta_2|\theta_1), \quad (10)$$

$$Q'(a_1 + d_2 - a_2) = r^b \int_{\theta_{\min}}^{\theta_{\max}} U'(r^d a_2 + \theta_2 l_2 - r^b d_2) dF(\theta_2|\theta_1). \quad (11)$$

Obviously, these two equations cannot hold simultaneously when the deposit rate differs from the borrowing rate.[11] This implies that an agent will not hold both assets and debt in period 2 in the absence of bankruptcy provisions. When there are either Chapter 7 or 13 bankruptcy provisions, however, the agent may hold both assets and debt simultaneously. The intuition is that although debt

[11] When the borrowing rate is the same as the deposit rate, the equilibrium condition pins down the agent's net asset position, $d_2 - a_2$.

requires paying an interest rate premium, the prospect of not having to pay the debt completely or part of it in the event of bankruptcy lowers the effective rate an agent pays on his or her debts. In other words, bankruptcy provisions encourage agents to borrow to save. Lehnert and Maki (1999) also obtain such a result in their paper through numerical simulation. They further argue that this theoretical result is corroborated by household behavior as documented in the Consumer Expenditure Survey.[12]

Result 2 *With either Chapter 7 or Chapter 13 bankruptcy provisions, an agent may simultaneously hold low-return assets and high-interest debt; with no bankruptcy provisions, however, an agent will hold only one of the two instruments.*

With regard to labor effort, equation (9) suggests that an agent's period 2 work effort is a decreasing function of the probability of filing for Chapter 13. This is a direct result of a Chapter 13 provision: Those who file under 13 lose part of their income to their creditors.

Result 3 *The work effort that an agent exerts in period 2 decreases as the probability of bankruptcy filing under Chapter 13 increases.*

The effects of changes of a_1 and θ_1 on agents' portfolio and bankruptcy decisions are more complicated. When a_1 increases, the agent does not need to borrow as much from period 2 to increase consumption in period 1 since now he or she can consume more period 1 assets. The reduced debt relative to asset holdings will in turn increase the debt repayment region. Hence, agents will take advantage of bankruptcy provisions less often.

An improved period 2 labor income prospect (larger θ_1)implies that the agent would like to borrow more to increase consumption in period 1. A better labor income prospect in the sense of first-order stochastic dominance also means that the agent repays loans more often given the amount borrowed. In other words, the agent now enjoys the bankruptcy provision less often, which will reduce borrowing in period 1. In this case, whether the agent will borrow more when θ_1 is higher depends on the net effect of the two forces. Similarly, bankruptcy and labor effort decisions will be affected by opposing forces.

Example 2 *The agent's utility takes the following functional form:* $log(c_1) + E_{\theta_2|\theta_1}[c_2 - S \cdot 1(x = 7 \text{ or } 13)] - 1.25 * l_2^2$. $F(\theta_2) = (\theta_2)^{\theta_1}$, *where* θ_1, $\theta_2 \in [0, 1]$. *Other parameter values are summarized in Table 2.*

[12] The Consumer Expenditure Survey (CE) is an annual survey of about 5,500 households conducted by the Bureau of Labor Statistics. Participants are surveyed four times over the course of the year and are asked about their expenditures, assets, liabilities, and incomes.

Table 2 Parameter Values for Example 2

Parameter	Values
bankruptcy cost (S)	2.06
Chapter 7 asset exemption level (E)	0.43
deposit rate (r^d)	1.00
borrowing rate (r^b)	1.06
portion of income garnished under Chapter 13 (ρ)	0.65

Table 3 Simulation Results of Example 2

Parameters		Simulation Results					
a_1	θ_1	a_2	l_2	d_2	Chapter 13 (percent)	Chapter 7 (percent)	Repayment (percent)
0.114	1.00	0.4985	0.5963	0.3937	0.2677	0.0715	0.6608
0.120	1.00	0.4993	0.5779	0.3734	0.2855	0.0177	0.6968
0.114	0.99	0.4976	0.5948	0.3926	0.2684	0.0737	0.6579

Table 3 summarizes results of the three experiments. According to these experiments, agents with higher period 1 assets save more, borrow less, and repay debt in period 2 more often. In the event of bankruptcy, they file for Chapter 13 more often in order to protect their assets. As a result, they put in less labor effort. Agents with inferior period 2 labor income prospects save less and borrow even less, and repay debt in period 2 less often. In the event that they become bankrupt, they file for Chapter 7 more often because they hold fewer assets in period 2. Consequently, they exert more labor effort.

Financial Intermediation and Credit Market Equilibrium

Thus far analysis has focused on a single agent, taking borrowing and lending rates as given. I now introduce financial intermediation and define a credit market equilibrium. There is a single creditor in this economy. Without loss of generality, I call this creditor a bank. Agents deposit with the bank at rate r^d and borrow from the bank at rate r^b. Note that borrowers are charged the same rate irrespective of their assets and debt position. The implicit assumption is that the lender does not observe borrowers' assets and cannot set prices according to loan size. This is a simplifying assumption. Alternatively, one could make the lending rate a function of the borrower's assets or loan size, or even a signal regarding his or her period two labor income shock.

The bank makes zero profit under the assumption of competitive financial intermediation. Let $G(a_1, \theta_1)$ denote the ex ante distribution of agents over period 1 wealth and period 2 labor income shock signals, and let there be a measure 1 of agents in the economy. The zero profit restriction translates into

$$\int_{a_1,\theta_1} [\int_{\theta_{\min}}^{\overline{\theta}} \rho\theta_2 l_2 F(\theta_2|\theta_1) + \int_{\overline{\theta}}^{\overline{\overline{\theta}}} \max[0, r^d a_2 - E] dF(\theta_2|\theta_1) + \int_{\overline{\overline{\theta}}}^{\theta_{\max}} r^b d_2 dF(\theta_2|\theta_1)] dG(a_1, \theta_1) = r^d \int_{a_1,\theta_1} \int_{\theta_2} a_2 dF(\theta_2|\theta_1), \quad (12)$$

where the first term on the left-hand side of equation (12) is the wage garnished from agents who file for bankruptcy under Chapter 13, the second term is the assets obtained from agents who file for bankruptcy under Chapter 7, and the third term is the loan repayment from those who do not default. The right-hand side of the equation is the cost of deposits. Obviously, the greater the number of agents filing for bankruptcy, the higher the borrowing rate the bank will charge.

A general equilibrium of our economy consists of a pair of interest rates (r^d, r^b) and a set of decisions (a_2, d_2, l_2, x) such that given the interest rates, (1) agents make decisions that maximize their expected utility, (2) the bank breaks even, and (3) capital markets clear as follows,

$$\int_{a_1,\theta_1} a_2(a_1, \theta_1) dG(a_1, \theta_1) = \int_{a_1,\theta_1} d_2(a_1, \theta_1) dG(a_1, \theta_1). \quad (13)$$

The left-hand side of equation (13) is the deposits the bank collects from the agents, and the right-hand side is the loans that the bank makes.

3. POLICY EXPERIMENTS

Most of the current proposals on bankruptcy reform center on three policy instruments: the bankruptcy exemption, the percentage of wage income that can be garnished by creditors in the event of bankruptcy, and the financial profiles of agents who must file for bankruptcy under Chapter 13. The National Bankruptcy Review Commission, for example, recommends large increases in bankruptcy exemptions. The Gekas bill goes in the opposite direction by forcing debtors in bankruptcy whose income is above the median to use 100 percent of their postbankruptcy earnings above a predetermined level to repay debt. In this section, I analyze the implications of each of the proposed bankruptcy law changes. In particular, I ask how these changes affect agents' repayment and bankruptcy chapter choices, their ex ante portfolio decisions, and their labor decisions and welfare.

Figure 3 Changes of Bankruptcy Regions due to an Increase in Exemption

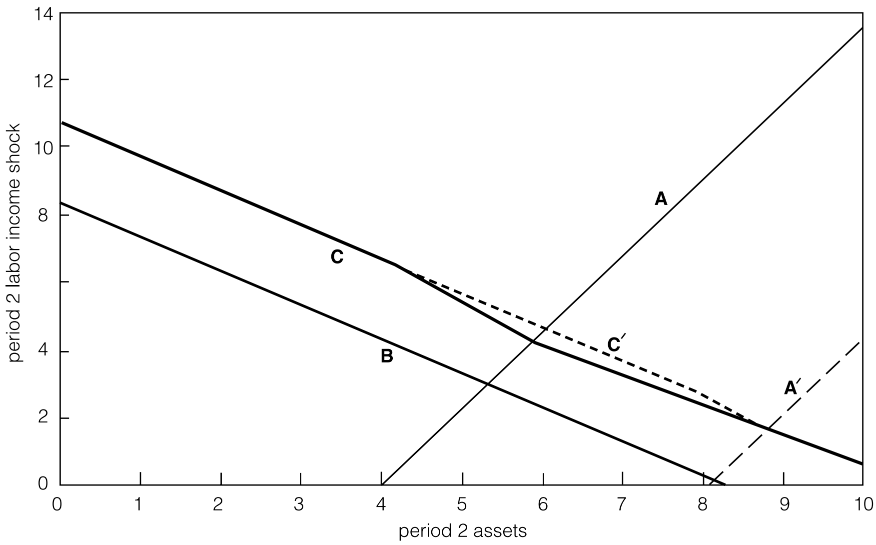

Bankruptcy Exemption

I start with the period 2 decision problem. Holding interest rates fixed, suppose the bankruptcy exemption in the model economy is increased from E_1 to E_2. There are two immediate effects. First, Chapter 7 bankruptcy is now more attractive than Chapter 13, though only for agents with medium assets. Agents whose assets are below $\rho\theta_2 l_2 + E_1$ will file for Chapter 7 under the old exemption level, and agents whose assets are above $\rho\theta_2 l_2 + E_2$ will not find the increase sufficient for them to change their bankruptcy chapters. Only those with assets between these two cutoff levels and whose labor income shocks are not high enough to make debt repayment more attractive will switch from Chapter 13 to Chapter 7.

Second, agents with assets between E_1 and $\rho\theta_2 l_2 + E_2$ may also benefit from the increase in exemption by filing for Chapter 7 rather than paying off their debt. To demonstrate, recall that agents are indifferent between filing for Chapter 7 and repaying debt when $U(\min\{r^d a_2, E\} + \theta_2 l_2\}) - S = U(r^d a_2 + \theta_2 l_2 - r^b d_2)$. The exemption level affects this equality only if agents' assets are above this level. Moreover, if agents' assets are too high, either Chapter 13 or repayment will be more attractive. Figure 3 depicts changes in bankruptcy regions when E is increased from 4 to 8 using Example 1. All the lines carry the same interpretation as those in Example 1. The dotted lines correspond to those under the new exemption level.

In period 1, because of the increased benefits of filing for bankruptcy under Chapter 7, agents will save less and/or borrow more so that in period 2 they have a higher chance of filing for bankruptcy under Chapter 7. This can be seen in equation (8). The threshold $\overline{\overline{\theta}}$ increases with the increase in E; for the equation to continue to hold, the agent needs to hold more debt relative to assets. The general equilibrium effect of the increase in debt and decrease in saving is that the deposit rate increases. In addition, the increase in default (particularly Chapter 7 default) increases the borrowing rate. This increase in borrowing rate will dampen the decrease in saving since agents will have to rely more on saving to smooth their consumption. An increase in the deposit rate clearly benefits those who save, and an increase in the borrowing rate hurts those who borrow. The implication is that increasing the exemption will benefit the rich and hurt the poor, especially those with good labor prospects.

Result 4 *Given assets and debt, agents with medium assets and low labor income shocks will benefit from an increase in the exemption by switching from Chapter 13 to Chapter 7. Agents with medium assets and medium income shocks will benefit from the increase in the exemption if they switch from repaying their debt to filing for Chapter 7. Ex ante, agents will exert more labor effort. General equilibrium price changes may dampen these results.*

Wage Garnishment

I next discuss the effects of changes in the percentage of wage income that can be garnished by creditors in the event of Chapter 13 bankruptcy. This wage garnishment is captured by the parameter ρ in the model.

Changes in ρ affect two types of marginal borrowers: those who are at the margin of filing for Chapter 7 bankruptcy as opposed to Chapter 13 and those who are at the margin of paying off their debts as opposed to filing for Chapter 13. A reduction in ρ will induce more agents in the first group to file for Chapter 13. The second group of borrowers repays debts because the utility from doing so is higher than filing for Chapter 13, i.e., $U(r^d a_2 + \theta_2 l_2 - r^b d_2) \geq U(r^d a_2 + (1-\rho)\theta_2 l_2) - S$. A smaller ρ will increase the value of filing for Chapter 13 and, hence, will make agents less likely to pay off debts. Again, the benefits of a smaller ρ accrue more to agents with better labor income shocks. Furthermore, since only agents with relatively greater assets will consider filing for Chapter 13, a reduction in wage garnishing will most benefit rich people with relatively high labor income shock. Figure 4 plots changes in bankruptcy regions when I reduce ρ from 0.45 to 0.25 in Example 1.

In period 1, holding interest rates constant, marginal agents in the first group will increase their savings, while those in the second group will increase

Figure 4 Changes of Bankruptcy Regions due to a Reduction in Wage Garnishment

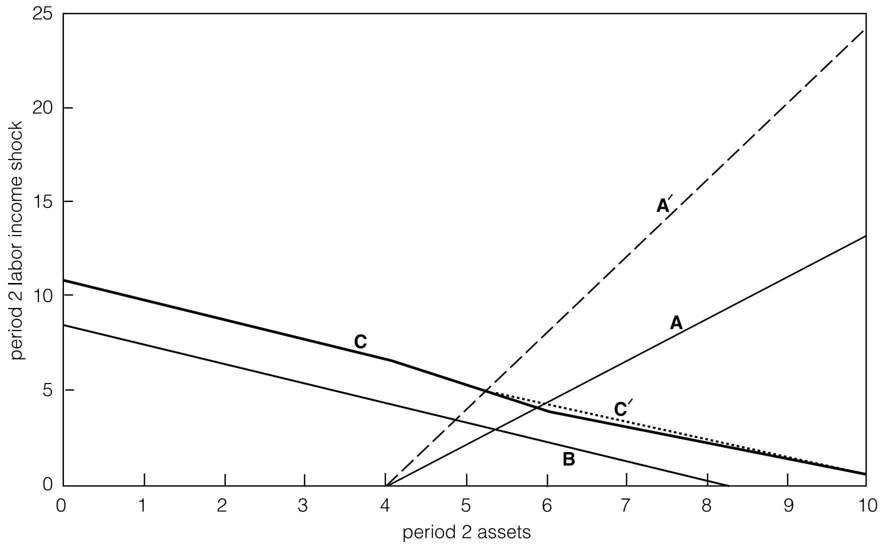

their borrowing to take advantage of the now more beneficial Chapter 13 bankruptcy provision. Agents will also increase their labor effort since now a smaller fraction of their labor income will be lost to creditors. If there is excess aggregate saving in the economy, the deposit rate will drop so that the aggregate saving equals aggregate borrowing in the new equilibrium. If there is excess aggregate borrowing, the deposit will rise to reach a new equilibrium. The general equilibrium effect of a reduction of ρ on the borrowing rate is twofold. On one hand, more agents will default, and the increased default risk will cause the borrowing rate to increase. On the other hand, with more Chapter 13 and fewer Chapter 7 filings the lender will be able to collect more, which will tend to reduce the borrowing rate.

Result 5 *Given their assets and debt positions, agents with good labor income shocks and high assets will benefit from a reduction in wage garnishment under Chapter 13. The benefits come from either switching from Chapter 7 to Chapter 13, or switching from repaying debts to filing for Chapter 13. Ex ante, agents respond to the changed incentives by either saving more or borrowing more and by exerting higher labor effort.*

Figure 5 Changes in Bankruptcy Regions due to an Income Mandate

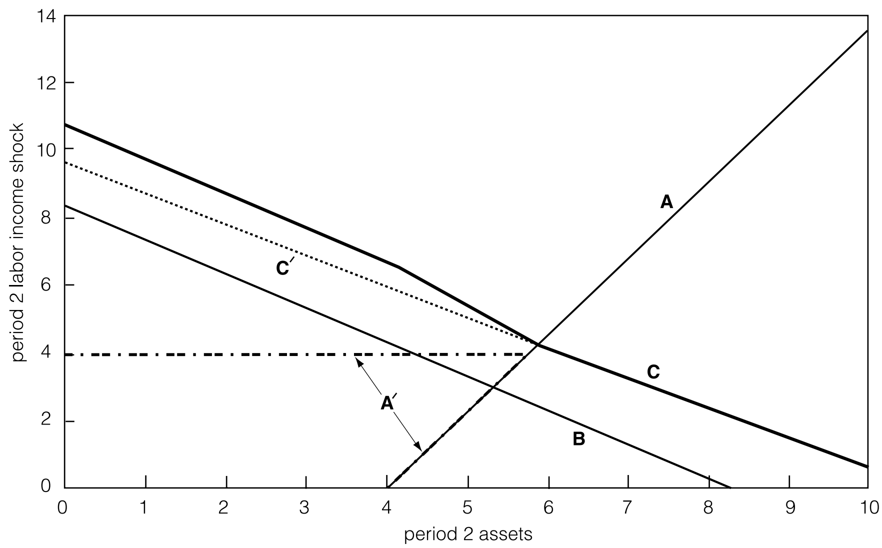

Income Mandate

Under an income mandate provision, if your income is above a certain cutoff level, then you cannot file for bankruptcy under Chapter 7. The cutoff is often determined according to the population income distribution. Let us suppose that all the other provisions remain unchanged, including the exemption level and the percentage of labor income that will be collected by creditors under Chapter 13.

The implementation of such an income mandate corresponds to setting a labor income shock threshold in our economy. Agents with a period 2 income shock above this level can file for bankruptcy only under Chapter 13. The effects of the income mandate are straightforward. It affects only those with the lowest amount of assets and medium labor income shocks. Those with more assets will file under Chapter 13 even without the mandate, and those with good labor income shocks will always repay their debts. Figure 5 depicts the imposition of a labor income shock cutoff equal to 4 in Example 1. The repayment line for low assets agent shifts down because of the income mandate. The reason for the downward shift is that the benefits from filing from bankruptcy (Chapter 13 under the new rule) are fewer than before (Chapter 7). As a result, agents are more likely to pay off their debt.

In period 1, agents realize they will have to file for Chapter 13 more often than in period 2; they will therefore put in less labor effort and will borrow less as a result. The reduced demand will drive down the deposit rate. The borrowing rate will also come down since the borrowing premium will be lower. More agents will file for Chapter 13, and more agents will pay off their debts.

