Economic Quarterly—Volume 99, Number 1—First Quarter 2013—Pages 45–81

Saving for Retirement with Job Loss Risk

Borys Grochulski and Yuzhe Zhang

In this article, we study optimal saving and consumption decisions. The optimal saving problem is among the most basic questions in economics and finance. How does one best decide on what portion of income they should consume now and what portion they should save for their future consumption needs? One important aspect of this question concerns saving for retirement. What is an optimal plan for saving enough to be able to retire? In particular, how does this plan depend on the risk of losing one's job? How much more should one save if the risk of becoming jobless increases?

Our primary objective in this article is to review several important results from the general theory of optimal consumption and saving decisions, as well as provide some novel analysis of the problem of saving for retirement in particular. The problem of optimal timing of retirement is most conveniently studied in a continuous-time framework, which we employ for our analysis. Our secondary objective is to provide an accessible exposition of the techniques useful in solving continuous-time models of the type we examine.

The basic framework economists have used to study the intertemporal tradeoff between current and future consumption has the following structure. An economic agent earns a stream of labor income that can change stochastically over time. At each point in time, the agent allocates his labor income to either current consumption or to savings. The agent's preferences over consumption streams are represented by a concave utility function, i.e., the agent is averse to fluctuations in his

Grochulski is a senior economist with the Richmond Fed. Zhang is an assistant professor with Texas A&M University. The authors would like to thank Arantxa Jarque, Marianna Kudlyak, Andrew Owens, and Ned Prescott for their helpful comments. The views expressed in this article are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System. E-mail: borys.grochulski@rich.frb.org; yuzhe-zhang@econmail.tamu.edu.

consumption. The duration of the agent's career is in the basic model approximated by infinity, i.e., the agent earns labor income and consumes indefinitely into the future. The portion of his labor income that the agent does not immediately consume adds to his financial wealth. In the basic model, there is only one asset in which all of the agent's financial wealth is invested. The asset pays off a riskless rate of return equal to the agent's intertemporal rate of time preference—the rate of return with which the agent's optimal consumption path absent all uncertainty would be constant forever.

The model we study in this article extends this basic framework by adding to the optimal consumption and saving decision a labor supply decision operating on the extensive margin, meaning we allow the agent to stop working. If the agent quits, he loses his labor income but gains leisure. The decision to quit is irreversible, so quitting means retiring. In retirement, the agent lives off of his savings and enjoys leisure. As in the basic model, the agent remains infinitely lived in our analysis.

For tractability and ease of exposition, we assume in our model a particularly simple stochastic structure for the agent's labor income process. The agent earns a constant stream of labor income for as long as he is not fired. If he is fired, he earns nothing and cannot go back to working ever again. Thus, being fired is in our model equivalent to being sent to involuntary retirement. The observed time path of the agent's labor income in our model is thus constant, at some positive level, until the agent either is fired or quits. Afterward, it is also constant at the level of zero.

Ljungqvist and Sargent (2004, Ch. 16) review the solution to the optimal consumption and saving problem in the basic framework with income fluctuating stochastically but without retirement. The main property of the optimal consumption plan is unbounded growth of financial wealth and consumption: Provided that the labor income process does not settle down in the long run (rather, it remains sufficiently stochastic), in almost all possible resolutions of uncertainty, the amount of financial wealth the agent holds and the amount the agent consumes grow over time without bound. When we allow for endogenous retirement, this property of the optimal wealth accumulation and consumption plan no longer holds. In all possible resolutions of uncertainty, wealth and consumption converge to a finite limit.

The intuition behind this result is simple. We show that the agent's optimal retirement plan takes the form of a wealth threshold rule: The agent retires as soon his accumulated financial wealth reaches a certain threshold. With this rule, wealth will not grow without bound prior to retirement. With finite wealth and no labor income after retirement, the agent's optimal consumption also remains bounded in the long run.

In fact, consumption is constant and equal to the amount of interest income generated by the agent's wealth in retirement.

The dynamics of consumption and savings are in our model as follows. Wealth and consumption increase monotonically over time for as long as the agent does not involuntarily lose his job. If the agent is fired, his wealth accumulation is stopped and his consumption jumps downward. If the agent reaches the voluntary retirement threshold, his wealth and consumption reach their permanent, retirement levels smoothly. For any level of financial wealth the agent starts out with, we compute the planned duration of the agent's career, i.e., his time to planned retirement. Agents with lower initial wealth retire later.

We provide several comparative statics results. We show how the agent's optimal path of wealth accumulation and consumption prior to retirement depends on the risk of losing his job, on the value of leisure he obtains in retirement, and on the level of the rate of return paid by the asset in which the agent invests his savings. Higher job loss risk implies the agent saves more, consumes less, and retires faster. Lower utility of leisure implies the agent saves less, consumes more, and retires later. When the interest rate is higher, the agent retires with lower wealth and generally consumes more prior to retirement. In solving for the agent's optimal retirement rule, we discuss the option value of postponing retirement.

In addition, we discuss, in the context of our model, two standard properties of the solution to the optimal consumption and saving problem. We show that in the model with retirement, like in the standard model without retirement, the agent's marginal utility of consumption is a martingale, which means the conditional expected change in its value is always zero. We also review the result known as the permanent income hypothesis (PIH). Defined narrowly, PIH states that the agent chooses to consume exactly the income from his total wealth at all times. Total wealth consists of both financial and human wealth, where human wealth is defined as the expected present value of all labor income the agent is to earn in the future. With quadratic preferences, PIH holds in the standard model without retirement. We show that adding an endogenous retirement decision to the model does not overturn PIH.

We provide an elementary-level discussion of all dynamic optimization techniques involved in the analysis of our continuous-time model, thus making it accessible to a broad audience.

Related Literature

Our study is related to the literature on optimal consumption and saving decisions with fluctuating income and incomplete markets, and to the literature on the optimal timing of retirement.

The vast literature on the optimal saving problem with fluctuating income is summarized in Ljungqvist and Sargent (2004, Ch. 16). Classic studies of this problem, which include Friedman (1957), Bewley (1977), and Hall (1978), take the agent's stochastic income process as exogenous, which means they abstract from retirement. Chamberlain and Wilson (2000) allow for stochastic changes to the interest rate and show under weak conditions that optimal consumption diverges with probability one. Marcet, Obiols-Homs, and Weil (2007) extend the classic framework by including the agent's labor supply decision along the intensive margin. They show that with endogenous labor income the result of divergence of almost all consumption paths does not hold due to a wealth effect suppressing the agent's labor supply and thus eventually eliminating fluctuations in the agent's income. Our analysis is similar but allows for changes in labor supply along the extensive margin, i.e., it incorporates the retirement decision.

Similar to our analysis, Ljungqvist and Sargent (forthcoming) study an optimal consumption and saving problem with endogenous retirement. They focus on the impact of the curvature of the life cycle income profile on savings and the timing of retirement in a finite-horizon model in which all income shocks are unanticipated. Our model assumes a flat income profile in an infinite-horizon model in which the agent anticipates the risk in his income and responds to it.

Kingston (2000) and Farhi and Panageas (2007) study the optimal retirement timing decision combined with the problem of optimal saving and asset allocation prior to retirement, where available assets are one risky and one riskless asset, as in Merton (1971). They show that the option to delay retirement lets agents take on more risk than they would have chosen otherwise. In particular, Farhi and Panageas (2007) show that investors close to retirement may find it optimal to invest more heavily in stocks than those whose retirement is far off in the future. Our analysis is different as we do not consider a portfolio allocation problem in this article. Rather, we assume an incomplete market structure in which the riskless asset is the only vehicle for saving and wealth accumulation, as in the classic models of optimal consumption and saving decisions.

Our article is organized as follows. Section 1 presents our model. Section 2 discusses the optimal consumption pattern after retirement. Section 3 describes the optimal timing of retirement. Sections 4 and 5 study consumption and wealth accumulation prior to retirement.

Sections 6 and 7 provide comparative statics results with respect to several parameters of the model, with particular attention given to the job loss hazard parameter. Section 8 concludes. Appendix A contains proofs. Appendixes B and C discuss two extensions of the model.

1. MODEL

We will study the following partial equilibrium model in continuous time with a single agent. The agent consumes a single consumption good and leisure. The agent is initially employed. When employed, the agent earns a flow of labor income of $y > 0$ units of the consumption good per unit of time. The agent also consumes a flow of leisure of $l_W > 0$ units per unit of time. If the agent is not working, his labor income is zero but his flow of leisure is $l_R > l_W$. The agent's preferences over deterministic paths of consumption and leisure are represented by a standard utility function

$$\int_0^\infty e^{-rt} U(c_t, l_t) dt,$$

where c_t is consumption, $l_t \in \{l_W, l_R\}$ is leisure, and $r > 0$ is the agent's intertemporal rate of time preference.

While employed, the agent faces the risk of losing his job. If he loses his job, he never works again, which effectively means that losing one's job represents in our model involuntary retirement. The job loss shock arrives stochastically with a constant hazard rate $\lambda > 0$. That is, for any date t at which the agent is employed and for any $s > 0$, the probability that the agent will have not lost his job by date $t + s$ is $e^{-\lambda s}$.

In addition to losing his job involuntarily, the agent can quit. In this case, as well, the separation from employment is permanent, i.e., quitting means retiring. If he retires, the agent gives up the flow of labor income y and gains the flow of extra leisure $l_R - l_W > 0$.

At each point in time, the agent decides how much of his current income to consume and how much to save. There is only one asset in which the agent can invest his savings. It is a riskless asset with a constant rate of return equal to the agent's rate of time preference r. Denote the amount of the riskless asset held by the agent at date t, i.e., the agent's financial wealth at t, by W_t.

