



Theses and Dissertations

2013-06-19

Investment-Consumption with a Randomly Terminating Income

James Benjamin Taylor Jr.
Brigham Young University - Provo

Follow this and additional works at: <https://scholarsarchive.byu.edu/etd>



Part of the [Mathematics Commons](#)

BYU ScholarsArchive Citation

Taylor, James Benjamin Jr., "Investment-Consumption with a Randomly Terminating Income" (2013). *Theses and Dissertations*. 4069.
<https://scholarsarchive.byu.edu/etd/4069>

This Dissertation is brought to you for free and open access by BYU ScholarsArchive. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of BYU ScholarsArchive. For more information, please contact scholarsarchive@byu.edu, ellen_amatangelo@byu.edu.

Investment-Consumption with a Randomly Terminating Income

James B. Taylor, Jr.

A dissertation submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

Jeffrey Humpherys, Chair
Scott Glasgow
Christopher Grant
Craig Merrill
Dennis Tolley

Department of Mathematics
Brigham Young University
June 2013

Copyright © 2013 James B. Taylor, Jr.

All Rights Reserved

ABSTRACT

Investment-Consumption with a Randomly Terminating Income

James B. Taylor, Jr.
Department of Mathematics, BYU
Doctor of Philosophy

We develop a stochastic control model for an investor's optimal investment and consumption over an uncertain planning horizon when the investor is endowed with a defaultable income stream. The distributions of the random time of default and the random terminal time are prescribed by deterministic hazard rates, and the investor makes investments in a standard financial market with a bond and a stock, modeled by geometric Brownian motion. In addition, the investor purchases insurance against both default and the terminal date, the default insurance serving as a proxy for the investor's disutility for default. We approximate the original continuous-time problem with a sequence of discrete-time Markov chain control problems by applying dynamic programming and the Markov chain approximation. We demonstrate how the problem can be solved numerically through a logarithmic transformation of the investor's wealth variable, even when the utilities are CRRA with large risk aversion parameter. The model and computational approach are applied to a retiree's optimal annuity decision in the presence of default risk, and we demonstrate that default risk can lead a retiree to annuitize significantly smaller proportions of savings, even when a portion of the defaulted annuity can be recovered, than is traditionally considered optimal by the retirement-finance community. Hence, we show that credit risk may play an important role in resolving the annuity puzzle.

Keywords: annuity puzzle, random endowment

ACKNOWLEDGMENTS

*Our doubts are traitors,
And make us lose the good we oft might win,
Through fearing to attempt.*

– William Shakespeare, *Measure for Measure*

There are several people who helped make this dissertation possible, to whom I express gratitude. First, I am grateful to Jeffrey Humpherys for introducing me to the question that gave rise to this work as well as for his support, and I am grateful to my advisement committee for their efforts in reviewing the manuscript and offering helpful comments. I am grateful to David Babbel and Craig Merrill for their insights into the annuity puzzle and for including me in their study of it. I am grateful to my grandmothers, Janice and Sandra, for their perseverance, and I am grateful to my parents, James and Jacque, for the opportunities they provided me and for their love. I am grateful to my children, Candace, Joshua, Levi, and Luke, for their patience while I worked on “my problem”, and lastly, I am grateful to my wife, Natalie, for everything.

CONTENTS

Contents	iv
List of Tables	vi
List of Figures	vii
1 Introduction	1
2 Mathematical Background	7
2.1 Diffusions	7
2.2 Stochastic Control Theory	12
2.3 Markov Chain Approximation	16
2.4 Richard's Model	21
3 Investment Consumption with Randomly Terminating Income: Formulation	36
3.1 Probabilistic Framework	37
3.2 Control Variables	38
3.3 Wealth Processes	41
3.4 Cost-Functionals and Control Problem	46
3.5 Motivation	46
3.6 Markov Chain Approximation	52
4 Application: Life Annuities	61
4.1 Model Calibration	66
4.2 Computation and Discussion	67
5 Conclusion	70
5.1 Summary of Key Results	70

5.2 Future Research	72
A Proofs and Derivations	76
A.1 Derivation of (2.39)	76
A.2 Proof of Lemma 3.28	77
A.3 Derivation of (3.38)	80
A.4 Derivation of (3.41)	81
Bibliography	84

LIST OF TABLES

4.1	Summary of Market and Risk Parameters	67
4.2	Optimal Annuitization as Percentage of Wealth, by Recovery Limit	68

LIST OF FIGURES

4.1 Annuity Value, Recovery, and Loss at Default	64
--	----

1. INTRODUCTION

If people do not believe that mathematics is simple, it is only because they do not realize how complicated life is.

– John Von Neumann

Life annuities are important financial instruments that help retirees hedge against various risks they face during retirement, including especially longevity risk, or the risk of outliving savings and other accumulated assets. The academic literature on life annuities has catalogued the benefits that annuities offer to retirees over other investment opportunities, and indeed, several authors, including [21] and [3], have suggested that retirees might optimally allocate their savings for the retirement period by placing much if not all of their savings into annuities. Nevertheless, the relatively small historical and current size of the annuity market suggests that retirees do not annuitize, for whatever reason, near levels considered optimal by researchers. This discrepancy between optimal “rational” levels of annuitization, on the one hand and realized levels of annuitization, on the other hand, is often referred to in academic financial economics as the *annuity puzzle*.

Many authors have studied the annuity puzzle and have proposed resolutions. For instance, some authors, including [10], suggest that a retiree’s desire to leave a bequest at death could significantly reduce the optimal level of annuitization since no annuitized wealth can be recuperated by heirs at the time of death. Some authors, including [3], have noted that the possibility of sudden, sharp declines in health could greatly reduce the optimal level of annuitization. Other authors have offered other additional factors that could drive retirees away from full annuitization. (See [1] for a more comprehensive list). Certainly, many of the proposed resolutions do impact the annuity decision to some degree and perhaps combine in effect. It is noteworthy that many of the proposed explanations can be categorized as aversions to one of two types of risk: liquidity risk - the inability to access capital when it is needed — and credit or counterparty risk — the risk that the counterparty in a contract,

such as an insurance company offering an annuity, will default on their financial obligations. Nevertheless, in spite of numerous attempted resolutions of the annuity puzzle, there is a gap in the literature for quantitative demonstrations of how aversion to liquidity risk, credit risk, or any other risk can induce rational agents to annuitize at the level observed in real markets.

In this dissertation, we develop a stochastic control model and more specifically, an optimal investment-consumption model in order to study the effect of default risk on an investor's optimal asset allocation given that one of the assets provides a constant income stream until a random date but may default. The model admits several interesting financial and economic applications, though we focus primarily on its application to the optimal annuity decision in the presence of default risk.

An overview of the model, applied to the annuity problem for concreteness, is as follows. Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ be a filtered probability space with a filtration $\mathbb{F} = (\mathcal{F}_t)$ that satisfies the usual hypothesis and is rich enough to support a Brownian motion, $W = (W_t)$, a random time τ_M , corresponding to a retiree's time of death, and a random time τ_D , corresponding to the time of default of the annuity. The Brownian motion and random times are assumed to be mutually independent. The retiree has some savings, x , at retirement a portion of which can be used to purchase an annuity. The retiree then receives a continuous stream of income from the annuity at constant rate α until the earlier of the investor's uncertain time of death and the annuity's random time of default. There may be some guaranty program so that upon default, the investor only loses that portion of the annuity not ensured by the guaranty. Throughout retirement, the retiree must select a consumption rate, $c \geq 0$, and divide any capital between a risk-free asset S_0 and a risky asset S_1 . The capital invested in S_0 is denoted π_0 , and the capital invested in S_1 is denoted π . The retiree may borrow against the future income stream given that life insurance p_M is purchased against her uncertain time of death and default insurance p_D called a credit default swap, is purchased against the uncertain time of default, in part to guarantee any repayment. The debt may never

exceed the expected remaining income from the endowment. The decision variables, return on investments, and expected remaining income combine to define the investor's total wealth process X which is constrained below by the discounted future income before default and constrained below by any remaining, non-defaultable income after default.

The sudden change in the dynamics of the total wealth as well as the jump in the borrowing limit at the time of default necessitates the construction of a coupled pair of optimal control problems. The post-default cost functional J_2 is defined as the retiree's expected utility for consumption until death and utility for bequest at death. The post-default optimization problem is then to maximize J_2 over the decision variables:

Problem (D). Given the retiree's total wealth x at the time of default $t = \tau_D$ maximize the retiree's expected utility for bequest and consumption over the retiree's remaining lifetime,

$$D(t, x) = \sup_{(c, \pi, p_M)} J_2(t, x, c, \pi, p_M).$$

Additionally, determine the optimal policies, (c^*, π^*, p_M^*) , if they exist.

Importantly, D is itself a utility function, the post-default utility.

The pre-default cost functional J_1 consists of the retiree's expected utility for consumption until the time of default, the retiree's utility for bequest if the retiree dies before default occurs, and the retiree's post-default utility for consumption and bequest if default occurs during the retiree's lifetime. The pre-default optimization problem, and the main problem of this dissertation, is to optimize the pre-default cost functional for a given annuity rate α over the decision variables:

Problem (V). Given a fixed initial wealth x and endowment income rate α , maximize the pre-default cost functional, J_1 , over the admissible decision variables,

$$V(x; \alpha) = \sup_{(c, \pi, p_M, p_D)} J_1(x, c, \pi, p_M, p_D). \quad (1.1)$$

Additionally, determine the optimal policies, $(c^*, \pi^*, p_M^*, p_D^*)$, if they exist.

The investment-consumption model outlined above belongs to the class of investment consumption models that feature a random endowment. An early example of such a model is [14], Richard's investigation of optimal life insurance purchasing. More recent studies include [4], [7], and [19]. In particular, Davis and Vellekoop [19], citing a lack of explicit results in the random endowment literature, seek a closed-form solution for the optimal control policies in a simple optimal consumption model with a randomly terminating income. However, the duality techniques employed by Davis and Vellekoop produce a duality gap and therefore fail to provide an explicit solution. Furthermore, they note difficulty in obtaining even numerical approximations to the solution when working with utility functions that become singular near subsistence levels of consumption and wealth, such as the constant relative risk-aversion (CRRA) utility functions with negative risk-aversion parameter. Unfortunately, in the expected-utility paradigm, such utility functions are the most commonly used by researchers to model the risk-averse investment behavior encountered in real markets. One of the important contributions of this dissertation is the introduction of techniques that overcome many of the difficulties in working with a randomly terminating income and singular utility functions.

Our investment-consumption model shares many features with Richard's life insurance model. In his original paper, Richard obtains a closed-form solution for the value function and optimal policies using the dynamic programming principle and solving a non-linear partial differential equation. Essential to Richard's solution was an unconstrained set of decision variables. For instance, the financial agent could freely buy and sell life insurance, which does not reflect how a small investor can participate in the insurance market. Pliska and Ye [13] revisited Richard's model, imposing constraints more in line with real markets. Pliska and Ye employed the Markov chain approximation (MCA) in order to solve the constrained version of Richard's model numerically. MCA is a powerful tool for solving control problems that produces a sequence of discrete-time problems which can be solved numerically and whose solutions converge to the solution of the original, continuous-time control problem.

We similarly employ MCA to obtain a sequence of discrete-time problems that approximate the solutions to the continuous-time control problems.

The second important contribution of this dissertation is in application to a retiree's preference for annuities. Once $V(x; \alpha)$ is known for a given level of income rate α , it is straightforward to optimize value function V over income rates and thus determine the optimal allocation of wealth into the risky endowment. In particular, let $A(\alpha)$ be the price of a life annuity as a function of the annuity income rate α . Then the optimal annuity problem is:

Problem (A). Given a retiree with total savings x at retirement, find the retiree's optimal level of annuitization, α^* ,

$$\alpha^* = \arg \max_{\alpha} \{V(x; \alpha) \mid 0 \leq A(\alpha) \leq x\} \quad (1.2)$$

Our model permits the division of the purchased annuity into a defaultable portion and a non-defaultable portion. This is significant because in practice, a retiree does not necessarily lose the entire future annuity income in the event of a default. First, state authorities assume control of the insolvent company's assets, and a retiree can recuperate losses during liquidation of the assets. Second, every state has a limited guaranty program that replaces the retiree's annuity income, either directly or through purchase of a new annuity on the retiree's behalf. Our model allows us to treat the guaranteed annuity as non-defaultable and only that portion of the annuity beyond the guaranty limits as defaultable. As a base case, we suppose the full annuity is defaultable and demonstrate that risk of default in whole significantly reduces a retiree's appetite for life annuities. While the base case may seem unrealistic, in fact, regulations limit the ability of annuity sellers to advertise the existence of state guaranty programs so that indeed, many retirees may not account for annuity guaranty in their decision-making.

In addition to the case of total default, we also analyze the retiree's optimal annuity decision as the relative portions of defaultable and non-defaultable annuity are adjusted.

We find that the optimal annuity level moves to full annuitization as the guaranty program covers an increasingly greater percentage of the annuity. This suggests that there may be benefits to social welfare in making retirees aware of state guaranty programs as it could provide retirees greater confidence in annualizing savings and therefore, more effectively funding retirement.

The thesis proceeds as follows. In Chapter 2, we review the mathematical theory used to develop and solve the investment and consumption model, including basic theory of diffusions, basic stochastic control theory, and a short introduction to the Markov chain approximation. We illustrate the theory by providing a somewhat careful review of Richard's model including its solution through MCA. In Chapter 3, we develop the new investment consumption model, which incorporates a randomly terminating income over an uncertain time horizon. Also in Chapter 3 we apply the Markov chain approximation to obtain a sequence of discrete-time problems whose solutions converge to the solution of problems (D) and (V). Chapter 4 presents the application of our investment consumption model to problem (A). Finally, in Chapter 5, we summarize our results and discuss future research related to our model, including three very interesting and important problems.

2. MATHEMATICAL BACKGROUND

The one thing probabilists can do which analysts can't is stop — and they never forgive us for it.

– Sid Port, quoted in [15]

In this chapter, we review the basic mathematical notions and theory on which we will heavily rely in subsequent chapters. We begin by reviewing fundamentals of diffusion processes that are relevant to stochastic control. This is followed by a brief introduction to stochastic control problems and a survey of the necessary theory. Third, we provide a slightly more thorough introduction to the main technique used in solving the stochastic control models under consideration in subsequent chapters, the Markov Chain Approximation. We conclude the chapter by highlighting the use of the basic theory to solve a classic stochastic control problem that is foundational to our model. Note that we will not provide the most general version that can be had for each theorem but will specialize to the relevant cases. For example, we will only require one-dimensional Brownian motion, so all theorems are adjusted accordingly.

2.1 DIFFUSIONS

Underpinning the theory of stochastic control is the theory of probability and stochastic processes and in particular, the theory of diffusions. Surely, the stochastic process most widely employed in stochastic modeling is Brownian motion, a continuous stochastic process with independent, stationary increments that are normally distributed. Standard references for the construction of Brownian motion and its basic properties, including the construction of the Ito integral are [8, 15, 16], from which much of the presentation in this section is drawn.

2.1.1 Stochastic Differential Equations. Stochastic differential equations (SDE) are ordinary differential equations plus a random noise term, prescribed mathematically by an Ito integral. Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ be a fixed, filtered probability space where the filtration $\mathbb{F} = (\mathcal{F}_t)$ satisfies the usual hypotheses,

(i) $\mathcal{F}_t = \bigcap_{s>t} \mathcal{F}_s$.

(ii) \mathcal{F}_0 contains the \mathbb{P} -null sets.

The first condition means, intuitively, that there are no jumps in the available information. The second condition means that it is common knowledge which events are certain not to happen, right from the beginning.

Remark. The usual hypotheses are a standard assumption for filtrations in stochastic control theory. Filtrations satisfying the usual hypotheses have several nice properties. For instance, if \mathbb{F} satisfies the usual hypotheses, τ is an \mathbb{F} -stopping time if and only if τ is an \mathbb{F} -optional time. These kinds of results are helpful when controlling a stopped process. In the model we develop below, random times play an important role, and the usual hypotheses will be very useful.

Let $b : [0, \infty) \times \mathbb{R} \times \Omega \rightarrow \mathbb{R}$ and $\sigma : [0, \infty) \times \mathbb{R} \times \Omega \rightarrow \mathbb{R}$. We require that

(i) $b(\cdot, \cdot, \omega), \sigma(\cdot, \cdot, \omega)$ are $[0, T] \times \mathbb{R}$ -Borel functions, almost surely- ω ;

(ii) $b(\cdot, x, \cdot), \sigma(\cdot, x, \cdot)$ are \mathbb{F} -progressively measurable, for all $x \in \mathbb{R}$.

We refer to b and σ as the drift and volatility, respectively.

Remark. In the classic theory of diffusions, stochastic differential equations are defined by deterministic drift and volatility functions, in which case it is sufficient that the drift and volatility are *adapted* to \mathbb{F} . However, the stochastic differential equations of interest in stochastic control define *controlled diffusions*, for which b and σ explicitly depend on ω through the control parameter. In the case where the drift and volatility depend explicitly on ω , it is essential that we assume they are \mathbb{F} -*progressively measurable*. A progressively measurable

process $u : [0, T] \times \Omega \rightarrow \mathbb{R}$ is a process such that for each $t \in [0, T]$, u is $\mathcal{B}([0, t]) \times \mathcal{F}_t$ -measurable, which is strictly stronger than being adapted. Progressive measurability is important for two reasons. First, progressive measurability is a technical condition so that expectations of integrals are well-defined and can be evaluated by Fubini's Theorem. Equally important for the models developed below, though, is that progressive measurability ensures that stopped processes are measurable, which will be important in the stochastic control models we develop below.