Result 6 *Given assets and debt positions, an income mandate for Chapter 13 filing hurts agents with few assets and medium labor income shocks. Ex ante in period 1, agents will not work as hard and will also reduce their borrowing. In general equilibrium, both deposit rates and borrowing rates will come down.*

To summarize my policy discussion, a reduction in assets exemption level under Chapter 7 benefits agents with medium assets and medium labor income shock. Ex ante, these agents save less and borrow more; however, they also work harder. The general equilibrium price effects are likely to dampen these results. A reduction in the percentage of wages that can be garnished under Chapter 13 benefits agents with high assets and medium labor income shock. These agents switch to Chapter 13 bankruptcy from either Chapter 7 bankruptcy or repayment. Ex ante, these agents work less hard. Whether the interest rate increases or not depends on the net changes of the total increase in saving and the total increase in borrowing. Finally, an income mandate hurts agents with medium labor income shocks and few assets. Ex ante, these agents will not work as hard and will also reduce their borrowing.

Two points need to be made before I conclude. First, a two-period model has been chosen to keep the analysis relatively simple and tractable. A proper analysis of efficiency losses and distributional concerns would obviously require a fully dynamic model of bankruptcy choice. The extension of the current model to an infinite horizon makes the bankruptcy cost at least partly endogenous (see “The Economic Environment,” above, for a discussion of the bankruptcy cost); the bankruptcy decisions, therefore, become truly dynamic. I speculate that most of the qualitative results of this article will survive this extension.

A second point to bear in mind is that I assume that the portion of income garnished under Chapter 13 is constant, while in practice it often depends on the debtor’s income as well as his or her debts. Relaxation of this assumption clearly would make Chapter 13 more attractive to those with fewer debts and higher labor income shocks.

4. CONCLUSION

The recent surge in U.S. consumer bankruptcy filings has prompted many reform proposals. At the center of these proposals is the issue of consumer bankruptcy choices, specifically whether agents should be encouraged to choose Chapter 13 over Chapter 7.

I have used a simple theoretical model with uninsurable labor income to investigate two sets of issues: What are the financial profiles of those who repay their debts and those who file for bankruptcy under one chapter or the other? and, What are the policy implications of the current reform proposals, including efficiency and distributional concerns? With respect to the first question, I confirm the general view that agents with relatively greater assets prefer Chapter 13, while those with relatively high labor income prefer Chapter 7. I also find that bankruptcy provisions tend to encourage "borrow to save" behavior; that is, some agents will hold low return assets and high risk-premium debt simultaneously. Furthermore, agents with a higher probability of filing for Chapter 13 will exert less labor effort.

I conducted three policy experiments: changes in the bankruptcy exemption under Chapter 7, changes in the percentage of labor income that can be garnished under Chapter 13, and the implementation of a labor income mandate for Chapter 13 filings. The experiments show that a reduction in the asset exemption level, an increase in the percentage of labor income that can be obtained by creditors, and the implementation of a labor income mandate will all encourage Chapter 13 bankruptcy filings over Chapter 7 and the repayment of debt. The efficiency cost of these changes is that agents will exert less effort, causing total output (labor income in my model) to drop. In terms of income and wealth distribution effects, three conclusions emerge. Changes in the asset exemption level affect those with medium assets and medium labor income. Changes in wage garnishment affect those with high assets and medium labor income shock. The implementation of an income mandate affects those with few assets and medium labor income.

REFERENCES

Adler, Barry, Ben Polak, and Alan Schwartz. "Regulating Consumer Bankruptcy: A Theoretical Inquiry," *The Journal of Legal Studies,* vol. 24 (June 2000), pp. 585–614.

Aiyagari, S. Rao. "Macroeconomics with Frictions," Federal Reserve Bank of Minneapolis *Quarterly Review,* vol. 21 (Summer 1997), pp. 28–36.

Athreya, Kartik. "Welfare Implications of the Bankruptcy Reform Act of 1998." Manuscript. Federal Reserve Bank of Richmond, 2000.

Brunstad, G. Eric, Jr. "Bankruptcy and the Problems of Economic Futility: A Theory on the Unique Role of Bankruptcy Law," *The Business Lawyer,* vol. 55 (February 2000), pp. 499–591.

Domowitz, Ian, and Robert L. Sartain. “Determinants of the Consumer Bankruptcy Decision,” *The Journal of Finance,* vol. 54 (February 1999), pp. 403–20.

Lehnert, A., and D. M. Maki. “The Great American Debtor: A Model of Household Consumption, Portfolio Choice and Bankruptcy.” Manuscript. Board of Governors of the Federal Reserve System, 1999.

Nelson, Jon P. “Consumer Bankruptcy and Chapter Choice: State Panel Evidence,” *Contemporary Economic Policy,* vol. 17 (1999), pp. 552–66.

Sullivan, Teresa A., Elizabeth Warren, and Lawrence Westbrook. *As We Forgive Our Debtors: Bankruptcy and Consumer Credit in America.* New York: Oxford University Press, 1989.

Wang, Hung-Jen, and Michelle J. White. “An Optimal Personal Bankruptcy Procedure and Proposed Reforms,” *Journal of Legal Studies*, vol. 29 (January 2000), pp. 255–86.

White, Michelle J. “Why Don’t More Households File for Bankruptcy?” Manuscript. University of Michigan, 1998.

Optimal Taxation in Infinitely-Lived Agent and Overlapping Generations Models: A Review

Andrés Erosa and Martin Gervais

The literature concerned with dynamic fiscal policy has evolved in two main directions over the last 20 years or so. On the one hand, there is a large literature on optimal taxation. In the context of a standard neoclassical growth model with infinitely-lived individuals, Chamley (1986) and Judd (1985) establish that an optimal income-tax policy entails taxing capital at confiscatory rates in the short run and setting capital income taxes equal to zero in the long run. Only labor income should be taxed in the long run. On the other hand, most applied work concerned with the dynamic impact of fiscal policy uses the life-cycle framework (Auerbach, Kotlikoff, and Skinner [1983], Auerbach and Kotlikoff [1987], and many others, surveyed in Kotlikoff [1998]). Unfortunately, the prescriptions that emanate from the former framework do not necessarily generalize to the latter.

This article reviews the basic results obtained under both the infinitely-lived agent model and the life-cycle model. The first section, which presents a nontechnical introduction to the optimal taxation literature, discusses the optimal taxation problem and the intuition behind the results obtained from both types of models. Section 2 more formally presents the results for the infinitely-lived agent model and Section 3 presents results for the life-cycle economy.

The review of the literature presented here complements that of Chari and Kehoe (1999) and Atkeson, Chari, and Kehoe (1999). The main focus of their

---

Andrés Erosa: Universitat Autònoma de Barcelona. Martin Gervais: Research Department, Federal Reserve Bank of Richmond. We would like to thank Kartik Athreya, Huberto Ennis, and Andreas Hornstein for helpful comments and suggestions. The views expressed in this article are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Richmond or the Federal Reserve System.

papers is on the infinitely-lived agent model. As the title of Atkeson, Chari, and Kehoe's article indicates, they emphasize that taxing capital is a bad idea. Their conclusion, especially for life-cycle economies, is based on very special cases. Our review suggests instead that there is no real consensus regarding the optimal tax on capital income. Rather, we demonstrate several empirically relevant settings in which optimal capital taxes are non-zero, both in the short run and in the long run.

1. A REVIEW OF OPTIMAL TAXATION

Statement of the Ramsey Problem

The problem of finding the optimal manner in which to finance a given stream of expenditures has a long tradition in public finance. The statement of the optimal taxation problem given here follows Ramsey's 1927 seminal paper, which formally recognized that individuals and firms react to changes in fiscal policy. When considering alternative fiscal policies, the government has to take into account that individuals and firms will behave in their own best interest, taking as given whichever fiscal policy the government has chosen. Each fiscal policy implies a feasible allocation of goods and factor services, along with prices, that fully reflects the optimal reaction of individuals and firms; that is, each fiscal policy implies a competitive equilibrium allocation.[1] Given a welfare criterion, which the government uses to evaluate different allocations, the Ramsey problem for the government is to pick the fiscal policy[2] that generates the competitive equilibrium allocation giving the highest value of the welfare criterion.

An equivalent way of formulating the Ramsey problem is to let the government pick an *allocation* directly—rather than a set of tax rates—but to restrict the set of allocations from which the government can choose. This set of allocations can be constructed as follows. Pick an arbitrary fiscal policy. Under this fiscal policy, the optimal behavior of individuals and firms generates a competitive equilibrium allocation. This allocation is one element in the set of allocations from which the government can choose. We refer to such an allocation as an *implementable allocation*: to implement this particular allocation as a competitive equilibrium, the government simply needs to choose the fiscal policy that generated it. By repeating this process for all possible fiscal policies, we can construct the set of all possible allocations that the government can implement. The resulting Ramsey problem then consists of picking, among all implementable allocations, the one that maximizes a

[1] Note, however, that many fiscal policies may generate the same competitive equilibrium. This is the case, in particular, when the government has more tax instruments than are needed to generate a particular allocation.

[2] Or one of them if there are many.

welfare criterion. In many situations, this alternative way of stating the Ramsey problem, referred to as the *primal approach* in the literature, turns out to be much more convenient than the dual problem of choosing tax rates.[3]

The Ramsey problem poses some additional challenges when its focus is on dynamic fiscal policies. Implicit in the statement of the problem is the following sequence of actions by the government, individuals, and firms. First, at initial date zero, the government announces a time path for the fiscal policy instruments. Taking this path of tax rates as given, individuals and firms then choose their paths of consumption, savings, leisure, and inputs in order to maximize utility and profits. When we get to period one, however, it is quite possible that the government will choose to revise its path of tax rates if given the opportunity to do so. Furthermore, individuals and firms will behave differently in period zero if they know that the government has an incentive to modify the path of tax rates in the future. This problem, known as the *time inconsistency* of policies, is particularly severe in infinitely-lived agent models, but it is also present in life-cycle economies.[4]

Infinitely-Lived Agent Models

Two central prescriptions emerge from the solution to the Ramsey problem in representative, infinitely-lived agent models. The first is that capital income should not be taxed in the long run. This result makes sense if we understand that a positive tax on the return from today's savings effectively makes consumption next period more expensive than consumption in the current period. In infinitely-lived agent models, then, a positive (and constant) tax on capital income in the steady state implies that the implicit tax rate of consumption in future periods increases without bound. On the other hand, the relevant elasticity of demand for consumption at all dates is constant.[5] Taxing dated consumption at different rates thus violates the general public finance principle that tax rates should be inversely proportional to demand elasticities. It follows that the capital income tax should be zero.

The second important aspect of optimal taxation in infinitely-lived agent models stems directly from the time inconsistency of optimal policies discussed above. Prior to date zero, which is the date when the government chooses the path of fiscal instruments, individuals presumably operate under

---

[3] On the primal approach, see Atkinson and Stiglitz (1980) and Lucas and Stokey (1983).

[4] The classic reference on time inconsistency of optimal plans is Kydland and Prescott (1977).

[5] In general equilibrium, the relevant elasticity does not have a readily recognizable representation. If the individuals' utility function is additively separable over time, then this elasticity depends on the intertemporal elasticity of substitution as well as some cross elasticity between consumption and leisure. In any case, the fact that both consumption and leisure are constant in steady state is sufficient to make this elasticity constant.

the assumption that the old fiscal policy will last forever. As far as the government is concerned, individuals' previous actions translate into the economy's initial conditions at date zero, as summarized by individuals' initial asset holdings (capital and government debt). Since these assets were accumulated in the past, at date zero individuals will supply their capital to firms regardless of the fiscal policy: this factor is inelastically supplied. As a result, taxing the return on these assets perfectly imitates a (nondistortionary) lump-sum tax. Without any restrictions on the size of that tax, it is efficient for the government to tax initial asset holdings at confiscatory rates.[6] In this way, the government can finance its stream of expenditures through the return on the levied capital and avoid distortionary taxes in the future. Indeed, if the return on that capital is sufficiently large to finance all future government expenditures, a Pareto optimal allocation is achieved because there is no need to ever use distortionary taxes.

The time inconsistency problem exists because when the government switches to a new fiscal policy, individuals are "surprised" and cannot react to the government's action. As such, the time inconsistency problem and the optimality of the front-loading policy are directly related. The former is not, however, confined to the initial switch in fiscal policy. As long as the path of taxes initially announced by the government involves distortionary taxes at some future date, the government has an incentive to redesign its original plan in order to take advantage of whatever lump-sum tax (capital levy) becomes available in the future. Economists have dealt with the general time inconsistency problem by assuming that the government has access to some commitment device, or a *commitment technology*, that allows the government to commit itself once and for all to the sequence of tax rates announced at date zero. In other words, the commitment technology prevents the government from revising the path of fiscal instruments over time. The optimality of confiscating initial holdings of financial assets, however, still remains an integral part of the solution to the Ramsey problem.

To avoid this arguably trivial solution to the optimal taxation problem, it is usually assumed that the government faces exogenous bounds on the size of feasible tax rates. For example, Chamley (1986) assumes that tax rates have to lie between zero and one. Chamley shows that the optimal policy under some assumptions, with respect to preferences, entails taxing capital income at the highest possible rate for a finite amount of time—while the lump-sum aspect of this tax outweighs its distortionary cost—and setting the capital income tax equal to zero thereafter.[7] Although this exogenous upper bound assumption

[6] Similarly, it is efficient for the government to renege on all government debt outstanding at date zero.

[7] In discrete time models, there is a period of transition during which the tax on capital income is strictly between zero and one. There is no such transition period in Chamley's original result since it was derived in a continuous time model.

may seem realistic, it is also completely arbitrary and has a pronounced impact on the solution to the optimal taxation problem: The higher the bound is, the more capital the government accumulates during the first few periods after the switch in fiscal policy and the lower the tax rate on labor income individuals have to face in the future, including the steady state. More generally, the size of the bound determines the magnitude of the welfare gains achieved by switching to the taxes prescribed by the Ramsey problem.

Life-Cycle Economies

The result developed above holds because the elasticity of consumption expenditures exhibits steady state constancy. In turn, this elasticity is constant precisely because consumption and leisure are themselves constant in steady state, which need not be the case in life-cycle economies.[8] In fact, one of the main reasons why economists use the life-cycle model is precisely because observed lifetime consumption and leisure profiles are not flat. Because the behavior of individuals has this life-cycle pattern, there is no reason to expect the relevant consumption elasticities to be constant.

It follows from this reasoning that consumption at different ages should be taxed at different rates. Alternatively, capital income should be taxed or subsidized at rates that depend on the age of the individual supplying the capital.[9] Through a similar argument, optimal labor income tax rates also vary with the age of the individual supplying labor. Although these arguments indicate that the relative capital and labor income tax rates should vary over the lifetime of individuals, they leave open the question of how to determine the level at which these tax rates should be set. We will return to this question below.

The choice of a welfare function to evaluate different implementable allocations is not as straightforward in life-cycle economies as it is in infinitely-lived agent models. The fact that standard infinitely-lived agent models are populated by a single representative individual dictates that the benevolent government or planner would use the representative agent's utility function as the welfare function. A life-cycle economy, however, involves many heterogeneous agents: each generation has (at least) a representative member, and a new generation is born every period. At a minimum, relative weights need to be assigned to each individual. It is usually assumed that these weights take

---

[8] This remains true even in infinitely-lived agent economies with heterogeneous consumers. See Judd (1985) for details.

[9] Recall that taxing consumption tomorrow more (less) than today is equivalent to taxing (subsidizing) capital income tomorrow.

the form of a discount factor, so that the planner places an ever decreasing weight on future generations.[10]

Irrespective of the precise form of these welfare weights, the impact of a capital levy is very different in a life-cycle environment than in an infinitely-lived agent economy, simply because a capital levy explicitly involves a redistribution between generations: The individuals on whom the burden of a capital levy falls are different from those who benefit from lower distortionary taxation in the future. For example, consider the impact of confiscating the assets of a (possibly retired) individual who, at date zero, is in his last period of life. Under this front loading policy, this individual's consumption would be very low (it may be zero) and so would his utility; his utility is not affected by the lower tax rates that future generations would face. Since this individual's utility has a positive weight in the welfare function, the value of the government's objective would be driven down considerably by the front loading policy. This is not to say that the government would not tax initial assets at all, but rather that the extent to which the government will do so is limited, at least relative to what is optimal in infinitely-lived agent models. Accordingly, there is no need to impose arbitrary bounds on feasible tax rates in life-cycle economies.

Recall that the level of the long-run labor income tax in the infinitely-lived agent model is a function of these exogenous bounds on feasible tax rates. The size of these bounds determines how much capital the government accumulates during the transition, and thus the tax revenue that needs to be collected in the long run and the tax rate on labor income. Since there is no need to impose such bounds in life-cycle models, what, then, determines the tax revenue that needs to be collected in the long run? The answer lies in the weights that the planner puts on different generations. In the usual case, where these weights are represented by a discount factor, the steady-state amount of government debt and the amount of tax collection are entirely driven by the size of the discount factor. A relatively low discount factor indicates that the government puts low weights on future generations relative to current generations. In such cases, the government will tend to have a relatively high amount of accumulated debt in the long run and will need to collect a relatively high amount of taxes. Similarly, a high discount factor implies low (or even negative) government debt and that a small amount of taxes is to be collected.

[10] Note that these weights have to get smaller over time—they have to converge to zero—for the welfare function to be well defined. For example, if all generations had the same weight, all (positive) allocations would give the same welfare value (infinity).

2. AN INFINITELY-LIVED AGENT ECONOMY

Consider an economy populated by a large number of identical individuals with infinite lives.[11] Each period, individuals are endowed with one unit of productive time. The representative individuals' preferences are ordered according to the following utility function

$$\sum_{t=0}^{\infty} \beta^t U(c_t, 1 - l_t). \tag{1}$$

Equation (1) expresses that in each period of their infinite lives, individuals care about consumption, denoted c_t, and leisure, $(1-l_t)$; the latter corresponds to the total endowment of time minus time devoted to work (l_t). The discount factor $1 > \beta > 0$ is used by individuals to discount utility in future periods to utility in the current period. We assume that the utility function U is strictly increasing in both arguments, is strictly concave, and satisfies the standard Inada conditions.[12] Two commonly used algebraic forms for the utility function will be considered in the text. The first is a utility function which is *separable* between consumption and leisure:

$$U(c, 1 - l) = \frac{c^{1-\sigma}}{1 - \sigma} + V(1 - l), \tag{2}$$

where the function V satisfies the above stated conditions and $1/\sigma$ is the intertemporal elasticity of substitution, which measures the degree to which individuals are willing to substitute consumption over time. The second functional form we consider is a *Cobb-Douglas* utility function:

$$U(c, 1 - l) = \frac{c^{1-\sigma}(1 - l)^{\eta}}{1 - \sigma}, \tag{3}$$

where $\eta = \theta(1 - \sigma)$. In equation (3), $1/\sigma$ has the same interpretation as it did under the separable utility function and θ measures the intensity of leisure in individuals' preferences.