With these assumptions, the law of motion for the agent's financial wealth W_t is as follows. While working, the financial wealth changes according to

$$dW_t = (rW_t + y - c_t)dt. \tag{1}$$

Thus, for example, if the agent were to consume exactly his labor income while working, i.e., if $c_t = y$, then his financial wealth would grow exponentially at the rate of interest r. When not working (i.e., in retirement), the agent's wealth follows

$$dW_t = (rW_t - c_t)dt. \tag{2}$$

The agent maximizes

$$\mathbb{E}\left[\int_0^{\min\{\tau,\tau_f\}} e^{-rt}U(c_t, l_W)dt + \int_{\min\{\tau,\tau_f\}}^{\infty} e^{-rt}U(c_t, l_R)dt\right], \tag{3}$$

where τ is the agent's planned, voluntary retirement time, τ_f is the time he is forced into involuntary retirement, and the expectation $\mathbb{E}$ is taken over the realizations of the involuntary job loss shock. In particular, we will take the utility function to be separable in consumption and leisure:

$$\begin{aligned} U(c, l_W) &= u(c), \\ U(c, l_R) &= u(c) + \psi, \end{aligned}$$

where u is strictly increasing and a strictly concave utility of consumption and $\psi \geq 0$ is the utility of the extra leisure the agent enjoys in retirement. In this specification, the agent's lifetime utility (3) can be more simply written as

$$\mathbb{E}\left[\int_0^{\infty} e^{-rt}u(c_t)dt + e^{-r\min\{\tau,\tau_f\}}\frac{\psi}{r}\right].$$

2. OPTIMAL SAVING AND CONSUMPTION IN RETIREMENT

We start by discussing the agent's optimal use of savings in retirement. Because the return on the financial wealth held by the agent is equal to the agent's rate of time preference and the agent faces no uncertainty in retirement, it is natural to guess that in retirement the agent will keep assets constant, $dW_t = 0$, and consume his capital income, i.e., the return rW_t at all t. Thus, the natural guess is that if the agent retires with assets W_t, the maximum present value of total lifetime utility he can obtain after retirement, denoted by $V(W_t)$, is

$$V(W_t) = \frac{1}{r}u(rW_t) + \frac{\psi}{r}. \tag{4}$$

In the remainder of this section, we will use a standard dynamic programming argument to confirm that this guess is correct. In the process, we will derive an optimality condition on the value function

V—known as the Bellman equation—that will be useful when we discuss the agent's optimal consumption and saving behavior prior to retirement in the next section.

Following the dynamic programming approach, we take a small time interval $[t, t+h)$ and assume that from time $t+h$ onward the agent will apply the optimal saving and consumption policy, which is not known to us as of now. Given this assumption, we seek an optimal consumption rate c within the time interval $[t, t+h)$. Because h is small, we can consider c to be constant over the interval $[t, t+h)$. The true optimal consumption rate to be applied at time t, c_t, will be obtained by taking the limit as h goes to zero.

Because the agent follows an optimal consumption plan after $t+h$, the total discounted value he will obtain as of time $t+h$ will be $V(W_{t+h})$, where W_{t+h} is the amount of financial wealth the agent holds at $t+h$. For a given consumption rate c to be applied in $[t, t+h)$, the total discounted utility value the agent obtains as of time t is

$$\int_0^h e^{-rs}\left(u(c)+\psi\right) ds + e^{-rh}V(W_{t+h}). \tag{5}$$

Because this plan is a feasible consumption plan for an agent with nonnegative wealth, the maximal utility value $V(W_t)$ must be at least as large as the value of this plan, so for any c it is true that

$$V(W_t) \geq \int_0^h e^{-rs}\left(u(c)+\psi\right) ds + e^{-rh}V(W_{t+h}).$$

When h becomes arbitrarily small, the maximized value of the right-hand side of this expression approaches the value on the left-hand side, which we can write as

$$V(W_t) = \max_c \left\{ \int_0^h e^{-rs}\left(u(c)+\psi\right) ds + e^{-rh}V(W_{t+h}) \right\} \tag{6}$$

with h approaching zero. Since h is very small, we can replace the expression on the right-hand side of (6) with its first-order approximation. For a function f differentiable at some point t, for small h, we can approximate $f(t+h)$ with $f(t)+f'(t)h$. In this approximation, the first of the two terms in (5) equals

$$0 + \left(u(c)+\psi\right) h,$$

and the second term equals

$$V(W_t) + \left(-rV(W_t) + V'(W_t)\frac{dW_t}{dt}\right) h.$$

The value in (5) is therefore approximated by

$$V(W_t) + \big(u(c)+\psi - rV(W_t) + V'(W_t)(rW_t - c)\big)\, h,$$

where we have used the law of motion for assets in retirement (2). With this approximation, we can thus write (6) as

$$V(W_t) = \max_c \left\{ V(W_t) + \left(u(c) + \psi - rV(W_t) + V'(W_t)(rW_t - c) \right) h \right\}. \tag{7}$$

Dividing by h and simplifying terms, we obtain the following condition for the value function V:

$$rV(W_t) = \max_c \left\{ u(c) + \psi + V'(W_t)(rW_t - c) \right\}. \tag{8}$$

We will refer to this condition as the Bellman equation for the value function V. This equation shows how the agent's total utility $V(W_t)$ (converted to flow units by multiplying it by r) depends on current utility and the change in financial wealth. Higher c will increase the current utility flow $u(c) + \psi$ at the cost of lower saving $rW_t - c$. The marginal value of wealth $V'(W_t)$ shows how costly a change in saving is to the agent in utility terms. In choosing the consumption rate c the agent optimally balances this tradeoff between his utility from current consumption and his utility from future wealth.

Next, by differentiating the Bellman equation (8), we will obtain the optimal consumption policy function. Note that the Envelope Theorem lets us treat c as a constant in this differentiation. Indeed, differentiation gives us

$$rV'(W_t) = V''(W_t)(rW_t - c_t) + V'(W_t)r,$$

which simplifies to

$$0 = V''(W_t)(rW_t - c_t). \tag{9}$$

Assuming the second derivative V'' is nonzero, we divide both sides by $V''(W_t)$ to obtain

$$c_t = rW_t. \tag{10}$$

This confirms our guess that the optimal consumption policy for the agent in retirement is to consume the interest income from his financial assets at all t. Using this policy in the law of motion for wealth in retirement, (2), we confirm that $dW_t = 0$ and so assets and consumption remain constant in retirement. Substituting constant consumption $c_{t+s} = rW_t$ into the agent's utility function at all times $t + s$ following the retirement date t leads to the value function (4), confirming the guess we made at the beginning of this section.

We will also note that the first-order condition for the maximum on the right-hand side of (8) is

$$u'(c_t) = V'(W_t).$$

This condition, along with the policy function (10), lets us determine the marginal value of wealth in retirement as

$$V'(W_t) = u'(rW_t). \tag{11}$$

Clearly, the same result can be obtained by differentiating (4) directly.[1]

3. OPTIMAL RETIREMENT DECISION

In this section, we show that the optimal retirement policy for the agent is a threshold policy: The agent retires when his wealth reaches a specific threshold level. At this threshold, the marginal value of income the agent can earn if he works is exactly matched by the value of the extra leisure the agent can get if he retires.

In our analysis of the optimal voluntary retirement rule, we will use one intuitive property of the optimal pre-retirement wealth accumulation path. Namely, that the optimal wealth accumulation path is non-decreasing, i.e., the agent actually does save for retirement. That the agent will choose an increasing wealth accumulation path $\{W_t; t \geq 0\}$ prior to retirement is very intuitive in our model because the agent's labor income process is non-increasing and the return on savings is equal to the agent's rate of time preference. It is clear from (1) that W_t decreases only if $c_t > rW_t + y$, i.e., when the agent consumes more than his capital income rW_t and labor income y combined. Doing so clearly cannot be optimal for the agent given the labor income process the agent faces. The agent earns constant labor income $y > 0$ when he works and has no labor income after he quits or loses his job. In order to smooth consumption, the agent will want to save at least a part of his labor income for as long as he works, i.e., will choose $c_t \leq rW_t + y$ prior to retirement. In Section 5, we will characterize precisely what portion of y will be saved at each point in time. For now, we will just state that $c_t > rW_t + y$ is never optimal for the agent, and thus W_t is at least weakly increasing over time.

We now move on to the agent's optimal retirement decision. We will analyze this decision in two steps. First, we will compare the agent's value from retiring now, i.e., at some given time t, with the value from retiring a little later, i.e., at $t + h$, for a small $h > 0$. Then, we will argue that if the agent prefers to retire at t rather than retire at $t + h$ for a small h, then he also prefers to retire at t over retiring at any future date, which means the agent's overall optimal retirement time is t.

[1] Further, differentiating (4) twice, we have $V''(W_t) = ru''(rW_t) < 0$, which justifies the assumption of nonzero V'' we made when we divided (9) by $V''(W_t)$.

As before, we will use the first-order approximation for payoffs at $t+h$. In addition, we will discretize the involuntary job loss shock by assuming that if the agent loses his job by time $t+h$, this loss will occur only at $t+h$ and not earlier. With h approaching zero, these approximations will be sufficiently precise.

Suppose then that the agent is employed and has financial wealth W_t as of some time t. As we know from the previous section, the agents' value of retiring now is $V(W_t)$. The value of postponing retirement by a small amount of time h, denoted here by $V^h(W_t)$, is

$$V^h(W_t) = \max_c \left\{ \int_0^h e^{-rs} u(c) ds + e^{-rh} V(W_{t+h}) \right\}$$

with wealth following (1) between t and $t+h$, as the agent keeps working between t and $t+h$. Note that it does not matter if at $t+h$ the job loss shock happens or does not happen, because the agent is retiring at $t+h$ anyway. Since h is small, we use the first-order approximation and express $V^h(W_t)$ as

$$\begin{aligned} V^h(W_t) &= \max_c \left\{ V(W_t) + \left(u(c) - rV(W_t) + V'(W_t) \frac{dW_t}{dt} \right) h \right\} \\ &= \max_c \left\{ V(W_t) + \big(u(c) - rV(W_t) + V'(W_t)(rW_t + y - c) \big) h \right\}, \end{aligned}$$

where the second line uses (1).