The SDE for the drift and volatility pair (b, σ) suppressing ω -dependence for simplicity, is the integral equation

$$X(t) = x + \int_0^t b(s, X(s)) ds + \int_0^t \sigma(s, X(s)) dW(s), \quad (2.1)$$

though it is often written formally in differential form as

$$\begin{cases} dX(t) &= b(t, X(t)) dt + \sigma(t, X(t)) dW(t), \\ X(0) &= x. \end{cases}$$

Definition 2.1 (Strong Solution of SDE). An \mathbb{F} -progressively measurable process X is a strong solution of the SDE (2.1) if

- (i) $\mathbb{P}[X(0) = x] = 1$,
- (ii) $\mathbb{P} \left[\int_0^t |b(s, X(s))| + \sigma^2(s, X(s)) ds < \infty \right] = 1$, for all $t \in [0, \infty)$, and
- (iii) (2.1) holds for all $t \in [0, \infty)$, almost surely.

The next theorem describes the conditions under which the existence and uniqueness of a strong solution to (2.1) is guaranteed.

Theorem 2.2 (Existence and Uniqueness). *Suppose that $b(t, \cdot), \sigma(t, \cdot)$ are globally Lipschitz-*

continuous in x for all t, ω , and satisfy the growth condition

$$|b(t, x, \omega)| + |\sigma(t, x, \omega)| \leq g(t, \omega) + K|x|$$

where K is a constant and g is a progressively measurable process such that

$$\mathbb{E} \left[\int_0^t g^2(s, \omega) ds \right] < \infty.$$

Then there exists a continuous, \mathbb{F} -progressively measurable process X that is a strong solution of (2.1).

Solutions to SDEs are called *diffusions*, or in our case, *controlled diffusions*. We enforce the hypotheses of Theorem 2.1 for the SDEs defined in our control models below.

2.1.2 Properties of Diffusions. One of the most useful properties of a diffusion is that it is a Markov process. Intuitively, the Markov property means that for any time t , the path of the diffusion X up until t does not influence the expected future path of X beyond the position of X at t . Mathematically, a Markov process is defined as follows

Definition 2.3 (Markov Process). A *Markov process* is an adapted process X on (Ω, \mathcal{F}) together with a family of probability measures $\{P^x\}_{x \in \mathbb{R}}$ on (Ω, \mathcal{F}) such that

- (i) $P^x[X(0) = x] = 1 \quad \forall x \in \mathbb{R}$;
- (ii) $P^x[X(t+s) \in B \mid \mathcal{F}_s] = P^x[X(t+s) \in B \mid X(s)], \quad \forall x \in \mathbb{R}, t, s \geq 0, B \in \mathcal{B}(\mathbb{R})$;
- (iii) $P^x[X(t+s) \in B \mid X(s) = y] = P^y[X(t) \in B], \quad \forall x \in \mathbb{R}, t, s \geq 0, B \in \mathcal{B}(\mathbb{R})$.

In fact, the progressively-measurable diffusions that we will encounter satisfy an even more impressive “memoryless” property known as the *strong Markov property*. The strong Markov property is similar to the Markov property except that we may even condition at stopping times.

A second essential property of diffusions is that they possess almost surely continuous sample paths, and a third is that smooth functions of diffusions satisfy a stochastic fundamental theorem of calculus known as Ito's formula (see [8], Theorem 5.3.3) :

Theorem 2.4 (Ito's Formula for a Diffusion). *Let $V : [0, \infty) \times \mathbb{R} \rightarrow \mathbb{R}$ be a continuously differentiable function, and let X be a solution to (2.1). Then for every $t \geq 0$, V satisfies*

$$\begin{aligned} V(t, X(t)) &= V(0, X(0)) + \int_0^t V_t(s, X(s)) ds + \int_0^t V_x(s, X(s)) dX(s) \\ &\quad + \int_0^t V_{xx}(s, X(s)) d\langle X(s) \rangle \\ &= V(0, X(0)) + \int_0^t (V_t(s, X(s)) + \mathcal{A}V(s, X(s))) ds \\ &\quad + \int_0^t V_x(s, X(s)) dW(s), \end{aligned} \tag{2.2}$$

where $\langle \cdot \rangle$ is the quadratic variation process, and the differential operator \mathcal{A} is

$$\mathcal{A}V(t, x) = V_x(t, x)b(x) + \frac{1}{2}V_{xx}(t, x)\sigma(x). \tag{2.3}$$

Using Ito's formula, the instantaneous mean rate of change of a smooth function of a diffusion can be computed (see [8], Problem 5.1.2).

Theorem 2.5 (Infinitesimal Generator). *Let X be a solution to the SDE (2.1), where b and σ are bounded and continuous functions. Then for each $f \in C^2(\mathbb{R})$,*

$$\lim_{h \downarrow 0} \mathbb{E}^{t,x} \left[\frac{V(t+h, X(t+h)) - V(t, x)}{h} \right] = \mathcal{A}V(x), \quad \forall x \in \mathbb{R}$$

Applying Theorem (2.5) to the functions $x \mapsto x$ and $x \mapsto x^2$, we obtain the local mean and local variance of a diffusion,

$$\mathbb{E}^x[X(t) - x] = b(x)t + o(t), \tag{2.4}$$

$$\mathbb{E}^x[(X(t) - x)^2] = \sigma^2(x)t + o(t), \tag{2.5}$$

respectively.

Controlled diffusions play a central role in the theory of stochastic control theory serving as the dynamical system to be controlled in some optimal manner. The aforementioned properties of diffusions lead to useful techniques for solving control problems such as dynamic programming and the Markov chain approximation, which are discussed below.

2.2 STOCHASTIC CONTROL THEORY

As researchers have recognized the importance of modeling variables that cannot be accounted for in deterministic ways, stochastic modeling has risen to prominence in varied applications, especially in financial and economic applications. Stochastic control theory combines classic control theory with the theory of diffusions. Throughout this section, we assume the same probability framework introduced in the previous section.

Stochastic control models consist of three basic ingredients:

- (i) a set of control variables, called “admissible” control (decision) variables,
- (ii) a controlled diffusion, and
- (iii) a cost functional, which provides a method for evaluating decision strategies by specifying a cost or reward based on the chosen decision and evolution of the diffusion.

The fundamental goal of a stochastic control problem is to optimize, in our case, maximize, the cost functional over the set of admissible controls, over some time interval called the *planning horizon*. We will assume in the sequel that the drift and volatility functions introduced above depend on ω only through a control variable u ,

$$b(t, x, \omega) = \tilde{b}(t, x, u(t, \omega))$$

$$\sigma(t, x, \omega) = \tilde{\sigma}(t, x, u(t, \omega)).$$

2.2.1 Admissible Controls. The first important component of any stochastic control model is the set of control, or decision, variables. Intuitively, admissible controls must

- only react to information as it becomes available;
- permit a well-defined, controlled diffusion;
- satisfy any explicit as well as implicit constraints; and
- lead to a well-defined cost functional.

The intuitive properties above give rise to the definition of the set of admissible controls:

Definition 2.6 (Admissible Controls). The set of *admissible controls*, \mathcal{U} , is the set of \mathbb{F} -progressively measurable processes, $u : [0, T] \times \Omega \rightarrow \mathbb{R}^d$, such that

- (i) the controlled SDE admits a unique solution, in which case it is called a *feasible* control;
- (ii) $u(t) \in \mathcal{V} \subseteq \mathbb{R}^d$ and $X(t) \in A \subseteq \mathbb{R}$, where \mathcal{V} and A are the constraint sets; and
- (iii) the cost functional is well-defined (see Definition 2.7 below.)

That u maps into \mathbb{R}^d is just to accommodate multiple, real-valued decision variables.

2.2.2 State Process and Cost Functional. The solution to the SDE comprises the second important component of a stochastic control model, which is the state variable, modeled as a controlled diffusion. The controls and state process give rise to the final important component of a stochastic control problem, the cost functional and associated value function. A cost functional, $J : [0, \infty) \times \mathbb{R} \times \mathcal{U} \rightarrow \mathbb{R}$, quantifies a cost or reward given the time at which the decision is made, the state of the controlled process, and the chosen control. A standard form for the cost functional, and the only one we are interested in, is defined by the expectation,

$$J(x, u) = \mathbb{E}^x \left[\int_0^T U(s, u(s)) ds + B(T, X(T)) \right], \quad (2.6)$$

where $U : [0, T] \times \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}$, and $B : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$. In the definition of the admissible controls, we required that the cost-functional is well defined. We can now supply meaning to that requirement.

Definition 2.7 (Well-defined Cost Functional). The cost functional J is well-defined for a feasible control u if

$$\int_0^T U(s, u(s)) ds + B(T, X(T)) \in L^1(\Omega, \mathcal{F}) \quad (2.7)$$

Associated with the cost functional is the value function, $V : \mathbb{R} \rightarrow \mathbb{R}$, which gives the optimized cost functional as a function of the starting point for the diffusion,

$$V(x) = \sup_{u \in \mathcal{U}} J(x, u). \quad (2.8)$$

Having defined the basic ingredients of a control problem, we now state the fundamental objective for the general formulation of stochastic control given above:

Problem (P). Given a set of admissible controls, \mathcal{U} , a controlled diffusion, X , and a cost functional J , determine a control policy $u^* \in \mathcal{U}$ for which

$$V(x) = J(x, u^*).$$

If it exists, the control u^* is called an optimal control policy.

Several useful techniques have been developed to help determine optimal control policies. Among these are dynamic programming approach, discussed next, and this is followed by the Markov chain approximation, which makes use of dynamic programming techniques.

2.2.3 Dynamic Programming and the HJB Equations. One of the important tools for solving optimal control problems is the dynamic programming principle (DPP), also known as Bellman's principle of optimality. Intuitively, because the trajectory of a diffusion depends on its present state but not the path taken to arrive at its present state,

the optimization problem is equivalent to optimally controlling the diffusion over some time interval, from t to t' , say, observing the state of the diffusion at t' , and then, optimally controlling the diffusion over the remaining time horizon.

A precise mathematical statement of the DPP requires an extension of the framework for control problems discussed above in order to accommodate conditioning on the state variable at various times, $X(t) = x$, and then optimizing over the remaining interval $[t, T]$. In particular, time dependence must be incorporated in the definition of the cost functional and value function, which can be done in a very natural way. Unfortunately, the entire underlying probability framework must be generalized to allow for changing (conditional) measures, $\mathbb{P}^{t, X(t)}$, as $X(t)$ evolves. That is, the probability measure under which the expectation in the cost functional (2.6) is computed constantly changes with the evolution of the diffusion, necessitating a weak formulation of the control problem under a family of probability spaces indexed by time and space. However, dynamic programming techniques will only play a formal role in solving the main control models below, and accordingly, we omit technical details pertaining to the weak formulation of the generic control problem, (P). We refer the interested reader to [24] for additional information.

We define the time dependent cost functional and value function by

$$\begin{aligned} J(t, x, u) &= \mathbb{E}^{t, x} \left[\int_t^T U(s, u(s)) ds + B(T, X(T)) \right], \\ V(t, x) &= \sup_{u \in \mathcal{U}} J(t, x, u). \end{aligned} \quad (2.9)$$

Theorem 2.8 (Dynamic Programming Principle). *For any $x \in \mathbb{R}$ and $0 \leq t \leq t' \leq T$,*

$$V(t, x) = \sup_{u \in \mathcal{U}} \mathbb{E}^{t, x} \left[\int_t^{t'} U(s, u(s)) ds + V(t', X(t')) \right]. \quad (2.10)$$

The DPP is useful because it permits the study of the average change in the value function over short time intervals. Suppose that the value function V is sufficiently smooth. Invoking Ito's Lemma and taking the time intervals to zero, the DPP leads to the so-called Hamilton-

Jacobi-Bellman (HJB) equation, which is a non-linear PDE satisfied by the value function, V . For instance, the HJB equation and terminal condition for the value function V defined above are

$$\begin{cases} V_t(t, x) + \sup_{u \in \mathcal{U}} \{b(t, x, u(t))V_x(t, x) + \frac{1}{2}\sigma^2(t, x, u(t))V_{xx}(t, x) + U(t, x)\} = 0 \\ V(T, x) = B(T, x). \end{cases}$$

In addition to providing a means for determining the value function V , the HJB equation helps determine the optimal control policies.

Remark. In stochastic control, a solution to the HJB equation provides a candidate for the value function and control policies. To rigorously establish that the candidates are indeed optimal requires a so-called verification theorem. In the application of dynamic programming in later chapters, we do not solve the HJB equation explicitly but instead use it to guide the construction of an approximating sequence of problems, as discussed in the next section.

2.3 MARKOV CHAIN APPROXIMATION

In this section, we present the Markov chain approximation (MCA), which is an approximation technique that can be used to solve a fairly large class of stochastic control problems. Essentially, MCA works by approximating the continuous-time controlled diffusion of a stochastic control problem with a sequence of discrete-time Markov chains (DTMC) that are locally consistent with the continuous-time diffusion, i.e., possess similar statistical characteristics over short time intervals to those of the diffusion. Associated with the sequence of DTMC is a sequence of discrete-time control problems that are essentially discrete-time analogues of the continuous-time control problem. Under certain conditions, it can be shown that the sequence of discrete-time value functions and the corresponding optimal controls converge to the value function and optimal controls of the continuous-time problem. In practice, the appropriate sequence of DTMC can often be determined by ex-

exploiting the Hamilton-Jacobi-Bellman (HJB) equations introduced in the previous section. The presentation in this section draws heavily from [9].

2.3.1 Local Consistency. Let $X = \{X(t)\}$ be a real-valued, continuous-time diffusion, and let $\xi^{\delta,h} = \{\xi_n^{\delta,h}\}_{n=1}^{\infty}$ be a discrete-time Markov chain (DTMC) on a state space $S_h \in \mathbb{R}$. The parameters δ and h determine the size of the discrete-time steps, $\Delta t^{\delta,h}$, such that $\sup \Delta t^{\delta,h} \rightarrow 0$ as $\delta, h \rightarrow 0$ but $\inf \Delta t^{\delta,h} > 0$ for each $\delta, h > 0$. We define the finite difference operator $\Delta \xi_n^{\delta,h} = \xi_{n+1}^{\delta,h} - \xi_n^{\delta,h}$. We denote by $\mathbb{E}_{\delta,h}^{x,n}$ the expectation operator with respect to the transition probabilities p^h conditioned on $\{\xi_n^{\delta,h} = x, u_n^{\delta,h}\}$.

Definition 2.9 (Locally Consistent DTMC). A DTMC $\xi^{\delta,h}$ is locally consistent with a diffusion X having drift b and volatility σ if

(i)

$$\mathbb{E}_{\delta,h}^{x,n}[\Delta \xi_n^{\delta,h}] = b(x, u_n^{\delta,h})\Delta t^{\delta,h} + o(\Delta t^{\delta,h}),$$

(ii)

$$\mathbb{E}_{\delta,h}^{x,n} \left[(\Delta \xi_n^{\delta,h} - \mathbb{E}_{\delta,h}^{x,n}[\Delta \xi_n^{\delta,h}])^2 \right] = \sigma^2(x, u_n^{\delta,h})\Delta t^{\delta,h} + o(\Delta t^{\delta,h}),$$

(iii) and

$$\lim_{\delta, h \downarrow 0} \sup_n \{|\xi_{n+1}^{\delta,h} - \xi_n^{\delta,h}|\} = 0.$$

Intuitively, local consistency means that over short time intervals, the mean and volatility of the DTMC are the same as that of the diffusion.

For the locally consistent DTMC $\xi^{\delta,h}$, we construct a discrete-time control problem that approximates the continuous-time control problem.

2.3.2 Convergence. The key mathematical theory justifying MCA is the proof that a DTMC, or at least an appropriate continuous-time interpolation of a DTMC, converges weakly to the diffusion with which it is locally consistent, and that therefore the value function and optimal control policies associated with the discrete-time control problem converge

to the continuous-time value function and optimal control policies. However, the main theorems pertaining to local consistency and convergence of MCA will not apply directly to the models we develop below. Therefore, we put off a more thorough discussion of convergence until Chapter 5.

In application, how to choose a locally consistent DTMC is not obvious a priori but can be determined by a careful study of the dynamics of the diffusion. An effective technique for exploiting the dynamics of the diffusion to construct the locally consistent DTMC is the finite difference method, discussed next.

2.3.3 Finite Difference Method. In a nutshell, the finite difference method converts the continuous-time dynamics of the continuous-time control problem as encoded in the Hamilton-Jacobi-Bellman (HJB) equation into discrete-time dynamics by discretizing the derivatives in the non-linear PDE. The discretized HJB is naturally associated with a discrete-time control problem, and the transition probabilities of the discrete-time Markov chain can be recovered from the finite difference equation. The resulting DTMC is naturally locally consistent with the controlled diffusion of the continuous-time problem. Below, we provide a simple example of the finite difference method in the case of the general stochastic control problem of the form (2.6).