Each period, individuals face the budget constraint

$$c_t + a_{t+1} = w_t l_t + (1 + r_t)a_t, \tag{4}$$

where w_t is the after-tax wage rate, r_t is the after-tax interest rate, and a_{t+1} is the amount of resources carried over from period t to period $t + 1$. Letting p_t be the Lagrange multiplier on the time-t budget constraint, the first order

---

[11] This section draws from Atkeson, Chari, and Kehoe (1999).

[12] The Inada conditions state that the marginal utility of consumption or leisure is very high (low) at very low (high) consumption levels, that is, $\lim_{c\to 0} U(c, 1-l) = \lim_{l\to 1} U(c, 1-l) = \infty$ and $\lim_{c\to\infty} U(c, 1-l) = 0$. Note that leisure time cannot exceed one since working time cannot be negative.

conditions for individuals are

$$\beta^t U_{c_t} - p_t = 0, \tag{5}$$

$$\beta^t U_{l_t} + p_t w_t = 0, \tag{6}$$

$$-p_t + (1 + r_{t+1}) p_{t+1} = 0, \tag{7}$$

where U_{c_t} and U_{l_t} denote the derivative of U with respect to c_t and l_t respectively, that is, the marginal utility of consumption and leisure. One could use these conditions to obtain the optimal consumption and leisure demands of individuals. Naturally, these demand functions would depend on the fiscal policy chosen by the government, and they would represent the reaction functions that the government takes into account when formulating a Ramsey problem. We show below that these first order conditions can be used not only to construct the budget constraint but also to construct a constraint that can be imposed on the government when formulating a Ramsey problem where the government chooses allocations rather than tax rates.

There is a single produced good in our economy that can be used as consumption (private or public) or capital. For the goods-producing sector of our economy, we assume that the input-output technology is represented by a neoclassical production function with constant returns to scale, $y_t = f(k_t, l_t)$, where y_t, k_t, and l_t denote the aggregate (or per capita) levels of output, capital, and labor, respectively. Profit maximization by firms implies that capital and labor services are paid their marginal products: before-tax prices of capital and labor in period t are given by $\hat{r}_t = f_{k_t} - \delta$, where $0 < \delta < 1$ is the depreciation rate of capital, and $\hat{w}_t = f_{l_t}$.

Feasibility requires that total (private and public) consumption plus investment be less than or equal to aggregate output

$$c_t + k_{t+1} - (1 - \delta)k_t + g_t \le y_t, \tag{8}$$

where g_t stands for date-t government consumption and all aggregate quantities are expressed in per capita terms.

To finance its given stream of expenditures, the government has access to a set of fiscal policy instruments and to a commitment technology to implement its fiscal policy. The set of instruments available to the government consists of government debt and proportional taxes on labor income and capital income.[13] The date-t tax rates on capital and labor services are denoted by τ_t^k and τ_t^w, respectively. In per capita terms, the government budget constraint at date t is given by

$$(1 + \hat{r}_t)b_t + g_t = b_{t+1} + (\hat{r}_t - r_t)a_t + (\hat{w}_t - w_t)l_t, \tag{9}$$

[13] In this framework, consumption taxes need to be ruled out to make the problem interesting since they can be used in conjunction with labor income taxes to perfectly imitate a levy of initial holdings of assets. See Chari and Kehoe (1999) for details.

where b_t represents government debt issued at date t, $w_t \equiv (1 - \tau_t^w)\hat{w}_t$, and $r_t \equiv (1 - \tau_t^k)\hat{r}_t$. Equation (9) expresses that the government pays its expenditures, which are composed of outstanding government debt payments (principal plus interest) and other government outlays, either by issuing new debt, by taxing interest income, or by taxing wage income.

In the spirit of Ramsey, the government takes individuals' optimizing behavior as given and chooses a fiscal policy to maximize a given welfare criterion. Since there is but a single representative agent in this economy, a natural way for the government to evaluate different allocations is to use the representative individual's utility function. If we let π denote a fiscal policy and denote $c_t(\pi)$ and $l_t(\pi)$ the solution to the consumer problem as a function of the fiscal policy, then the Ramsey problem is

$$\max_{\{\pi\}} \sum_{t=0}^{\infty} \beta^t U(c_t(\pi), 1 - l_t(\pi)), \tag{10}$$

subject to feasibility (8) and the government budget constraint (9) for all $t \geq 0$. This Ramsey problem corresponds to the dual. Note that the problem is fairly difficult to analyze because any tax instrument enters all demand functions. Given this difficulty, the primal approach is much more tractable.

The Primal Approach

To construct a Ramsey problem where the government chooses allocations rather than tax rates, we must restrict the set of allocations from which the government can choose. This set should include only allocations that are competitive equilibrium under some fiscal policy. To construct this set, we use the fact that for any given fiscal policy, the competitive equilibrium must satisfy the consumer's optimality conditions, including the budget constraint, as well as those of the firm. Using these optimality conditions, we can derive a condition that competitive equilibria must satisfy.

Our first step is to iterate on the budget constraint (4) to express this sequence of constraints as a single, present-value budget constraint

$$\sum_{t=0}^{\infty} \left[\prod_{j=1}^{t} \frac{1}{1 + r_j} \right] (c_t - w_t l_t) = (1 + r_0) a_0. \tag{11}$$

The term inside the square brackets is a shorthand to express the multiplication of many terms, $\frac{1}{1+r}$ in this case. Next, the consumer's first order conditions (5) through (7) imply

$$p_t/p_0 = \prod_{j=1}^{t} \frac{1}{1+r_j}, \tag{12}$$

$$w_t = \frac{U_{l_t}}{U_{c_t}}. \tag{13}$$

Using equations (5), (12), and (13) we can rewrite the present value budget constraint (11) as

$$\sum_{t=0}^{\infty} \beta^t (U_{c_t} c_t + U_{l_t} l_t) = A_0, \tag{14}$$

where $A_0 \equiv (1+r_0)U_{c0}a_0$, $r_0 = (1-\tau_0^k)\hat{r}_0$ and τ_0^k is taken to be fixed to make the problem interesting.

Equation (14) is referred to as the *implementability constraint*. It can be shown that any competitive equilibrium allocation has to satisfy this constraint, and that any feasible allocation satisfying the implementability constraint is a competitive equilibrium (see Chari and Kehoe [1999] for details). Imposing this constraint on the government's problem accomplishes exactly what we wanted: It ensures that any allocation picked by the government can be implemented as a competitive equilibrium.

We can now state the Ramsey problem in terms of allocations. This problem consists of maximizing welfare, given by the representative consumer's utility function (1), subject to feasibility (8) and the implementability constraint (14). Note that by Walras's law, if the individual's present value budget constraint (11) holds under a feasible allocation, then the government's budget constraint (9) is also satisfied. Let λ be the Lagrange multiplier on the implementability constraint (14) and define the pseudo-welfare function W to include the implementability constraint

$$W_t = U(c_t, 1-l_t) + \lambda(U_{c_t} c_t + U_{l_t} l_t). \tag{15}$$

The multiplier λ will be strictly positive if it is necessary for the government to use distortionary taxation. The term multiplying λ essentially gives a bonus to date-t allocations that bring in extra government revenues, thereby relieving other periods from distortionary taxation, and the same term imposes a penalty in the opposite situation. The Ramsey problem, in terms of allocations, is

$$\max_{\{c_t, l_t, k_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t W_t - \lambda A_0, \tag{16}$$

subject to feasibility (8). The form the above problem takes, the *primal*, is very similar to a first-best planning problem except that the pseudo-welfare function replaces the utility function.

Prescriptions

Using the primal allows us to characterize optimal fiscal policies. With few exceptions, our focus will be on the capital income tax. The first order conditions for an optimum imply[14]

$$-\frac{W_{l_t}}{W_{c_t}} = -\frac{U_{l_t}[1+\lambda(1+H_t^l)]}{U_{c_t}[1+\lambda(1+H_t^c)]} = f_{l_t}, \tag{17}$$

$$\frac{W_{c_t}}{W_{c_{t+1}}} = \frac{U_{c_t}[1+\lambda(1+H_t^c)]}{U_{c_{t+1}}[1+\lambda(1+H_{t+1}^c)]} = \beta(1+f_{k_{t+1}}-\delta), \tag{18}$$

for $t = 1, 2, \ldots,$ where

$$H_t^c \equiv \frac{U_{c_t,c_t}c_t + U_{l_t,c_t}l_t}{U_{c_t}}, \tag{19}$$

$$H_t^l \equiv \frac{U_{c_t,l_t}c_t + U_{l_t,l_t}l_t}{U_{l_t}}. \tag{20}$$

Equation (17) equates the government's marginal rate of substitution between consumption and leisure at date t to the marginal product of labor of that date. Similarly, equation (18) sets the government's marginal rate of substitution between consumption today and consumption tomorrow equal to the return on capital (net of depreciation). Note that the government's marginal rate of substitution, unlike that of an individual, takes the implementability constraint into account. Also, the government cares about before-tax prices, whereas individuals face after-tax prices.

Atkeson, Chari, and Kehoe (1999) refer to the terms H^c and H^l as *general equilibrium elasticities* since they capture the relevant distortions for setting the capital and labor income tax rates in general equilibrium. Notice that if $H_t^c = H_{t+1}^c$, then equation (18) implies that

$$\frac{U_{c_t}}{U_{c_{t+1}}} = \beta(1+f_{k_{t+1}}-\delta) = \beta(1+\hat{r}_{t+1}). \tag{21}$$

This condition is of particular interest, for when $H_t^c = H_{t+1}^c$, the tax on capital income in period $t+1$ is zero. To observe this result, notice that the consumer's first order conditions (5) and (7), which the Ramsey allocation has to satisfy, imply that

$$\frac{U_{c_t}}{U_{c_{t+1}}} = \beta(1+r_{t+1}) = \beta(1+(1-\tau_{t+1}^k)\hat{r}_{t+1}). \tag{22}$$

For both equations (21) and (22) to hold, τ_{t+1}^k must equal zero. Similarly, if $H_t^c = H_t^l$, the tax rate on labor income has to be equal to zero for the optimality

[14] Note that the first-order conditions at time zero are different from the above equations since consumption at date zero appears inside the term A_0.

condition in order for the labor decision of individuals ($-U_{l_t}/U_{c_t} = w_t$) to be compatible with that of the government (equation (17)). We have just demonstrated the following proposition:

Proposition 1 *In our infinitely-lived agent model, (i) the optimal tax rate on labor income at date t is different from zero unless $H_t^c = H_t^l$, and (ii) the optimal tax rate on capital income at date $t+1$ is different from zero unless $H_t^c = H_{t+1}^c$.*

Chamley's (1986) celebrated result on the optimality of not taxing capital in the long run follows directly from Proposition 1. Suppose that the Ramsey allocation converges to a steady state where consumption and leisure are constant by definition. It follows immediately that in such a steady state, the function H_t^c is also constant, which implies that the optimal capital income tax is zero in the long run.

Proposition 1 also implies that for utility functions that are separable in consumption and leisure (so that $U_{c,l} = 0$) and have a constant intertemporal elasticity of substitution (so that $U_{c,c}c/U_c$ is constant), the capital income tax should be zero in all but the first period. Of course, capital should be taxed at a confiscatory rate in the first period regardless of preferences since this tax perfectly imitates a lump-sum tax.

Proposition 2 *Under utility functions of the form given by* (2)*, the optimal capital income tax in our infinitely-lived agent model is zero for $t > 1$.*

Chamley also shows, under the same separable utility function, that imposing a bound on feasible tax rates implies the following for the optimal tax on capital income: The tax rate should be equal to the upper bound for a finite amount of time, after which it should be equal to zero. In discrete time models, there is a period between the two regimes where the tax rate is strictly between zero and the upper bound. The intuition for this result is that taxing capital income has two effects. While the capital income tax partially imitates a lump-sum tax because the initial stock of capital is given, it also introduces a distortion on the savings decision. As a result, the lump-sum aspect of the tax dominates for periods sufficiently close to date zero, and the distortionary aspect of the tax dominates thereafter.

The intuition for the above result remains intact under more general utility functions, in the sense that early taxation of capital income is preferred to later taxation. It is less clear, however, whether capital income should be taxed throughout the transition. In particular, for the Cobb-Douglas utility function (3), the optimal capital income tax is zero only in the long run.

3. A LIFE-CYCLE ECONOMY

We now consider a life-cycle economy.[15] This economy is similar to the infinitely-lived agent economy considered in Section 2, except it is populated by overlapping generations of individuals with finite lives. Individuals still make consumption and labor/leisure choices in each period so as to maximize their lifetime utility, and firms still operate a neoclassical production technology. The payments received by individuals on their factors (capital and labor) are subject to proportional taxes, which we now assume can be conditioned on age.

Individuals live $(J+1)$ periods, from age 0 to age J. At each time period, a new generation is born and is indexed by date of birth. At date zero, when the change in fiscal policy occurs, the generations alive are $-J, -J+1, \ldots, 0$. In order to take these initial generations into account in the following analysis, it will prove convenient to denote the age of individuals alive at date zero by $j_0(t)$. For all other generations we set $j_0(t) = 0$, so that for any generation t, $j_0(t) = \max\{-t, 0\}$. One can thus think of $j_0(t)$ as the first period of an individual's life affected by the date zero switch in fiscal policy. We let $\mu_j = 1/(J+1)$ represent the share of age-j individuals in the population. The labor productivity level of an age-j individual is denoted z_j.

We let $c_{t,j}$ and $l_{t,j}$, respectively, denote consumption and time devoted to work by an age-j individual who was born in period t. Note that $c_{t,j}$ and $l_{t,j}$ actually occur in period $(t+j)$. Similarly, the after-tax prices of labor and capital services are denoted $w_{t,j}$ and $r_{t,j}$, respectively. The problem faced by an individual born in period $t \geq -J$ is to maximize lifetime utility subject to a sequence of budget constraints:

$$U^t(\pi) \equiv \max \sum_{j=j_0(t)}^{J} \beta^{j-j_0(t)} U(c_{t,j}, 1-l_{t,j}), \tag{23}$$

$$\text{s.t.} \quad c_{t,j} + a_{t,j+1} = w_{t,j} z_j l_{t,j} + (1+r_{t,j}) a_{t,j}, \qquad j = j_0(t), \ldots, J. \tag{24}$$

This problem mimics the consumer's problem from Section 2, except for the need to index all variables by age. We denote $U^t(\pi)$ the indirect utility function, that is, the maximum lifetime utility obtained by an individual from generation t under fiscal policy π.

Let $p_{t,j}$ denote the Lagrange multiplier associated with the budget constraint (24) faced by an age-j individual born in period t. The necessary and sufficient conditions for a solution to the consumer's problem are given by (24)

[15] This section draws from Erosa and Gervais (2000).

and

$$\beta^{j-j_0(t)}U_{c_{t,j}} - p_{t,j} = 0, \tag{25}$$

$$\beta^{j-j_0(t)}U_{l_{t,j}} + p_{t,j}w_{t,j}z_j \leq 0, \quad \text{with equality if } l_{t,j} > 0, \tag{26}$$

$$-p_{t,j} + p_{t,j+1}(1 + r_{t,j+1}) = 0, \tag{27}$$

$$a_{t,J+1} = 0, \tag{28}$$

$j = j_0(t), \ldots, J$, where $U_{c_{t,j}}$ and $U_{l_{t,j}}$ denote the derivative of U with respect to $c_{t,j}$ and $l_{t,j}$, respectively.[16]

The feasibility constraint is still given by (8). However, the date-t aggregate levels of consumption and labor input—the latter being expressed in efficiency units—are now obtained by adding up the weighted consumption (or effective labor supply) of all individuals alive at date t, where the weights are given by the fraction of the population that each individual represents:

$$c_t = \sum_{j=0}^{J} \mu_j c_{t-j,j},$$

$$l_t = \sum_{j=0}^{J} \mu_j z_j l_{t-j,j}.$$

The set of instruments available to the government consists of government debt and proportional, age-dependent taxes on labor income and capital income.[17] The date-t tax rates on capital and labor services supplied by an age-j individual (born in period $(t - j)$) are denoted by $\tau^k_{t-j,j}$ and $\tau^w_{t-j,j}$, respectively. In per capita terms, the government budget constraint at date $t \geq 0$ is given by

$$(1 + \hat{r}_t)b_t + g_t = b_{t+1} + \sum_{j=0}^{J}(\hat{r}_t - r_{t-j,j})\mu_j a_{t-j,j} + \sum_{j=0}^{J}(\hat{w}_t - w_{t-j,j})\mu_j z_j l_{t-j,j}, \tag{29}$$

where $w_{t,j} \equiv (1 - \tau^w_{t,j})\hat{w}_{t+j}$ and $r_{t,j} \equiv (1 - \tau^k_{t,j})\hat{r}_{t+j}$. Equation (29) has exactly the same interpretation as equation (9) in the infinitely-lived agent model.

[16] We assumed in Section 2 that the labor supply was always between zero and one. In the context of a life-cycle economy, however, the labor supply can realistically hit a corner solution if labor productivity gets sufficiently small. For example, individuals may become less productive as they age and choose to retire.

[17] Recall that consumption taxes had to be ruled out in the infinitely-lived agent model because they could perfectly imitate a levy of initial holdings of assets. Because the government's incentive to confiscate initial holdings of assets is endogenously limited in life-cycle economies, we could allow for consumption taxes. However, it can be shown that a consumption tax would be a redundant instrument. See Erosa and Gervais (2000) for details.

Because of the presence of many heterogeneous individuals in this economy, the choice of a welfare function is not as straightforward here as it is in the infinitely-lived agent model. Below, we assume that social welfare is defined as the discounted sum of individual lifetime welfares (as in Samuelson [1968] and Atkinson and Sandmo [1980]). In other words, the government chooses a sequence of tax rates in order to maximize

$$\sum_{t=-J}^{\infty} \gamma^t U^t(\pi), \tag{30}$$

where $0 < \gamma < 1$ is the intergenerational discount factor and $U^t(\pi)$, as was defined earlier, is the indirect utility function of generation t as a function of the government tax policy. As was the case with the infinitely-lived agent model, it is much easier to characterize optimal fiscal policies using the primal approach.