Using the first-order approximation (7) for the value of retiring at t, $V(W_t)$, we have that postponing retirement by h is strictly preferred to retiring immediately, i.e., $V^h(W_t) > V(W_t)$, if and only if

$$\begin{aligned} & \max_c \left\{ V(W_t) + \big(u(c) - rV(W_t) + V'(W_t)(rW_t + y - c) \big) h \right\} \\ > \; & \max_c \left\{ V(W_t) + \big(u(c) + \psi - rV(W_t) + V'(W_t)(rW_t - c) \big) h \right\}. \end{aligned}$$

Dividing by h, simplifying, and taking terms that do not depend on c out of the maximization on each side, we have

$$\max_c \left\{ u(c) - V'(W_t) c \right\} + V'(W_t) y > \max_c \left\{ u(c) - V'(W_t) c \right\} + \psi.$$

Since the maximization problems on both sides of this inequality are the same, we simplify the above condition further to obtain

$$V'(W_t) y > \psi. \tag{12}$$

This says that whenever the utility flow from the additional leisure the agent can obtain by retiring is smaller than the utility he draws from the flow of his labor income, the agent will prefer to postpone retirement. From (11) we know that the agent's marginal value of wealth in retirement is $V'(W_t) = u'(rW_t)$. Thus, inequality (12) is

equivalent to

$$\frac{\psi}{y} < u'(rW_t). \tag{13}$$

Let u'^{-1} denote the inverse function of u'. Since the right-hand side of (13) is strictly decreasing in W_t, it is true that this inequality holds for all $W_t < W^*$, where the threshold value W^* is given by

$$W^* = \frac{1}{r} u'^{-1}\left(\frac{\psi}{y}\right). \tag{14}$$

This means that postponing retirement (by at least a small instant) is preferred at all wealth levels W_t strictly smaller than W^*. The agent thus will not retire voluntarily with wealth $W_t < W^*$. Intuitively, for as long as his wealth is below W^*, by continuing to postpone retirement, the agent obtains a larger current flow return (his labor income is more valuable than the leisure forgone to obtain it) and retains the option to retire later.

Now that we know the agent will not retire with wealth smaller than W^*, we should ask if the agent will choose to retire as soon as his wealth reaches W^*. We know already that the agent with wealth W_t equal to or larger than W^* prefers to retire at t over retiring a bit later. But what about the possibility of retiring much later? Does $W_t \geq W^*$ also mean that the agent prefers to retire at date t rather than at any future date $T > t$? The answer is yes because, as we argued earlier in this section, the time path of wealth the agent chooses is never decreasing. Indeed, suppose the agent's wealth as of t satisfies $W_t \geq W^*$, but he does not retire until some later date $T > t$. Because the path of wealth is non-decreasing, $W_s \geq W^*$ at all dates s in $t \leq s \leq T$. In particular, for a small $h > 0$, at date $s = T - h$, the agent's wealth is greater than or equal to W^*, so, by our previous argument, the agent prefers to retire at $T - h$ rather than wait until T. Because his wealth is not smaller than W^* at $T - 2h$, as well, the agent will prefer to retire at $T - 2h$ rather than at $T - h$. Extending this reasoning backward in time all the way back to date t shows that the agent's overall preferred retirement rule is to retire as soon as his wealth reaches W^*.

In sum, the optimal retirement rule takes on a threshold form. The agent chooses to postpone retirement for as long as his wealth is below the threshold W^* and retire immediately when his wealth reaches W^*. It is worth noting in (14) that the optimal wealth threshold W^* increases in labor income y, decreases in the value of leisure ψ, and does not depend on the intensity of the job loss risk λ. If $\psi = 0$, i.e., if working is not costly to the agent in terms of forgone leisure at all, then $W^* = \infty$. In this case, the agent never chooses to retire voluntarily.

The Option Value of Postponing Retirement

Because retirement is permanent in our model, when the agent retires he loses the option of working at later dates. The threshold retirement rule we derived tells us, however, that the value of this option is zero for the agent in the problem we study.

In general, a one-time, irreversible action has a positive option value for an agent if he is willing to forgo an immediate benefit that the action can produce in order to retain the option of taking the action in the future.[2] In our model, the agent retires as soon as the current flow return from doing so turns positive, i.e., when the value of the flow of leisure, ψ, becomes as large as the value of the flow of labor income y, $V'(W_t)y$. The agent is not willing to delay retirement beyond that point because once wealth reaches W^* the agent will continue to prefer the flow of leisure ψ over the flow of his labor income y at all future times in all possible realizations of uncertainty he faces. In fact, once the agent retires with wealth $W_t \geq W^*$, there is no realization of uncertainty in which he might want to go back to working, even if he could return.

The value of having the option to work in the future that the agent gives up by retiring would in our model be positive if the parameters determining the threshold wealth level W^* could change in a way that increases the value of working relative to the value of consuming leisure. In particular, the value of this option would be positive if the agent could receive a positive income shock increasing the level of his labor income y, or a taste shock decreasing the utility of leisure ψ, or a taste shock increasing the agent's marginal utility of consumption u', or a shock destroying a part of the agent's financial wealth W_t. In Appendix C, we discuss this point in more detail, focusing on the possibility of an increase in labor income y.

4. CONSUMPTION, SAVING, AND WEALTH ACCUMULATION PRIOR TO RETIREMENT

In this section, we study the agent's optimal saving and consumption decisions prior to retirement, i.e., when his wealth is strictly less than W^*. The guess-and-verify method we used earlier to solve for optimal consumption in retirement will not work here because wealth and consumption have nontrivial dynamics prior to retirement. In order to

[2] For example, it is often optimal for a business owner to keep her business open for some time after it begins to make a loss (a flow of negative profit). The option for the profits that the business may generate in the future has a value that keeps the owner from shutting the business down as soon as the current profit flow turns negative (see Leland [1994]). Pindyck (1991) discusses the option value of undertaking a one-time investment in a stochastic environment.

study these dynamics, we will derive intertemporal optimality conditions leading to a dynamic system in wealth and consumption. We will then use standard methods to analyze this system.

Bellman Equation

Let us denote by $J(W_t)$ the maximal discounted expected utility value a working agent can obtain given his wealth W_t. Since retiring immediately is optimal when $W_t \geq W^*$, we have $J(W_t) = V(W_t)$ for all $W_t \geq W^*$. Since not retiring is strictly preferred by the agent when $W_t < W^*$, we have $J(W_t) > V(W_t)$ for all $W_t < W^*$. We look now to learn more about $J(W_t)$ for $W_t < W^*$. We proceed by deriving the Bellman equation for J analogous to the Bellman equation for V we derived earlier.

Take a small $h > 0$ and assume that an agent who works at t and holds financial wealth W_t chooses to consume at some constant rate c inside the time interval $[t, t+h)$. In addition, assume that if the agent wants to quit inside $(t, t+h)$, he will do so only at $t+h$. Likewise, assume that if the agent loses his job involuntarily during this short period of time, this will happen only at the end of the period, i.e., at date $t+h$. As before, these assumptions will be innocuous when we take the limit with h going to zero. Following the dynamic programming approach, we suppose that from time $t+h$ onward the agent applies an optimal (to us yet unknown) consumption and saving policy. The total utility value the agent obtains by following this strategy with some fixed consumption rate c is

$$\int_t^{t+h} e^{-rs}u(c)ds + e^{-rh}\left[e^{-\lambda h}J(W_{t+h}) + (1-e^{-\lambda h})V(W_{t+h})\right]. \quad (15)$$

The term in square brackets represents the expectation of the value the agent will draw at time $t+h$. With probability $e^{-\lambda h}$ he does not lose his job as of $t+h$ and $J(W_{t+h})$ represents the continuation value he obtains at that time. With probability $1-e^{-\lambda h}$ he loses his job, and thus the continuation value he obtains at $t+h$ is the retirement value $V(W_{t+h})$.

With the optimal choice of c, the value in (15) approaches the overall maximal value the agent can obtain, $J(W_t)$, which we write as

$$\begin{aligned} J(W_t) \;=\; & \max_c \Big\{ \textstyle\int_t^{t+h} e^{-rs}u(c)ds \\ & + e^{-rh}\Big[e^{-\lambda h}J(W_{t+h}) + (1-e^{-\lambda h})V(W_{t+h})\Big]\Big\} \end{aligned} \quad (16)$$

with h approaching zero. Since h is very small, we can apply the first-order approximation to the value in (15) and write it as

$$J(W_t) + \big(u(c) - (r+\lambda)J(W_t) + J'(W_t)(rW_t + y - c) + \lambda V(W_t)\big)h.$$

Using this approximation in (16), we have

$$\begin{aligned} J(W_t) &= \max_c \big\{ J(W_t) + \big(u(c) - (r+\lambda)J(W_t) + \\ &\quad + J'(W_t)(rW_t + y - c) + \lambda V(W_t)\big)h \big\}. \end{aligned}$$

Dividing by h and simplifying terms, we get the Bellman equation for J:

$$(r+\lambda)J(W_t) = \max_c \big\{ u(c) + J'(W_t)(rW_t + y - c) + \lambda V(W_t) \big\}. \quad (17)$$

To compare it with the Bellman equation for V, (8), let us rewrite (17) as

$$rJ(W_t) = \max_c \big\{ u(c) + J'(W_t)(rW_t + y - c) \big\} - \lambda \left(J(W_t) - V(W_t) \right). \quad (18)$$

Bellman equations (8) and (18) differ in three ways. First, the trade-off between consumption and saving is different, as prior to retirement the agent earns the stream of income y. Second, the level of J is also influenced by the lower flow of leisure prior to retirement. These two differences are reflected in the expression inside the maximization with respect to c in (18). Third, (18) contains an extra term, $-\lambda \left(J(W_t) - V(W_t) \right)$, that reflects the possibility of the agent's involuntarily losing his job. In this term, λ is the intensity with which the agent loses his job and $J(W_t) - V(W_t)$ is the loss of value that occurs in that event.