Let X be a controlled diffusion satisfying the SDE

$$dX(t) = b(t, u(t)) dt + \sigma(t, u(t)) dW(t),$$

and let V be the value function of a stochastic control problem, given by

$$V(t, x) = \sup_{u \in \mathcal{U}} \mathbb{E}^{t,x} \left[\int_t^T U(t, u(s)) ds + B(T, X(T)) \right].$$

As stated in the previous section, the dynamic programming principle and Ito's lemma give

rise to the HJB equation for V ,

$$0 = \sup_{u \in \mathcal{U}} \left\{ V_t(t, x) + b(t, x, u(t))V_x(t, x) + \frac{1}{2}\sigma^2(t, x, u(t))V_{xx}(t, x) + U(t, u(t)) \right\}. \quad (2.11)$$

Let $\delta, h > 0$ be the discrete-time step size in time and space, respectively. The HJB equation is then discretized as follows:

$$V_t(t, x) = \frac{V^{\delta, h}(t + \delta, x) - V^{\delta, h}(t, x)}{\delta} + o(\delta), \quad (2.12)$$

$$b^+(t, x)V_x(t, x) = b^+(t, x) \frac{V^{\delta, h}(t + \delta, x + h) - V^{\delta, h}(t + \delta, x)}{h} + o(h), \quad (2.13)$$

$$b^-(t, x)V_x(t, x) = b^-(t, x) \frac{V^{\delta, h}(t + \delta, x) - V^{\delta, h}(t + \delta, x - h)}{h} + o(h), \quad (2.14)$$

$$V_{xx}(t, x) = \frac{V^{\delta, h}(t + \delta, x) + V^{\delta, h}(t + \delta, x - h) - 2V^{\delta, h}(t, x)}{h^2} + o(h^2), \quad (2.15)$$

$$(2.16)$$

where $b^+ = \max\{b, 0\}$, $b^- = \max\{-b, 0\}$, and $b = b^+ - b^-$.

Remark. V_x is discretized differently based on its coefficient due to the following observation. Since the drift b represents the local velocity of the diffusion, b^+ and b^- are the nonnegative and nonpositive components of local velocity, respectively. Therefore, b^+ pushes the diffusion located at x in a nonnegative direction, i.e., to x or $x+h$ while b^- pushes the diffusion located at x in a nonpositive direction, i.e., to x or $x-h$. For more details pertaining to the choice of discretization, see [9], page 93 and pp. 326–330.

Introducing the discretizations (2.12)-(2.15) into the HJB equation (2.11) yields the dif-

ference equation

$$\begin{aligned}
0 = & \sup_{u \in \mathcal{U}} \frac{V^{\delta,h}(t + \delta, x) - V^{\delta,h}(t, x)}{\delta} + b^+(t, x) \frac{V^{\delta,h}(t + \delta, x + h) - V^{\delta,h}(t + \delta, x)}{h} \\
& + b^-(t, x) \frac{V^{\delta,h}(t + \delta, x) - V^{\delta,h}(t + \delta, x - h)}{h} \\
& + \frac{1}{2} \sigma^2(t, x, u(t)) \frac{V^{\delta,h}(t + \delta, x + h) + V^{\delta,h}(t + \delta, x - h) - 2V^{\delta,h}(t + \delta, x)}{h^2} \\
& + U(t, x, u(t)),
\end{aligned}$$

which can be rearranged as

$$\begin{aligned}
V^{\delta,h}(t, x) = & P_{\delta,h}^{x,t}(x, x + h) V^{\delta,h}(t + \delta, x + h) + P_{\delta,h}^{x,t}(x, x) V^{\delta,h}(t + \delta, x) \\
& + P_{\delta,h}^{x,t}(x, x - h) V^{\delta,h}(t + \delta, x - h) + \delta U(t, u(t)),
\end{aligned} \tag{2.17}$$

where

$$P_{\delta,h}^{x,t}(z) = \begin{cases} \frac{\delta}{h} (b^+(t, x) + \frac{1}{2h} \sigma^2(t, x, u(t))) & \text{if } z = x + h \\ \frac{\delta}{h} (b^-(t, x) + \frac{1}{2h} \sigma^2(t, x, u(t))) & \text{if } z = x - h \\ 1 - P_{\delta,h}^{x,t}(x + h) + P_{\delta,h}^{x,t}(x - h) & \text{if } z = x. \end{cases}$$

We define a time increment $\Delta t^{\delta,h} = \delta$ and a DTMC $\xi^{\delta,h}$ on the discrete state space

$$\left\{ (t, x) = (k\delta, jh) : k, j \in \mathbb{N}, 0 \leq k \leq \frac{T}{\delta}, -k \leq j \leq k \right\} \tag{2.18}$$

by the transition probabilities

$$\mathbb{P}_{\delta,h}^{x,t}(z) = \begin{cases} P_{\delta,h}^{x,t}(z) & \text{if } z \in \{x, x + h, x - h\}, \\ 0 & \text{otherwise.} \end{cases}$$

Note that we may interpret the right-hand-side of (2.17) as the dynamic programming principle for discrete-time state process $\xi^{\delta,h}$ and discrete-time value function $V^{\delta,h}$, under the

transition probability $\mathbb{P}_{\delta,h}^{x,t}$,

$$V^{\delta,h}(t, x) = \mathbb{E}_{\delta,h}^{x,t}[V(t + \delta, \xi^{\delta,h})] + U(t, u(t))\Delta t^{\delta,h}.$$

Remark. It is crucial to the above interpretation, and indeed, to the success of the MCA algorithm, that $0 \leq P_{\delta,h}^{x,t} \leq 1$ and $\sum_z P_{\delta,h}^{x,t}(z) = 1$, for all $x \in \mathbb{R}$, $t \in [0, T]$ so that the $P_{\delta,h}^{x,t}$ may legitimately be interpreted as transition probabilities.

It is straightforward to establish that $\xi^{\delta,h}$ is locally consistent with the diffusion X .

At this point we need only solve the discrete-time control problem for $\xi^{\delta,h}$ and then take limits as $\delta, h \downarrow 0$ in order to obtain the optimal value function V and the associated optimal policies, u^* , if they exist. In practice, we simply approximate the value function and optimal policies by solving the discrete problem numerically on a sufficiently fine grid.

2.3.4 Limitations of the Markov Chain Approximation. In the literature, MCA has been successfully applied to various stochastic control models to approximate the value function and optimal control policies, including in financial and economic applications. However, difficulties sometimes arise in the application of MCA. As noted, the transition probabilities $P_t^{\delta,h}$ must satisfy $0 \leq P_t^{\delta,h} \leq 1$, $\sum_z P_t^{\delta,h}(z) = 1$. These conditions may fail when the transition probabilities depend explicitly on the state variable x . In addition, implementing MCA computationally can be difficult when the cost functional involves a singularity, i.e., $\lim_{u \downarrow 0} U(t, u) = -\infty$ or $\lim_{x \downarrow 0} B(t, x) = -\infty$. These challenges arise in the control models presented in the next section and in subsequent chapters, and we will demonstrate how a well-chosen change of variable can help overcome these difficulties.

2.4 RICHARD'S MODEL

Frequent applications of stochastic control models, known collectively as investment consumption models, arise in the fields of finance and economics. The basic set-up consists

of an investor with some capital who finds it necessary to spend, or consume, and invest over some planning horizon and who derives utility from the total consumption during the investment period as well as the liquid wealth at the end of the investment period. The control variables in investment consumption models are investment decisions, consumption decisions, and possibly other financial and economic decisions. The investor's wealth process, which is a controlled diffusion influenced by investment and consumption decisions, serves as the state variable. The cost functional is built from utility functions, which are a standard economic tool for quantifying an individual's behavior in the presence of risk. Merton [11] introduces the basic continuous-time investment-consumption models for which the wealth process is modeled as a diffusion. The literature devoted to extensions of Merton's basic model is vast. Among the many variations, one important extension endows the investor with an income that is risky in one way or another, such as the study by Richard [14]. The investment-consumption model studied in subsequent chapters involves a randomly terminating endowment income and shares many features with Richard's model. Therefore, we present Richard's model in some detail. We proceed by first developing the probabilistic and financial framework for Richard's model. We then define the key elements of control models discussed in Section 2.2 for Richard's model, including the controls, state variable, and cost functional. After identifying the basic ingredients of the control model, we apply the Markov chain approximation introduced in Section 2.3 by deriving and discretizing the Hamilton-Jacobi-Bellman equation for the value function. While we avoid many technicalities, some details are included because they provide intuition for the more complex model in the next chapter.

2.4.1 Decision Variables and Wealth Process. Let $(\Omega, \mathbb{F}, \mathcal{F}, \mathbb{P})$ be a fixed, filtered probability space with filtration $\mathbb{F} = (\mathcal{F}_t)$ that is rich enough to support a Brownian motion, W , and a random terminal time τ_M . The random time τ_M possesses conditional distributions

specified by the survival function,

$$\bar{F}_M(t | s) = \mathbb{P}[\tau_D > t | \tau_M \geq s] = \exp\left(-\int_s^t \lambda_M(u) du\right).$$

An investor is endowed with an initial wealth of x and receives an exogenous stream of income at a rate of $\alpha(t)$, which continues until the earlier of τ_M and a fixed time T , which provides an upper bound to the control problem time horizon. For concreteness, we will assume that the random time τ_M corresponds to the investor's uncertain time of death, although it could be used to model a number of random events. For each $t \in [0, \tau_M \wedge T)$, the investor chooses a consumption rate, $c(t)$ while investing any liquid capital in the financial market. The basic assets available for investment are a risk-free asset with constant rate of return r , $S_0(t) = \exp(rt)$, and a risky asset S_1 , modeled by geometric Brownian motion with constant mean rate of return μ and volatility σ ,

$$\frac{dS_1(t)}{S_1(t)} = \mu dt + \sigma dW(t).$$

The capital invested in the S_0 at time t is denoted by $\pi_0(t)$, and the capital invested in S_1 is denoted by $\pi(t)$. In addition to consumption and investment in the basic assets, the market offers instantaneous term life insurance, meaning that the investor can purchase an insurance policy at time t for $p_M(t)$ that pays out $\frac{p_M(t)}{\eta_M(t)}$ if $\tau_M = t$ but expires otherwise. Here, η_M is just a factor that determines the insurance payout rate.

Remark. A simple heuristic argument suggests that the actuarially fair payout rate per dollar of instantaneous life insurance is $\eta_M(t) = \lambda_M(t)$. Formally, the instantaneous payout per dollar is actuarially fair if

$$1 = \frac{1}{\eta_M(t)} \mathbb{P}^t[\tau_M \in dt] = \frac{1}{\eta_M(t)} \lambda_M(t) \bar{F}_M(t | t) = \frac{\lambda_M(t)}{\eta_M(t)}, \quad (2.19)$$

i.e.,

$$\eta_M(t) = \lambda_M(t). \quad (2.20)$$

Nonetheless, for reasons discussed below, we allow for $\eta_M \geq \lambda_M$.

The exogenous income, consumption, and insurance purchases define the investor's cumulative income process, Γ , by

$$\Gamma(t) = \int_0^t (\alpha(s) - c(s) - p_M(s)) ds, \quad t < \tau_M \wedge T. \quad (2.21)$$

Next, we define the investor's financial gains process G expressed in differential form as

$$\begin{aligned} dG(t) &= \pi_0 \frac{dS_0}{S_0} + \pi \frac{dS_1}{S_1} \\ &= \pi_0 r dt + \pi(\mu dt + \sigma dW(t)) \quad t < \tau_M \wedge T, \end{aligned} \quad (2.22)$$

The investor's Γ -financed liquid wealth, L , is defined as the sum of the cumulative income and the gains, subject to the condition that the liquid wealth is continuously reinvested in the risky and risk-free assets,

$$L(t) = \Gamma + G; \quad (2.23)$$

$$L(t) = \pi_0 + \pi. \quad (\Gamma\text{-financed}) \quad (2.24)$$

(2.21)- (2.24) yield the differential for L ,

$$dL(t) = (\alpha(t) - c(t) - p_M(t) + rL(t)) \quad (2.25)$$

$$+ (\mu - r)\pi(t) dt + \sigma\pi(t)dW(t), \quad t < \tau_M \wedge T \quad (2.26)$$

While the liquid wealth may become negative, i.e., the investor may borrow in net, the wealth is constrained by the expected, discounted remaining exogenous income, I , defined

by

$$I(t) = \int_t^T \exp\left(-\int_t^s (r + \eta_M(u)) du\right) \alpha(s) ds. \quad (2.27)$$

The borrowing constraint is then

$$L(t) \geq -I(t), \quad t < \tau_M \wedge T. \quad (2.28)$$

Intuitively, the borrowing limit represents the fact that no lender will loan to a borrower beyond what can be expected to be repaid. Note in equation (2.27) that the future income is discounted by η_M instead of λ_M . The actuarially fair value of the expected remaining income, on the other hand, is

$$\begin{aligned} & \mathbb{E}^t \left[\int_t^{\tau_M \wedge T} \exp(-r(s-t)) \alpha(s) ds \right] \\ &= \int_t^T \exp\left(-\int_t^s (r + \lambda_M(u)) du\right) \alpha(s) ds. \end{aligned}$$

Hence, the discount factor in (2.27) is not necessarily actuarially fair. One reason for this is that, in practice, insurers assume asymmetric information, i.e., that those purchasing life insurance face a more dire outlook on survival than the general population, a phenomenon known as self-selection. Therefore, we allow for insurers to discount by η_M instead of λ_M .

The differential for I is

$$\begin{aligned} dI(t) &= I'(t) dt \\ &= (-\alpha(t) + (r + \eta_M(t))I(t)) dt \end{aligned} \quad (2.29)$$

The state variable for Richard's investment and consumption model is the investor's total wealth, X , defined as the sum of the liquid wealth and the expected remaining wealth,

$$X = L + I. \quad (2.30)$$

Remark. In his original study, Richard used the liquid wealth L as the state variable and then later changed variables to incorporate the total wealth. It is more convenient for our purposes to begin with the total wealth as the state variable.

The differential of the total wealth is

$$\begin{aligned}
dX(t) &= dL(t) + db(t) \\
&= (\alpha(t) - c(t) - p_M(t) + r(X(t) - I(t)) + (\mu - r)\pi(t)) dt + \sigma\pi(t) dW(t) \\
&\quad + (-\alpha(t) + (r + \eta_M(t))I(t)) dt \\
&= (-c(t) + (-\eta_M(t)X(t) + \eta_M I(t) - p_M(t)) + (r + \eta_M(t))X(t) + (\mu - r)\pi(t)) dt \\
&\quad + \sigma\pi(t)dW(t) \\
&= (-c(t) - \eta_M(t) \zeta_M(t) + (r + \eta_M(t))X(t) + (\mu - r)\pi(t)) dt \\
&\quad + \sigma\pi(t)dW(t), \quad t < \tau_M \wedge T,
\end{aligned} \tag{2.31}$$

where

$$\zeta_M = X - I + \frac{p_M}{\eta_M}. \tag{2.32}$$

Intuitively, ζ_M is the investor's liquid wealth plus the insurance payout; that is, $\zeta_M(t)$ is the investor's terminal wealth in the event $\tau_M = t$. Thus, $\eta_M \zeta_M$ may be viewed as the cost of financing a desired terminal wealth.

By (2.28), X is constrained to be nonnegative,

$$X(t) \geq 0, \quad t < \tau_M \wedge T. \tag{2.33}$$

Now that the state variable and constraints have been specified, we can fully define the set of admissible controls. The purpose of the admissibility conditions is to guarantee the existence and uniqueness of a solution to the SDE for the total wealth, (2.31), as well as meet any imposed constraints. Note that in the context of the theory of controlled diffusions

presented earlier, if we let b and σ be the drift and volatility,

$$\begin{aligned} b(t, x, u_1, u_2, u_3) &= -u_1 - \eta_M(t)u_3 + (r + \eta_M(t))x + (\mu - r)u_3 \\ \sigma(t, x, u_1, u_2, u_3) &= \sigma u_3, \end{aligned}$$

then for any reasonably-behaved η_M , the SDE drift and volatility coefficients are uniformly Lipschitz on $[0, T] \times [0, \infty) \times \mathbb{R}_+^3$,

$$\begin{aligned} |b(t, x, u_1, u_2, u_3) - b(t, y, u_1, u_2, u_3)| &= |r + \eta_M(t)||x - y| \\ |\sigma(t, x, u_1, u_2, u_3) - \sigma(t, y, u_1, u_2, u_3)| &= 0, \end{aligned} \quad (2.34)$$

as required for the existence and uniqueness of a controlled diffusion.

Definition 2.10 (Admissible Controls). An admissible control triple (c, π, ζ_M) is a triple of progressively measurable processes such that

$$(i) \quad \int_0^T (|c(s)| + |\eta_M(s)\zeta_M(s)| + |\pi(s)| + \pi^2(s)) ds < \infty; \quad (2.35)$$

$$(ii) \quad c \geq 0;$$

$$(iii) \quad \zeta_M \geq X - I; \text{ and}$$

$$(iv) \quad X \geq 0.$$

Between the Lipschitz condition and (2.35), there exists a unique solution to the SDE (2.31) for any initial data $(t, x) \in [0, T] \times [0, \infty)$ and any set of admissible controls.

2.4.2 Utility Functions and Cost Functional. The standard cost functional for financial and economic applications, and the one that will be used in the models investigated here, is built around the concept of a utility function, which is defined next.

Definition 2.11. (Utility Function) A utility function $U : [0, \infty) \times (0, \infty) \rightarrow \mathbb{R}$ is continuously differentiable function for which

- (i) $U'(t, \cdot) > 0$ for each $t \in [0, \infty)$;
- (ii) $U''(t, \cdot) < 0$ for each $t \in [0, \infty)$; and
- (iii) $\lim_{x \rightarrow \infty} U'(t, x) = 0$ for each $t \in [0, \infty)$.