The Primal Approach

As before, we need to impose restrictions so that any allocation chosen by the government can be decentralized as a competitive equilibrium. In life-cycle models, each generation has its own implementability constraint. The implementability constraints are obtained by using the consumers' optimality conditions (25) through (27) and acknowledging the fact that factors are paid their marginal products to substitute out prices from consumers' budget constraints (24). After adding up these budget constraints, the resulting implementability constraint associated with the cohort born in period t is given by

$$\sum_{j=j_0(t)}^{J} \beta^{j-j_0(t)}(U_{c_{t,j}} c_{t,j} + U_{l_{t,j}} l_{t,j}) = A_{t,j_0(t)}, \tag{31}$$

where $A_{t,j_0(t)} \equiv U_{c_{t,j_0(t)}}(1 + r_{t,j_0(t)})a_{t,j_0(t)}$. It should be emphasized that implicit in this implementability constraint is the existence of age-dependent tax rates. Additional restrictions need to be imposed for an allocation to be implementable with age-independent taxes. In other words, the set of allocations from which the government can pick depends crucially on the instruments available to the government.

The Ramsey problem in this life-cycle economy consists of choosing an allocation to maximize the discounted sum of successive generations' utility, subject to each generation's implementability constraint (31) as well as the feasibility constraint (8) for $t = 0, \ldots$:

$$\max_{\left\{\{c_{t,j}, l_{t,j}\}_{j=j_0(t)}^{J}, k_{t+J+1}\right\}_{t=-J}^{\infty}} \sum_{t=-J}^{\infty} \gamma^t W_t.$$

The pseudo-welfare function W_t is defined as in Section 2 to include generation t's implementability constraint in addition to its lifetime utility. If we let $\gamma^t \lambda_t$ be the Lagrange multiplier associated with generation t's implementability constraint (31), then the function W_t is defined as

$$W_t = \sum_{j=j_0(t)}^{J} \beta^{j-j_0(t)} \left[U\left(c_{t,j}, 1-l_{t,j}\right) + \lambda_t (U_{c_{t,j}} c_{t,j} + U_{l_{t,j}} l_{t,j})\right] - \lambda_t A_{t,j_0(t)}. \quad (32)$$

As was the case in the infinitely-lived agent model, the government budget constraint (29) has been omitted from the Ramsey problem since it has to hold by Walras's law.

Prescriptions

In this section we show that the solution to the Ramsey problem generally features non-zero tax rates on labor and capital income and demonstrate how these rates vary with age. In particular, and in contrast with infinitely-lived agent models, if the Ramsey allocation converges to a steady state solution, optimal capital income taxes will in general differ from zero even in that steady state. Although the main results of this section hold more generally, we will restrict attention to steady states for ease of exposition.

Let $\gamma^t \phi_t$ denote the Lagrange multiplier associated with the time-t feasibility constraint (8). The steady state solution is characterized by the following equations:

$$1/\gamma = 1 - \delta + f_k, \quad (33)$$

$$-\frac{W_{l_j}}{W_{c_j}} = -\frac{U_{l_j}[1+\lambda(1+H_j^l)]}{U_{c_j}[1+\lambda(1+H_j^c)]} \leq z_j f_l, \quad \text{with equality if } l_j > 0, \quad (34)$$

$$\frac{W_{c_j}}{W_{c_{j+1}}} = \frac{U_{c_j}[1+\lambda(1+H_j^c)]}{\beta U_{c_{j+1}}[1+\lambda(1+H_{j+1}^c)]} = f_k + 1 - \delta, \quad (35)$$

where

$$H_j^c = \frac{U_{c_j,c_j} c_j + U_{l_j,c_j} l_j}{U_{c_j}}, \quad (36)$$

$$H_j^l = \frac{U_{c_j,l_j} c_j + U_{l_j,l_j} l_j}{U_{l_j}}, \quad (37)$$

as well as the feasibility and implementability constraints (8) and (31).

The first order condition with respect to capital, equation (33), implies that the solution to this Ramsey problem has the modified golden rule property: The marginal product of capital (net of depreciation) equals the discount rate applied to successive generations $(1/\gamma - 1)$ (see Samuelson [1968]). Equation (34) equates the government's marginal rate of substitution between consumption and leisure of an age-j individual to the effective marginal product of labor of that same individual. Similarly, equation (35) sets the government's marginal rate of substitution between consumption today and consumption tomorrow equal to the return on capital (net of depreciation).

We now derive necessary conditions under which the Ramsey allocation features zero taxation of either labor or capital income. Any optimal fiscal policy has to satisfy the consumer's optimality conditions, and we derive the necessary conditions by comparing the optimality conditions from the consumer's problem to those of the Ramsey problem.

Combining the consumer's first order conditions for consumption (25) and labor (26), and applying them to the nontrivial case of positive labor supply, we obtain

$$-\frac{U_{l_j}}{U_{c_j}} = z_j w_j = z_j \hat{w}(1 - \tau_j^w), \tag{38}$$

which corresponds to the usual optimality condition that the marginal rate of substitution between labor and consumption be equal to the relative price of labor faced by the consumer. Compare equation (38) to its analogue from the Ramsey problem (34). Using the fact that $\hat{w} = f_l$, the tax rate on labor income for an age-j individual is given by

$$\tau_j^w = \frac{\lambda(H_j^l - H_j^c)}{1 + \lambda + \lambda H_j^l}. \tag{39}$$

Since λ is generally different from zero, this tax rate on labor income will be equal to zero only if $H_j^l = H_j^c$.

The same logic applies to the tax rate on capital income. For this case, consider the consumer's first order condition for consumption (25) at age j and $j+1$. Using the consumer's first order condition for asset holdings (27), we get

$$\frac{U_{c_j}}{\beta U_{c_{j+1}}} = 1 + r_{j+1} = 1 + (1 - \tau_{j+1}^k)\hat{r}. \tag{40}$$

Equation (40) corresponds to the usual intertemporal condition that sets the marginal rate of substitution between consumption today and consumption tomorrow equal to the relative price of the same commodities, which is equal to the gross interest rate. The government's counterpart of (40) is given by equation (35). Using these two equations and the fact that $\hat{r} = f_k - \delta$, we obtain

$$\frac{1+\hat{r}}{1+(1-\tau^k_{j+1})\hat{r}} = \frac{1+\lambda+\lambda H^c_j}{1+\lambda+\lambda H^c_{j+1}}, \tag{41}$$

which implies that the tax rate on capital income is different from zero unless $H^c_j = H^c_{j+1}$. These results are summarized in the following proposition.

Proposition 3 *In our life-cycle economy, (i) the optimal tax rate on labor income at age j is different from zero unless $H^l_j = H^c_j$, and (ii) the optimal tax rate on capital income at age $j+1$ is different from zero unless $H^c_j = H^c_{j+1}$.*

Proposition 3 is very similar to Proposition 1, but with age of individuals replacing the time period. Essentially, prescriptions that hold in the transition of infinitely-lived agent models hold in the steady state of life-cycle economies. Since zero-capital income tax is merely a special case in the transition of infinitely-lived agent models, we should expect the same prescription in the steady state of life-cycle economies.

The intuition as to why the celebrated Chamley-Judd zero-capital tax result does not extend to life-cycle economies should be clear. Since consumption and leisure are constant in the steady state of infinitely-lived agent models, H^c is constant; thus, zero-capital income taxation is optimal regardless of the form of the utility function. In contrast, consumption and leisure are generally not constant over an individual's lifetime in life-cycle models, even in steady state. There is in fact no reason to expect $H^c_j = H^c_{j+1}$, and consequently capital income taxes will generally not be equal to zero in the long run. Obviously, if the economy is specified so that individuals' behavior features no life-cycle elements, i.e., labor supply and consumption are independent of age, then optimal taxation works as in infinitely-lived agent models in the sense that capital income is not taxed.

Proposition 4 *In our life-cycle model, if (i) $z_j = z$, $j = 0, \ldots, J$ and (ii) $\gamma = \beta$, then it is not optimal to tax capital income in the long run.*

The proof of Proposition 4 is fairly intuitive. From the first order condition with respect to capital (33), when $\gamma = \beta$, the steady state return on capital coincides with individuals' rate of time preference, i.e., $f_k - \delta = \hat{r} = 1/\gamma - 1 = 1/\beta - 1$. The consumer's optimization conditions (25) through (27) then imply that, given a constant productivity profile, consumption and leisure do not depend on age in steady state. In turn, the absence of life-cycle behavior implies that the function H^c_j does not depend on age, which, following Proposition 3, is sufficient for the optimal capital income tax rate to be zero in steady state.

Proposition 4 is a generalization of a result in Atkeson, Chari, and Kehoe (1999) that they use to prescribe zero-capital income taxation in overlapping generations economies. It should be noted, however, that the conditions stated in Proposition 4 have empirically unappealing implications. In particular, they imply that individuals consume their labor earnings period by period. As a result, individuals do not accumulate any assets and the entire stock of capital is owned by the government.

Optimal capital income taxes are also zero when preferences are such that uniform commodity taxation over the lifetime of individuals is optimal. The separable utility function (2) is one form under which the capital income tax is zero. Because this utility function is separable in consumption and leisure (so that $U_{c,l} = 0$) and has a constant intertemporal elasticity of substitution (so that $U_{c,c}c/U_c$ is constant), the general equilibrium elasticity (the function H^c) is constant. Relative to Proposition 2, the general equilibrium elasticity here is not only independent of time but of age as well.

Proposition 5 *Under utility functions of the form given by* (2)*, the optimal capital income tax in our life-cycle model is zero for* $t > 1$.

Proposition 5, combined with Proposition 2, sounds like good news. If we were confident that individuals' preferences were reasonably well represented by utility functions that are separable in consumption and leisure, then we would be reasonably confident that prescribing zero-capital income tax was the right thing to do. There is, however, an important caveat: The result in Proposition 5 relies heavily on the government's ability to age-condition tax rates. If the government were constrained to use tax rates that are independent of age, then the optimal capital income tax would no longer be zero.[18] Furthermore, applied work in public finance is usually conducted with utility functions that are not separable in consumption and leisure.[19]

Under the Cobb-Douglas utility function, it is straightforward to show that optimal capital income taxes in this case are zero in the long run only under very restrictive conditions, as stated in Proposition 4. The principles guiding the optimal manner in which to tax capital over the lifetime of individuals are stated in the following proposition:

Proposition 6 *For utility function of the form given by* (3)*, the tax rate on capital income at age* $j + 1$ *is positive if and only if* $l_{j+1} < l_j$.

[18] Although showing this result is beyond the scope of this paper, the intuition is that the government will use a non-zero capital income tax to imitate the optimal age-dependent labor income tax (see Erosa and Gervais [2000] for details). Alvarez, Burbidge, Farrell, and Palmer (1992) derive a similar result in a partial equilibrium setting. This type of finding is reminiscent of results in Stiglitz (1987) and Jones, Manuelli, and Rossi (1997), where the government taxes capital income when it is constrained to use tax rates that are independent of individuals' skill levels.

[19] Auerbach and Kotlikoff (1987), for example, use a nested CES utility function with intertemporal elasticity of substitution equal to 0.25 and intratemporal elasticity of substitution equal to 0.8. Both elasticities would have to be equal to unity for the CES utility function to be separable in consumption and leisure.

Proof. The proof follows directly from the definitions of H_j^c (equation (36)) under utility function (3). Since $H_j^c = -\sigma - \eta/(1 - l_j)$ we can rewrite equation (41) as

$$\frac{1+\hat{r}}{1+r_{j+1}} = \frac{1+\lambda+\lambda(-\sigma-\eta/(1-l_j))}{1+\lambda+\lambda(-\sigma-\eta/(1-l_{j+1}))}. \tag{42}$$

Notice that $\tau_{j+1}^k > 0$ if and only if

$$\frac{1+\hat{r}}{1+r_{j+1}} = \frac{1+\hat{r}}{1+(1-\tau_{j+1}^k)\hat{r}} > 1. \tag{43}$$

From equations (42) and (43), we obtain that $\tau_{j+1}^k > 0$ if and only if $l_{j+1} < lj$. ■

By taxing (subsidizing) capital, the government makes consumption and leisure in the future more (less) expensive than today. Proposition 6 suggests the government uses capital income taxes to smooth individuals' leisure and consumption profiles over their lifetimes. Under Cobb-Douglas utility, the share of consumption in an individual's total expenditures is constant, so that consumption and leisure always move together over time. If consumption and leisure are high tomorrow relative to today, then the government will tax the return on today's savings at a positive rate tomorrow. By doing so, the government gives individuals an incentive to consume more and save less today, and thus to consume less tomorrow.

An implication of the principle of optimal taxation developed in Proposition 6 is that capital income should not be taxed during retirement. This follows directly from the fact that labor supply is constant during retirement. Notice, however, that leisure time during retirement is taxed indirectly because the return on savings is taxed prior to retirement.

4. CONCLUSION

We have reviewed optimal taxation in both an infinitely-lived agent model and a life-cycle economy. Our review shows that there is no consensus regarding the optimal tax on capital income. Although the optimal capital income tax is invariably zero in the long-run equilibrium of infinitely-lived agent models, the conditions under which that is the case in life-cycle economies are very stringent. Even under a separable utility function, the capital income tax will only be zero if the government has access to age-dependent labor income taxes. Furthermore, both models suggest that capital income should be taxed at non-zero rates during the transition unless individuals have separable preferences between consumption and leisure. Thus, the strong conclusions and recommendations of much of this literature must be treated with caution.

REFERENCES

Alvarez, Yvette, John Burbidge, Ted Farrell, and Leigh Palmer. "Optimal Taxation in a Life Cycle Model," *The Canadian Journal of Economics*, vol. XXV (February 1992), pp. 111–22.

Atkeson, Andrew, V.V. Chari, and Patrick J. Kehoe. "Taxing Capital Income: A Bad Idea," Federal Reserve Bank of Minneapolis *Quarterly Review*, vol. 23 (Summer 1999), pp. 3–17.

Atkinson, Anthony B., and Agnar Sandmo. "Welfare Implications of the Taxation of Savings." *The Economic Journal*, vol. 90 (September 1980), pp. 529–49.

______________, and Joseph E. Stiglitz. *Lectures on Public Economics*. New York: McGraw-Hill, 1980.

Auerbach, Alan J., and Laurence J. Kotlikoff. *Dynamic Fiscal Policy*. Cambridge: Cambridge University Press, 1987.

______________, and Jonathan Skinner. "The Efficiency Gains from Dynamic Tax Reform," *International Economic Review*, vol. 24 (February 1983), pp. 81–100.

Chamley, Christophe. "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives," *Econometrica*, vol. 54 (May 1986), pp. 607–22.

Chari, V. V., and Patrick J. Kehoe. "Optimal Fiscal and Monetary Policy," in J.B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics, Vol. 1*. Amsterdam: North-Holland, 1999, pp. 1671–745.

Erosa, Andrés, and Martin Gervais. "Optimal Taxation in Life-Cycle Economies," Federal Reserve Bank of Richmond Working Paper 00-2, September 2000.

Jones, Larry E., Rodolfo E. Manuelli, and Peter E. Rossi. "On the Optimal Taxation of Capital Income," *Journal of Economic Theory*, vol. 73 (March 1997), 93–117.

Judd, Kenneth L. "Redistributive Taxation in a Simple Perfect Foresight Model," *Journal of Public Economics*, vol. 28 (October 1985), pp. 59–83.

Kotlikoff, Laurence J. "The A-K Model: Its Past, Present, and Future." NBER Working Paper 6684, August 1998.

Kydland, Finn E., and Edward C. Prescott. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy*, vol. 85 (June 1977), pp. 473–91.

Lucas, Robert E., Jr., and Nancy L. Stokey. "Optimal Fiscal and Monetary Policy in an Economy without Capital," *Journal of Monetary Economics*, vol. 12 (July 1983), pp. 55–93.

Ramsey, Frank P. "A Contribution to the Theory of Taxation," *Economic Journal*, vol. 37 (1927), pp. 47–61.

Samuelson, Paul A. "The Two-Part Golden Rule Deduced as the Asymptotic Turnpike of Catenary Motions." *Western Economic Journal*, vol. 6 (March 1968), pp. 85–89.

Stiglitz, Joseph E. "Pareto Efficient and Optimal Taxation and the New New Welfare Economics," in A.J. Auerbach and M. Feldstein, eds., *Handbook of Public Economics, Vol. 2*. Amsterdam: North-Holland, 1987, pp. 991–1042.

The Wealth Effect in Empirical Life-Cycle Aggregate Consumption Equations

Yash P. Mehra

This article presents an empirical model of U.S. consumer spending that relates consumption to labor income and household wealth. This specification is consistent with the life-cycle hypothesis of saving first popularized in the 1960s by Ando, Modigliani, and their cohorts.[1] My analysis here extends the previous research in several directions. First, I examine the dynamic relationship between consumption, income, and wealth using cointegration and error correction methodology. In previous research, the traditional life-cycle model has often been examined using either levels or first differences of these variables. While the use of differences does avoid the pitfall of spurious correlation due to common trending series, it tends to lead to the omission of the long-run equilibrium (cointegrating) relationships that may exist among levels of these variables. In fact, Gali (1990) goes so far as to present a theoretical life-cycle model that generates a common trend in aggregate consumption, labor income, and wealth. Therefore, my empirical work here tests for the presence of a long-run equilibrium (cointegrating) relationship between the level of aggregate consumer spending and its economic determinants such as labor income and wealth. I then examine the short-run dynamic relationship among these variables using an error correction specification proposed in Davidson et al. (1978).

The present article investigates whether wealth has predictive content for future consumption. If it does, then changes in wealth may lead to changes in

The author wishes to thank Huberto Ennis, Pierre Sarte, and Roy Webb for many useful suggestions. The views expressed in this paper are those of the author and do not necessarily reflect the views of the Federal Reserve Bank of Richmond or the Federal Reserve System.

1 See, for example, Ando and Modigliani (1963) and Modigliani and Brumberg (1979).

consumer spending.[2] I also examine whether consumer spending is sensitive to stock market wealth. The 1990s saw an enormous increase in household wealth generated by the rising value of household stock holdings. This increase has generated considerable interest among policymakers in identifying the magnitude of the stock market wealth effect on consumption. For example, in his recent testimony before the U.S. Congress, Chairman Alan Greenspan has stated that wealth-induced consumption growth has partly been responsible for generating aggregate demand in excess of potential. The Chairman says that the wealth effect may have added on average 1 1/2 to 2 percentage points to annual growth rate of real GDP over the last few years.[3] The empirical work here addresses the wealth effect by considering aggregate consumption equations that relate consumption directly to equity wealth.