Euler Equation

As before, we use the envelope and first-order conditions associated with the Bellman equation. Using the Envelope Theorem in differentiation of the Bellman equation (17) yields

$$(r+\lambda)J'(W_t) = J''(W_t)(rW_t + y - c_t) + J'(W_t)r + \lambda V'(W_t).$$

Simplifying terms and rearranging, we get

$$\lambda \left(J'(W_t) - V'(W_t) \right) = J''(W_t) \left(rW_t + y - c_t \right). \quad (19)$$

Unlike in the post-retirement problem we studied earlier, in the pre-retirement problem the envelope condition (19) does not by itself determine the optimal consumption rule. However, it gives us an important intertemporal optimality condition for consumption known as the Euler

equation. To derive it, we use the chain rule to express the time derivative of $J'(W_t)$ as

$$\begin{aligned}\frac{dJ'(W_t)}{dt} &= J''(W_t)\frac{dW_t}{dt} \\ &= J''(W_t)(rW_t + y - c_t),\end{aligned}$$

where the second equality uses (1). This lets us write (19) as

$$\frac{dJ'(W_t)}{dt} = \lambda\left(J'(W_t) - V'(W_t)\right).$$

Next, we use the first-order condition in the maximization problem in the Bellman equation (17),

$$u'(c_t) = J'(W_t), \tag{20}$$

to write the above as

$$\frac{du'(c_t)}{dt} = \lambda\left(u'(c_t) - V'(W_t)\right).$$

Finally, we use (11) to eliminate V' from the above equation and express it purely in terms of the marginal utility of consumption:

$$\frac{du'(c_t)}{dt} = \lambda\left(u'(c_t) - u'(rW_t)\right). \tag{21}$$

This is the Euler equation for consumption prior to retirement. It shows how the marginal utility of consumption changes along an optimal path of consumption and financial wealth accumulation prior to retirement.[3]

Martingale Property

Before we use the Euler equation to study optimal consumption and asset accumulation, let us discuss an implication of the Euler equation known as the martingale property of marginal utility. As studied by Hall (1978) and many others, (21) implies that at all times prior to retirement the expected change in marginal utility of consumption is zero, i.e., marginal utility of consumption is a so-called martingale.[4] In discrete-time models that are most commonly used in the literature, the Euler equation takes the familiar form of $u'(c_t) = \mathbb{E}_t\left[u'(c_{t+1})\right]$ at all t, where $\mathbb{E}_t\left[\cdot\right]$ is the conditional expectation operator. In discrete time, it is thus easy to see that the expected change in u' is zero. In continuous time, the martingale property is slightly less self-evident but can still be seen as follows.

[3] Note that, trivially, the Euler equation also holds after retirement.

[4] Because consumption is constant in retirement, marginal utility of consumption is trivially also a martingale after the agent retires, voluntarily or not.

Take a small $h > 0$ and a date t at which the agent is not retired. In the time interval $[t, t+h)$, the agent will be hit with the job loss shock with probability $1-e^{-\lambda h}$, which will cause his marginal utility at $t+h$ to change (jump) by $u'(rW_{t+h})-u'(c_{t+h})$. With probability $e^{-\lambda h}$, the agent will not lose his job, in which case the change in his marginal utility over the time interval $[t, t+h)$ will be simply $u'(c_{t+h}) - u'(c_t)$. Marginal utility is a martingale when the average (i.e., expected) value of these two changes is zero, i.e., when

$$(1 - e^{-\lambda h}) \left(u'(rW_{t+h}) - u'(c_{t+h})\right) + e^{-\lambda h} \left(u'(c_{t+h}) - u'(c_t)\right) = 0.$$

Rearranging this condition, we have

$$\begin{aligned} u'(c_{t+h}) - u'(c_t) &= (e^{\lambda h} - 1) \left(u'(c_{t+h}) - u'(rW_{t+h})\right) \\ &= \lambda h \left(u'(c_{t+h}) - u'(rW_{t+h})\right), \end{aligned}$$

where the second equality uses the linear approximation $e^{\lambda h} = 1 + \lambda h$. Dividing by h and taking formally the limit as $h \to 0$, we get the Euler equation (21). Thus, the Euler equation (21) says exactly that the time trend $du'(c_t)/dt$ in marginal utility along the path that consumption follows conditional on the job loss shock not occurring is the negative of the jump in marginal utility that occurs if the agent loses his job, $u'(rW_t)-u'(c_t)$, times the intensity of the job loss λ. This trend exactly offsets the jump-induced change in marginal utility, making the overall expected change in marginal utility zero, i.e., marginal utility indeed is a martingale.[5]

Dynamic Analysis

As we saw earlier, consumption and financial wealth have trivial dynamics in retirement: Both remain constant over time. Prior to retirement, however, wealth and consumption do change over time. We will now use the Euler equation (21) and the law of motion for wealth (1) to study the dynamics of wealth and consumption prior to retirement. To do this, we use the chain rule

$$\frac{du'(c_t)}{dt} = u''(c_t)\frac{dc_t}{dt}$$

and the strict concavity of u, implying $u'' \neq 0$, to express the Euler equation (21) as

$$\frac{dc_t}{dt} = \lambda \frac{u'(c_t) - u'(rW_t)}{u''(c_t)}. \tag{22}$$

[5] In this respect, the marginal utility process is in our model similar to a compensated Poisson process. See Problem 1.3.4 in Karatzas and Shreve (1997).

Figure 1 Phase Diagram for (W,c)

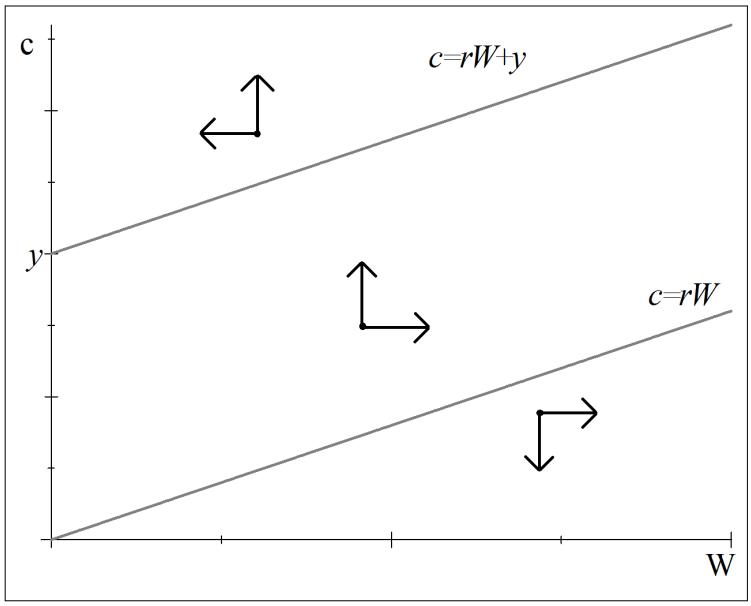

Together with the law of motion for financial wealth W_t, (1), this gives us a dynamic system describing the evolution of consumption and wealth prior to retirement. In the rest of this subsection, we will study qualitative properties of this system. We will use a phase diagram to describe the shape of the time paths in the plane (W, c) that satisfy the differential equations (1) and (22). Any such path is called a solution to the system (1), (22), and there are an infinite number of them (every point in the domain for (W, c) belongs to one solution). The solutions represent all paths of consumption and wealth accumulation the agent might want to follow while working that are consistent with intertemporal optimization. That is, any path that is not a solution to (1), (22) is not optimal for the agent. In order to select the optimal path from among all solutions to this system, a boundary condition is needed. In standard infinite-horizon analysis, the transversality condition serves this role. In our model with endogenous retirement, this condition will be provided by the optimal voluntary retirement rule we obtained in the previous section.

The phase diagram for the system of differential equations (1), (22) is shown in Figure 1. It provides a graphical representation of the

directions in which the system (W_t, c_t) moves along all possible solution paths. In Figure 1, these directions are marked by horizontal and vertical arrows. These arrows are determined as follows.

From (22) we see that along a solution path consumption will increase over time, i.e., $dc_t/dt > 0$, if and only if $u'(c_t) - u'(rW_t) < 0$, i.e., if and only if $c_t > rW_t$. The line $c = rW$ therefore divides the state space (W, c) into two regions: one in which consumption grows over time (the region above this line), and one in which it decreases over time (the region below it). Similarly, we have from (1) that W_t grows over time if and only if $c_t < rW_t + y$. Therefore, the line $c = rW + y$ divides the state space (W, c) into a region of wealth growth (below this line) and a region of wealth decline (above it). Since the two lines are parallel, we see that there are three regions in the state space (W, c) differentiated by distinct dynamic properties of the system (W_t, c_t). Above the line $c = rW + y$, wealth declines and consumption grows. In the band $rW < c < rW + y$, both wealth and consumption grow. Below the line $c = rW$, wealth grows and consumption declines.

The qualitative conclusions we can obtain from the phase diagram are as follows. Inside the band $rW < c < rW + y$, solution paths increase in both the c and the W direction and fall into one of the following three types. Paths of the first type will reach the upper straight line $c = rW + y$, where they bend backward as W_t begins to decrease while c_t continues to increase. Paths of the second type will reach the lower straight line $c = rW$, where they bend downward with c_t declining and W_t continuing to increase. Note that none of the paths of the first or second type return to the band $rW < c < rW + y$ once they leave it. Paths of the third type will stay inside the band $rW < c < rW + y$ forever.

Further characterization of the solution paths can be obtained analytically in the special case in which the Euler equation (22) is linear or numerically in other cases. In the remainder of this article, we will focus on the case with a linear Euler equation and discuss analytical solutions. As we will see in the next section, the Euler equation is linear when the utility function u is quadratic. In Appendix B, we briefly discuss how the results change for other utility functions, in particular for preferences exhibiting constant relative risk aversion (CRRA).

5. EXACT SOLUTION WITH LINEAR EULER EQUATION

We specialize the utility function to

$$u(c) = -\frac{1}{2}(c - B)^2$$

and restrict its domain to $c \leq B$. Under this specification, marginal utility is linear and therefore the Euler equation (22) is linear in consumption as well as in wealth:

$$\frac{dc_t}{dt} = \lambda \left(c_t - rW_t \right). \tag{23}$$

The system of differential equations (1), (23) is now linear and can be solved in closed form. In particular, we have the following lemma providing analytical expressions for all solution paths of the system (1), (23).

Lemma 1 *Let $\{(W_t, c_t); t \geq 0\}$ be a solution path. If there exists τ such that $c_\tau = rW_\tau + y$, then*

$$c_t = rW_t + \frac{ry}{r+\lambda}\left(1 - e^{-(r+\lambda)(\tau - t)}\right) + ye^{-(r+\lambda)(\tau - t)}. \tag{24}$$

If there exists τ such that $c_\tau = rW_\tau$, then

$$c_t = rW_t + \frac{ry}{r+\lambda}\left(1 - e^{-(r+\lambda)(\tau - t)}\right). \tag{25}$$

Otherwise,

$$c_t = rW_t + \frac{ry}{r+\lambda}. \tag{26}$$

Proof. In Appendix A. ■

Figure 2 plots several sample solution paths $\{(W_t, c_t); t \geq 0\}$ of the three types given in the above lemma. Solution paths (24) bend backward with wealth declining over time at all dates $t > \tau$, where τ is such that $c_\tau = rW_\tau + y$. None of these solution paths will be optimal for the agent because, as we saw earlier, it is never optimal for the agent to see his financial wealth decrease while he is saving for retirement. Along all solution paths (25) and (26) wealth is increasing. These paths, therefore, are our candidates for the optimal path of consumption and saving prior to retirement.