For a given admissible consumption, portfolio, and insurance decision, a standard cost functional is the expected total utility for consumption throughout the interval of uncertain length and bequest at the random terminal date, conditioned on the investor's survival to time t and wealth at time t , $X(t) = x$, by

$$\mathbb{E}^{t,x} \left[\int_t^{\tau_M \wedge T} U(s, c(s)) ds + B(\tau_M, \zeta_M(\tau_M)) \right]. \quad (2.36)$$

However, by the independence of the Brownian motion and the random time, (2.36) is equivalent to

$$\begin{aligned} & \mathbb{E}^{t,x} \left[\int_t^{\tau_M \wedge T} U(s, c(s)) ds + B(\tau_M, \zeta_M(\tau_M)) \right] \\ &= \mathbb{E}^{t,x} \left[\int_t^T U(s, c(s)) \mathbf{1}_{s < \tau_M} ds + B(\tau_M, \zeta_M(\tau_M)) \mathbf{1}_{\tau_M < T} + B(T, \zeta_M(T)) \mathbf{1}_{T < \tau_M} \right] \\ &= \mathbb{E}^{t,x} \left[\int_t^T U(s, c(s)) \bar{F}_M(s | t) ds + \int_t^T B(s, \zeta_M(s)) f_M(s | t) ds + B(T, \zeta_M(T)) \bar{F}_M(T | t) \right] \\ &= \mathbb{E}^{t,x} \left[\int_t^T (U(s, c(s)) + \lambda_M(s) B(s, \zeta_M(s))) \bar{F}_M(s | t) ds + B(T, \zeta_M(T)) \bar{F}_M(T | t) \right]. \end{aligned}$$

In the third and fourth lines above, the expectation is with respect to the the measure associated with Brownian motion only.

Thus, we define the cost functional given the investor's current wealth, $X(t) = x$, by

$$J(t, x, c, \pi, \zeta_M) = \mathbb{E}^{t,x} \left[\int_t^T (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}_M(s | t) ds + B(T, \zeta_M(T)) \bar{F}_M(T | t) \right], \quad (2.37)$$

and the associated value function by

$$V(t, x) = \sup_{(c, \pi, \zeta_M) \in \mathcal{U}} J(t, x, c, \pi, \zeta_M). \quad (2.38)$$

We can now state the main problem for Richard's model:

Problem (R). For each $t, x \in [0, T] \times [0, \infty)$, find an optimal consumption, portfolio, and insurance policies, $(c^*, \pi^*, \zeta_M^*) \in U$ such that

$$V(t, x) = J(t, x, c^*, \pi^*, \zeta_M^*),$$

if one exists.

In his original study, Richard did not constrain the life insurance variable to be nonnegative, as we have, instead allowing the investor to both buy *and* sell life insurance. Richard formally obtained an explicit solution to the PDE associated with the unconstrained problem. However, if life insurance is constrained so that the investor can merely *purchase* insurance, i.e. $p_M \geq 0$, the HJB for V cannot be solved explicitly. For the constrained problem, Pliska and Ye [13] demonstrate how to use the Markov chain approximation to approximate the solution numerically. We explore their approach next.

2.4.3 Markov Chain Approximation with Logarithmic Transformation. In Section 2.3, we introduced a powerful technique for solving stochastic control problems, the Markov chain approximation. We also described an algorithm that generates a sequence of discrete-time control problems that approximate the continuous-time control problem called

the finite difference method. We now apply these techniques to Richard's model. According to the finite difference method, we begin by obtaining the HJB equation for V , which is derived in the appendix:

$$0 = \sup_{c, \pi, \zeta_M \in \mathcal{U}} \{ -\lambda_M(t)V(t, x) + V_t(t, x) + b(t, x)V_x(t, x) + \pi^2(t)\sigma^2V_{xx}(t, x) + U(t, c(t)) + \lambda_M(t)B(t, \zeta_M(t)) \}, \quad 0 \leq t < T. \quad (2.39)$$

In addition to (2.39), V satisfies a terminal condition and a boundary condition. At $t = T$, since $\bar{F}(T|T) = 1$, we have $V(T, x) = B(T, \zeta_M(T))$. However, we note that $I(T) = 0$, as there is no income beyond T . Moreover, it may reasonably be expected that $p_M(t) \downarrow 0$ as $t \downarrow T$. Indeed, in Richard's unconstrained problem he showed that the insurance control p_M was proportional to the percentage of the total wealth represented by future income. As there is no future income at T , we should have $p_M(T) = 0$. Hence, $\zeta_M(T) = X(T)$. Therefore, the terminal condition is

$$V(T, x) = B(T, x). \quad (2.40)$$

The boundary condition follows from the constraint on the total wealth X so that zero wealth is an absorbing state from which there can be no further consumption or gains. Thus,

$$\lim_{x \downarrow 0} V(t, x) = J(t, 0, 0, 0, 0) \quad (2.41)$$

We do not intend to solve equation (2.39) directly but simply let it guide the choice of discrete-time Markov Chain (DTMC). Instead, we discretize the continuous-time HJB equation in order to discover the transition probabilities for the discrete-time Markov Chain by way of a discrete-time dynamic programming relation. However, the direct discretization of (2.39) does not lend such an interpretation. To highlight the difficulties encountered in

applying the finite difference method directly, consider the finite difference approximations:

$$V_t(t, x) = \frac{V^{\delta, h}(t + \delta, x) - V^{\delta, h}(t, x)}{\delta} + o(\delta), \quad (2.42)$$

$$b^+(t, x, u)V_x(t, x) = b^+(t, x, u) \frac{V^{\delta, h}(t + \delta, x + h) - V^{\delta, h}(t + \delta, x)}{h} + o(h), \quad (2.43)$$

$$b^-(t, x, u)V_x(t, x) = b^-(t, x, u) \frac{V^{\delta, h}(t + \delta, x) - V^{\delta, h}(t + \delta, x - h)}{h} + o(h), \quad (2.44)$$

$$V_{xx}(t, x) = \frac{V^{\delta, h}(t + \delta, x) + V^{\delta, h}(t, x - h) - 2V^{\delta, h}(t, x)}{h^2} + o(h), \quad (2.45)$$

where

$$b^+(t, x) = (r + \eta_M(t))x + \pi(t)(\mu - r) \quad (2.46)$$

$$b^-(t, x) = c(t) + \eta_M(t)\zeta_M(t). \quad (2.47)$$

That $b^+, b^- \geq 0$ requires some justification, but suppose for now that it is true. Substituting (2.42) – (2.45) into (2.39) results in the difference equation

$$\begin{aligned} V^{\delta, h}(t, x) = \sup_{c, \pi, \zeta_M \in \mathcal{U}} \{ & P_{\delta, h}^{t, x}(x + h)V^{\delta, h}(t + \delta, x + h) + P_{\delta, h}^{t, x}(x)V^{\delta, h}(t + \delta, x) \\ & + P_{\delta, h}^{t, x}(x - h)V^{\delta, h}(t + \delta, x - h) + \delta(U(t, c) + \lambda_M(t)B(t, \zeta_M)) \}, \end{aligned} \quad (2.48)$$

where

$$P_{\delta, h}^{t, x}(z) = \begin{cases} \frac{\delta}{h} \left(b^+(t, x) + \frac{\sigma^2 \pi^2(t)}{2h} \right) & \text{if } z = x + h \\ \frac{\delta}{h} \left(b^-(t, x) + \frac{\sigma^2 \pi^2(t)}{2h} \right) & \text{if } z = x - h \\ 1 - P_{\delta, h}^{t, x}(x + h) - P_{\delta, h}^{t, x}(x - h) & \text{if } z = x \end{cases}$$

(2.48) presents two main barriers to MCA:

- (i) The transition $P_{\delta, h}^{t, x}(x + h)$ depends *explicitly* on the wealth variable x . Because x is unbounded, it is difficult, if not impossible, to choose discretization parameters δ and h sufficiently small so that $P_{\delta, h}^{t, x}$ is a true probability measure on the wealth state space.

(ii) In application, the most commonly used utility functions are the CRRA utilities,

$$U(t, z) = \exp(-\rho t) \frac{z^{1-\gamma}}{1-\gamma},$$

and a CRRA utility for terminal wealth, or bequest,

$$B(t, z) = \exp(-\rho t) \frac{z^{1-\gamma}}{1-\gamma}.$$

Note in particular that the boundary condition becomes

$$\lim_{x \downarrow 0} V(t, x) = -\infty. \quad (2.49)$$

For any reasonable choice of discretization parameters, the wealth variable transitions very close to the boundary $x = 0$ with positive probability, but the singular boundary condition at $x = 0$ makes numerical computation practically impossible.

Fortunately, a simple change of variable overcomes both of these difficulties. Following Ye [13], [22], we make the logarithmic transformation

$$y = \log(x), \quad (2.50)$$

so $\exp(y)$ is just the total wealth. We also define a value function in terms of u so that it agrees with the original value function,

$$\widehat{V}(t, y) = V(t, \exp(y)) = V(t, x), \quad (2.51)$$

and change control variables by

$$\widehat{c} = \exp(-y)c, \quad (2.52)$$

$$\widehat{\pi} = \exp(-y)\pi, \quad (2.53)$$

$$\widehat{\zeta}_M = \exp(-y)\zeta_M. \quad (2.54)$$

Then

$$b(t, x(y)) = \exp(y)(r + \eta_M(t) + \widehat{\pi}(t)(\mu - r) - \widehat{c}(t) - \eta_M(t)\widehat{\zeta}_M(t)). \quad (2.55)$$

$$(2.56)$$

We define the set of admissible controls, $\widehat{\mathcal{U}}$, for \widehat{c} , $\widehat{\pi}$, and $\widehat{\zeta}_M$, analogous to the definition for the admissible set \mathcal{U} way. Derivatives of V in x and \widehat{V} in y are related by

$$V_t(t, x) = \widehat{V}_t(t, y) \quad (2.57)$$

$$V_x(t, x) = \exp(-y)\widehat{V}_y(t, y) \quad (2.58)$$

$$V_{xx}(t, x) = \exp(-2y)(\widehat{V}_{yy}(t, y) - \widehat{V}_u(t, y)). \quad (2.59)$$

Substituting (2.57)-(2.59) into (2.39) and rearranging, we obtain

$$0 = \sup_{c, \pi, \zeta_M \in \widehat{\mathcal{U}}} \left\{ -\lambda_M(t)\widehat{V}(t, y) + \widehat{V}_t(t, y) + \widehat{b}(t, y)\widehat{V}_y(t, y) + \frac{\sigma^2\widehat{\pi}^2(t)}{2}\widehat{V}_{yy}(t, y) + U(t, \exp(y)\widehat{c}) + \lambda_M(t)B(t, \exp(y)\widehat{\zeta}_M(t)) \right\}, \quad 0 \leq t < T, \quad (2.60)$$

where

$$\widehat{b}(t, y) = r + \eta_M(t) + \widehat{\pi}(t)(\mu - r) - \widehat{c}(t) - \eta_M(t)\widehat{\zeta}_M(t) - \frac{\sigma^2\widehat{\pi}^2(t)}{2}. \quad (2.61)$$

In addition to (2.60), the terminal condition and boundary conditions are, respectively,

$$\widehat{V}(T, y) = B(T, \exp(y)) \quad (2.62)$$

$$\lim_{y \downarrow 0} V(t, y) = -\infty \quad (2.63)$$

We define \widehat{b}^+ and \widehat{b}^- by

$$b^+(t, y) = r + \eta_M(t) + \widehat{\pi}(t)(\mu - r), \quad (2.64)$$

$$b^-(t, y) = \widehat{c}(t) + \eta_M(t)\widehat{\zeta}_M(t) + \frac{\sigma^2 \widehat{\pi}^2(t)}{2}, \quad (2.65)$$

and then discretize as before, giving the discretized HJB equation for \widehat{V} ,

$$\begin{aligned} \widehat{V}(t, y) = & \sup_{\widehat{c}, \widehat{\pi}, \widehat{\zeta}_M \in \widehat{U}} \widehat{P}_{\delta, h}^{t, y}(y)(y, y + h)\widehat{V}^h(t + \delta, y + h) + \widehat{P}_{\delta, h}^{t, y}(y)\widehat{V}^h(t + \delta, y) \\ & + \widehat{P}_{\delta, h}^{t, y}(y)\widehat{V}^h(t + \delta, y - h) + \delta(U(t, \exp(y))\widehat{c}(t) + \lambda_M(t)B(t, \exp(y))\widehat{\zeta}_M)) \end{aligned} \quad (2.66)$$

where

$$\widehat{P}_{\delta, h}^{t, y}(z) = \begin{cases} \frac{\delta}{h} \left(\widehat{b}^+(t, y) + \frac{\sigma^2 \pi^2(t)}{2h} \right) & \text{if } z = y + h \\ \frac{\delta}{h} \left(\widehat{b}^-(t, y) + \frac{\sigma^2 \pi^2(t)}{2h} \right) & \text{if } z = y - h \\ 1 - \widehat{P}_{\delta, h}^{t, y}(y + h) - \widehat{P}_{\delta, h}^{t, y}(y - h) & \text{if } z = y \end{cases}$$

At this point, it is possible, in principle, to fix the parameters δ, h and then use discrete dynamic programming, iterating backward from the terminal condition, to solve the discretized problem numerically. For a sufficiently fine time-wealth grid, the numerical solution to the discrete-time problem provides a good approximation to the continuous-time value function as well as the continuous-time optimal policies. The transformed difference equation for \widehat{V} , (2.66) exhibits two important features that ensure the success of a numerical implementation. First, the explicit dependence on the transformed wealth variable y has been eliminated so that it is readily possible to choose δ and h such that $P_{\delta, h}^{t, y}$ is a legiti-

mate transition probability. Second, the boundary condition has been pushed to $y = -\infty$. Numerical solutions to (2.66) can be obtained by numerical dynamic programming over a finite, trinomial tree. Even for fine grid spacings, the logarithmic transformation attenuates the rate at which the singular boundary approaches $-\infty$ enough so that the computational method succeeds.

According to the theory for MCA, local consistency is sufficient to guarantee convergence, which Ye claimed. However, the hypotheses of the theorems cited by Ye are not satisfied in Richard's model. Moreover, the DTMC is not locally consistent with the continuous-time wealth process. Nonetheless, Ye compared numerical solutions obtained via MCA with Richard's closed-form solutions in the unconstrained case and found apparent convergence. We will address the convergence of Ye's application of MCA to Richard's model further in Chapter 5, but for now, the important issue of convergence will be left unresolved.

A few points concerning the solution to Richard's problem are worth noting. First, in the unconstrained case, the the value function V takes the functional form

$$V(t, x) = a(t) \frac{x^{1-\gamma}}{1-\gamma}, \quad (2.67)$$

The function a is significant only in that it depends on time alone, not on the wealth. Hence, V has the form of a CRRA utility in total wealth, and the optimal value function is influenced by the total wealth but not by how that wealth is divided between liquid and illiquid wealth. Second, both Richard and Ye, in the constrained case, confirm that indeed, $p_M^* \rightarrow 0$ as $t \uparrow T$. More specifically, they find that the greater the proportion of total wealth is illiquid, the more insurance is desired. This is important because it confirms the intuition that life insurance is a hedge against lost future income, and the costs of life insurance can be viewed as a proxy for aversion to an early death.

3. INVESTMENT CONSUMPTION WITH RANDOMLY TERMINATING INCOME: FORMULATION

This problem has proved remarkably hard to crack.

– Mark H. A. Davis, speaking about randomly terminating income.

In this chapter, we formulate an investment-consumption model, that consists of an investor with some initial capital who receives an endowment income that may default and who invests and consumes over a planning horizon with an unknown terminal date. In the previous chapter, we introduced Richard's model. The model developed below can be viewed as an extension of that model. Like Richard, we permit consumption, investment, and insurance purchasing in a simple financial market over an uncertain planning horizon.

The main difference between that model and ours lies in the features of the endowment income. In Richard's model, the endowment income is received continuously throughout the entire planning interval. In our model, an endowment income is also received continuously, but the endowment may default *before* the end of the planning interval. In the event that the endowment defaults before the end of the planning interval, the investor must nonetheless continue to consume and invest until the uncertain terminal date. Thus, whereas Richard's model consists of a single random time beyond which no decisions are necessary, our model consists of independent random times, one signaling default of the investor's income and the other signalling the end of the need for control. Like Richard, we allow hedging against the uncertain terminal time by permitting the purchase of terminal date-contingent insurance. In addition, we also allow hedging against the uncertain default time by permitting the purchase of default-contingent insurance. In real financial markets, contracts that protect against default of an income stream are called *credit default swaps*. In a normal credit default swap, if an obligator defaults on payments owed, a third party — the seller of the credit default swap — assumes the insolvent party's obligation to the obligee. Thus, a credit default swap insures the obligee against default. In our model, rather than replace the

investor's income stream with another income stream as is done in real markets, the credit default swap simply pays out a lump sum upon default and therefore operates more like the terminal date-contingent insurance of Richard's model.

An interesting feature of our model is that the dynamics of the investor's total wealth, which serves as the state variable of the control problem, fundamentally change at the time of default. This change in dynamics around an independent random time necessitates the development of two control problems — one associated with the investor's decisions prior to default, and one associated with the investor's decisions after default.

We proceed by first introducing the probabilistic framework for the model. We then define the basic financial market and prescribe the investor's decision variables, which include consumption, investment, and insurance purchasing. Once the decision variables are introduced, we define the investor's liquid wealth and illiquid wealth. These together determine the investor's total wealth, which serves as the state variable for the model. Constraints on the investor's total wealth lead to the set of admissible controls. Finally, an expected-utility cost functional is defined and the main optimal control problem stated.