Poterba (2000) has suggested that the marginal propensity to consume out of wealth may have declined in the 1990s. According to Poterba, the main reason for this decline is the growing importance of equities in household wealth. Since the number of households holding equities is still lower than the number holding many other kinds of assets, and since a growing part of equity investments are held in tax-favored retirement accounts, the marginal propensity to consume out of equity wealth may be small. Furthermore, the marginal propensity to consume out of total wealth may appear to decline if the recent increase in household wealth reflects the growing importance of equities, as has been the case in the 1990s.[4] In order to see whether the relationship between consumption and wealth has changed during the 1990s, I estimate consumption equations over the full sample period 1959Q1 to 2000Q2 as well as two other subperiods ending in the early 1990s, 1959Q1 to 1990Q2 and 1959Q1 to 1995Q2.

The empirical results that are presented here show that aggregate consumer spending is cointegrated with labor income and wealth over the sample period 1959Q1 to 2000Q2, indicating the existence of a long-term equilibrium relationship between consumption and its economic determinants, such

---

[2] The other recent work that has applied cointegration techniques to consumption equations is in Ludvigson and Steindel (1999).

[3] In his testimony Chairman Alan Greenspan notes: "For some time now, the growth of aggregate demand has exceeded the expansion of production potential. . . . A key element in this disparity has been the very rapid growth of consumption resulting from the effects on spending of the remarkable rise in household wealth. . . . Historical evidence suggests that perhaps three to four cents of every additional dollar of stock market wealth eventually is reflected in increased consumer purchases. . . . [D]omestic demand growth, influenced importantly by the wealth effect on consumer spending, has been running 1-1/2 to 2 percentage points at an annual rate in excess of even the higher, productivity-driven growth in potential supply since late 1997." (Greenspan 2000a, p. 2; Greenspan 2000b, pp. 1–4).

[4] The other suggested factor that could contribute to a lower marginal propensity to consume out of wealth is the falling cost of leaving bequests. Poterba (2000) points out that estate tax reform has been a very active topic of congressional debate in recent years, with numerous proposals calling for elimination of the "death tax." These tax changes raise the attractiveness of leaving bequests, inducing households with high net worth to reduce their current marginal propensity to consume out of wealth.

as income and wealth. The coefficients that appear on income and wealth variables in the estimated cointegrating regression are statistically significant and measure the long-term responses of consumption to income and wealth. The results thus indicate wealth has a significant effect on consumption.

In the short-term consumption equations estimated here, the error correction variable appears with an expected negatively signed coefficient and is statistically significant, indicating the presence of lags in the response of consumption to income and wealth. This result also shows that wealth has predictive content for future consumption. Hence we may conclude that a persistent decline in equity wealth can lead to lower consumer spending.

My results do indicate that the long-term marginal propensity to consume out of wealth declined somewhat during the 1990s, as suggested in Poterba (2000). However, point estimates of the long-term marginal propensity to consume out of wealth have remained in a narrow range of 0.03 to 0.04, indicating a $1 increase in equity values raises consumption by 3 to 4 cents. The long-term marginal propensity to consume out of equity wealth is small, but even with such relatively low estimates of the marginal propensity to consume out of wealth, the consumption effect of the 1990s stock market boom is substantial. Estimates that I derived using the aggregate consumption equation indicate that the equity wealth effect may have added on average about 1 percentage point to the annual growth rate of real GDP over 1995 to 1999, as noted in Greenspan's testimony (2000a, p. 2).

During 2000, household equity wealth declined about 18 percent, after rising at an average annual rate of 18 percent per year during the preceding five-year period (1995–1999). The consumption equations estimated here indicate that the decline in equity wealth is likely to depress consumer spending. Hence, the growth rate of consumer spending in the near term is likely to fall below the robust 4 percent yearly growth rate observed during the preceding five-year period.

1. THE MODEL AND THE METHOD

An Aggregate Consumption Function

The aggregate consumption function studied is given in (1a).

$$\tilde{C}_t = a_0 + a_1 Y_t + a_2 Y^e_{t+k} + a_3 W_t, \tag{1a}$$

where $\tilde{C}_t$ is current planned consumption, Y_t is actual current-period labor income, W_t is actual current-period wealth, and Y^e_{t+k} is average anticipated future labor income over the earning span (k) of the working age population. Equation (1a) simply states that aggregate planned consumption depends upon the anticipated value of lifetime resources, which equals current and anticipated future labor income and current value of financial assets. This consumption function can be derived from the well-known life-cycle hypothesis of saving popularized by Ando and Modigliani (1963). Their analysis begins at the level of an individual consumer, whose utility depends upon his or her own consumption in current and future periods. The individual is then assumed to maximize his or her utility subject to available resources, these resources being the sum of current and discounted future earnings over the individual's lifetime. As a result of this maximization, the current consumption of the individual can be expressed as a function of his or her resources and the rate of return on capital with parameters depending upon his or her age. The individual consumption functions are then aggregated to arrive at a consumption function that is linear in income and wealth variables, as in (1a). More recently, Gali (1990) has shown that this type of aggregate consumption function can also be derived from the dynamic optimizing behavior of consumers with finite horizons and life-cycle savings. His model[5] goes one step further in establishing the presence of a common upward trend in aggregate consumption, labor income, and nonhuman wealth.

The consumption function in (1a) identifies income and wealth as the main economic determinants of aggregate, planned consumption. However, actual consumer spending in a given time period may not equal planned spending for a number of reasons. The first is the presence of adjustment costs. For example, individuals may not be able to adjust within each period their spending on housing services, given large searching, moving, and finance costs. Also, if there is considerable habit persistence in consumption behavior, then individuals may adjust their spending slowly to bring it in line with the level suggested by the economic determinants in (1a). Another reason may be the presence of liquidity-constrained individuals, who cannot smooth their consumption by borrowing against their future income due to capital market restrictions. For these individuals, actual consumer spending may be more closely tied to current labor income than is suggested by equation (1a) (Campbell and Mankiw 1989).

Given these considerations, I estimate the consumption function allowing adjustment lags in consumer spending. In particular, I postulate that actual consumer spending adjusts to the planned level with the following error correction dynamic specification (Davidson et al. 1978).

$$\Delta C_t = b_0 - b_1(C_{t-1} - \tilde{C}_{t-1}) + b_2 \Delta \tilde{C}_t + \sum_{s=1}^{k} b_{3s} \Delta C_{t-s} + \mu_t \qquad (2)$$

where C_t is actual consumer spending, μ is a disturbance term, and other variables are defined as before. In this specification, changes in current period

[5] The life-cycle model of aggregate consumption and savings in Gali (1990) is a discrete-time, quadratic-utility, open-economy version of the overlapping generations framework.

consumption depend upon changes in current period planned consumption, the gap between the last period's actual and planned consumption, and lagged actual consumption. The disturbance term μ in (2) captures the short-run influences of unanticipated shocks to actual consumer spending. The magnitude of the coefficient b_1 measures how rapidly consumers close the gap between their actual and planned consumption within each period. The larger the magnitude of b_1 (in absolute), the more rapid the adjustment. If we substitute (1b) into (2), we get the short-term dynamic consumption equation (3):

$$\begin{aligned}\Delta C_t &= b_0 - b_1(C_{t-1} - a_0 - a_1 Y_{t-1} - a_2 Y^e_{t+k-1} - a_3 W_{t-1}) \\ &\quad + b_2(a_1 \Delta Y_t + a_2 \Delta Y^e_{t+k} + a_3 \Delta W_t) + \sum_{s=1}^{k} b_{3s} \Delta C_{t-s} + \mu_t \qquad (3)\end{aligned}$$

The key feature of equation (3) is that consumption depends upon levels and first differences of the determinants of planned consumption, namely labor income and wealth.

Estimation Issues and Data Properties

The empirical estimation of equation (1a) or (3) raises several issues. One issue is that expected future labor income is not directly observable. As a result of the presence of information lags, even current-period income and wealth may not be observable (Goodfriend 1986). The simple procedure I follow here is to assume that expected future labor income is proportional to expected current labor income, meaning current and expected future income move together (Ando and Modigliani 1965). Furthermore, I also assume that current-period values of income and wealth are unknown and that planned consumption therefore depends upon their anticipated values. It is assumed that consumers form expectations about their current-period income and wealth by taking into account information known to them in period $t-1$, i.e., expectations of income and wealth are rational. Under these assumptions, one may have:

$$C_t = d_0 + d_1 Y^e_t + d_2 Y^e_{t+k} + d_3 W^e_t \qquad (1b)$$

$$Y^e_{t+k} = \beta Y^e_t;\ \beta > 0; \qquad (4)$$

$$Y_t = Y^e_t + \varepsilon_{1t} = E(Y_t/I_{t-1}) + \varepsilon_{1t} \qquad (5a)$$

$$W_t = W^e_t + \varepsilon_{2t} = E(W_t/I_{t-1}) + \varepsilon_{2t} \qquad (5b)$$

where Y_t^e is anticipated current-period labor income, W_t^e is anticipated current-period wealth, E is the expectations operator, I_{t-1} is the information set used in forming expectations of current-period income and wealth, and ε_1 and ε_2 are forecast errors assumed to be uncorrelated with time $t-1$ information. Equation (1b) is similar to equation (1a) except that it makes aggregate planned consumption depend upon anticipated current and future income and wealth variables. Equation (4) states the simplifying assumption that expected future income is proportional to current-period income. Equation (5) relates actual current-period income and wealth variables to their forecasts based on past information summarized in I_{t-1}.

If we substitute (4), (5), and (1b) into (2), we get (6), which is a short-term consumption equation containing current and lagged income and wealth variables.

$$\begin{aligned} \Delta C_t &= b_0 - b_1(C_{t-1} - d_0 - (d_1 + d_2\beta)Y_{t-1} - d_3 W_{t-1}) \\ &\quad + b_2((d_1 + d_2\beta)\Delta Y_t + d_3 \Delta W_t) + \sum_{s=1}^{k} b_{3s}\Delta C_{t-s} + v_t \end{aligned} \tag{6}$$

where

$$\begin{aligned} v_t &= u_t - b_2(d_1 + d_2\beta)\varepsilon_{1t} - b_2 d_3 \varepsilon_{2t} + (b_2 - b_1)(d_1 + d_2\beta)\varepsilon_{1t-1} \\ &\quad + (b_2 - b_1)d_3\varepsilon_{2t-1}. \end{aligned}$$

Short-term consumption equation (6) is similar to (3) in that it contains the levels as well as first differences of income and wealth variables. However, it differs in that the disturbance term v is now serially correlated. As can be seen, the disturbance term v_t in (6) is a linear combination of current and past values of three disturbance terms μ, ε_1, and ε_2. It can be verified that v_t is serially correlated, even if μ, ε_1, and ε_2 are not. In fact, v_t has first-order moving average serial correlation under the maintained assumption that the disturbance term μ and forecast errors ε_1 and ε_2 are not serially correlated.

Ordinary least squares are likely to provide biased estimates of the coefficients of the short-term consumption equation (6) because the disturbance term v_t is correlated with current-period income and wealth variables.[6] However, it can be verified that v_t is not correlated with period $t-2$ information used by consumers to forecast current-period income and wealth variables, as in (7):

[6] This correlation arises because the composite error term v_t consists of forecast errors ε_1 and ε_2, which are correlated with current-period income and wealth variables as indicated in equations (5a) and (5b).

$$E(v_t/I_{t-2}) = 0. \tag{7}$$

That suggests equation (6) can be estimated by instrumental variables, using variables in the information set I_{t-2} as instruments.[7] In particular, I estimate equation (6), using Hansen's (1982) generalized method of moments estimator (GMM). Under the identifying assumptions in (7), this procedure produces efficient instrumental variables estimates. Furthermore, the procedure generates a test of identifying restrictions used in estimating the consumption equation.[8]

Another key feature of short-term consumption equation (6) is that it contains lagged levels of consumption, income, and wealth variables, along with their first differences. If these variables have unit roots, then estimation of (6) would not yield consistent estimates of income and wealth parameters unless consumption were cointegrated with income and wealth variables (Engle and Granger 1987). I therefore examine the stochastic properties of data, investigating in particular whether there exists a cointegrating (equilibrium) relationship between consumption and its determinants, such as income and wealth.[9]

The investigation of the presence of a cointegrating relationship between consumption, labor income, and wealth is of interest for another reason. Namely, aggregate consumption equations in this article are estimated under the assumption that expected future labor income is proportional to current-period labor income. This assumption implies that expected future labor income will share the trend in current-period labor income. In the short run, however, expected future income may deviate from current. The presence of those short-term deviations implies the disturbance term that contains the omitted variable, namely expected future income, will be correlated with current income and wealth variables included in these regressions. This complication, however, has no effect on long-term estimates of income and wealth elasticities

[7] The extra lag in the instruments also helps meet several other potential objections. First, Goodfriend (1986) has noted that aggregate variables are not in individuals' information sets contemporaneously because of delays in government publication of aggregate statistics. Since such delays are typically no more than a few months, lagging instruments an extra quarter largely avoids this problem. Second, it has been suggested that those goods labeled as nondurable in the National Income Accounts are in fact durable. Durability would introduce a first order moving average term into the change in consumer expenditure (Mankiw 1982); this would not affect the procedure in the present article that uses twice-lagged instruments.

[8] The GMM procedure generates estimates of the coefficients under the identifying restrictions given in equation (7) and is thus more efficient than the standard instrumental variables procedure. It also yields a test of the identifying restrictions, enabling one to test the model adequacy.

[9] The traditional life-cycle model in Ando and Modigliani (1965) simply implies that aggregate consumption is linearly related to labor income and wealth. It says nothing about the cointegration properties of these variables. In contrast, the theoretical life-cycle model in Gali (1990) implies that consumption, income, and wealth variables may share a common trend. But whether this theoretical implication is consistent with actual data still needs to be tested, so one must test for the presence of cointegration.

that are derived using consumption equations involving cointegrated variables. The intuition behind this result is that if consumption is cointegrated with current income and wealth variables, then the residuals will be stationary and hence will have no effect on estimates of coefficients that capture correlation among trending variables. For that reason, aggregate consumption equations estimated using only current-period trending labor income and wealth variables would provide consistent estimates of long-term coefficients.[10]

Definition and Measurement of Variables

The aggregate consumption equations (1a) and (6) relate consumption to labor income and wealth. Following Ando and Modigliani (1965), I identify the income effect on consumption by including labor income (net of taxes) and identify the wealth effect by including net worth of households in the aggregate consumption function. In these specifications, wealth affects consumption directly through its market value, which provides a source of purchasing power used to iron out fluctuations in income arising from transitory developments as well as from the normal life cycle.

In order to estimate the effect of the recent stock market boom on consumption, I also consider the specification that relates consumption directly to equity wealth. As is now widely known, most of the recent increase in household wealth has been associated with the recent explosion in equity values, which are held by fewer households than many other assets. Consequently, the marginal propensity to consume out of stock market wealth may be smaller than that to consume out of total wealth.

The consumption variable implicit in standard theories of consumer behavior used to derive equation (1a) is measured as a flow rather than a stock. Expenditures on durable consumer goods are not included in the measure of consumption because they represent additions and replacements to the asset stock, whose short-term dynamics may be quite different. Hence, consumption equations have generally been estimated with consumption measured either as household spending on nondurable goods and services alone or on that item plus the (imputed) flow from the stock of durable consumer goods. Both approaches should yield similar qualitative inferences about the underlying determinants of consumption, and I have chosen to follow the first approach. However, since there is considerable interest in identifying the effect of recent stock market wealth on *actual* total consumer spending, I also estimate consumption equations with consumption measured as total consumer spending, including expenditures on durable consumer goods. I estimate only the long-term consumption equations, though, because identifying assumptions made

---

[10] This will not necessarily hold for short-term consumption equations.

Figure 1 Time Series: Real, Per Capita, and in Logs

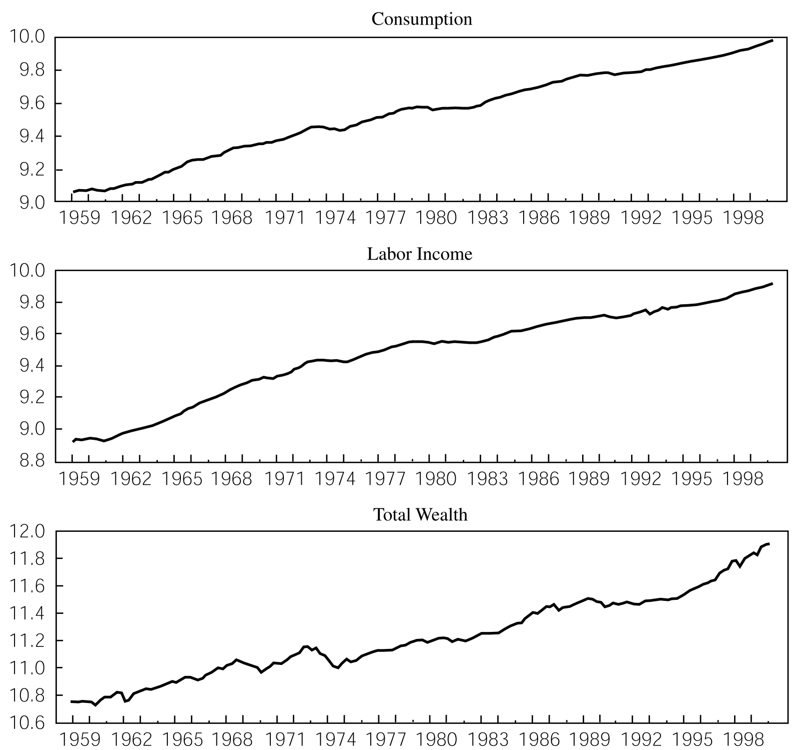

here to estimate short-term consumption equations may not be valid if the measure of consumption used includes expenditures on the stock of durable goods.[11]

To estimate consumption equations, I use standard, U.S. quarterly time series data over 1959Q1 to 2000Q2. Consumption is measured as per capita consumption of nondurables and services, in 1996 dollars (rC).[12] Labor income is measured as disposable labor income per capita, in 1996 dollars ($rDLY$). Total wealth is per capita net worth of households, in 1996 dollars

[11] If consumption includes durable goods, then the disturbance term in short-term consumption equations may follow a more complicated serial correlation pattern than the one assumed in equation (6) of the text. This has no effect, however, on estimates of the long-term consumption equations that involve trending variables. Mankiw (1982) explores some other implications of including durable goods in consumption.

[12] The consumption series is scaled up so that its sample mean matches the sample mean of total consumption. This adjustment matters if one wants to predict the level of consumer spending using consumption equations that involve measures of total labor income and wealth.