The Optimal Accumulation Path

As we saw in Section 3, the agent's retirement decision is determined by a simple wealth threshold rule. The agent retires as soon as his wealth reaches the level W^*. The threshold W^* depends on the parameters r, ψ, and y, as shown in (14). At retirement (and forever after), the agent's optimal level of consumption is $c_t = rW^*$. The Euler equation (23) and the wealth accumulation equation (1) tell us that prior to retirement the agent follows one of the non-backward-bending paths depicted in Figure 2. For a given value of W^*, which one of these paths will the agent follow?

Figure 2 Solution Paths with Quadratic Utility

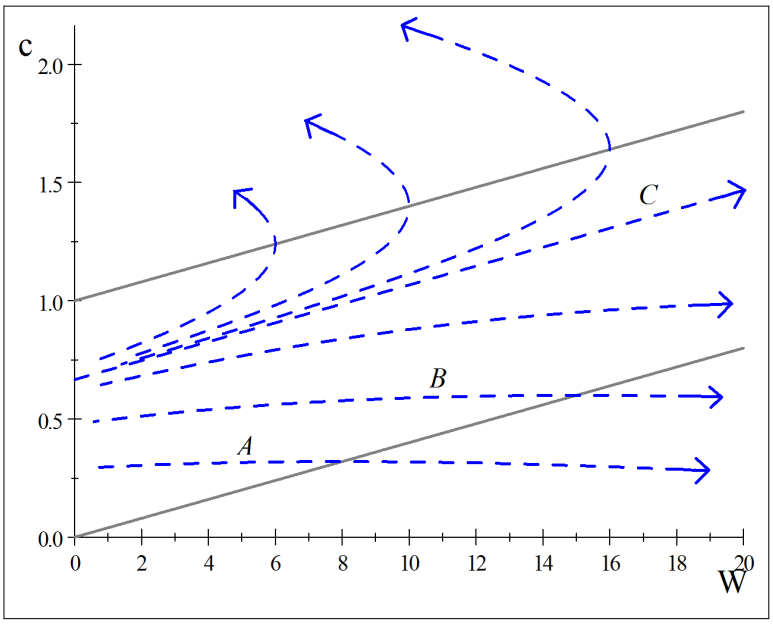

Notes: Parameters used in this plot: $y = 1$, $r = 0.04$, $\lambda = 0.02$.

Since the agent's wealth and consumption remain constant in retirement, in Figure 2 the evolution of wealth and consumption after voluntary retirement is represented by a single point for each threshold value W^*. That point is (W^*, rW^*). Thus, once the agent retires, the time path of his wealth and consumption is absorbed at (W^*, rW^*). It is easy to see in Figure 2 that for each value $W^* \geq 0$ there is a unique solution path $\{(W_t, c_t); t \geq 0\}$ leading to the point (W^*, rW^*). That path is the optimal path for the agent whose retirement wealth threshold is W^*. Why this path? Because all other paths would imply a jump in consumption at retirement, which the agent wants to avoid. Because his utility function is concave, the agent prefers a smooth consumption path with no jump at retirement. The level of consumption in voluntary retirement, rW^*, thus determines the optimal accumulation path the agent follows prior to retirement. It is the single path that intersects the line $c = rW$ at $W = W^*$.

For example, the solution path labelled A in Figure 2 crosses the line $c = rW$ at $W = 8$. Thus, this solution path is optimal for the agent whose desired retirement wealth is $W^* = 8$. Likewise, the solution path B is optimal for the agent whose desired retirement wealth is $W^* = 15$. The solution path C follows a straight line parallel to the line $c = rW$ and therefore never crosses it. This solution path is optimal for the agent whose retirement threshold is $W^* = \infty$, which means the agent plans to never retire voluntarily. From the formula (14) we see that $W^* = \infty$ when $\psi = 0$, i.e., when the agent does not at all value the extra leisure he can obtain by retiring. Since the value of the extra leisure is zero for this agent, it is natural that he never chooses to retire. This is the case studied in standard infinite-horizon models of optimal saving and consumption decisions, e.g., Ljungqvist and Sargent (2004, Ch. 16).

Note that the argument implying that the agent's preferred path of wealth and consumption is the one that leads to the point (W^*, rW^*) does not use the assumption of quadratic preferences. Rather, this argument is based on the agent's preference for smooth consumption, and so it applies to any concave utility function u. Thus, although the shape of the optimal path of wealth accumulation and consumption $\{(W_t, c_t); t \geq 0\}$ will in general not be the same as that presented in Figure 2, it will be true for any concave utility function that the optimal accumulation path is the unique solution path that leads to the point (W^*, rW^*).[6]

Also, the phase diagram in Figure 1, which works for any concave u, shows that the only way for a solution path to approach (W^*, rW^*) is through the middle band $rW < c < rW + y$ of the state space (W, c), where W_t and c_t are both strictly increasing over time. This confirms the validity of the assumption we made in Section 2 about wealth following an increasing path prior to retirement.

Planned Retirement and Optimal Saving Rate

Figure 2 provides a clear illustration of the following point. When the option to retire is added to the standard, infinite-lived-agent model of optimal consumption and saving, the model's prediction on the optimal amount of saving unambiguously increases.

In its textbook version (see Ljungqvist and Sargent [2004]), the standard model of optimal consumption and saving decisions abstracts

[6] See Appendix B for the case of the CRRA utility function.

from retirement. Labor income fluctuates stochastically, but the agent does not have an option to retire and end his flow of labor income altogether. The model we consider in this article assumes a particularly simple form of stochastic income fluctuations (labor income is a step function initially positive and jumping down to zero at a random date τ_f), but allows for endogenous retirement.

In the special case with $\psi = 0$, our model is a version of the standard model with no retirement. As we saw earlier, when $\psi = 0$ the agent never retires voluntarily, so his labor income effectively follows an exogenous process, as in the standard model. From (26) we have in that case that the optimal fraction of labor income y to be saved by the agent is

$$\begin{aligned} \frac{dW_t/dt}{y} &= \frac{rW_t + y - c_t}{y} \\ &= 1 - \frac{r}{r+\lambda} \\ &= \frac{\lambda}{r+\lambda}. \end{aligned}$$

Thus, the standard model without retirement would predict $\frac{\lambda}{r+\lambda}$ as the agent's optimal rate of saving out of labor income. With positive utility of leisure, $\psi > 0$, our model predicts voluntary retirement in finite time τ as well as a higher optimal rate of saving prior to retirement. From (25) we have

$$\begin{aligned} \frac{dW_t/dt}{y} &= \frac{rW_t + y - c_t}{y} \\ &= 1 - \frac{r}{r+\lambda}\left(1 - e^{-(r+\lambda)(\tau - t)}\right) \\ &> \frac{\lambda}{r+\lambda}, \end{aligned}$$

where the strict inequality follows from the agent's time to retirement being finite, i.e., $\tau - t < \infty$. Given that people do save for retirement in reality, models that disregard retirement underpredict the optimal rate of saving. Figure 2 shows this very clearly: The infinite-horizon solution path that runs parallel to the line $c = rW$ is everywhere above all solution paths that cross this line.

That the optimal saving rate should be higher when agents save for retirement is of course very intuitive. With retirement, the agent's labor income is more front-loaded relative to the case without retirement. To

smooth this front-loading out, the agent saves more. Our analysis lets us see this point clearly in Figure 2.[7]

Time to Retirement

As we see in (25), the agent's optimal consumption at time t depends on the agent's current wealth W_t and the amount of time left before his planned retirement, $\tau - t$. The agent's target retirement wealth level W^* is given in (14). But how do we find the agent's target retirement time τ?

From the law of motion for wealth prior to retirement, (1), we have

$$W_\tau = W_t + \int_0^{\tau-t} dW_{t+s} = W_t + \int_0^{\tau-t} (y - (c_{t+s} - rW_{t+s}))ds.$$

Using the retirement condition $W_\tau = W^*$ and the consumption rule (25) we have

$$\begin{aligned} W^* &= W_t + \int_0^{\tau-t} \left(1 - \frac{r}{r+\lambda}(1 - e^{-(r+\lambda)s})\right) y ds \\ &= W_t + \frac{\lambda}{r+\lambda}(\tau - t)\, y + \frac{r}{(r+\lambda)^2}\Big(1 - e^{-(r+\lambda)(\tau-t)}\Big) y. \quad (27) \end{aligned}$$

For any given values for r, λ, y, W^*, and W_t, this condition can be solved for the agent's planned time to retirement $\tau - t$. Because the right-hand side of (27) is increasing in both $\tau - t$ and W_t, the time to retirement is decreasing in current wealth.[8]

In sum, the dynamics of consumption and wealth accumulation are as follows. The agent determines his target retirement wealth level W^*, as in (14). Then the agent follows the unique wealth accumulation and consumption path $\{(W_t, c_t); t \geq 0\}$ in Figure 2 that leads to the point (W^*, rW^*). How far away from the retirement point the agent starts on this path depends on his initial wealth W_0. Unless he loses his job before reaching wealth W^*, the agent retires voluntarily as soon as his wealth attains W^*. After retirement, he consumes at the constant rate rW^* and his financial wealth remains constant at W^*. Thus, the solution path the agent follows in Figure 2 is absorbed at the point (W^*, rW^*). If the agent is forced into involuntary retirement at some date $\tau_f < \tau$, i.e., when his wealth is $W_{\tau_f} < W^*$, his consumption jumps down at τ_f from his preferred accumulation path to the point $(W_{\tau_f}, rW_{\tau_f})$, and is absorbed there. That is, consumption stays constant in retirement at

[7] The same is true in the case of CRRA preferences we discuss in Appendix B. See Figure 7.

[8] This also confirms that wealth grows over time while the agent is working.

the level $rW_{\tau_f} < rW^*$ and financial wealth stays constant at $W_{\tau_f} < W^*$.

Permanent Income Hypothesis

It is well known, see Ljungqvist and Sargent (2004, Ch. 16), that with quadratic preferences the optimal saving and consumption rule satisfies the permanent income hypothesis (PIH). Under PIH, it is optimal for the agent at each point in time to consume simply the income from his total wealth, where total wealth includes financial wealth and human capital. Human capital of an agent is defined as the present value of all the labor income that the agent is yet to earn. Thus, permanent income has two components: the income from currently held financial wealth and the income from currently held human capital.