3.1 PROBABILISTIC FRAMEWORK

Let $(\Omega, \mathbb{F}, \mathcal{F}, \mathbb{P})$ be a probability space upon which is placed a Brownian motion, W , and two random times, τ_M and τ_D . Furthermore, let $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ be the completion of the filtration generated by W , $\mathbf{1}_{\tau_M \leq t}$, and $\mathbf{1}_{\tau_D \leq t}$. Hence, the events $\{\tau_M \leq t\}$ and $\{\tau_D \leq t\}$ are \mathcal{F}_t -measurable for all t . We assume that the Brownian motion and random times are mutually independent. The distributions for the random times τ_M and τ_D are prescribed by survival functions,

$$\bar{F}_M(t | s) = \mathbb{P}[\tau_M > t | \tau_M > s] = \exp\left(-\int_s^t \lambda_M(u) du\right) \quad (3.1)$$

$$\bar{F}_D(t | s) = \mathbb{P}[\tau_D > t | \tau_D > s] = \exp\left(-\int_s^t \lambda_D(u) du\right). \quad (3.2)$$

That is, $(\bar{F}_M(\cdot | s))_{s \geq 0}$ and $(\bar{F}_D(\cdot | s))_{s \geq 0}$ define families of conditional distribution functions for τ_M and τ_D , respectively. The conditional densities for τ_M and τ_D follow from differentiation,

$$\bar{f}_M(t | s) = -\lambda_M(t) \bar{F}_M(t | s) \quad (3.3)$$

$$\bar{f}_D(t | s) = -\lambda_D(t) \bar{F}_D(t | s) \quad (3.4)$$

By independence, the conditional distributions for the minimal random time, $\tau_M \wedge \tau_D$, are

$$\bar{F}(t | s) = \mathbb{P}[\tau_M \wedge \tau_D > t | \tau_M \wedge \tau_D > s] = \bar{F}_M(t | s) \bar{F}_D(t | s), \quad (3.5)$$

and the densities are

$$\bar{f}(t | s) = (\lambda_M(t) + \lambda_D(t)) \bar{F}(t | s).$$

Remark. That the distribution of a random time can be specified by a hazard rate λ requires only the benign assumption that the probability of the random time occurring on sets of Lebesgue measure zero is zero, a natural assumption for mortality and default. Indeed, if $\mathbb{P}[\tau \leq t]$ is absolutely continuous with respect to Lebesgue measure, m , then

$$\mathbb{P}[\tau > t] = \exp\left(-\int_0^t \lambda(u) dm(u)\right),$$

where

$$\lambda(t) = \frac{d\mathbb{P}(t)/dm(t)}{1 - \mathbb{P}[\tau \leq t]}.$$

3.2 CONTROL VARIABLES

With the probability framework fixed, we now begin defining pieces of the investment consumption model. An investor with initial capital x must make certain financial decisions over a time interval, or planning horizon, of uncertain length. The random variable τ_M ,

which will represent a retiree's uncertain time of death in a later chapter, corresponds to the uncertain terminal date. The investor is only concerned with making decisions over the interval $[0, \tau_M \wedge T)$, where T is a fixed upper bound for the planning horizon.

3.2.1 Consumption and Portfolio. The investor must decide how much to consume at each moment throughout the planning interval. Mathematically, the investor's consumption process is an \mathbb{F} -progressively measurable process $c \geq 0$. The investor must also make investment decisions throughout the planning interval. The financial market consists of a risk-free asset S_0 such as a risk-free bond, and a risky asset S_1 such as a stock. The evolutions of these assets are defined by the SDE

$$dS_0(t) = rS_0(t) dt, \quad (3.6)$$

$$dS_1(t) = \mu S_1(t) dt + \sigma S_1(t) dW(t), \quad (3.7)$$

where, $r > 0$ is the constant risk-free rate of return, $\mu > r$ is the constant mean rate of return on the risky asset, and $\sigma > 0$ is the volatility.

Remark. The financial parameters r, μ , and σ are taken to be constants, but in fact nearly every development that follows would work equally well with deterministic, continuous functions.

An investor's investment, or portfolio decisions, π_0, π , are \mathbb{F} -progressively measurable processes such that the process π_0 corresponds to the dollar amount invested in the bond, and π corresponds to the dollar amount invested in the stock. We will see below that there is no need to track π_0 separately.

3.2.2 Insurance. The investor receives a stream of income from an exogeneous endowment at a rate $\alpha(t)$ continuously throughout the planning horizon until the earlier of the time of default τ_D and the end of the planning interval, $\tau_M \wedge \tau_D \wedge T$. The assumption that the τ_D is independent of the Brownian motion and hence, the financial market, does not completely

reflect the relationship between stock market performance and the default probability of other financial instruments, but in some situations it serves as a reasonable approximation. Nevertheless, we assume independence for simplicity and tractability. A portion of the endowment income is considered non-defaultable, either because of the financial stability of its source or because of some institutional guaranty, and it pays an income rate of $\alpha_{ND}(t)$. The defaultable portion of the income rate is denoted $\alpha_D(t)$. That is,

$$\alpha(t) = \begin{cases} \alpha_D(t) + \alpha_{ND}(t) & \text{if } t \leq \tau_D, \\ \alpha_{ND}(t) & \text{if } t > \tau_D, \end{cases}$$

The investor may hedge against default of the endowment income by purchasing τ_D -contingent insurance, i.e. an instantaneous credit default swap. The investor's life insurance premium process is an \mathbb{F} -progressively measurable process p_M . We constrain p_M to be non-negative; that is, we require that the investor can purchase but not sell insurance. We will follow conventions in the literature by writing the per dollar payout rate for insurance contract as $\frac{1}{\eta_D(t)}$; that is, dollar credit default swap contract purchased for $p_D(t)$ is awarded $\frac{p_D(t)}{\eta_D(t)}$ if $\tau_D = t$ but expires otherwise. A heuristic argument suggests that the actuarially fair payout rate per dollar of instantaneous default insurance is $\eta_D(t) = \lambda_D(t)$. Formally, the instantaneous payout per dollar is actuarially fair if

$$1 = \frac{1}{\eta_D(t)} \mathbb{P}^t[\tau_D \in dt] = \frac{1}{\eta_D(t)} \lambda_D(t) \bar{F}_D(t | t) = \frac{\lambda_D(t)}{\eta_D(t)},$$

i.e.,

$$\eta_D(t) = \lambda_D(t).$$

Typically, however, insurance contracts are not actuarially fairly priced, owing to asymmetric information between the buyer and seller. Whether an individual purchases insurance or not depends fundamentally on that individual's subjective beliefs about the probability of some contingency. Therefore, those who insure against an event tend to see that event realized

more frequently than actuarial models calibrated to the general population might imply. One way that insurers compensate for this phenomenon is by loading the actuarially fair hazard rates. In the case of a credit default swap, the insurer assumes that the income stream defaults more frequently for insurance purchasers than the population at large, so the hazard rate is loaded. For example, we could take $\eta_D(t) = k_D \lambda_D(t)$, where $k_D \geq 1$ is the loading factor. Of course, financial intermediaries also require a profit, though we generally ignore transaction costs. Either way, we take $\eta_D \geq \lambda_D$.

The investor may also hedge against the uncertain terminal date by purchasing instantaneous term τ_M -contingent insurance. For concreteness, we will refer to this insurance as *life insurance*. The investor's life insurance premium process is an \mathbb{F} -adapted process p_M that pays $\frac{p_M(t)}{\eta_M(t)}$ if $\tau_M = t$ but expires otherwise. As with the credit default swap, the actuarially fair payout rate is $\eta_M(t) = \lambda_M(t)$, but we allow for $\eta_M \geq \lambda_M$.

3.3 WEALTH PROCESS

In this section, we describe the interplay between the investor's decisions, returns on investment to determine the investor's wealth, culminating in the definition of the state variable for our investment consumption model, the investor's total wealth, as well as the set of admissible controls.

First, the investor's cumulative income process, $\Gamma(t)$ is defined by

$$\Gamma(t) = x + \int_0^t (\alpha_D(s) \mathbf{1}_{s < \tau_D} + \alpha_{ND}(s) - c(s) - p_M(s) - p_D(s)) ds, \quad 0 \leq t \leq \tau_M \wedge T. \quad (3.8)$$

The investor's capital gains process, G , is the return on investments, given in differential form by

$$\begin{aligned} dG(t) &= \pi_0(t) \frac{dS_0(t)}{S_0(t)} + \pi(t) \frac{dS_1(t)}{S_1(t)} \\ &= \pi_0(t) r dt + \pi(t)(\mu dt + \sigma dW(t)), \quad 0 \leq t \leq \tau_M \wedge T. \end{aligned} \quad (3.9)$$

The investor's Γ -financed liquid wealth, L , is defined as the sum of the cumulative income and the gains, subject to the condition that the liquid wealth is always continuously reinvested in the stock and bond,

$$L(t) = \Gamma + G; \quad (3.10)$$

$$L(t) = \pi_0 + \pi. \quad (3.11)$$

From (3.11), we see that the risk-free portfolio process, π_0 , is completely determined by the liquid wealth and the risky portfolio process, π , and therefore, it will not need to be tracked separately from π in the sequel. Combining (3.10) and (3.11) yields the differential for L ,

$$\begin{aligned} dL(t) &= d\Gamma(t) + dG(t) \\ &= (\alpha(t) - c(t) - p_M(t) + rL(t) + (\mu - r)\pi(t)) dt + \sigma\pi(t) dW(t), \quad 0 \leq t \leq \tau_M \wedge T. \end{aligned} \quad (3.12)$$

3.3.1 Borrowing Constraint. While the liquid wealth may become negative, i.e., the investor may borrow in net, the liquid wealth is constrained from below by the expected remaining discounted endowment income, I , defined by

$$I(t) = \begin{cases} I_D(t) + I_{ND}(t) & \text{if } t \leq \tau_D, \\ I_{ND}(t) & \text{if } t > \tau_D, \end{cases}$$

where I_D and I_{ND} are the expected remaining discounted, defaultable and non-defaultable endowment incomes, respectively,

$$I_D(t) = \int_t^T \exp\left(-\int_t^s (r + \eta_M(u) + \eta_D(u)) du\right) \alpha_D(s) ds, \quad (3.13)$$

$$I_{ND}(t) = \int_t^T \exp\left(-\int_t^s (r + \eta_M(u)) du\right) \alpha_{ND}(s) ds. \quad (3.14)$$

Observe that the non-defaultable income is discounted by the time-value of money and terminal-time hazard rate while the defaultable income is discounted by both these and the default rate. Further note in (3.13) and (3.14) that the future income is discounted by η_M instead of λ_M and η_D instead of λ_D so that I is not the actuarially fair expected remaining income. Indeed, the actuarially fair expected remaining incomes are

$$\begin{aligned} & \mathbb{E}^t \left[\int_t^{\tau_M \wedge T \wedge \tau_D} \exp(-r(s-t)) \alpha_D(s) ds \right] \\ &= \int_t^T \exp \left(- \int_t^s (r + \lambda_M(u) + \lambda_D(u)) du \right) \alpha_D(s) ds. \end{aligned}$$

and

$$\begin{aligned} & \mathbb{E}^t \left[\int_t^{\tau_M \wedge T} \exp(-r(s-t)) \alpha_{ND}(s) ds \right] \\ &= \int_t^T \exp \left(- \int_t^s (r + \lambda_M(u)) du \right) \alpha_{ND}(s) ds. \end{aligned}$$

The reason for the discrepancy, as discussed above for insurance purchasing, is that in practice, the investor's future cash flows are discounted above the actuarially fair rate due to adverse selection.

The borrowing constraint for the liquid wealth is thus

$$L(t) \geq \begin{cases} -I(t) & \text{if } t \leq \tau_D, \\ -I_{ND}(t) & \text{if } t > \tau_D. \end{cases}$$

The differential for I is

$$dI(t) = \begin{cases} (-\alpha_D(t) - \alpha_{ND}(t) + (r + \eta_M(t) + \eta_D(t))I_D(t) + (r + \eta_M(t))I_{ND}(t)) dt & \text{if } t \leq \tau_D, \\ (-\alpha_{ND}(t) + (r + \eta_M(t))I_{ND}(t)) dt & \text{if } t > \tau_D. \end{cases}$$

The state variable for our investment and consumption model is the investor's total wealth, X , defined as the sum of the liquid wealth and the expected remaining wealth,

$$X = L + I, \quad 0 \leq t \leq \tau_M \wedge T. \quad (3.15)$$

We will find it convenient to define separately the pre- and post-default total wealth, X_1 and X_2 , respectively, i.e.,

$$X(t) = X_1(t)\mathbf{1}(t \leq \tau_D) + X_2(t)\mathbf{1}(t > \tau_D), \quad 0 \leq t \leq \tau_M \wedge T. \quad (3.16)$$

The differential of the total wealth before default is, after simplification,

$$\begin{aligned} dX_1(t) &= dL(t) + dI(t) \\ &= [-c(t) - \eta_M(t)\zeta_M(t) - \eta_D(t)\zeta_D(t) + (r + \eta_M(t) + \eta_D(t))X_1(t) + (\mu - r)\pi(t)] dt \\ &\quad + \sigma\pi(t) dW(t), \quad 0 \leq t \leq \tau_M \wedge T \wedge \tau_D, \end{aligned} \quad (3.17)$$

and after default,

$$\begin{aligned} dX_2(t) &= dL(t) + dI_{ND}(t) \\ &= [-c(t) - \eta_M(t)\zeta_M(t) + (r + \eta_M(t))X_2(t) + (\mu - r)\pi(t)] dt \\ &\quad + \sigma\pi(t) dW(t), \quad \tau_D \leq t \leq \tau_M \wedge T, \end{aligned} \quad (3.18)$$

where

$$\zeta_M = X - I_D - I_{ND} + \frac{p_M}{\eta_M}, \quad (3.19)$$

$$\zeta_D = X - I_D + \frac{p_D}{\eta_D}. \quad (3.20)$$

Intuitively, ζ_M is the investor's total wealth plus the insurance payout, less the unrealized future income; that is, $\zeta_M(t)$ is the investor's terminal wealth given death occurs at t . Furthermore, ζ_D can be interpreted as the investor's total wealth plus the default payout, less the defaulted income; that is, $\zeta_D(t)$ is the investor's total wealth at the time of default. The insurance variables p_M and p_D can easily be recovered from ζ_M and ζ_D , and we will find it convenient to work with ζ_M and ζ_D . The no-short selling constraints on insurance, $p_M, p_D \geq 0$, lead to constraints on ζ_M and ζ_D , namely,

$$\zeta_M(t) \geq X(t) - I(t) \quad (3.21)$$

$$\zeta_D(t) \geq X(t) - I_D(t). \quad (3.22)$$

The constraints (3.21) , (3.22) imply the non-negativity of ζ_M and ζ_D . Of course, there is no need for default insurance post-default, so we automatically take $\zeta_D(t) = 0$ for $t > \tau_D$.

By the constraint on the liquid wealth, X_1 and X_2 are each constrained to be nonnegative,

$$X_1, X_2 \geq 0. \quad (3.23)$$

When necessary to avoid confusion, we will denote the pre-default decision variables by $c_1, \pi_1, \zeta_{M_1}, \zeta_{D_1}$ and the post-default decision variables by $c_2, \pi_2, \zeta_{M_2}, \zeta_{D_2}$.

3.3.2 Admissible Controls. Having defined the investor's total wealth and introduced the full set of constraints, we are now prepared to define the set of admissible controls. Let $\mathbb{G} = (\mathcal{G}_t)$ be the completion of the filtration generated by Brownian-motion. Note that this filtration is automatically right-continuous, so it satisfies the usual hypotheses.

Definition 3.1 (Pre-default Admissible Controls). The set of pre-default admissible controls, \mathcal{U}_1 , is the collection of \mathbb{G} -progressively measurable consumption, portfolio, and insurance processes, c, π, ζ_M, ζ_D , such that

$$(i) \int_0^T |c(s)| + |\pi(s)| + |\eta_M(s)\zeta_M(s)| + |\eta_D(s)\zeta_D(s)| + \pi^2(s) ds < \infty;$$

(ii) ζ_M, ζ_D satisfy (3.21)-(3.22); and

(iii) $X_1 \geq 0$.

Definition 3.2 (Post-default Admissible Controls). The set of post-default admissible controls, \mathcal{U}_2 , is the collection of \mathbb{G} -progressively measurable consumption, portfolio, and insurance processes, c, π, ζ_M , such that

(i) $\int_0^T |c(s)| + |\pi(s)| + |\eta_M(s)\zeta_M(s)| + \pi^2(s) ds < \infty$;

(ii) ζ_M satisfies (3.21); and

(iii) $X_2 \geq 0$.

Note that the drift and volatility for the pre-default wealth and post-default wealth are uniformly Lipschitz continuous for each fixed control. Therefore, condition (ii) in the respective definitions of admissible controls is sufficient to guarantee existence and uniqueness of \mathbb{G} -progressively measurable wealth processes, X_1 and X_2 . Also, notice that the pre-default and post-default controls are completely ignorant of the default time τ_D .

We now complete the investment-consumption model by defining the cost functional.

3.4 COST-FUNCTIONALS AND CONTROL PROBLEM

The cost functional for our model is actually a pair of coupled cost functionals that treat the investor's pre-default decisions and post-default decisions independently. We first motivate the definition of the cost functionals.

3.5 MOTIVATION

In Definition 2.11, we introduced utility functions, which led to a natural cost functional for investment-consumption problems. For our model, the natural cost functional possesses the

same generic form as for Richard's problem,

$$J(t, x, c, \pi, \zeta_M, \zeta_D) = \mathbb{E}^{t,x} \left[\int_t^{\tau_M \wedge T} U(s, c(s)) ds + B(\tau_M, \zeta_M(\tau_M)) \right], \quad (3.24)$$

with associated value function

$$V(t, x) = \sup_{c, \pi, \zeta_M, \zeta_D} J(t, x, c, \pi, \zeta_M, \zeta_D). \quad (3.25)$$

Here, $\mathbb{E}^{t,x}$ is the expectation conditioned on $X(t) = x$ as well as $\tau_M, \tau_D > t$.