(rNW). Household stock market wealth is measured as per capita corporate equities held by them, in 1996 dollars (rEQ).[13]

Accordingly, the aggregate consumption equations investigated take the following forms.

$$rC_t = r\tilde{C} + \varepsilon_t = f_0 + f_1 rDLY_t + f_2 rNW_t + \varepsilon_t \tag{8a}$$

$$\begin{aligned} \Delta rC_t &= g_0 - \lambda(rC_{t-1} - f_0 - f_1 rDLY_{t-1} - f_2 rNW_{t-1}) \\ &\quad + g_1 \Delta rDLY_t + g_2 \Delta rNW_t + \sum_{s=1}^{k} g_{3s} \Delta rC_{t-s} + v_t \end{aligned} \tag{8b}$$

$$rC_t = r\tilde{C} + \varepsilon_t = f_0 + f_1 rDLY_t + f_{21} rEQ_t + f_{22} rNWO_t + \varepsilon_t \tag{9a}$$

$$\begin{aligned} \Delta rC_t &= g_0 - \lambda(rC_{t-1} - f_0 - f_1 rDLY_{t-1} - f_{21} rEQ_{t-1} \\ &\quad - f_{22} rNWO_{t-1}) + g_1 \Delta rDLY_t + g_{21} \Delta rEQ_t \\ &\quad + g_{22} \Delta rNWO_t + \sum_{s=1}^{k} g_{3s} \Delta rC_{t-s} + v_t \end{aligned} \tag{9b}$$

where NWO is nonequity net worth and other variables are defined as before. Equations (8a) and (9a) are the long-term consumption equations that relate the level of actual consumer spending to determinants of the level of planned consumer spending. Equation (8a) relates consumption to labor income and total wealth. Equation (9a) is similar to (8a) except that total wealth is decomposed into equity and nonequity components. The estimated coefficients that appear on $rDLY$, rNW, rEQ, and $rNWO$ in (8a) and (9a) measure long-run marginal propensities to consume out of income and wealth variables. Equations (8b) and (9b) are the short-term consumption equations that relate changes in actual consumer spending to changes in income and

[13] As indicated above, consumption is measured either as total personal consumption expenditure or as expenditure on nondurable goods and services. Labor income is defined as wages and salaries + transfer payments + other labor income − personal contributions for social insurance − taxes. Taxes are defined as [wages and salaries/ (wages and salaries + proprietors' income + rental income + personal dividends + personal interest income)] personal tax and non-tax payments. Total wealth is measured by household net worth, and stock market wealth by direct household holdings of corporate equity. The quarterly data on income and consumption are from the Department of Commerce, Bureau of Economic Analysis, and the data on household net worth and corporate equities are from DRI. The latter are based on flow of funds data. The wealth component of net worth includes direct household corporate holdings, mutual funds holdings, holdings of private and public pension plans, personal trusts, insurance companies, and other nonfinancial assets. Household equity wealth is, however, direct holdings of corporate equity. The nominal labor income is deflated using the deflator for personal disposable income, and nominal wealth variables using the deflator for personal consumption expenditure.

wealth variables, allowing for the presence of lags in the adjustment of actual consumption to planned consumption.

The consumption equations discussed above are linear in levels of variables. The coefficients that appear in these equations measure the effects on consumption of a dollar increase in income and wealth variables. This specification is appropriate if these series—consumption, income, and wealth—follow homoscedastic linear processes in levels, with or without unit roots. However, aggregate time-series data on these variables appear to be closer to linear in logs of variables than levels. This can be seen in Figure 1, which charts time series for logs of real (per capita) consumption, labor income, and total wealth. In view of the apparent superiority of log modification, the consumption equations here are estimated using logs of variables, so that estimated coefficients are elasticities. The implied level responses are then backed out using the consumption-income and consumption-wealth ratios evaluated at their sample mean values. (Hereafter, lowercase letters denote log variables.)[14]

2. EMPIRICAL RESULTS

Unit Roots Test Results

Figure 1 clearly indicates that the time series for (the logs of) consumption, labor income, and wealth are nonstationary. In order to test whether these variables are stationary around a linear trend or have stochastic trends, I perform tests for the presence of unit roots. Since unit-root tests are sensitive to the presence of a linear trend, I first investigate whether these series possess a linear trend at all. The presence of a linear trend is investigated by running the following regression,

$$\Delta x_t = c_0 + \sum_{s=1}^{k} c_{1s}\Delta x_{t-s} + \eta_t, \tag{10}$$

where x stands for the pertinent variable. The variable x has a linear trend if the t-statistic for the coefficient that appears on the constant in (10) is large. Panel A in Table 1 presents t-statistics for rc, $rdly$, rnw, req, and $rnwo$. The lag length k is chosen by the Akaike (1973) information criterion (AIC). Those test results indicate that the real consumer spending, labor income, and nonequity net worth (rc, $rdly$, $rnwo$) have a linear trend, whereas other remaining variables do not. The results further indicate that unit root tests

[14] Campbell and Mankiw (1990) and Ludvigson and Steindel (1999) also estimate aggregate time-series consumption equations that are linear in logs of variables.

Table 1 Tests for Trend and Unit Roots

Series	Panel A t-statistic	Panel B Dickey-Fuller Lag	t-test	Phillips-Perron Z_t
x	(1)	(2.1)	(2.2)	(3)
rc	3.8*	4	−2.6	−1.9
$rdly$	4.2*	1	−1.3	−1.2
rnw	0.8	0	−0.9	−1.3
req	−1.4	0	−0.5	−0.2
$rnwo$	1.7*	4	−2.6	−1.7
Δrc		4	−5.4*	−9.3*
$\Delta rdly$		0	−13.3*	−13.4*
Δrnw		0	−11.6*	−11.7*
Δreq		0	−11.8*	−11.8*
$\Delta rnwo$		0	−10.1*	−10.9*

* Indicates significance at the 5 percent level.

Notes: All variables are in their natural logarithms and on a per capita basis: rc is real consumer spending on nondurable goods and services; $rdly$ is real disposable labor income; rnw is real household net worth; req is real value of corporate equities held by households; $rnwo$ is real household net worth excluding corporate equities; and Δ is the first difference operator.

The t-statistic in column (1) tests the null hypothesis that the constant term is zero in a least-square regression of Δx on a constant, linear trend, and its own eight lagged values. If this t-value is large, then the series x has a linear trend. The t-test in column (2.2) is the standard augmented Dickey-Fuller test of a unit root. The optimal lag length used in performing the augmented Dickey-Fuller test is selected by the Akaike information criterion and is reported in column (2.1). The statistic Z_t in column (3) is the standard Phillips-Perron test of a unit root. The Phillips-Perron test reported allows for the presence of serial correlation and heterogeneity in the disturbance term. The unit root tests reported above allow for the presence of a linear trend in the pertinent series.

discussed below should be performed allowing for the presence of a linear trend, in at least some of these variables.

The presence of unit roots is investigated using Dickey-Fuller (DF) and Phillips-Perron (PP) tests.[15] The DF test is performed with the following regression.

$$x_t = d_0 + d_1 TR_t + d_3 x_{t-1} + \sum_{s=1}^{k} d_{4s} \Delta x_{t-s} + \eta_t \qquad (11)$$

where TR is a linear trend. The variable x has a unit root if $d_3 = 1$ in (11). The DF t-statistic tests the hypothesis $d_3 - 1 = 0$. The DF test above relies on the

[15] See Dickey and Fuller (1979) and Phillips and Perron (1988).

Table 2 Residual Based Tests for Cointegration

Cointegrating Regression Regress rc on	Optimal Lag*	Dickey-Fuller t-test	Phillips-Ouliaris Z_t	Critical Values 5%	10%
$(C, TR, rdly, rnw)$	1	−3.34	−4.30	−4.16	−3.84
$(C, TR, rdly, req, rnwo)$	1	−4.37	−4.53	−4.49	−4.08

*Optimal lag is chosen using the Akaike (1973) information criterion.

Notes: All variables are defined as in Table 1. Dickey-Fuller and Phillips-Ouliaris test statistics are applied to the fitted residuals from the cointegrating regressions reported above in Table 2. The optimal lag length chosen in implementing the Dickey-Fuller test is selected by the Akaike information criterion. Eight lags are used in constructing the covariance matrix used to construct the Phillips-Perron test. Critical values reported above are from Tables IIb and IIc in Phillips and Ouliaris (1990).

assumption that the disturbance term is a finite order autoregressive process (AR). The PP test, however, does not rely on a finite order AR, but instead employs a correction for general order serial correlation and heteroskedasticity, based in part on the spectral representation of the disturbance sequence at frequency zero. The PP test statistic Z_t tests $d_3 - 1 = 0$, and it has the same limiting distribution as DF.

Panel B in Table 1 presents DF and PP tests for the presence of a unit root in variables rc, $rdly$, rnw, req, and $rnwo$. In the performance of the DF test, the lag k chosen by AIC is one. Eight autocovariance terms are used in the performance of the PP test.[16] As can be seen in Panel B of Table 1, the DF test results are consistent with the presence of a unit root, suggesting that these variables are not stationary around a linear trend. Panel B also presents DF and PP tests for the presence of a unit root in first differences of the same variables. Those test statistics indicate that first differences do not have a unit root, implying that these variables follow a first order integrated ($I(1)$) process. The unit-root test results thus suggest that we can carry tests for cointegration and thereby assess whether real consumer spending is cointegrated with its determinants, such as labor income and wealth.

Test Results for Cointegration

Table 2 reports test statistics corresponding to the Phillips-Ouliaris (1990) residual-based cointegration tests. I apply DF and PP unit root tests to the

[16] PP unit root tests using four autocovariance terms yield similar results.

Table 3 Dynamic OLS Estimates of the Aggregate Consumption Equation

Panel A: $rc_t = f_0 + f_1 rdly_t + f_2 rnw_t + f_3 TR_t$

	Elasticities		Implied Level Response Coefficients	
Sample Period	f_1	f_2	$rDLY$	rNW
1960Q2–2000Q2	0.51(32.1)	0.14(4.1)	0.62	0.03
1960Q2–1995Q2	0.49(24.9)	0.23(7.5)	0.60	0.04
1960Q2–1990Q2	0.48(20.0)	0.22(7.5)	0.57	0.04

Panel B: $rc_t = f_0 + f_1 rdly_t + f_{21} req_t + f_{22} rnwo_t + f_3 TR_t$

	Elasticities			Implied Level Response Coefficients		
Sample Period	f_1	f_{21}	f_{22}	$rDLY$	rEQ	$rNWO$
1960Q2–2000Q2	0.49(15.6)	0.02(4.1)	0.15(3.6)	0.59	0.03	0.03
1960Q2–1995Q2	0.43(14.7)	0.03(6.1)	0.21(7.1)	0.52	0.04	0.05
1960Q2–1990Q2	0.45 (15.1)	0.03(5.9)	0.12(1.4)	0.54	0.04	0.03

Notes: The coefficients with t-values in parentheses reported above are elasticities, estimated using dynamic ordinary least squares (Stock and Watson 1993). The t-values have been corrected for the presence of serial correlation. The consumption equations are estimated including four leads and lags of the first differences of right-hand-side explanatory variables $rdly$, rnw, req, and $rnwo$.

Implied level response coefficients are backed out by multiplying estimated elasticities with their relevant consumption-income or consumption-wealth ratios evaluated at their respective sample means. A level response coefficient measures the effect of a *dollar* increase in the relevant variable on consumption.

The uppercase variables are in levels and lowercase in their logs. All variables are defined as in Table 1.

residuals of long-term consumption equations (8a) and (9a). If real consumer spending is cointegrated with income and wealth variables, then the error term that appears in consumption equations is stationary, and one may reject the hypothesis of a unit root in ε_t. As the table shows, DF and PP tests generally reject the null hypothesis of unit root in ε_t, indicating that real consumer spending is cointegrated with labor income and wealth specified alternatively in (8a) and (9a). This is also shown in Figures 2 and 3, which chart actual consumption, the values predicted by the cointegrating regression (which are a measure of planned consumption), and the residuals from the cointegrating regression (which are a measure of the gap between actual and planned consumption). Figure 2 uses the aggregate consumption equation with total wealth and Figure 3 uses it with equity wealth (see Panels A and B, Table 3). As is evident in the figures, actual consumption moves with planned

Figure 2

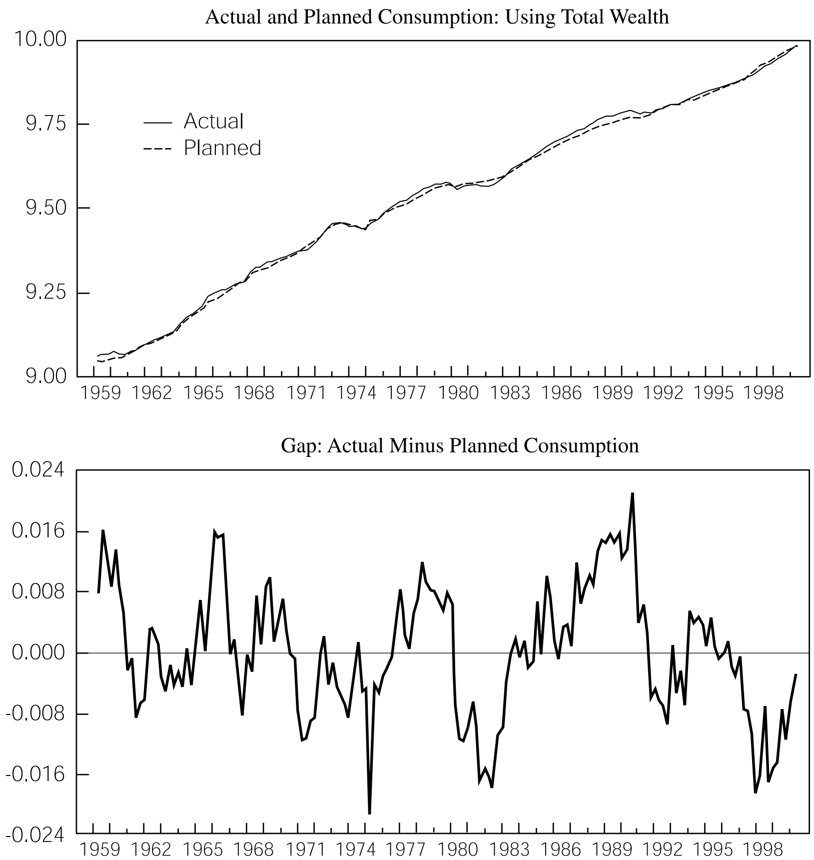

consumption over time and the gap between them is stationary during the sample period studied.

Long-term Consumption Equations: Marginal Propensities to Consume out of Income and Wealth

The cointegration test discussed above implies that consumption equations (8a) or (9a) with variables in log levels can provide reliable inferences about the long-term influences of income and wealth on consumption. These consumption equations can be estimated superconsistently by ordinary least squares, despite the fact that the expected future labor income variable is not explicitly included in these equations. However, statistical inferences cannot be carried out using the conventional standard errors since the resulting parameter esti-

Figure 3

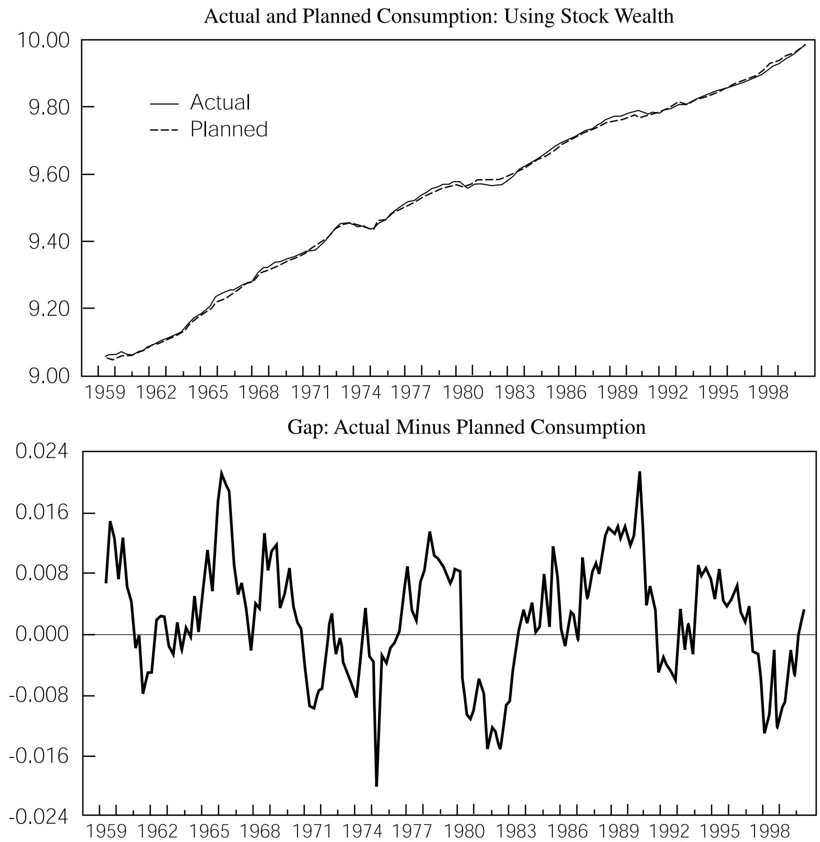

mates have nonstandard distributions. I therefore estimate consumption equations here using Stock and Watson's (1993) dynamic ordinary least squares (DOLS), which includes leads and lags of the differences in the right-hand-side variables as additional regressors.[17] Table 3 presents the dynamic OLS estimates of the aggregate consumption equation; Panel A reports estimates for consumption equation (8a) and Panel B does so for consumption equation (9a). I present estimates for the full sample period 1960Q2 to 2000Q2 as well as for two other subperiods ending in the 1990s, 1960Q2 to 1990Q4 and 1960Q2 to 1995Q4.[18]

[17] The standard errors reported have been corrected for the presence of serial correlation and heteroskedasticity.

[18] The estimation period begins in 1960Q2, prior observations being used to allow for lagged values of income and wealth variables in consumption equations.