That PIH holds in our model is most clearly evident in (26), i.e., in the case with $\psi = 0$ in which the agent never retires voluntarily. If the agent's stock of financial wealth is W_t, his permanent income from it is rW_t because, as we saw earlier, if the agent consumes rW_t, he never depletes his financial wealth and therefore is able to maintain this consumption forever. If the agent is working at t, the expected present value of his future labor income is

$$\mathbb{E}\left[\int_0^{\tau_f} e^{-rs} y ds\right] = \frac{y}{r+\lambda}.$$

Permanent income from human capital $y/(r+\lambda)$ is $ry/(r+\lambda)$ because this is the perpetual flow equivalent of stock $y/(r+\lambda)$. According to PIH, with financial wealth W_t and with human capital $y/(r+\lambda)$, the agent's consumption at t should be $r(W_t + \frac{y}{r+\lambda})$, which it is, as we see in (26).

The agent's optimal rule for consumption and saving obeys PIH also when he chooses to voluntarily retire at a future date τ. In this case, the agent's human capital as of $t < \tau$ is

$$\mathbb{E}\left[\int_0^{\min\{\tau_f,\tau\}} e^{-rs} y ds\right] = \frac{y}{r+\lambda}\left(1 - e^{-(r+\lambda)(\tau-t)}\right). \tag{28}$$

Thus, (25) is consistent with PIH because the agent in this case as well consumes exactly the return on his financial and human capital at all times. Note that the value in (28) is less than $y/(r+\lambda)$ because a part of expected future income is lost due to the agent's planned retirement at τ. The closer t is to τ, the lower the agent's human capital. Because the agent saves at all $t < \tau$, however, his financial wealth W_t grows as t gets closer to τ. It fact, financial wealth grows faster than human capital declines, and so the agent's permanent consumption increases

over time for as long as the agent does not lose his job. As wealth approaches W^*, the agent's human capital goes down to zero smoothly, and his consumption increases smoothly to rW^*. If the agent loses his job involuntarily at some date τ_f before his financial wealth reaches W^*, the agent's human capital discontinuously jumps down to zero and his permanent consumption jumps down to just the return on his financial assets, rW_{τ_f}.

6. RETIREMENT SAVING AND THE JOB LOSS RISK

In this section, we study the dependence of the optimal consumption and saving plan on the job loss rate λ.

Proposition 1 *At any $W_t < W^*$, the larger the job loss intensity λ, the lower consumption c_t, the higher the wealth accumulation rate dW_t/dt, and the shorter the time to planned retirement $\tau - t$. If $\lambda \to \infty$, then $c_t \to rW_t$, $dW_t/dt \to y$, and $\tau - t \to (W^* - W_t)/y$.*

Proof. In Appendix A. ■

This proposition shows that if we compare two agents identical in all respects (same wealth, same income) except for the job loss rate λ, the agent with larger job loss risk will consume less and save more than the other agent. Intuitively, the agent with higher λ holds less human capital than the agent whose λ is lower. The labor income flow rate y, the same for both agents, therefore, is higher relative to total wealth for the agent with higher λ, and so he will save a larger portion of y than the other agent. In other words, labor income y is less permanent for the agent with higher λ, so intertemporal consumption smoothing implies he will save more. Figure 3 illustrates this point by plotting optimal paths for consumption and wealth for several values of λ.

This comparative statics result can be interpreted as showing the agent's response to a completely unanticipated shock to the job loss risk the agent faces in our model. Under the parametrization used in Figure 3, if $\lambda = 0.02$, the agent will follow the highest of the three accumulation paths plotted in that figure. If at some point prior to retirement the intensity parameter λ jumps to 0.1, the agent will switch at that point to the lowest of the three paths. This means that his saving rate will increase and consumption will decrease without any change to his current income. This example illustrates a response of

Figure 3 Optimal Accumulation Paths for Three Values of the Job Loss Rate

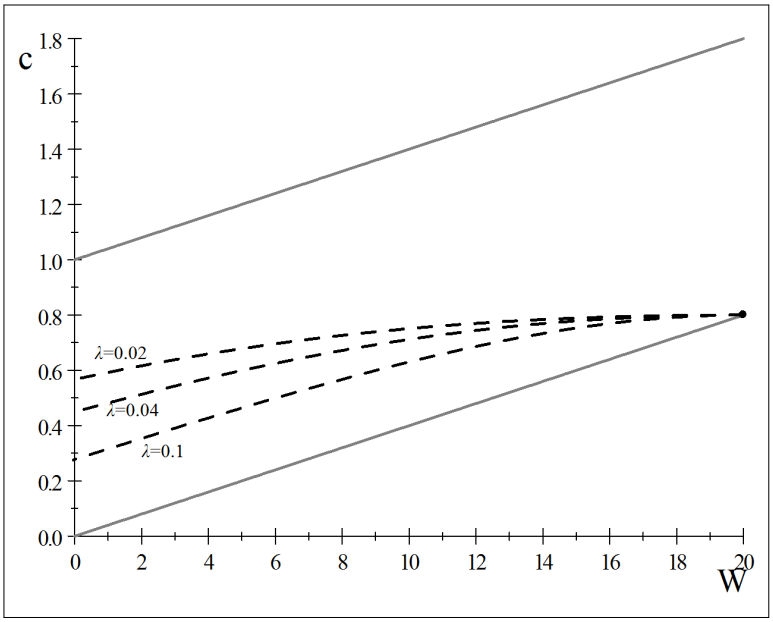

Notes: Other parameters used in this plot: $r = 0.04$, $y = 1$, ψ such that $W^* = 20$.

optimal consumption to a change in the expectations the agent holds about the future.[9]

If λ is very large, then, as Proposition 1 shows, the agent saves close to 100 percent of his labor income y and consumes close to rW_t. This again is intuitive, as when λ is large, the agent's human capital is close to zero and financial wealth constitutes the bulk of his total wealth. The level of permanent consumption he can afford is thus close to the level he could maintain if he had lost his job already, which with assets W_t is exactly rW_t.

[9] The discussion in this paragraph assumes that the agent does not anticipate that λ could jump, and that once it does jump, the agent firmly expects it to never jump again. Clearly, this is an oversimplification. We can expect, however, that our conclusion here continues to hold when the jumps in λ are anticipated. That is, although the shape of the accumulation paths in Figure 3 must be adjusted, we expect consumption to decline when λ increases in a model in which changes in λ are anticipated by the agent.

Figure 4 Dependence of Time to Retirement on the Job Loss Rate

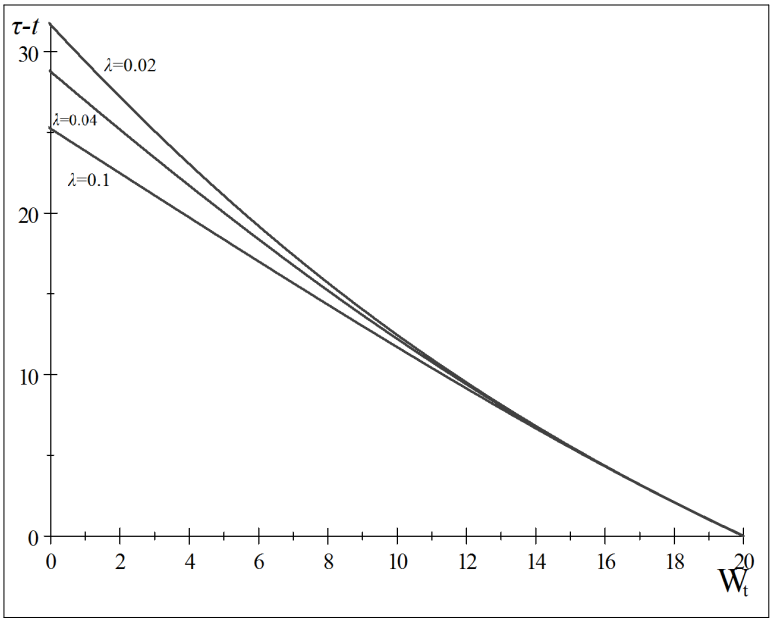

Notes: Parameters used in this plot are the same as those in Figure 3.

Proposition 1 also shows that conditional on not losing the job before the planned retirement date τ, the agent with larger λ will reach his desired retirement wealth level W^* faster. Note that the theoretical limit with $\lambda \to \infty$ of the time to retirement, $(W^* - W_t)/y$, is consistent with the agent saving 100 percent of his labor income and living only off his asset income already before retirement.

Figure 4 plots the planned time to retirement $\tau - t$ against wealth W_t for several values of λ. In the example presented in that figure, we have $r = 0.04$, which makes one unit of time correspond roughly to one year. Annual labor income y is normalized to 1, and $W^* = 20$, which means that the agent wants to retire as soon as his stock of wealth reaches the equivalent of 20 years of labor income. With $\lambda = 0.02$, meaning the event of involuntary and permanent job loss on average occurs once in 50 years, the agent who starts out with zero initial wealth plans to retire after about 32 years. With $\lambda = 0.04$, i.e., when

involuntary retirement is a once-in-a-quarter-century event, the agent plans to retire after roughly 28.5 years. With $\lambda = 0.01$, the permanent job loss shock becomes a once-in-a-decade event in expectation. In this case, the agent plans to retire after 25 years. These numbers illustrate the fact that the agent can only partially insure himself against the permanent job loss shock in our model. With $\lambda = 0.02$, the probability that an agent with zero wealth reaches voluntary retirement is $e^{-0.02*32}$, which equals roughly 53 percent. For an agent with the same initial wealth but with $\lambda = 0.1$, this chance is only $e^{-0.1*25}$, i.e., about 8 percent.

Differentiating with respect to λ the expressions for optimal consumption in (25) and (26), it is easy to check that the response of c_t to a given change in λ is stronger the longer the agent's planned time to retirement $\tau - t$. In particular, the response of consumption to changes in the job loss risk is the strongest in the case of $\psi = 0$, where the agent plans to never retire voluntarily. This result is very intuitive given that fast planned retirement means human capital is a small portion of the agent's total wealth.

7. ADDITIONAL COMPARATIVE STATICS RESULTS

With closed-form solution for the optimal path of saving and consumption, we can provide several additional comparative statics results.