We develop J by first partitioning the planning horizon around τ_M . Noting that there is no utility for consumption after $\tau_M \wedge T$, the argument of the expectation in (3.24) becomes

$$\begin{aligned} \int_t^{\tau_M \wedge T} U(s, c(s)) ds + B(\tau_M, \zeta_M(\tau_M)) &= \int_t^{\tau_M \wedge T \wedge \tau_D} U(s, c(s)) ds \\ &+ \int_{\tau_D}^{\tau_M \wedge T} U(s, c(s)) \mathbf{1}_{\tau_D < \tau_M \wedge T} ds + B(\tau_M, \zeta_M(\tau_M)) \mathbf{1}_{\tau_D < \tau_M \wedge T} + B(\tau_M, \zeta_M(\tau_M)) \mathbf{1}_{\tau_D \geq \tau_M \wedge T}. \end{aligned} \quad (3.26)$$

Assume for the moment that the strong Markov property for X can be applied to (3.26) for the \mathbb{F} -stopping time τ_D , i.e., for all Borel-measurable f ,

$$\mathbb{E}^{t,x} [f(X(\tau_D + s)) | \mathcal{F}_{\tau_D}] = \mathbb{E}^{\tau_D, X(\tau_D)} [f(X(s))]. \quad (3.27)$$

The strong Markov property for X at τ_D provides the following lemma, a proof of which is in the appendix.

Lemma 3.3.

$$\begin{aligned} V(t, x) &= \sup_{c, \pi, \zeta_M, \zeta_D} \mathbb{E}^{t,x} \left[\int_t^{\tau_M \wedge T \wedge \tau_D} U(s, c(s)) ds + B(\tau_M, \zeta_M(\tau_M)) \mathbf{1}_{\tau_D \geq \tau_M \wedge T} \right. \\ &\quad \left. + D(\tau_D, X(\tau_D)) \mathbf{1}_{\tau_D < \tau_M \wedge T} \right], \end{aligned} \quad (3.28)$$

where

$$D(t, x) = \sup_{c, \pi, \zeta_M} \mathbb{E}^{t, x} \left[\int_t^{T \wedge \tau_M} U(s, c(s)) + B(\tau_M, \zeta_M(\tau_M)) \right]. \quad (3.29)$$

Remark. The expectation in the definition of D above is conditioned on $\tau_D = t, \tau_M > t, X(\tau_D) = x$, as in (3.27).

At this point, the usual path forward is

- (i) apply the expectation operator in (3.28) for τ_D and τ_M independently of W to eliminate explicit dependence on random times;
- (ii) apply the dynamic programming principle to the value function, V ; and
- (iii) derive the Hamilton-Jacobi-Bellman equation and apply the Markov chain approximation.

We proceed with (i), at least formally. Then we have

$$\begin{aligned} V(t, x) &= \sup_{c, \pi, \zeta_M, \zeta_D} \mathbb{E}^{t, x} \left[\int_t^{\tau_M \wedge T \wedge \tau_D} U(s, c(s)) ds + B(\tau_M, \zeta_M(\tau_M)) \mathbf{1}_{\tau_D \geq \tau_M \wedge T, \tau_M < T} \right. \\ &\quad \left. + B(T, \zeta_M(T)) \mathbf{1}_{\tau_D \geq \tau_M \wedge T, \tau_M \geq T} + D(\tau_D, X(\tau_D)) \mathbf{1}_{\tau_D < \tau_M \wedge T} \right], \\ &= \sup_{c, \pi, \zeta_M, \zeta_D} \mathbb{E}^{t, x} \left[\int_t^{\tau_M \wedge T \wedge \tau_D} U(s, c(s)) ds + B(\tau_M, \zeta_M(\tau_M)) \mathbf{1}_{\tau_M < T \wedge \tau_D} \right. \\ &\quad \left. + B(T, X(T)) \mathbf{1}_{T < \tau_M \wedge \tau_D} + D(\tau_D, X(\tau_D)) \mathbf{1}_{\tau_D < \tau_M \wedge T} \right], \end{aligned} \quad (3.30)$$

where we have used the fact that if the investor hasn't died by time T , they must accept the utility for their terminal wealth at T as well as the fact that $p_M(T) = 0$.

Suppose now that we apply the expectation operator in (3.30) independently to τ_M and

τ_D . Then

$$\begin{aligned}
V(t, x) &= \sup_{c, \pi, \zeta_M, \zeta_D} \mathbb{E}^{t, x} \left[\int_t^T U(s, c(s)) \bar{F}(s | t) ds + \int_t^T B(s, \zeta_M(s)) \mathbf{1}_{s < \tau_D} f_M(s | t) \right. \\
&\quad \left. + B(T, X(T)) \bar{F}(T | t) + \int_t^T D(s, X(s)) \mathbf{1}_{s < \tau_M} f_D(s | t) \right], \\
&= \sup_{c, \pi, \zeta_M, \zeta_D} \mathbb{E}^{t, x} \left[\int_t^T (U(s, c(s)) + \lambda_M(s) B(s, \zeta_M(s)) + \lambda_D(s) D(s, X(s))) \bar{F}(s | t) ds \right. \\
&\quad \left. + B(T, X(T)) \bar{F}(T | t) \right]. \tag{3.31}
\end{aligned}$$

(3.31) can be interpreted as rather than worrying about precisely when the random times occur, the investor may equivalently view death and default as happening continuously throughout the interval, but with the utilities for terminal wealth and post-default optimal decisions scaled down by the respective hazard rate.

It is, unfortunately, difficult to justify the application of Fubini's theorem in the preceding given that the controls and state variable are not independent of the random times. There are also difficult technical issues pertaining to the the appropriate filtration for the application of the dynamic programming principle to the value function, V . However, the formal manipulations above suggest a pair of cost functionals for the pre-default wealth and post-default wealth, respectively, to which dynamic programming techniques readily apply. These are defined next.

3.5.1 Post-default Cost Functional. The investor's cost functional, post-default, is the map $J_2 : [0, T] \times [0, \infty) \times \mathcal{U}_2$ defined by

$$\begin{aligned}
J_2(t, x, c, \pi, \zeta_M) &= \mathbb{E}^{t, x} \left[\int_t^T (U(s, c(s)) + \lambda_M(s) B(s, \zeta_M(s))) \bar{F}_M(s | t) ds \right. \\
&\quad \left. + B(T, X(T)) \bar{F}_M(T | t) \right], \tag{3.32}
\end{aligned}$$

where $\mathbb{E}^{t, x}$ is the expectation operator conditioned only on $X_2(t) = x$.

The post-default value function, D , is defined by

$$D(t, x) = \sup_{(c, \pi, \zeta_M) \in \mathcal{U}_2} J_2(t, x, c, \pi, \zeta_M). \quad (3.33)$$

Intuitively, D optimizes the investor's utility over the decision variables from the time of default to the end of the planning horizon for each time of default. Given D , it remains only to optimize the investor's decision-making until default, which is the purpose of the pre-default cost functional.

3.5.2 Pre-default Cost Functional. The pre-default cost functional is the map $J_1 : [0, T] \times [0, \infty) \times \mathcal{U}_1$ defined by

$$J_1(t, x, c, \pi, \zeta_M, \zeta_D) = \mathbb{E}^{t,x} \left[\int_t^T (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s)) + \lambda_D(s)D(s, \zeta_D(s))) \bar{F}(s|t) + B(T, X(T))\bar{F}_M(T|t) \right], \quad (3.34)$$

where $\mathbb{E}^{t,x}$ is the expectation conditioned only on $X_1(t) = x$. The pre-default value function, V , is defined by

$$V(t, x) = \sup_{(c, \pi, \zeta_M, \zeta_D) \in \mathcal{U}_1} J_1(t, x, c, \pi, \zeta_M, \zeta_D). \quad (3.35)$$

The primary control problem is then

Problem (V). Find pre-default admissible controls $c_1^*, \pi_1^*, \zeta_{M_1}^*, \zeta_{D_1}^*$, if they exist, such that

$$V(t, x) = J_1(t, x, c_1^*, \pi_1^*, \zeta_{M_1}^*, \zeta_{D_1}^*). \quad (3.36)$$

In this case, the maximizing set of control policies are *post-default optimal*.

Of course, a pre-requisite to solving Problem (V) is solving the sub-problem

Problem (D). Find post-default admissible controls, $c_2^*, \pi_2^*, \zeta_{M_2}^*$, if they exist, such that

$$D(t, x) = J_2(t, x, c_2^*, \pi_2^*, \zeta_{M_2}^*). \quad (3.37)$$

In this case, the maximizing set of control policies are *pre-default optimal*.

The optimal policies for the whole planning horizon are “glued” together from the optimal pre-default and post-default controls, if they exist:

Definition 3.4 (Optimal Control Policies). Given a set of pre-default optimal controls, $(c_1, \pi_1, \zeta_{M_1}, \zeta_{D_1})$ and a set of post-default optimal controls, $(c_2, \pi_2, \zeta_{M_2})$, an *optimal control policy*, $(c, \pi, \zeta_M, \zeta_D)$, over the full planning horizon is the \mathbb{F} -progressively measurable process defined by

$$(c(t), \pi(t), \zeta_M(t), \zeta_D(t)) = \begin{cases} (c_1(t), \pi_1(t), \zeta_{M_1}(t), \zeta_{D_1}(t)) & \text{if } t \in [0, \tau_D \wedge \tau_M \wedge T], \\ (c_2(t), \pi_2(t), \zeta_{M_2}(t), 0) & \text{if } t \in (\tau_D, \tau_M \wedge T]. \end{cases}$$

The advantage to treating the optimal control problem before default and after default separately is that dynamic programming techniques and the Markov chain approximation can now be employed to approximate the value function and optimal controls.

Now that the admissible controls, state variables, and cost functionals have been defined, the investment consumption model is complete. In the next section, we present a technique for approximating, numerically, solutions to Problems (D) and (V), but first, we make a few observations. In our model, the key decision variable both mathematically and practically is the default insurance. Without default insurance, the investor’s borrowing limit is no longer the whole expected future income but only the non-defaultable income since without some form of insurance, there remains the possibility that default could occur and leave the investor unable to fulfill obligations made against the defaultable income. Mathematically, this requires the imposition of a moving boundary condition, and it is unclear how to proceed in solving the associated control problem. The addition of default insurance to the model allows the investor to guarantee debts, so the investor’s borrowing limit can jump at the time of default without creating the potential for insolvency. Hence, default insurance allows for a smooth transition from the pre-default problem to the post-default problem.

Pragmatically, default insurance can be seen as a proxy for the investor's aversion to default. Because the investor is averse to the loss of the remaining income at default, insurance purchases are made as a hedge. Both Richard [14] and Ye [13] showed that the optimal amount of life insurance purchased by an investor is directly proportional to the size of future income relative to wealth. That is, the greater the percentage of wealth that is tied up in future income and hence would be lost through an early death, the more insurance is purchased. Similarly, we expect that the amount of default insurance purchased should be proportional to the percentage of total wealth that would be lost upon default, so part of the cost of an optimal default insurance policy is the price the investor is willing to pay in order to avoid loss of capital through default.

3.6 MARKOV CHAIN APPROXIMATION

In this section, we first derive the HJB equations for D and V . Recall that the HJB equations can be discretized and then interpreted as the dynamic programming principle for a discrete-time control problem. However, as we encountered in Richard's model in Section 2.4, direct discretization of the HJB equations is neither amenable to the discrete dynamic programming interpretation nor to numerical implementation. Therefore, after obtaining the HJB equations, we apply Ye's logarithmic transformation, and the finite difference method is then applied to the transformed HJB.

3.6.1 Post-Default HJB Equation. We first present the HJB equation for the post-default utility D , which serves as a sort of subproblem for V . Standard dynamic programming techniques, details of which are provided in the appendix, show that any solution to Problem (D) satisfies

$$0 = \sup_{c, \pi, \zeta_M \in \mathcal{U}_2} \left\{ -\lambda_M(t)D(t, x) + D_t(t, x) + b_2(t, x)D_x(t, x) + \pi^2(t)\sigma^2 D_{xx}(t, x) + U(t, c(t)) + \lambda_M(t)B(t, \zeta_M(t)) \right\}, \quad (3.38)$$

where

$$b_2(t, x) = -c(t) - \eta_M(t)\zeta_M(t) + (r + \eta_M(t))x + (\mu - r)\pi(t).$$

D also satisfies a terminal condition and a boundary condition. The terminal condition follows directly from (3.32)

$$D(T, x) = B(T, x), \quad x \in (0, \infty). \quad (3.39)$$

The boundary condition follows from the constraint for X_2 . If the investor ever reaches the debt limit, all future funds are allocated to debt repayment, so there can be no further consumption.

$$\lim_{x \downarrow 0} D(t, x) = -\infty, \quad t \in [0, T]. \quad (3.40)$$

3.6.2 Pre-Default HJB Equation. Next, we provide the HJB equation for V . Again, the details of the derivation are in the appendix.

$$\begin{aligned} 0 = \sup_{c, \pi, \zeta_M, \zeta_D \in \mathcal{U}_1} & -(\lambda_M(t) + \lambda_D(t))V(t, x) + V_t(t, x) + b_1(t, x)V_x(t, x) + \pi^2(t)\sigma^2V_{xx}(t, x) \\ & + U(t, c(t)) + \lambda_M(t)B(t, \zeta_M(t)) + \lambda_D(t)D(t, \zeta_D(t)), \end{aligned} \quad (3.41)$$

where

$$b_1(t, x) = -c(t) - \eta_M(t)\zeta_M(t) - \eta_D(t)\zeta_D(t) + (r + \eta_M(t) + \eta_D(t))x + (\mu - r)\pi(t).$$

The terminal and boundary conditions for V similar to those for D ,

$$\begin{cases} V(T, x) = B(T, x) & x \in (0, \infty) \\ \lim_{x \downarrow 0} V(t, x) = -\infty & t \in [0, T]. \end{cases}$$

We reiterate that the purpose of the HJB equations are only to guide the choice of

discrete-time control problem and discrete-time Markov chain. To rigorously establish that solutions to (3.38) and (3.41) together with their respective terminal and boundary conditions are solutions to the optimal control problems (D) and (V), respectively, would require a so-called verification theorem, which we do not pursue here.

3.6.3 Logarithmic Transformation. Just as for the HJB equation derived in Section 2.4, (3.38) and (3.41) present difficulties for applying MCA. First, note that that in both equations, the coefficient of V_x depends explicitly on x , so direct discretization does not yield a discrete-time dynamic programming interpretation. Second, the singular boundary condition at $x = 0$ poses a serious challenge for a numerical implementation. We overcome both of the difficulties by following Ye in making a logarithmic transformation in order to obtain an equivalent partial differential equation that can be appropriately discretized and that moves the boundary out of the computational region.

The logarithmic transformation is accomplished changing the total wealth variable of the HJB equations by

$$y = \log(x). \quad (3.42)$$

When necessary we will denote the pre- and post-default transformed variables as y_1 and y_2 , respectively. Furthermore, we define the transformed value functions, \widehat{V} and \widehat{D} by

$$\widehat{V}(t, y) = V(t, \exp(y)) = V(t, x) \quad (3.43)$$

$$\widehat{D}(t, y) = D(t, \exp(y)) = D(t, x). \quad (3.44)$$

Hence, \widehat{V} and \widehat{D} equal the original value functions V and D so that the study of the transformed problem directly yields the results of the original problem. Derivatives of D and \widehat{D}

are related by

$$\begin{aligned}
 D_t(t, x) &= \widehat{D}_t(t, y) \\
 D_x(t, x) &= \exp(-y)\widehat{D}_y(t, y) \\
 D_{xx}(t, x) &= \exp(-2y)(\widehat{D}_{yy}(t, y) - \widehat{D}_y(t, y)),
 \end{aligned}$$

and similarly for V

$$\begin{aligned}
 V_t(t, x) &= \widehat{V}_t(t, y) \\
 V_x(t, x) &= \exp(-y)\widehat{V}_y(t, y) \\
 V_{xx}(t, x) &= \exp(-2y)(\widehat{V}_{yy}(t, y) - \widehat{V}_y(t, y)).
 \end{aligned}$$

Finally, we transform the control variables,

$$\widehat{c} = \exp(-y)c, \quad (3.45)$$

$$\widehat{\pi} = \exp(-y)\pi, \quad (3.46)$$

$$\widehat{\zeta}_M = \exp(-y)\zeta_M, \quad (3.47)$$

$$\widehat{\zeta}_D = \exp(-y)\zeta_D, \quad (3.48)$$

and note that because the transformation is smooth and monotonically increasing, a policy is optimal in the transformed variables if and only if the corresponding untransformed policy is optimal. The transformed control variables are interpreted as the percentage of the investor's total wealth allocated to the various decisions, rather than the dollar value.