If we focus on full sample estimates, we can see that labor income and wealth variables have theoretically expected signs and are significant in estimated consumption equations. The point estimate of the long-term elasticity of consumption with respect to labor income is 0.51. The point estimate of the implied level response is 0.62, suggesting that the consumption effect of a dollar increase in labor income is about 62 cents.[19] The point estimate of the long-term elasticity of consumption with respect to total wealth is 0.14. The point estimate of the implied level response is 0.03, indicating that a $1 increase in wealth raises consumer spending by 3 cents. The estimates also show that the elasticity of consumption is substantially smaller with respect to equity wealth than with respect to nonequity wealth, 0.02 versus 0.15 (see estimates in Panel B, Table 3). The implied level responses, though, do not differ and both equal 0.03, indicating that a $1 increase in equity values raises consumer spending by 3 cents. These estimates of the marginal propensities to consume out of income and wealth variables are not out of line with estimates in some other recent studies. For example, in 1996 estimates from the FRB/US quarterly model placed the marginal propensity to consume at 0.03 for stock wealth and 0.075 for other net wealth (Brayton and Tinsley 1996).

Estimates for two other subperiods ending in the early 1990s suggest similar inferences about aggregate consumption behavior. Income and wealth variables remain significant in estimated consumption equations. The elasticity of consumption is smaller with respect to equity wealth than with respect to nonequity wealth, but the implied level responses are not very different from each other and did not change much during the 1990s. In particular, the point estimate of the marginal propensity to consume out of equity values is 0.04 for these two subperiods, close to 0.03 estimated using the full sample period.

Poterba (2000) has suggested that the marginal propensity to consume out of wealth may have declined in the 1990s. Much of the rise in total wealth reflects the growing importance of equities. And, since the marginal propensity to consume out of equity wealth is smaller than that for wealth as a whole, the rising share of equities in total wealth pulls down the overall marginal propensity to consume in the latter. The estimates here support Poterba's conjecture (see Table 3). The parameter that measures the elasticity of consumption with respect to total wealth does show a decline during the 1990s: its point estimate declines to 0.14 in 2000 from 0.22 in 1990. However, the parameter that measures the elasticity of consumption with respect to equity values did not decline much during that decade; its point estimate of 0.03 in 2000 is close to 0.04 in 1990. Together these estimates suggest that

[19] As shown in Gali (1990, p. 439), the presence of finite horizon and life cycle savings implies the marginal propensity to consume out of labor income will be less than one. However, in the absence of life-cycle savings as in an infinite-horizon consumption model, the marginal propensity to consume is unity. The exact magnitude of the income coefficient in a life-cycle model, however, depends upon, among other things, the age structure of population and the relative distribution of income and net worth over the different age groups (Ando and Modigliani 1965).

Table 4 GMM Estimates of the Short-term Aggregate Consumption Equation

Panel A: $\Delta rc_t = g_0 + \lambda(rc_{t-1} - \tilde{rc}_{t-1}) + g_1 \Delta rdly_t + g_2 \Delta rnw_t + \sum_{s=1}^{3} g_{3s} \Delta rc_{t-s}$
where $\tilde{rc}_t = f_0 + f_1 rdly_t + f_2 rnw_t + f_3 TR_t$

Sample Period	λ	g_1	g_2	$\sum_{s=1}^{3} g_{3s}$	SER	*J-test*(S)
1961Q2–2000Q2	−0.15	0.33	0.06	0.39	0.00390	29.4
	(4.5)	(4.4)	(1.7)	(4.7)		(0.49)
1961Q2–1995Q2	−0.24	0.30	0.08	0.41	0.00393	28.1
	(5.2)	(4.4)	(2.8)	(5.3)		(0.56)
1961Q2–1990Q2	−0.26	0.29	0.11	0.39	0.00405	23.1
	(5.3)	(4.6)	(3.7)	(4.8)		(0.81)

Panel B: $\Delta rc_t = g_0 + \lambda(rc_{t-1} - \tilde{rc}_{t-1}) + g_1 \Delta rdly_t + g_{21} \Delta req_t + g_{22} \Delta rnwo_t + \sum_{s=1}^{3} g_{3s} \Delta rc_{t-s}$
where $\tilde{rc}_t = f_0 + f_1 rdly_t + f_{21} req_t + f_{22} rnwo_t + f_3 TR_t$

Sample Period	λ	g_1	g_{21}	g_{22}	$\sum_{s=1}^{3} g_{3s}$	SER	*J-test*(S)
1961Q2–2000Q2	−0.17	0.32	0.01	0.10	0.33	0.00533	21.3
	(4.2)	(3.6)	(2.0)	(1.5)	(2.8)		(0.32)
1961Q2–1995Q2	−0.24	0.21	0.01	0.15	0.38	0.00413	19.7
	(5.1)	(2.9)	(2.0)	(2.1)	(3.6)		(0.41)
1961Q2–1990Q2	−0.26	0.26	0.01	0.10	0.39	0.00419	18.6
	(5.1)	(3.3)	(2.4)	(1.4)	(3.5)		(0.48)

Notes: The coefficients (with t-values in parentheses) reported above are GMM estimates of the short-term aggregate consumption equation. The instruments used consist of a constant, eight twice-lagged values of changes in real consumer spending Δrc_{t-2}, real disposable labor income $\Delta rdly_{t-2}$, real disposable property income, real household net worth ΔrNW_{t-2} (or its two components), and short-term nominal interest rate, and one twice-lagged value of the error correction variable $rc_{t-2} - \tilde{rc}_{t-2}$. SER is the standard error of the regression, and *J-test* (with Significance level in parentheses) is the test of overidentifying restrictions used in estimating the consumption equation and is distributed Chi-squared χ^2 (Hansen 1982).

most of the decline observed in the elasticity of consumption with respect to total wealth in the 1990s may just reflect the changing mix of wealth stocks during this decade.

Short-term Consumption Equations

Table 4 reports estimates of short-term consumption equations (8b) and (9b). Consumption equations are estimated for the full sample period as well as for two subperiods ending in the early 1990s. If we focus on full sample estimates, we can see that all estimated coefficients appear with theoretically expected signs and are statistically significant. In particular, the estimated coefficient that appears on the lagged value of the error correction variable is negative, indicating the presence of adjustment lags in consumer spending.[20] The estimated coefficients that appear on current period income and wealth variables are statistically significant, showing that consumer spending does respond to current period changes in income and wealth. The estimates also indicate that the short-term elasticity of consumption with respect to wealth differs across wealth stocks. The short-term elasticity is smaller with respect to changes in equity values than with respect to changes in total or nonequity net worth.

If we examine subperiod estimates, we reach similar conclusions about the nature of short-term consumption behavior. Income and wealth variables continue to be significant in estimated consumption equations (see Table 4). The point estimate of the short-term elasticity is substantially smaller with respect to changes in the market value of equities than with respect to changes in nonequity wealth. The estimates show too that short-term elasticities of consumption with respect to changes in wealth variables did not change much during the 1990s. In particular, the point-estimate of the short-term elasticity to consume out of current-period equity values remained around 0.01 during much of that decade.[21]

Quantifying the Stock Market Wealth Effect on Consumption

The empirical work in the above sections finds that the long-term marginal propensity to consume out of equity values appears to be quite small, with a $1 increase in the market value of corporate equities raising consumer spending only by 3 to 4 cents. However, the increase in household wealth generated by the recent explosion in equity values has been large. In order to quantify the effect of the recent stock market boom on consumer spending, I now derive estimates of consumption growth that could be attributed to increase in equity values. I focus on the 1990s and derive estimates of stock-market-induced consumption growth, using the following long-term consumption equation:

[20] The result that the error correction variable is significant in short-term consumption equations implies that past income, stock market wealth, and other wealth are useful in predicting current period changes in consumption. This finding is similar in nature to the one in Hall (1978), who finds past changes in stock prices are significant in predicting changes in current consumption.

[21] GMM estimation of short-term consumption equations use twice-lagged values of instruments. I get qualitatively similar results if instead one-period lagged instruments are used in estimation.

$$rc_t = f_0 + f_1 rdly_t + f_{21} req_t + f_{22} rnwo_t + \varepsilon_t.$$

In particular, estimates of equity-induced consumption growth over one-year horizons are calculated using the component $f_{21}\Delta req_t$. Since one is interested in quantifying the effect of stock market wealth on *actual* total consumer spending, the long-term consumption equation is re-estimated using total consumer spending, not just spending on nondurable goods and services. As indicated before, the definition of consumption used in the life-cycle hypothesis of saving should include consumption of the stock of consumer durable goods, not expenditures on their purchase.

I first present evidence in Table 5 indicating that long-term consumption equations estimated using total consumer spending provide reasonable estimates of long-term elasticities of consumption with respect to income and wealth variables. The test for cointegration continues to indicate the presence of an equilibrium relationship between consumer spending and its economic determinants, such as labor income and wealth (see Table 5). The estimates of long-term elasticities indicate that income and wealth variables have significant effects on real consumer spending. In particular, the point estimate of the long-term marginal propensity to consume out of equity values is small and has remained in a narrow 0.04 to 0.06 range in the 1990s. That estimated range is quite close to the range generated using consumer spending on nondurable goods and services.

Table 6 presents the quantitative estimate of wealth-induced consumption growth. It makes clear that the part of consumption growth that can be explained by a rise in equity values has not been trivial. Between 1990 and 1999, stock-market-induced consumption growth ranged from 0.6 to 2.1 percentage points per year. Over 1995 to 1999, real consumer spending increased at an average annual rate of about 4.0 percent per year, and the part that can be explained by stock market wealth averaged 1.5 percent per year. This wealth effect would represent an increment to the growth rate of real GDP of about 1 percentage point per year. The wealth effect is even stronger if we consider the effect on consumption of increase in total wealth, not just equity values. The estimates in Table 6 indicate that total wealth effect may have added to the growth rate of real GDP about 2 percentage points per year over 1995 to 1999. These calculations illustrate that even with relatively low estimates of the marginal propensity to consume out of stock market wealth, the consumption effect of the 1990s stock market boom is substantial.

3. CONCLUDING OBSERVATIONS

I have used cointegration and error correction methodology to estimate aggregate consumption equations that relate consumer spending either to la-

Table 5 Results Using Total Real Consumer Spending

Panel A: Residual Based Tests for Cointegration

Regress rC on	Optimal Lag*	Dickey-Fuller t-test	Phillips-Ouliaris Z_t	Critical Values 5%	10%
$(C, TR, rdly, rnw)$	1	−3.54	−4.65	−4.16	−3.84
$(C, TR, rdly, req, rnwo)$	1	−4.60	−4.81	−4.49	−4.08

Panel B: Dynamic OLS Estimates of the Aggregate Consumption Equation

$$rc_t = f_0 + f_1 rdly_t + f_2 rnw_t + f_3 TR_t$$

Sample Period	Elasticities f_1	f_2	Implied Level Response Coefficients $rDLY$	rNW
1960Q2–2000Q2	0.54	0.21	0.65	0.04
	(23.6)	(5.4)		
1960Q2–1995Q2	0.52	0.31	0.62	0.06
	(18.3)	(6.8)		
1960Q2–1990Q2	0.50	0.29	0.59	0.05
	(14.9)	(6.6)		

$$rc_t = f_0 + f_1 rdly_t + f_{21} req_t + f_{22} rnwo_t + f_3 TR_t$$

Sample Period	Elasticities f_1	f_{21}	f_{22}	Implied Level Response Coefficients $rDLY$	rEQ	$rNWO$
1960Q2–2000Q2	0.48	0.03	0.22	0.58	0.04	0.05
	(11.1)	(3.9)	(3.9)			
1960Q2–1995Q2	0.41	0.04	0.28	0.62	0.06	0.05
	(9.8)	(5.1)	(6.8)			
1960Q2–1990Q2	0.43	0.04	0.23	0.51	0.06	0.05
	(11.2)	(5.2)	(1.7)			

*See Table 2.

Notes: The estimates reported above use total consumer spending as the measure of aggregate consumption. For other details see Notes in Tables 1, 2, and 3.

bor income and total net worth or to labor income, corporate equities, and nonequity net worth. The results indicate that while wealth has a significant effect on consumer spending, the long-term marginal propensity to consume out of wealth is small. The results also show that the long-term marginal propensity to consume out of equity wealth did not change very much during the 1990s, with its point estimates staying in a narrow 0.03 to 0.04 range. But even with relatively low estimates of the marginal propensity to consume out of equity wealth, the consumption effects of the 1990s stock market boom

Table 6 Wealth-Induced Consumption Growth

Year (Q4 to Q4)	Actual Consumption Growth	Predicted Consumption Growth	Stock-Market-Induced Cons. Growth	Total Wealth-Induced Cons. Growth
(1)	(2)	(3)	(4)	(5)
1990	0.6	0.8	0.6	−0.1
1991	0.4	2.4	2.1	1.8
1992	4.2	4.9	1.4	0.9
1993	3.3	3.0	1.5	1.4
1994	3.5	2.3	0.8	1.1
1995	2.7	2.6	1.7	2.8
1996	3.1	3.4	1.3	2.4
1997	4.0	5.1	1.6	3.2
			(1.5*, 1.0**)	(2.8*, 1.9**)
1998	4.8	5.9	1.2	2.8
1999	5.4	5.5	1.5	2.8

*Mean value of equity wealth-induced consumption growth over 1995 to 1999.

**Mean value of the Increment to the Growth Rate of Real GDP over 1995 to 1999. It is simply two-thirds of wealth-induced consumption growth.

Notes: The predicted values in columns (3), (4), and (5) are generated using rolling regression estimates of the long-term consumption equation $rc_t = f_0 + f_1 rdly_t + f_{21} req_t + f_{22} rnwo_t + f_3 TR_t$ over sample periods that all begin in 1960Q2 but end in the year shown in column (1). The stock-market-induced consumption growth is $f_{21} req_t$ and total wealth-induced consumption growth is $f_{21} req_t + f_{22} rnwo_t$. Actual and predicted values are annualized rates of growth of total real consumer spending over Q4 to Q4 periods ending in the year shown.

are substantial. The estimates in this article indicate that the wealth effect may have added on average 1 to 2 percentage points per year to the growth rate of real GDP in the second half of the 1990s, which is in keeping with Greenspan's testimony (2000a,b).

The short-term consumption equations estimated here indicate that consumption does respond to current-period changes in income and wealth. However, consumption is also correlated with lagged values of income and wealth variables; this result implies that short-term swings in household wealth generated by changing equity values could lead to short-term swings in consumer spending.

The empirical work here covers the sample period 1959Q1 to 2000Q2. The data for the third and fourth quarters of 2000 are now available and indicate that equity wealth during that year declined about 18 percent, after rising at an average annual rate of 18 percent per year during the preceding five-year period (1995–1999). The main prediction of the consumption equations that I have estimated here is that this decline in equity wealth is likely to depress consumer spending. The results indicate that the growth rate of consumer spending in the short term is likely to fall below the robust 4 percent yearly growth rate observed during the preceding five-year period.

The empirical work discussed in this article supports the presence of a significant wealth effect on consumption. However, some caveats are in order. These results indicate that estimates of the wealth elasticity were not stable during the 1990s. Also, the short-term consumption equations are estimated including variables suggested by the life-cycle model. The empirical work leaves open the question of whether consumption may be influenced by some additional factors such as consumer confidence, energy prices, and short-term interest rates.

REFERENCES

Akaike, H. “Information Theory and the Extension of the Maximum Likelihood Principle,” in *Second International Symposium on Information Theory,* B.N. Petrov and F. Csaki, eds., Budapest: Akademia Kiado, 1973.

Ando, Albert, and Franco Modigliani. “The ‘Life Cycle’ Hypothesis of Saving: Aggregate Implications and Tests,” *American Economic Review,* vol. 53 (March 1963), pp. 55–84.

Brayton, Flint, and Eileen Mauskopf. “Structure and Uses of the MPS Quarterly Econometric Model of the United States,” *Federal Reserve Bulletin,* vol. 73 (February 1987), pp. 93–109.

Brayton, Flint, and Peter Tinsley, eds. “A Guide to the FRB/US: A Macroeconomic Model of the United States,” Finance and Economics Discussion Working Paper 42, Board of Governors of the Federal Reserve System, 1996.

Campbell, John Y., and Gregory Mankiw. “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence,” NBER Working Paper 2924, May 1990.

Davidson, James E. H., David F. Hendry, Frank Srba, and Stephen Yeo. “Econometric Modelling of the Aggregate Time-Series Relationship Between Consumers’s Expenditure and Income in the United Kingdom,” *The Economic Journal,* vol. 88 (December 1978), pp. 661–92.

Dickey, D. A., and W. A. Fuller. “Distribution of the Estimators for Autoregressive Time Series with a Unit Root,” *Wayne Journal of the American Statistical Association,* vol. 74 (June 1979), pp. 427–31.

Engle, Robert F., and C. W. J. Granger. “Co-Integration and Error Correction: Representation, Estimation, and Testing.” *Econometrica*, vol. 55 (March 1987), pp. 251–76.

Gali, Jordi. “Finite Horizons, Life-Cycle Savings, and Time-Series Evidence on Consumption,” *Journal of Monetary Economics*, vol. 26 (December 1990), pp. 433–52.

Goodfriend, Marvin. “Information-Aggregation Bias: The Case of Consumption.” Mimeo, Federal Reserve Bank of Richmond, 1986.

Greenspan, Alan. Testimony before the Committee on Banking and Financial Services of the U.S. House of Representatives, February 17, 2000a.

__________. Testimony before the Committee on Banking, Housing, and Urban Affairs of the U.S. Senate, July 20, 2000b.

Hall, Robert E. “Stochastic Implications of the Life Cycle–Permanent Income Hypothesis: Theory and Evidence,” *Journal of Political Economy*, vol. 86 (December 1978), pp. 971–87.

Hansen, Lars Peter. “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica*, vol. 50 (July 1982), pp. 1029–54.

Ludvigson, Sydney, and Charles Steindel. “How Important Is the Stock Market Effect on Consumption?” Federal Reserve Bank of New York *Economic Policy Review*, vol. 5 (July 1999), pp. 29–52.

Mankiw, N. Gregory. “Hall’s Consumption Hypothesis and Durable Goods,” *Journal of Monetary Economics*, vol. 10 (Novermber 1982), pp. 417–25.

Phillips, Peter C. B., and S. Ouliaris. “Asymptotic Properties of Residual Based Tests for Cointegration,” *Econometrica*, vol. 58 (January 1990), pp. 165–93.

__________, and Pierre Perron. “Testing for a Unit Root in Time Series Regression,” *Biometrika*. vol. 75 (June 1988), pp 335–46.

Poterba, James M. “Stock Market Wealth and Consumption,” *Journal of Economic Perspectives*, vol. 14 (Spring 2000), pp. 99–118.

Stock, James H., and Mark W. Watson. “A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems,” *Econometrica*, vol. 61 (July 1993), pp. 783–820.