We saw already in Figure 2 how the optimal path of consumption and wealth accumulation depends on the parameter ψ. In (14), higher leisure utility ψ implies a lower retirement threshold W^*. In Figure 2, we see that lower W^* means faster retirement with a higher saving rate along the optimal accumulation path.

We can also examine how consumption, saving, and the retirement decision depend on the level of labor income y. We know from (14) that the retirement threshold wealth level W^* is increasing in y. Using (25), it is not hard to show that if two agents have the same financial wealth W_t and face the same job loss rate λ, the agent with higher labor income y will consume more and retire later. The numerical example given in Figure 5 illustrates this point. In that figure, paths leading to lower retirement points are everywhere below those leading to higher retirement wealth thresholds. Those higher paths correspond to higher labor income y earned during employment.

Finally, we examine how the solution to the agent's optimal consumption, saving, and retirement problem depends on the real interest rate r. Dashed lines in Figure 6 show three accumulation paths, each optimal at a different level of r. That the retirement wealth threshold

Figure 5 Consumption and Wealth Accumulation for Three Different Values of Labor Income

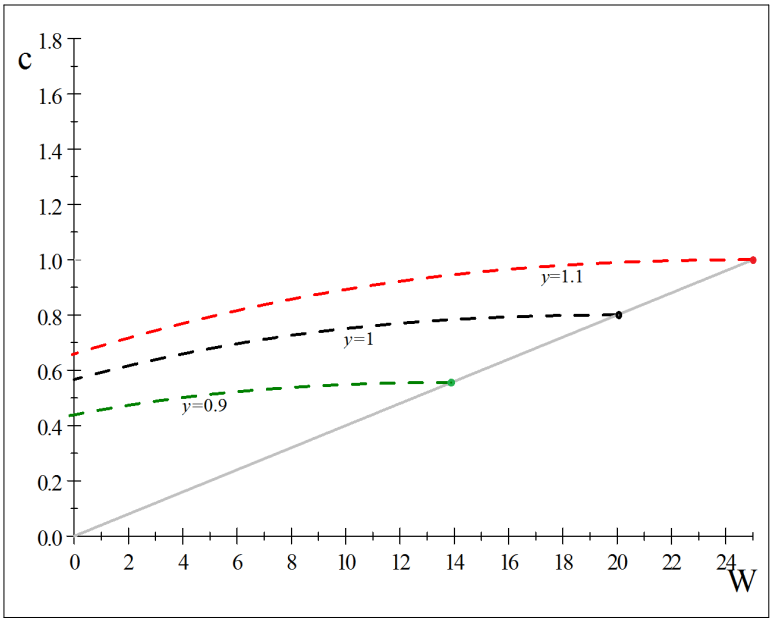

Notes: Other parameters as in Figure 2.

W^* is lower at higher r can be seen from the terminal points of the accumulation paths in Figure 6, or directly from the formula for W^* given in (14). Since the marginal value of wealth in retirement $u'(rW_t)$ decreases in r, it is intuitive that when r is higher the agent chooses to give up labor income y in return for utility ψ earlier, i.e., at a lower wealth threshold. Figure 6 shows that prior to retirement, at higher r both the agent's consumption c_t and his interest income rW_t are higher. Interest income is represented in Figure 6 by the straight lines rW connecting the origin to the terminal points of the optimal accumulation paths. How the wealth accumulation rate $dW_t/dt = rW_t - c_t + y$ depends on r is determined by the magnitudes of c_t and rW_t. In fact, the rate of wealth accumulation is increasing in r at high levels of wealth W_t and decreasing in r at low levels of W_t. For example, at $W_t = 12$, the vertical distance between (any two) dashed lines (representing c_t) is smaller than the vertical distance between the solid lines

Figure 6 Optimal Consumption and Wealth Accumulation Paths at Three Levels of the Interest Rate

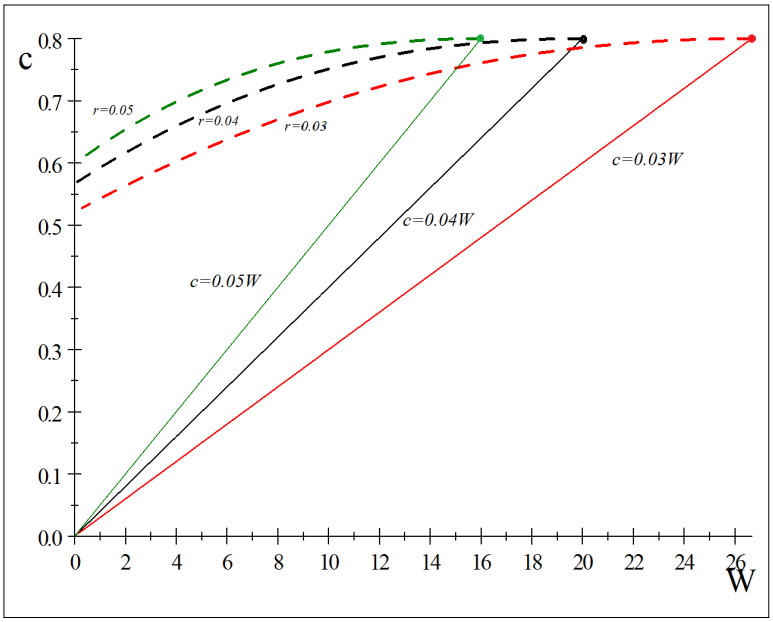

(representing rW_t). At $W_t = 1$, the opposite is true. The cumulative effect of these differences on the agent's wealth is positive. With some algebra that we omit here, it can be shown that the agent retires faster when r is higher. That is, for any given W_t the agent's time to planned retirement, $\tau - t$, is shorter the higher the interest rate r.

8. CONCLUSION

This article studies optimal consumption and saving decisions in an infinite-horizon model that allows for endogenous retirement. Relative to the standard model with no retirement, the optimal saving rate is higher. An increase in the job loss risk decreases consumption, even without the actual job loss occurring. Accounting for retirement subdues the magnitude of the response in consumption to changes in the job loss risk. These results may be important for quantitative analyses of observed consumption and saving decisions.

The strong assumption we make on the shape of the agent's profile of labor income lets us abstract in this article from borrowing constraints. Since his income can only decrease, the agent never wants to borrow in our model, so the no-borrowing constraint is natural in our analysis, and it never binds. Increasing and hump-shaped paths of income are standard in life-cycle models. Incorporating such paths into our model would require an extension of our analysis accounting for the possibility of binding borrowing constraints.

Our analysis of optimal saving for and timing of retirement can be extended to study other types of actions for which savings are important. For instance, due to down-payment requirements, the optimal timing of a house purchase by a household will depend on the financial wealth of the household. Our analysis in this article can be adapted to study jointly the saving decisions and the optimal timing of this purchase.

APPENDIX: APPENDIX A

Proof of Lemma 1

Multiplying (1) by r and subtracting it from (23) we obtain a linear differential equation

$$\frac{d(c_s - rW_s)}{ds} = (r+\lambda)(c_s - rW_s) - ry,$$

which, with the notation $z_s = c_s - rW_s$, we can write more compactly as

$$\frac{dz_s}{ds} = (r+\lambda)z_s - ry. \tag{29}$$

The solution to this equation is standard. Differentiating $z_s e^{-(r+\lambda)s}$, we have

$$d\left(z_s e^{-(r+\lambda)s}\right) = dz_s e^{-(r+\lambda)s} - (r+\lambda)z_s e^{-(r+\lambda)s}ds = -rye^{-(r+\lambda)s}ds,$$

where the second equality uses (29). Integrating from t to τ and solving for z_t yields

$$z_t = \frac{ry}{r+\lambda}\left(1 - e^{-(r+\lambda)(\tau-t)}\right) + z_\tau e^{-(r+\lambda)(\tau-t)}.$$

Writing the boundary condition $c_\tau = rW_\tau + y$ as $z_\tau = y$ and using it in the above general solution gives us (24). With the boundary condition $c_\tau = rW_\tau$, we have $z_\tau = 0$, which gives us (25). For (26), we take a limit of (25) with $\tau \to \infty$.

Proof of Proposition 1

Write (25) as

$$c_t - rW_t = \frac{r}{r+\lambda} y \left(1 - e^{-(r+\lambda)(\tau - t)}\right)$$

and note that the right-hand side of this equality is decreasing in λ and goes to zero as $\lambda \to \infty$. This proves the proposition's conclusions about c_t and, using (1), dW_t/dt. Next, write (27) as

$$W^* - W_t = \frac{\lambda}{r+\lambda} (\tau - t)\, y + \frac{r}{(r+\lambda)^2} \left(1 - e^{-(r+\lambda)(\tau - t)}\right) y$$

and check that the right-hand side of this equality is strictly increasing with respect to both λ and $\tau - t$. Because W^* does not depend on λ, the left-hand side is constant. Thus, the time to retirement $\tau - t$ must decrease when λ increases to keep the right-hand side constant.

APPENDIX: APPENDIX B

Figure 7 provides the analog of Figure 2 for a nonquadratic utility function u. In particular, this figure depicts numerically computed solution paths to the system of differential equations (1)–(22) for constant relative risk aversion (CRRA) preference represented by the utility function u of the form

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}.$$

Qualitatively, these graphs are similar to one another for all values of $\gamma > 0$.

Our analysis determining the optimal accumulation path for a given voluntary retirement wealth threshold W^* from Section 5 is unchanged. The main difference between CRRA preferences and quadratic preferences is that the permanent income hypothesis does not hold under CRRA preferences. With CRRA preferences, agents have the so-called precautionary motive for saving, which is absent under quadratic preferences. When the precautionary saving motive is present, the agent will increase the amount he saves in response to an increase in the riskiness of his income process, holding his expected income constant. (See Ljungqvist and Sargent [2004] for a general discussion of precautionary savings.)