Substituting the transformed derivatives and transformed variables into (3.38) and (3.41)

and simplifying, we obtain

$$0 = \sup_{c, \pi, \zeta_M \in \mathcal{U}_2} \left\{ -\lambda_M(t) \widehat{D}(t, y) + \widehat{D}_t(t, y) + \widehat{b}_2(t, y) \widehat{D}_y(t, y) + \frac{\widehat{\pi}^2(t) \sigma^2}{2} \widehat{D}_{yy}(t, y) + U(t, \exp(y) \widehat{c}) + \lambda_M(t) B(t, \exp(y) \widehat{\zeta}_M(t)) \right\}, \quad (3.49)$$

$$0 = \sup_{c, \pi, \zeta_M, \zeta_D \in \mathcal{U}_1} \left\{ -(\lambda_M(t) + \lambda_D(t)) \widehat{V}(t, y) + \widehat{V}_t(t, y) + \widehat{b}_1(t, y) \widehat{V}_y(t, y) + \frac{\widehat{\pi}^2(t) \sigma^2}{2} \widehat{V}_{yy}(t, y) + U(t, \exp(y) \widehat{c}) + \lambda_M(t) B(t, \exp(y) \widehat{\zeta}_M(t)) + \lambda_D(t) \widehat{D}(t, y + \log(\widehat{\zeta}_D(t))) \right\}, \quad (3.50)$$

where

$$\widehat{b}_2(t, y) = -\widehat{c}(t) - \eta_M \widehat{\zeta}_M(t) + r + \eta_M(t) + \widehat{\pi}(t)(\mu - r) - \frac{\widehat{\pi}^2(t) \sigma^2}{2},$$

$$\widehat{b}_1(t, y) = -\widehat{c}(t) - \eta_M \widehat{\zeta}_M(t) - \eta_D(t) \widehat{\zeta}_D(t) + r + \eta_M(t) + \eta_D(t) + \widehat{\pi}(t)(\mu - r) - \frac{\widehat{\pi}^2(t) \sigma^2}{2}.$$

The transformed terminal and boundary conditions are

$$\widehat{D}(T, y) = B(T, \exp(y)) \quad \text{and} \quad \widehat{D}_{y \downarrow -\infty}(t, y) = -\infty, \quad (3.51)$$

$$\widehat{V}(T, y) = B(T, \exp(y)) \quad \text{and} \quad \widehat{V}_{y \downarrow -\infty}(t, y) = -\infty. \quad (3.52)$$

Note that the HJB equation for \widehat{V} involves \widehat{D} instead of D , which will be convenient for optimization purposes as it makes use of the post- and pre-default log wealth variables relationship,

$$y_2 = y_1 + \log(\widehat{\zeta}_D^1).$$

With the HJB equations in hand, we continue with the finite difference method by discretizing the continuous-time PDE to get discrete-time difference relations.

3.6.4 Discretization. As discussed in Section 2.3, the discretization must be chosen carefully to ensure that the finite difference relations can be interpreted as the expectation of a discrete-time process. Let δ be the discrete time step size, and let h be the discrete

wealth step size. Finite difference formulas are chosen as follows:

$$\widehat{D}_t(t, y) = \frac{\widehat{D}^{\delta, h}(t + \delta, y) - \widehat{D}^{\delta, h}(t, y)}{\delta} + o(\delta) \quad (3.53)$$

$$\widehat{b}_2^+(t, y)\widehat{D}_y(t, y) = \widehat{b}_2^+(t, y)\frac{\widehat{D}^{\delta, h}(t + \delta, y + h) - \widehat{D}^{\delta, h}(t + \delta, y)}{h} + o(h), \quad (3.54)$$

$$\widehat{b}_2^-(t, y)\widehat{D}_y(t, y) = \widehat{b}_1^-(t, y)\frac{\widehat{D}^{\delta, h}(t + \delta, y) - \widehat{D}^{\delta, h}(t + \delta, y - h)}{h} + o(h), \quad (3.55)$$

$$\widehat{D}_{yy}(t, y) = \frac{\widehat{D}^{\delta, h}(t + \delta, y) + \widehat{D}^{\delta, h}(t, y - h) - 2\widehat{D}^{\delta, h}(t, y)}{h^2} + o(h), \quad (3.56)$$

where

$$\widehat{b}_2^+(t, y) = r + \eta_M(t) + \widehat{\pi}(t)(\mu - r) \quad (3.57)$$

$$\widehat{b}_2^-(t, y) = \widehat{c}(t) + \eta_M(t)\widehat{\zeta}_M(t) + \frac{\sigma^2\widehat{\pi}(t)^2}{2}, \quad (3.58)$$

and

$$\widehat{V}_t(t, y) = \frac{\widehat{V}^{\delta, h}(t + \delta, y) - \widehat{V}^{\delta, h}(t, y)}{\delta} + o(\delta), \quad (3.59)$$

$$\widehat{b}_1^+(t, y)\widehat{V}_y(t, y) = \widehat{b}_1^+(t, y)\frac{\widehat{V}^{\delta, h}(t + \delta, y + h) - \widehat{V}^{\delta, h}(t + \delta, y)}{h} + o(h), \quad (3.60)$$

$$\widehat{b}_1^-(t, y)\widehat{V}_y(t, y) = \widehat{b}_1^-(t, y)\frac{\widehat{V}^{\delta, h}(t + \delta, y) - \widehat{V}^{\delta, h}(t + \delta, y - h)}{h} + o(h), \quad (3.61)$$

$$\widehat{V}_{yy}(t, y) = \frac{\widehat{V}^{\delta, h}(t + \delta, y) + \widehat{V}^{\delta, h}(t, y - h) - 2\widehat{V}^{\delta, h}(t, y)}{h^2} + o(h), \quad (3.62)$$

where

$$\widehat{b}_1^+(t, y) = r + \eta_M(t) + \eta_D(t) + \widehat{\pi}(t)(\mu - r) \quad (3.63)$$

$$\widehat{b}_1^-(t, y) = \widehat{c}(t) + \eta_M(t)\widehat{\zeta}_M(t) + \eta_D(t)\widehat{\zeta}_D(t) + \frac{\sigma^2\widehat{\pi}(t)^2}{2}. \quad (3.64)$$

Inserting the finite difference formulas (3.53) - (3.56), (3.59) - (3.62) into the HJB equations (3.49) and (3.51), we obtain for \widehat{D} ,

$$\begin{aligned} \widehat{D}^{\delta,h}(t, y) = & \frac{1}{1 + \delta\lambda_M(t)} \sup_{\widehat{c}, \widehat{\pi}, \widehat{\zeta}_M} \left\{ Q_{\delta,h}^{t,y}(y+h) \widehat{D}^{\delta,h}h(t+\delta, y+h) + Q_{\delta,h}^{t,y}(y) \widehat{D}^{\delta,h}(t+\delta, y) \right. \\ & \left. + Q_{\delta,h}^{t,y}(y-h) \widehat{D}^{\delta,h}(t+\delta, y-h) + \delta U(t, \exp(y)\widehat{c}) + \delta\lambda_M(t)B(t, \exp(y)\widehat{\zeta}_M) \right\}, \end{aligned} \quad (3.65)$$

where

$$Q_{\delta,h}^{t,y}(z) = \begin{cases} \frac{\delta}{h} \left(b_2^+(t, y) + \frac{\widehat{\pi}^2(t)\sigma^2}{2h} \right) & \text{if } z = y+h \\ \frac{\delta}{h} \left(b_2^-(t, y) + \frac{\widehat{\pi}^2(t)\sigma^2}{2} \right) & \text{if } z = y-h \\ 1 - Q_{\delta,h}^{t,y}(y+h) - Q_{\delta,h}^{t,y}(y-h) & \text{if } z = y \end{cases}$$

Similarly for \widehat{V} ,

$$\begin{aligned} V^{\delta,h}(t, y) = & \frac{1}{1 + \delta(\lambda_M(t) + \lambda_D(t))} \sup_{\widehat{c}, \widehat{\pi}, \widehat{\zeta}_M} \left\{ P_{\delta,h}^{t,y}(y+h) \widehat{V}^{\delta,h}(t+\delta, y+h) + P_{\delta,h}^{t,y}(y) \widehat{V}^{\delta,h}(t+\delta, y) \right. \\ & \left. + P_{\delta,h}^{t,y}(y-h) \widehat{V}^{\delta,h}(t+\delta, y-h) + \delta U(t, \exp(y)\widehat{c}) + \delta\lambda_M(t)B(t, \exp(y)\widehat{\zeta}_M) \right. \\ & \left. + \delta\lambda_D(t) \widehat{D}(t, y + \log(\widehat{\zeta}_D(t))) \right\}, \end{aligned} \quad (3.66)$$

where

$$P_{\delta,h}^{t,y}(z) = \begin{cases} \frac{\delta}{h} \left(b_1^+(t, y) + \frac{\widehat{\pi}^2(t)\sigma^2}{2h} \right) & \text{if } z = y+h \\ \frac{\delta}{h} \left(b_1^-(t, y) + \frac{\widehat{\pi}^2(t)\sigma^2}{2h} \right) & \text{if } z = y-h \\ 1 - P_{\delta,h}^{t,y}(y+h) - P_{\delta,h}^{t,y}(y-h) & \text{if } z = y \end{cases}$$

The transition probabilities $Q_{\delta,h}^{t,y}$ and $P_{\delta,h}^{t,y}$, for appropriate choices of δ and h , constitute legitimate transition probabilities. As noted above, the transformed control variables are each a percentage of total wealth, and we may reasonably expect them to be relatively small percentages. Moreover, the consumption and insurance policies are constrained to be nonnegative. In addition, the financial parameters in real markets are relatively small, as

well. Hence, we have $0 \leq Q_{\delta,h}^{t,y}(z), P_{\delta,h}^{t,y}(z) \leq 1$ for all t, z .

3.6.5 Numerical Algorithm. The difference relations (3.65) and (3.66), along with the respective terminal conditions, through backward iterations, can be employed to numerically approximate \widehat{V} and \widehat{D} as well as the optimal control policies. For an initial wealth y_0 and temporal and spatial parameters δ and h , respectively, we generate a recombining trinomial tree that originates from $(t, y) = (0, y_0)$. The post-default utility, \widehat{D} must be approximated first because its values are required in the finite difference relation for \widehat{V} .

In the optimization step of the numerical algorithm for \widehat{D} , the optimal policies \widehat{c}_2^* , $\widehat{\pi}_2^*$, and $\widehat{\zeta}^* M_2$ are obtained from

$$\widehat{c}_2^*(t, y) = \arg \max_{\widehat{c} \geq 0} \left\{ \delta U(t, \exp(y)\widehat{c}(t)) - \widehat{c}(t) \frac{\widehat{D}^{\delta,h}(t + \delta, y) - \widehat{D}^{\delta,h}(t + \delta, y - h)}{h} \right\}, \quad (3.67)$$

$$\begin{aligned} \widehat{\pi}_2^*(t, y) = \arg \max_{\widehat{\pi}} \left\{ \widehat{\pi}(t)(\mu - r) \frac{\widehat{D}(t + \delta, y + h) - \widehat{D}(t + \delta, y)}{h} \right. \\ \left. + \widehat{\pi}^2(t) \left(\frac{\widehat{D}^{\delta,h}(t + \delta, y + h) + \widehat{D}^{\delta,h}(t + \delta, y - h) - 2\widehat{V}(t + \delta, y)}{h^2} \right. \right. \\ \left. \left. - \frac{\widehat{D}^{\delta,h}(t + \delta, y) - \widehat{D}(t + \delta, y - h)}{h} \right) \right\}, \quad (3.68) \end{aligned}$$

$$\begin{aligned} \widehat{\zeta}_M^* M_2(t, y) = \arg \max_{\widehat{\zeta}_M \geq 1 - \exp(-y)I(t)} \left\{ \delta B(t, \exp(y)\widehat{\zeta}_M(t)) \right. \\ \left. - \eta_M(t)\widehat{\zeta}_M(t) \frac{\widehat{D}^{\delta,h}(t + \delta, y) - \widehat{D}^{\delta,h}(t + \delta, y - h)}{h} \right\}. \quad (3.69) \end{aligned}$$

Since the functional forms of U and B are known and $\widehat{D}(t + \delta, y)$ is known for all y , analytic methods can be used to find the optimal policies. Indeed, for CRRA utilities, the RHS of (3.67)-(3.69) either have a unique local and global maximum or else no maximum exists, so either the optimal policies are at the global maximum, if feasible, or else the constraint is tight. After the optimization is performed at each discretized $t = k\delta$, we update $\widehat{D}(k\delta, y)$ over the discretized wealth grid via the difference formula (3.65). Proceeding iteratively backward, we obtain a numerical approximation to the solution of Problem (D).

Once \widehat{D} is approximated numerically, \widehat{V} can also be approximated. The algorithm for \widehat{V} is virtually the same as for \widehat{D} , though with a complication at the optimization step. The optimization for the control policies \widehat{c}_1 , $\widehat{\pi}_1$, and ζ_{M_1} is once again straightforward. However, the optimization for ζ_{D_1} requires some care. The optimal policy $\zeta_{D_1}^*$ is obtained from

$$\zeta_{D_1}^*(t, y) = \arg \max_{\widehat{\zeta}_D \geq 1 - \exp(-y)I_D(t)} \left\{ \delta \widehat{D}^{\delta, h}(t, y + \widehat{\zeta}_D(t)) - \eta_D(t) \widehat{\zeta}_D(t) \frac{\widehat{V}^{\delta, h}(t + \delta, y) - \widehat{V}^{\delta, h}(t + \delta, y - h)}{h} \right\} \quad (3.70)$$

Although the RHS of (3.70) is analogous to the RHS of (3.69), \widehat{D} is known only numerically. If a feasible optimal policy exists, then the optimization step is straightforward. Otherwise, we must rely on the fact that \widehat{D} is strictly increasing because it is a utility, so the constraint must be tight. In that case, we interpolate over the numerical data for \widehat{D} to the boundary in $\widehat{\zeta}_D$. Proceeding iteratively backward, we obtain a numerical approximation to the solution of Problem (V).

As a final comment, in Section 2.4, we noted that the standard convergence proofs for MCA do not apply directly to Richard's model, and similarly, they cannot be directly applied to the model developed in this chapter. We postpone further discussion of convergence until Chapter 5.

4. APPLICATION: LIFE ANNUITIES

And Methuselah lived after he begat Lamech seven hundred eighty and two years, and begat sons and daughters. And all the days of Methuselah were nine hundred sixty and nine years: and he died.

– Genesis 5:27

One of the most important concerns of financial researchers and policy makers in the United States is optimal financing of retirement. Indeed, in 2003 a leading academic on pension finance and investment strategy, Zvi Bodie, wrote, “For many people, the most important goal of financial planning is an adequate retirement income.” [2] In the past, employers offered retirement pensions that served to provide for the needs of retirees. However, in recent years, employers have been modifying or eliminating their pension funds, placing the responsibility for retirement financing on the individuals. Unfortunately, many retirees are not well-prepared to identify the investment vehicles that best match their needs and preferences for risk from among the many complex financial instruments available today. Writes Bodie, “From a social welfare perspective, this development might actually be a step backward. Risk is being transferred to those who are least qualified to manage it [W]e seem to expect people to choose an appropriate mix of stocks, bonds, and cash after reading a brochure published by an investment company. Some people are likely to make serious mistakes.”

One of the most effective financial products available to help individuals fund retirement is a life annuity. Retirees seek to maintain a reasonably comfortable lifestyle throughout retirement on the one hand while also seeking to provide a legacy to any heirs. If retirees consume too aggressively in order to accomplish the former, they risk outliving their savings. If they consume too conservatively to accomplish the later, or simply to avoid outliving their savings, they may unnecessarily forego a more comfortable lifestyle and increased happiness. As discussed in the introduction, researchers have demonstrated under a variety of assump-

tions about financial market behavior as well as individual attitudes towards risk that one of the most efficient means by which retirees can maintain a reasonably comfortable lifestyle and still leave an acceptable legacy is by placing a significant portion of their accumulated wealth in annuity contracts upon retirement. For a financial market consisting of a risk-free bond and an annuity contract, Yaari [21] demonstrates that if a retiree has no bequest motive, full annuitization is optimal. Indeed, in Yaari's simple model, life annuities don't simply outperform other investment opportunities in expectation but dominate other investments with probability one! In a more recent study, Davidoff, Brown, and Diamond [3] demonstrate under a wide array of more realistic assumptions that retirees should annuitize a substantial proportion of their savings.

In spite of the consensus view held by academic researchers, participation in the private annuity market in the United States remains relatively low . A number of commentators have suggested factors that might lead retirees to annuitize less than academic models propose. An incomplete list of these includes:

- (i) health shocks,
- (ii) social security crowding out private annuitization,
- (iii) premiums above actuarially fair prices,
- (iv) high-profile failures of insurance companies,
- (v) irrevocability,
- (vi) imperfect information, and
- (vii) erosion of purchasing power due to inflation.

For a more complete list, see [1].

The risk of health shocks, for instance, is a risk that retirees face when purchasing annuities that is not included in standard models. In the typical life-cycle approach to

modeling a retiree's uncertain time of death, survival functions with deterministic hazard rates are commonly used. Thus, in such models, although the time of death is a random variable, the known survival distribution is fully incorporated into the actuarial pricing of the annuity contract so that the depreciation in value of the annuity contract is continuous and deterministic throughout the life of the retiree. In reality, however, a retiree's survival probability changes stochastically upon certain health events, such as diagnosis of cancer or other serious illnesses. In that case, the value of the annuity to the retiree can realize a significant downward jump. Moreover, and perhaps more importantly since the re-sell value of an annuity is not usually a major factor, in the event of a health shock, a retiree may encounter a sudden, increased need for liquid wealth in order to meet medical expenses. However, an annuity contract is often either irrevocable or equipped with steep cancellation fees. A retiree's inability to liquidate an annuity without substantial loss is a financial risk known as *liquidity risk*.

In addition to liquidity risk, retirees also face counterparty, or credit, risk when purchasing an annuity. Credit risk is the risk annuitants bear due to the financial soundness, or lack thereof, of the institution that provides the annuity contract, resulting in the potential for annuity default. Not only do drops in the credit rating of the annuity seller result in annuity value depreciation, but more importantly, a leap to insolvency can result in a significant loss of invested capital. Given the unattractiveness to retirees of re-entering the labor market should a capital loss necessitate it, credit risk could impact the annuitization decision. Moreover, historical failures of insurance companies as well as recent high-profile failures of large financial institutions generally, may tend to heighten awareness of and aversion to the counterparty risk inherent in annuity contracts.