Using the Federal Funds Futures Market to Predict Monetary Policy Actions

Raymond E. Owens and Roy H. Webb

Changes in interest rates directly affect anyone who borrows or lends. A benchmark interest rate is the federal funds rate, the monetary policy instrument of the Federal Reserve System ("The Fed"). The federal funds rate serves as an anchor for the financial system, and other interest rates key off its current level and expected changes in it. Accurate predictions of changes in the federal funds rate are, therefore, of great value to persons engaged in a wide variety of business activities.

Forecasting interest rates during the last few decades has been especially difficult. Over that period, the economy has been rocked by a number of macroeconomic shocks that have had substantial impacts on interest rates. Equally difficult for analysts has been the challenge of accurately anticipating monetary policy actions in a setting in which the monetary policy process has been opaque. In recent decades, monetary policy goals have been numerous and on occasion contradictory, and policy has generally followed discretion rather than a set of clear, consistent rules. Clarity has also been limited by institutional practices that have provided incomplete information on monetary policy decisions to the public. Prior to 1994, for example, the Federal Open Market Committee (FOMC) used an operating procedure that targeted borrowed reserves and yielded a federal funds rate objective that was difficult to elucidate even well after the fact (Cook, 1989). In addition, during that period the FOMC chose not to immediately reveal its policy decision or its inclination regarding near-term future policy actions at the conclusion of its meetings, leaving financial market participants to guess the action taken.

This article originally appeared in *Business Economics, The Journal of the National Association for Business Economics* (vol. 36, April 2001), pp. 44–48. The views and opinions expressed in this article are solely those of the authors and do not necessarily represent those of the Federal Reserve Bank of Richmond or of the Federal Reserve System.

Beyond these factors, at least until October 1988, the specific sources of changes in short-term interest rate forecasts also were often difficult to identify because financial market forecasters often relied on the yields on short-term Treasury securities as their benchmark for short-term interest rates. Although changes in these rates were often affected by anticipated Fed policy actions, interest rate movements were also affected by changes in expected inflation, Treasury refunding plans, and other variables. These factors could lead to a highly variable spread between rates on short Treasuries and those on federal funds. As a result, a change in interest rates could arise from sources other than monetary policy actions, and no independent means was usually available to decompose the change into the impacts from the individual factors. This situation changed to some extent in 1988 when the Chicago Board of Trade (CBOT) began trading 30-day Federal Funds Futures, a contract based on the average monthly federal funds interest rate, the Fed's monetary policy instrument. This contract has been widely interpreted as an unbiased forecast of the expected interest rate on federal funds and has been considered a useful tool in identifying the impact of anticipated changes in monetary policy on interest rates. Of course, this estimate does not necessarily move in lockstep with expected movements in interest rates on other short-term securities, because of the other factors often embedded in those rates.

In this article, we review the development and basic mechanics of the Federal Funds Futures market. Following this description, we show that efforts to assess the usefulness of this market as a predictor of subsequent Fed monetary policy actions have generally supported the value of this tool. Our new look at the market emphasizes that the Federal Funds Futures market provides a valuable forecasting tool to the public at a nearly zero cost—namely an unbiased, reasonably accurate forecast of the future federal funds rate changes by the FOMC.

1. THE FEDERAL FUNDS FUTURES MARKET

Federal Funds Futures contracts began trading on the floor of the Chicago Board of Trade in October 1988. This event signaled the beginning of essentially public, market-based forecasts of future interest rates on federal funds. The traditional price-discovery mechanism of futures markets thus began to provide outside observers with the basic knowledge needed to construct informed forecasts of FOMC target changes. There are several steps involved in processing the market quotes, however, and at this point it will be helpful to review the specifics of the contract.

The contract traded is, of course, a well-defined instrument, and identifying changes in the federal funds rate embedded in the contract prices requires some simple arithmetic. First, though, are the basics of the contract. Federal Funds Futures contracts are traded for the current month and for future months—effectively about six or seven months out. The contracts are for the interest paid on a principal amount of $5 million of overnight federal funds held for thirty days and are priced on the basis of 100 minus the overnight federal funds rate for the delivery month. A 7.25 yield, for example, equals a price of 100 minus 7.25, or 92.75. For settlement purposes, the contract is to be compared to the average daily federal funds effective rate as reported by the Federal Reserve Bank of New York.

An additional feature of the contracts is that their pricing information is widely available in a timely fashion. The closing prices from the previous day's trading are quoted in the financial pages of most major newspapers. Moreover, nearly real time quotes are available on the Internet, with the CBOT's website being a reliable source.[1]

In the four months of each year in which there is no meeting of the FOMC—and assuming inter-meeting changes in the funds rate are not anticipated—the contracts' prices represent the expected federal funds target rate previously announced by the FOMC, after accounting for small deviations such as "misses" by the Fed's trading desk or special liquidity premiums that may exist in the market. In these months, the estimate of the federal funds rate should differ from the actual rate only by the misses. In contrast, for each of the eight months in which the FOMC meets, calculating the expected federal funds rate is slightly more complicated. In these months, the expected average for the period represents a weighted average of the federal funds rate before the FOMC meeting and the rate expected after the meeting. When rates are expressed in percentages, this is equivalent to:

$$i^f_{t,h} = \frac{k i^e_{t+h} + (m-k) i^{\hat{e}}_{t+h}}{m} \qquad (1)$$

where $i^f_{t,h}$ is the Federal Funds Futures contract rate at time t for h periods ahead, i^e_{t+h} is the expected federal funds rate leading up to the FOMC meeting k days into the month, $i^{\hat{e}}_{t+h}$ is the estimate of the funds rate after the meeting, and there are m days in the month of the FOMC meeting.

The expected federal funds rate after the FOMC meeting can be derived as:

$$i^{\hat{e}}_{t+h} = \frac{m i^f_{t,h} - k i^e_{t+h}}{m-k} \qquad (2)$$

This expected federal funds rate can be interpreted as a forecast of the of the FOMC's target rate subsequent to the meeting. It is often useful to convert this forecast to an anticipated probability that the FOMC changes its target rate. We can derive that anticipated probability by adding assumptions that

[1] The CBOT's quotes (with a ten minute delay) may be found at www.cbot.com/cbot/quotes/fin_futures/0,1860,FF,00.html.

we believe are generally realistic. First, assume that the FOMC changes rates only at scheduled meetings. That has been a good assumption since 1994—the FOMC has changed its target between meetings on only four occasions (April 18, 1994, October 15, 1998, January 3, 2001, and April 19, 2001). Second, assume that the FOMC chooses between no change in its target and a change of amount delta. Then by definition (and suppressing the subscripts for clarity):

$$i^{\hat{e}} = p(i^e + \Delta \mathrm{i}^T) + (1 - p)i^e \tag{3}$$

where $\Delta \mathrm{i}^T$ is the expected change in the target rate and p is the anticipated probability that the FOMC changes its target. This can be solved for p, yielding

$$\mathrm{P} = \frac{100(i^{\hat{e}} - i^e)}{\Delta \mathrm{i}^T} \tag{4}$$

where, again, the subscripts are suppressed for clarity. This calculation thus extracts the probability of a target change that is implied by the futures quote.

2. MONETARY POLICY TRANSPARENCY

The previous section showed that forecasters can mechanically derive the expected funds rate and the probability of a change in the federal funds rate from information in the Federal Funds Futures contract prices. But a significant issue regarding forecasts of federal funds rate changes that has so far been ignored is the degree of transparency in the monetary policy process. Because the Fed uses the federal funds rate as its primary monetary policy instrument, forecasting the federal funds rate is nearly equivalent to forecasting the decisions of the FOMC. The clearer the rules that govern monetary policy, the easier it is to forecast the federal funds rate.

The conduct of monetary policy by the Fed is not fully transparent. Part of the lack of transparency rests on the basic approach of relying on discretion rather than on a set of fixed rules. Other sources of opaqueness arise when information about monetary policy decisions is not promptly made available to the public. From an efficiency standpoint, this policy approach has been the focus of much debate among monetary economists (see, for example, Goodfriend, 1996). From a forecasting standpoint, however, the lack of transparency poses a challenge. Forecasters must accept some degree of opaqueness as a given; but, in recent years, the process has become somewhat more transparent.

A quantum increase in transparency occurred in February 1994, when the FOMC began publicly announcing its decision regarding the federal funds rate target immediately following the conclusion of their meeting. This information reduced the uncertainty surrounding the federal funds rate in the period following meetings, and greatly assisted forecasters.

While this change in procedure improved forecasters' accuracy, it marked a point of departure in efforts to assess the reliability of the Federal Funds Futures contract prices as predictors of the federal funds rate. It is widely understood that substantial changes in the economic environment or in the policy regime can markedly reduce the value of pre-change data in gauging subsequent activity. In this case, the 1994 change meant that the track record of forecasts using data from before 1994 could not be used to ascertain forecast reliability going forward. As a result, we focus only on information in the post-1994 period. In the next section, the role and reliability of this information will be assessed.

3. EMPIRICAL RESULTS

This section examines how well futures prices predict policy actions by the FOMC. We choose to limit our focus to policy actions made at the second FOMC meeting in 1994 and later. At its first meeting in 1994 the FOMC shattered its precedent in two ways. First, the committee explicitly announced that it had a target for the Federal funds rate; previously, obscure language such as "degrees of reserve pressure" had amounted to a code for funds rate changes in the FOMC's records of policy actions and other publications. Second, the FOMC announced its funds rate target on the afternoon of February 4, only a few hours after the decision was made. Previously no information was announced until several weeks after an FOMC meeting.

This move toward greater transparency by the FOMC would be expected to improve the precision of forecasts of future policy moves, and thus increase the efficiency of the Federal Funds Futures market. Söderström (1999) has documented substantial differences in the performance of the market before 1994 and after. While the earlier period is of undoubted historical interest, the later period is more relevant for practitioners who would like to extract information from futures market prices.

An important question to ask of forecasts is whether they are unbiased predictors. Thus, we first examine whether the forecast extracted from futures prices accurately predicts the policy action taken by the FOMC thirty days later. Specifically, we estimate the following regression equation

$$\Delta i_t^T = \alpha + \beta(i_{t-30}^f - i_{t-30}^T) + \epsilon_{\mathrm{t}} \qquad (5)$$

where i_t^T is the FOMC's target for the federal funds rate at the end of date t, Δ is the difference operator, i_{t-30}^f is the value of the federal funds rate target at date t anticipated by market participants thirty days earlier, α and β are parameters to be estimated, and ϵ_{t} is an error term assumed to be white noise. Unlike other studies, we do *not* use monthly average data. The use of monthly averages introduces a variety of influences on the effective funds rate that obscures the focus of our study, the predictability of FOMC target changes. For example, seasonal reserve demands can introduce large movements into the effective funds rate at the end of calendar quarters. In addition, the timing of an FOMC meeting in a given month will alter the effective forecast horizon between months. Furthermore, since 1994 the FOMC began to make most of its target rate changes at scheduled meetings, thereby removing much uncertainty of the timing of possible changes. To bypass these and other complicating factors, our dependent variable is recorded for each FOMC meeting, and we therefore have eight observations each year. If the forecasts are unbiased, then $\hat{\alpha} = 0$ and $\hat{\beta} = 1$.

The forecasts did not display significant bias at the conventional five percent level. Our sample period covers March 1994 to January 2001, a period in which there were fifty-six FOMC meetings. The OLS estimate of α is -4.29, with an estimated standard error of 2.75, and the estimate of β is 0.89, with an estimated standard error of 0.12. The F value for testing the joint restrictions $\hat{\alpha} = 0$ and $\hat{\beta} = 1$ is 2.58, and thus the unbiasedness hypothesis is not rejected at the five percent level; however, the hypothesis would be rejected at the ten percent level. No serial correlation of the residuals was apparent. The market prediction picked up a large portion of the changes in the actual funds rate, which is indicated by an R^2 statistic of .49. However, the average market prediction was larger than the average actual change, as indicated by the rejection of unbiasedness at the ten percent level.

In seventeen of the eighteen times the FOMC changed its target, the predicted change had the same sign as the actual change; on the other occasion the predicted change was zero. Thus, in this relatively small sample the market accurately predicted the direction of change.

Looking at the tendency toward overpredicting target changes, it could be useful to know whether there was a tendency toward overpredicting the frequency of changes. A quick look is suggestive. We first calculated the implicit probability of a target change from the market prices as follows. First, if the predicted change was at least twenty-five basis points, the implicit probability of a target change was set to unity. Otherwise, the absolute value of the predicted change was divided by twenty-five (basis points) to calculate the implicit probability of a target change. Over the sample period the FOMC changed its target at about one third of the meetings considered. Yet our estimate of the implicit probability of a move averaged 0.52, well above the observed frequency of .32.

More formal analysis confirms this finding. We used probit analysis to estimate

$$\mathrm{I}\Delta i_{\mathrm{t}}^{T} = \alpha + \beta \, \mathrm{Pr} \, \Delta i_{\mathrm{t}-30}^{T} + e_t \tag{6}$$

where $\mathrm{I}\Delta i_{\mathrm{t}}^{T}$ is an indicator variable that takes the value of one if the FOMC moves at its meeting at date t, and zero if it chooses not to move. On three occasions in the sample period, the FOMC changed its target between meetings; on those occasions, we set the value to one at the next meeting. Pr $\Delta i_{\mathrm{t}-30}^{T}$ is the implicit probability that the FOMC will change its funds rate target in the next thirty days. Once again, if the implicit probability is an unbiased estimate of the observed frequency of target changes, then $\hat{\alpha} = 0$ and $\hat{\beta} = 1$.

The probit estimate of α was -1.37, with an estimated standard error of 0.36, and the estimate of β was 1.67, with an estimated standard error of 0.56. The chi-squared statistic for testing the joint hypothesis that $\hat{\alpha} = 0$ and $\hat{\beta} = 1$ was 30, and thus the hypothesis is rejected at conventional significance levels.

Thus, the implicit probability of a target change significantly overpredicts the frequency of changes thirty days ahead. However, the R^2 statistic for equation (4) was reasonably large at .49, which compares favorably with other variables that have been used to estimate interest rate changes. Also, the market accurately predicts the direction of target changes. Remembering that we have only fifty-six observations, the marginal rejection of unbiasedness may be a small sample phenomenon that would not be expected to persist (see Webb, 1987).

4. COMPARISON WITH OTHER RESULTS

There has not been a consensus in the literature on whether federal funds rate futures prices are biased predictors of FOMC target changes. Robertson and Thornton (1997), for example, studied monthly averaged data from 1988 to 1997 and found a significant bias in one-month-ahead predictions. Our method of analysis differs from theirs in several important respects. First, they used pre-1994 data. In addition, their counterpart to equation (4) does not contain a slope coefficient. Also, Robertson and Thornton use monthly average data for the effective funds rate in their empirical work, whereas we use the FOMC's target on a particular day. The use of monthly averages introduces seasonal effects, most importantly end-of-quarter spikes in the funds rate due to balance sheet window dressing by financial institutions. Also, during the post-1994 period the FOMC mostly changed its target at scheduled meetings, which varied considerably in their timing within a month. Thus, when examining a one-month-ahead forecast, the effective horizon of the forecast would vary considerably. For example, in 1995 the Fed met on February 1, but in March they met on the 28th. Based on the market price on the last day of the previous month, the forecast horizon would be one day and twenty-eight days, respectively. And in four months each year there is no scheduled FOMC meeting. These factors create predictable errors in the regression analysis.

Söderström (1999), however, reached a different conclusion. Also using monthly averaged data, his analysis documented the major difference between pre-1994 data and more recent data. He also recognized the importance of seasonal effects on monthly-averaged data and used dummy variables to attempt to adjust for their average effects. He found that, including the dummy variables, the market's prediction of the effective federal funds rate was unbiased when made immediately before an FOMC meeting. A potential shortcoming of his approach, in contrast to the one presented in this paper, remains the limitation of monthly averaged data for forecasting FOMC meeting outcomes.

As an example of the power of the market forecast, consider November 17, 2000. The FOMC had issued a press release following its meeting on November 15 in which it stated its belief that the balance of risks was tilted toward conditions generating inflationary pressures. Many analysts interpret such a statement to imply that the FOMC believes the next change in its funds rate target is more likely to be an increase than a decrease. However, the market quote for November 17 was 93.53, which implies that the contract's funds rate for December was 6.47 percent. Using equations (2) and (4) above, market participants placed a twenty-seven percent implicit probability on a twenty-five basis point decrease in the federal funds rate at the next FOMC meeting on December 17. While the FOMC did not change the funds rate at the December meeting, they lowered the funds rate fifty basis points early in January 2001, and another fifty basis points later that month at the scheduled FOMC meeting. Thus the market recognized the direction of the next move mid-November, and conveyed that information to market observers.

5. CONCLUSION

The federal funds rate plays a key role in the financial and economic environment facing individuals and businesses. Accurately forecasting the rate can be valuable but has often been very difficult. This paper describes two important innovations in forecasting the funds rate—the development of the Federal Funds Futures market in 1988 and a substantial improvement in the transparency of monetary policy in 1994. The paper then assesses the impact of these innovations on forecasters' ability of anticipate changes in the funds rate from 1994 to 2000.

We use the information from the Federal Funds Futures market 30 days in advance of FOMC meetings to gauge market participants' views of the likelihood and magnitude of FOMC target rate changes. We found that futures market prices were unbiased predictors of target rate changes when evaluated at the usual five percent level. At the ten percent level, though, the hypothesis of no bias was rejected. Despite that marginal bias, market forecasts are valuable. They accurately predict the direction of target changes and are a means to enhance the prospective accuracy of market forecasts. We believe that these developments give readers the ability to forecast the federal funds rate at least as well as highly paid Fed watchers did in the not-too-distant past.

REFERENCES

Chicago Board of Trade. 1997. *Insights Into Pricing the CBOT Federal Funds Futures Contract.*

Cook, Timothy. 1989. “Determinants of the Federal Funds Rate: 1979–1982.” *Economic Review,* Federal Reserve Bank of Richmond. Volume 1, pp. 3–19.

Goodfriend, Marvin. 1996. “Monetary Policy Comes of Age: A 20th Century Odyssey.” *1996 Annual Report.* Federal Reserve Bank of Richmond. Pp. 3–25.

Robertson, John C., and Daniel L. Thornton. 1997. “Using Federal Funds Futures Rates to Predict Fed Actions.” *Economic Review.* Federal Reserve Bank of St. Louis. Volume 5, pp. 45–53.

Söderström, Ulf. 1999. “Predicting Monetary Policy Using Federal Funds Prices.” Working Paper Number 307. Stockholm School of Economics and Sveriges Riksbank, Stockholm, Sweden.

Webb, Roy. 1987. “The Irrelevance of Tests for Bias in Series of Macroeconomic Forecasts.” *Economic Review.* Federal Reserve Bank of Richmond. Volume 1, pp. 3–9.