In Figure 7, precautionary savings are best seen by comparing the solution path labelled C with the dotted line labelled PIH. The

Figure 7 Solution Paths for CRRA Preferences

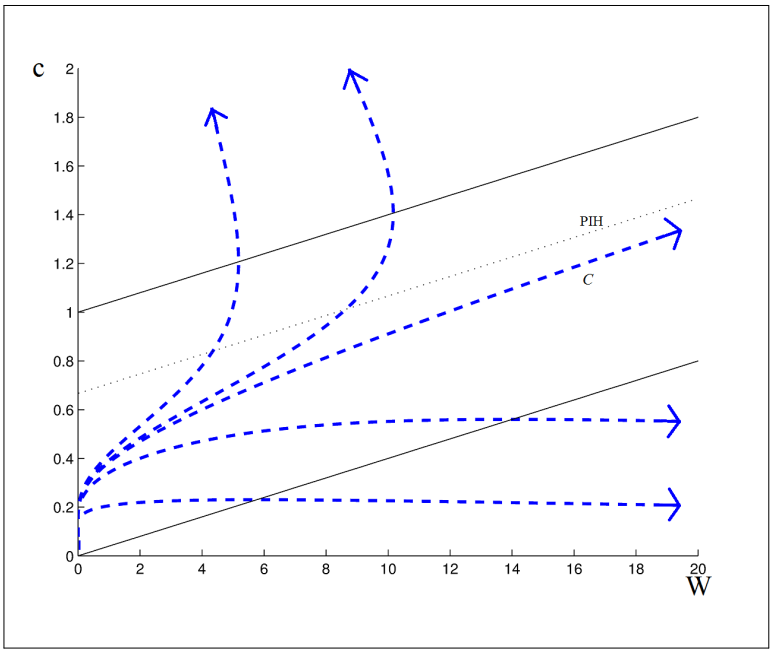

solution path C in Figure 7 is analogous to the solution path C in Figure 2. It is the single solution path that never leaves the middle band of the graph bounded by the lines $c = rW$ and $c = rW + y$. It is the optimal solution path under CRRA preferences for an agent whose $\psi = 0$, i.e., an agent who never chooses to retire voluntarily. The line labelled PIH in Figure 7 is the solution path that would be optimal for that agent if he did not have a precautionary saving motive (i.e., it is an exact replica of the solution path C from Figure 2). At any level of wealth W_t, the vertical distance in Figure 7 between line PIH and the solution path C measures precautionary saving of the agent with CRRA preferences. As we see, precautionary saving is positive at all wealth levels and its magnitude decreases in W_t. In fact, solution C converges to line PIH as $W_t \to \infty$.

As in Figure 2, each solution path crossing the line $c = rW$ is an optimal accumulation path for an agent whose value of the leisure preference parameter ψ is strictly positive. In these cases, as well, precautionary saving can be seen by comparing corresponding solution

paths in Figures 2 and 7. For any given voluntary retirement wealth threshold $W^* > 0$, the solution path leading to the retirement point $(W, c) = (W^*, rW^*)$ will in Figure 2 be strictly above the path leading to the same path in Figure 7. The vertical distance between these two paths will represent precautionary saving. All solution paths in Figure 7 converge to zero consumption when wealth goes to zero, while in Figure 2 they do not. Comparing solutions with voluntary retirement in finite time, as in the case of no voluntary retirement, we thus see that the precautionary saving motive is the strongest at very low wealth levels.

APPENDIX: APPENDIX C

In this appendix, we discuss an extension of our model in which the option value of delaying retirement is positive.

Let us add a positive labor income shock to our model. That is, instead of assuming that at all times prior to retirement the agent's labor income is constant, suppose it can increase from y to $\bar{y} > y$. Suppose this upward jump arrives with Poisson intensity $\sigma > 0$. Also, let's assume the job loss shock is independent of the level of income and, as before, it arrives with Poisson intensity $\lambda > 0$.

We will show that with this positive income shock, the agent with income y will not choose to retire as soon as his wealth reaches the threshold W^* but rather will prefer to keep working. The reason why the agent prefers to keep working is that postponing retirement has a positive option value when there is a chance that his labor income increases in the future.

Let $\bar{J}(W_t)$ be the maximal utility value the agent can obtain when his income is already high, i.e., $\bar{y}$. Because once it hits $\bar{y}$ income stays constant until retirement; our previous analysis applies: The agent whose income is $\bar{y}$ will want to retire exactly when his wealth hits the threshold

$$\bar{W}^* = \frac{1}{r} u'^{-1}\left(\frac{\psi}{\bar{y}}\right) > \frac{1}{r} u'^{-1}\left(\frac{\psi}{y}\right) = W^*.$$

We will show, however, that with low income y the agent will not want to retire as soon as his wealth reaches W^*. That is, the retirement rule with wealth threshold level W^* that we obtained in Section 3 is no longer optimal for the agent.

Consider the following strategy for an agent whose labor income is low, y, and whose wealth is W_t. Suppose the agent works over a small time interval $[t, t+h)$. By time $t+h$, three things can happen. The agent loses his job, gets a promotion, or neither. Suppose the agent behaves optimally after a promotion thus obtaining in that event the value $\bar{J}(W_{t+h})$. In the event of the job loss, he behaves optimally in retirement and so he obtains $V(W_{t+h})$. If neither promotion nor job loss happen, suppose the agent retires voluntarily at $t+h$, thus obtaining the value $V(W_{t+h})$ in this event as well. Thus, the agent's strategy is to postpone retirement by h and see if he gets a promotion. If he does not, he quits. Denote by $\bar{V}^h(W_t)$ the value that this strategy gives the agent as of date t.

We proceed analogously to Section 4. We have

$$\begin{aligned}\bar{V}^h(W_t) &= \max_c \Big\{ \textstyle\int_0^h e^{-rs} u(c) ds \\ &\quad + e^{-rh} \Big(e^{-\lambda h}(1-e^{-\sigma h})\bar{J}(W_{t+h}) + e^{-\lambda h} e^{-\sigma h} V(W_{t+h}) \\ &\quad + \left(1 - e^{-\sigma h}\right) V(W_{t+h}) \Big) \Big\}\end{aligned}$$

with wealth following (1) between t and $t+h$, as the agent works between t and $t+h$. Because h is small, we use the first-order approximation and express $\bar{V}^h(W_t)$ as

$$\begin{aligned}\bar{V}^h(W_t) &= \max_c \big\{ V(W_t) + \big(u(c) + \sigma \bar{J}(W_t) - (r+\sigma) V(W_t) \\ &\quad + V'(W_t) \tfrac{dW_t}{dt} \big) h \big\}.\end{aligned}$$

Next, we compare this value to the value of retiring immediately at t, which we know to be $V(W_t)$. We have that the value of postponing retirement by at least h is strictly preferred to retiring immediately, i.e., $\bar{V}^h(W_t) > V(W_t)$, if and only if

$$\begin{aligned}&\max_c \big\{ V(W_t) + \big(u(c) + \sigma \bar{J}(W_t) - (r+\sigma) V(W_t) \\ &\quad + V'(W_t)(rW_t + y - c) \big) h \big\} \\ > \;& \max_c \big\{ V(W_t) + \big(u(c) + \psi - rV(W_t) + V'(W_t)(rW_t - c) \big) h \big\}.\end{aligned}$$

Dividing by h, simplifying terms, and removing the identical maximization problems with respect to c on both sides of this condition simplifies it to

$$\sigma \left(\bar{J}(W_t) - V(W_t) \right) + V'(W_t) y > \psi.$$

Now we note that $\bar{J}(W^*) - V(W^*) > 0$ because with high labor income $\bar{y}$ the agent only wants to retire with wealth $\bar{W}^* > W^*$ and not earlier. By definition of W^*, we have $V'(W^*)y = \psi$. Therefore,

$$\sigma \left(\bar{J}(W^*) - V(W^*) \right) + V'(W^*) y > \psi.$$

This means that with low income y and wealth W^*, the agent prefers to postpone retirement. The reason for this is that the term $\sigma\left(\bar{J}(W^*) - V(W^*)\right)$ is strictly positive. This term represents the option value of delaying retirement. For as long as the agent is not retired, he has a chance to see his labor income increase, in which case he would prefer to continue working until his wealth reaches $\bar{W}^*$. Because retirement is permanent, by retiring with wealth $W^* < \bar{W}^*$, the agent closes this possibility to himself or, in other words, gives up this option. By delaying retirement, he keeps this option open.

By continuity, the above condition holds in the neighborhood of W^*, i.e., also for some wealth $W_t > W^*$. At that wealth level we have $V'(W_t)y < \psi$, i.e., in terms of his current payoff the agent would be strictly better off to retire immediately. He does not choose to do so, however, because the option value $\sigma\left(\bar{J}(W_t) - V(W_t)\right)$ is larger than the payoff from retiring $\psi - V'(W_t)y$.

REFERENCES

Bewley, Truman. 1977. "The Permanent Income Hypothesis: A Theoretical Formulation." *Journal of Economic Theory* 16 (December): 252–92.

Chamberlain, Gary, and Charles A. Wilson. 2000. "Optimal Intertemporal Consumption Under Uncertainty." *Review of Economic Dynamics* 3 (July): 365–95.

Farhi, Emmanuel, and Stavros Panageas. 2007. "Saving and Investing for Early Retirement: A Theoretical Analysis." *Journal of Financial Economics* 83 (January): 87–121.

Friedman, Milton. 1957. *A Theory of the Consumption Function.* Princeton, N.J.: Princeton University Press.

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

Karatzas, Ioannis, and Steven E. Shreve. 1997. *Brownian Motion and Stochastic Calculus, Second Edition.* New York: Springer Science+Business Media, Inc.

Kingston, Geoffrey H. 2000. "Efficient Timing of Retirement." *Review of Economic Dynamics* 3 (October): 831–40.

Marcet, Albert, Francesc Obiols-Homs, and Philippe Weil. 2007. "Incomplete Markets, Labor Supply and Capital Accumulation." *Journal of Monetary Economics* 54 (November): 2,621–35.

Merton, Robert C. 1971. "Optimum Consumption and Portfolio Rules in a Continuous-Time Model." *Journal of Economic Theory* 3 (December): 373–413.

Leland, Hayne E. 1994. "Corporate Debt Value, Bond Covenants, and Optimal Capital Structure." *Journal of Finance* 49 (September): 1,213–52.

Ljungqvist, Lars, and Thomas J. Sargent. 2004. *Recursive Macroeconomic Theory, Second Edition*. Cambridge, Mass.: The MIT Press.

Ljungqvist, Lars, and Thomas J. Sargent. Forthcoming. "Career Length: Effects of Curvature of Earnings Profiles, Earnings Shocks, Taxes, and Social Security." *Review of Economic Dynamics*.

Pindyck, Robert S. 1991. "Irreversibility, Uncertainty, and Investment." *Journal of Economic Literature* 29 (September): 1,110–48.