Although annuity default is a real and important risk assumed by retirees, institutional safety nets exist to help alleviate the effects of default. First, upon default, an insurance company's assets are seized by authorities who then act to liquidate the company's assets and divide the proceeds as fairly as possible among the insurance company's obligees. Second,

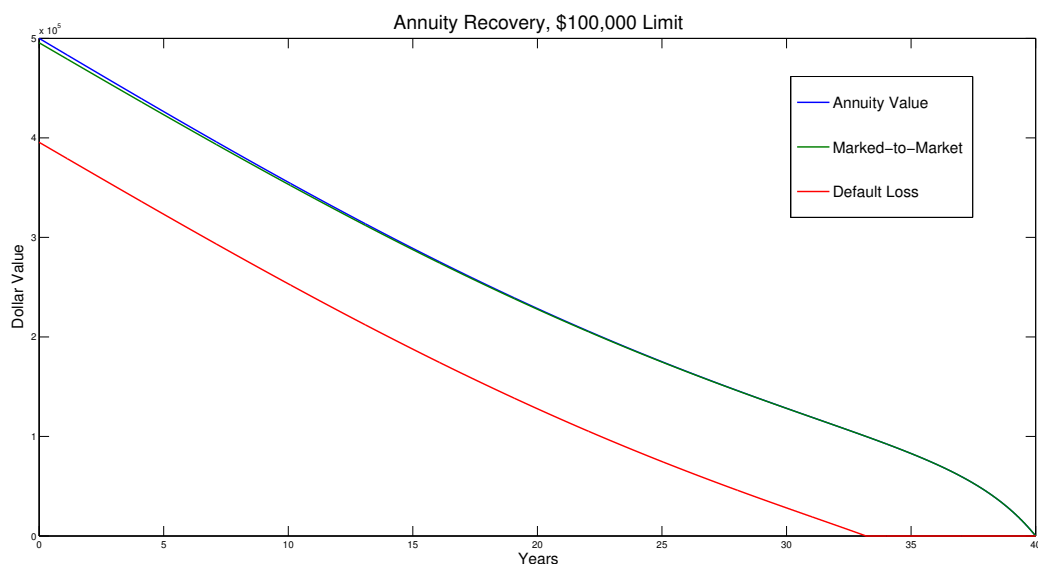


Figure 4.1: Annuity Value, Recovery, and Loss at Default

individual states each sponsor life and health insurance guaranty programs that help ensure annuity contracts, although the guaranty is limited — [1] includes a table that shows state-by-state coverage limits. The most common caps on state coverage are around \$100-200,000. On a national level, the National Organization of Life and Health Guaranty Associations (NOLHGA) works together with the state guaranty associations to provide oversight and assistance in case of the failure of an insurance company that operates in several states. Through litigation and government sponsored programs, retirees may recuperate a significant proportion of their losses in case of annuity default.

As an example of default with recovery, suppose an average 65-year old Utahn entering retirement purchases a simple life annuity from an insurance company for \$500,000. In return, the insurance company pays, continuously, at a constant annual rate of about \$31,000 until either the retiree dies or reaches 105 years of age. Further suppose that at any given time, the probability that the insurance company will default within a year, given that is presently operational, is about 0.1%. In Figure 4.1, we plot the nominal value of the contract, the marked-to-market value given the possibility of default and the recovery limit of \$200,000, and the retiree’s loss on default.

The small discrepancy between the nominal value of the contract and the marked-to-market value owes to the fact that the insurance company ignores its own default probability in the price. Notice that a little after year 33, or when the retiree reaches age 88, the value of the remaining annuity income reaches \$100,000 so that the contract is henceforth fully ensured. Thus, the loss from default is zero for the final 17 years of the contract, should the retiree survive.

It seems reasonable that a guaranty from a state or federal government that partially ensures annuity contracts could help eliminate some of a retiree's aversion to institutional credit risk and thereby attenuate the affect of credit risk on the retiree's optimal annuitization decision. However, as discussed in [1], due to various regulations, annuity sellers cannot advertise the existence of guaranty programs, so retirees do not necessarily account for recovery in their decision-making.

Below, we employ the model developed in the previous chapter to investigate a retiree's optimal annuity decision. There are, of course, a number of variations on the annuity contract available to retirees, but we restrict our attention to simple annuities, in which a lump-sum is paid at retirement in return for a continuously paid, constant rate of income until either the retiree dies or reaches some T years of age. In the next section, we specify parameters for the model, relying on typical financial parameters and actuarial tables.

After specifying the parameters, we study a 65-year-old retiree's optimal annuity decision in two main cases. First, we study a retiree's optimal level of annuitization, for varying insurer credit risk, that results in a total loss to the retiree upon default; that is, the retiree believes that none of a defaulted annuity can be recovered. This case will be called the *no recovery* case, and it serves to provide a lower bound for optimal annuitization by considering the worse case scenario.

Following the worse case of no recovery, we study a retiree's optimal annuity decision for varying credit risk using typical recovery limits of \$100,000–\$300,000. This case is called the *partial recovery* case.

4.1 MODEL CALIBRATION

As a first step in applying the investment consumption model of Chapter 3 to optimal annuitization, it is necessary to determine “realistic” parameters for the financial market, the utilities, and the random time probabilities. We draw heavily from [1].

4.1.1 Financial Market Parameters. Recent yields on US Treasury Inflation-protected Securities (TIPS), an inflation-adjusted, risk-free asset, have been well below 1%. Historic inflation-adjusted interest rates, however, are typically higher, and the annuity puzzle has persisted through several regime changes in the risk-free rate. Therefore, we take $r = .02$, which is a more historical value. In finance, the market mean rate of return is implied by the equity risk premium, $\mu - r$. Recent estimations of the market risk premium shows a premium of around 0.03, so we take $\mu = 0.05$. Finally, a standard metric for calibrating market volatility, the Chicago Board of Exchange Volatility Index, was around 20 in January 2013, or in other words, $\sigma = 0.2$.

4.1.2 CRRA Utility Parameters. As discussed in Section 2.4, the standard choice of utility for inter-temporal consumption and bequest is the CRRA utility,

$$U(t, z) = \exp(-\rho t) \frac{z^{1-\gamma}}{1-\gamma}.$$

The parameter γ represents a retiree’s risk aversion, with larger γ implying greater aversion to risk. The parameter ρ is a retiree’s subjective time-preference, or patience. While some authors have commented on the virtues of the CRRA utility while others have noted its flaws, we choose it mainly because it is the most widely used utility and because it is tractable. Common choices for γ and ρ are 4 and 0.03, respectively.

4.1.3 Mortality Probability. A common mortality model in actuarial practice is the Gompertz-Makeham model, which according to [1], takes as a functional form for the hazard

rate

$$\lambda_D(t) = \frac{1}{d} \exp\left(\frac{R + t - c}{d}\right),$$

where R is the retirement age and c and d are parameters.

Babbel and Merrill calibrated the Gompertz-Makeham model against actuarial tables, finding a good fit with $c = 87.98$ and $d = 11.19$.

4.1.4 Default Probability. The Municipal Bond Fairness Act of 2008 [12] included historical corporate bond default rates by credit rating, over different agency ratings. The data included in that report suggest a historic default rate of about 0.5% for AAA-rated bonds, which is Standard & Poor's highest credit rating, and 1% for AA-rated companies. Since most life insurance companies possess high credit ratings, we investigate default hazard rates from 0 to 3%, i.e., $\lambda_D \in [0, 0.04]$.

4.1.5 Summary of Parameters. Table 4.1 summarizes the financial market and utility parameters chosen for the numerical study.

Table 4.1: Summary of Market and Risk Parameters

r	μ	σ	γ	ρ	λ_D	T
.02	.05	.20	4	.03	[0, 0.04]	105

4.2 COMPUTATION AND DISCUSSION

We apply the calibrated model to the annuity puzzle as follows. Suppose a retiree has $x = \$500,000$ in savings, and wishes to annuitize A dollars. The insurance company prices the annuity actuarially fairly, i.e., $\eta_M = \lambda_M$, except that it does not include its own credit risk when computing the annuity price. Ignoring transaction costs, the price of an annuity paying rate α is then

$$A(\alpha) = \alpha \int_0^T \exp\left(-\int_t^s (r + \eta_M(u)) du\right) ds. \quad (4.1)$$

The maximum income rate that can be purchased is

$$\alpha_{\max} = \frac{x}{\int_0^T \exp\left(-\int_t^s (r + \eta_M(u)) du\right) ds}. \quad (4.2)$$

Then for each $\alpha \in [0, \alpha_{\max}]$, we approximate $V(x; \alpha)$ numerically as described in Section 3.6.5. After $V(x; \alpha)$ is known, the optimal annuity rate is

$$\alpha^* = \arg \max_{\alpha} \{V(x; \alpha)\},$$

and the optimal level of annuitization as percentage of wealth is

$$\theta = \frac{A(\alpha^*)}{x}$$

4.2.1 Computational Details. The numerical method was implemented in Matlab R2012b. The machines on which it was implemented were 2012 Mac Pro computers with 2×2.4 GHz 6-Core Intel Xeon processors and 64 GB 1333 MHz memory cards. For a given default rate, the problem was parallelized in the annuity rate α over the 12 processors. With $\delta = 0.01$, $h = 0.02$, and a somewhat coarse α grid step of 1500, or 8% of α_{\max} , computing time for a fixed default rate was around 12 hours.

4.2.2 Numerical Results. Table 4.2 shows the optimal fraction of initial wealth annuitized, θ , without guaranty as well as for guaranty limits of \$100,000, \$200,000, and \$300,000.

Table 4.2: Optimal Annuitization as Percentage of Wealth, by Recovery Limit

	$\lambda_D = 0$.01	.02	.03	.04
No Recovery	1	0.34	0.24	0.19	0.15
100K Limit	1	0.53	0.43	0.39	0.34
200K Limit	1	0.68	0.58	0.53	0.47
300K Limit	1	0.82	0.73	0.63	0.63

Table 4.2 shows that, as expected, full annuitization is optimal when there is no chance for default (Column 1). When there is no guaranty, θ drops to 0.34 for a hazard rate of

just 0.01. Hence, even for an insurance company with the highest credit rating it can be optimal for a retiree to annuitize as little as 34% of their savings. Table 4.2 also shows that, as expected, the guaranty attenuates somewhat the affect of default on the optimal recovery rate.

Numerical studies also demonstrate that just as for life insurance in Richard's model, the optimal purchase of default insurance is a function of the ratio of illiquid wealth to liquid wealth. Take for example, the no recovery case $\lambda_D = .01$, and suppose that the retiree annuitizes the optimal 34% of total savings, and suppose that the retiree's overall wealth doesn't significantly decrease from the initial \$500,000 over the course of retirement. Then we find that the optimal proportion of wealth, $\frac{p_D^*(t)}{X(t)}$, spent of credit default swaps is virtually nothing. Alternatively, consider the same retiree but for whom total wealth plummets due to repeated losses on investments. In the latter case, the retiree becomes willing to spend an increasing amount of capital to insure the annuity contract. Thus, at retirement, the retiree optimal level of annuitization is one for which the costs of insuring against default is zero except under unlikely, unfavorable outcomes.

Similarly, we find that life insurance purchases increase with increasing size of remaining annuity income relative to the total wealth, suggesting that retirees who purchase an annuity only to see their remaining savings disappear find it attractive to purchase insurance to provide an adequate bequest in case they do not live long enough to recuperate the annuitized wealth.

5. CONCLUSION

I have had my results for a long time: but I do not yet know how I am to arrive at them.

– Carl Friedrich Gauss

In this chapter, we summarize the key contributions of this dissertation, including the investment consumption model developed in Chapter 3 and its application to the optimal annuity problem in Chapter 4. Following that, we outline several directions for future research, including a discussion of the convergence problem as well as extensions of our investment consumption model and applications.

5.1 SUMMARY OF KEY RESULTS

The first significant contribution of this dissertation is the development of an investment consumption model that permits the study of optimal investment behavior in the presence of a risky, illiquid income stream. This was accomplished by incorporating a random time of default into a standard investment consumption model. However, the sudden change in the investor's wealth dynamics at the time of default presents a significant challenge to the analysis of standard investment consumption models, as has been noted in previous attempts to study such a model.

We have overcome the difficulty of dealing with sudden change in dynamics, firstly, by including a credit default swap to hedge help the investor against default. Mathematically, the credit default swap circumvents the problems associated with large jumps in the investor's wealth at the time of default by permitting a smooth transition between the investor's wealth before and after default. As a modeling tool, moreover, the default insurance is very useful because it serves as a proxy for the investor's disutility for default. Capital spent purchasing default insurance could have instead either been consumed or invested towards the terminal wealth. Moreover, borrowing against the illiquid income stream is more expensive when the

income may default, which further reduces the utility for consumption and terminal wealth.

In addition to serving as a proxy for disutility towards default, the credit default swap enabled us to define a pre-default cost functional and a separate post-default cost functional, coupled through the post-default value function, or the default utility. Because each of the cost functionals was defined over a deterministic interval with decision processes adapted only to the filtration generated by Brownian motion, we were able to apply dynamic programming techniques to obtain optimal pre-default and post-default decisions. The pre-default and post-default decisions were then “glued” together, using measurable selection, at the time of default.

The second way by which we have overcome the usual difficulties encountered with the application of numerical dynamic programming techniques, and specifically, the Markov chain approximation, is by applying a logarithmic transformation on the total wealth. As illustrated in Section 3.6.5, the logarithmic transformation of the wealth variable and control variables simultaneously removed explicit dependence on the wealth variable from the derivative coefficients in the discretized HJB equation, allowing for a probabilistic interpretation and removed computational difficulties of working near wealth boundary.

A particularly nice feature of our model is its ability to cope with different income streams that carry distinct rates of default. This, for instance, permitted investigation of an income stream with a limited guaranty. In fact, a straightforward extension of our model allows for separate sources of income that may default at different times and with different frequency.

The second significant contribution of this dissertation is the application of the investment consumption model to the optimal annuity problem in the presence of credit risk. We have demonstrated that the potential for annuity default leads to costs sufficient to drive the retiree’s optimal annuitization level far below what conventional economic theory would suggest. Relatively low annuitization remains optimal even when a retiree believes that guaranty programs will help replace a portion of the lost annuity. Hence, while there are certainly a number of important factors that contribute to the so-called annuity puzzle,

credit risk certainly stands out as a prominent player.

5.2 FUTURE RESEARCH

The investment model developed throughout this dissertation leads naturally to several other problems, which are detailed below.

5.2.1 MCA Convergence. In his dissertation as well as subsequent papers on MCA with a logarithmic transformation and its application to Richard's model, Ye [13], claims that the standard verification of local consistency is sufficient to establish the convergence of MCA, although he does not even verify the local consistency. In [23], Ye demonstrates the local consistency necessary for convergence on a model much simpler than Richard's. However, the theorems relating to local consistency that Ye claims as justification for the convergence of his approach are not satisfied in their application to Richard's model or even to the simpler model of [23]. Every proof of convergence of MCA of which we are aware, including the ones cited directly by Ye, require boundedness of the utility functions. However, the CRRA utilities employed by Ye as well as in this paper are manifestly *un*-bounded.

The lack of a proof of convergence for MCA in the case of unbounded utility functions represents an important gap in the theory, since in any useful application, the utility is either a CRRA or other unbounded utility function. We seek a proof of convergence for MCA for CRRA utilities, which would establish convergence in Ye's application of MCA as well as our own.

Not only does the theory not apply to Ye's application of MCA to Richard's model as it does not apply to ours, but the sequence of discrete time Markov chains suggested the finite difference method are not locally consistent with the controlled diffusion of total wealth!

The locally consistent DTMC paradigm developed by Kushner is extremely useful, but it is just a sufficient condition, not a necessary one. In [5], the authors prove the convergence of a Markov chain approximation scheme applied to a singular control problem, but their

proof does not rely on local consistency. Instead, they rely on the existence and uniqueness of a viscosity solution for the continuous-time HJB equation, and they then show that the solutions of a sequence of approximating, discrete-time control problems converge to the viscosity solution. In spite of the fact that local consistency cannot be applied to Ye's MCA approach, which our approach is based on, Ye provided convincing evidence that his method does converge to the solution of the control problem by applying his method to control problems with well-known explicit solutions. Ye showed numerically that by choosing sufficiently small discretization parameters, the error between the closed-form solutions and the numerical approximations is very small.

Problem. Provide a convergence proof for the MCA technique used in obtaining solutions in Ye's model as well as, of course, our own.

5.2.2 Regime-switching Models. In order to more completely model important market variables, basic investment consumption models are sometimes extended to include parameters that switch between a finite set of states as a Markov chain. Such models are often called *regime-switching* models in financial and economic applications. The theory of diffusions has accordingly been extended to permit a drift and volatility that jumps between finitely many states (see [20]), and studies of investment consumption with regime-switching parameters have been made, such as in [18]. A general version of Ito's formula permits the application of dynamic programming techniques to regime-switching models, and even the Markov chain approximation can be modified to handle regime-switching, as discussed in [17]. Two interesting applications that can be studied by combining our model with regime-switching include the optimal annuity decision with health shocks, and callable bonds.

5.2.3 Optimal Annuities with Health Shocks. As discussed in Chapter 4, one of the factors believed to significantly affect a retiree's annuity decision is the possibility of a health shock, or a stochastic change in a retiree's survival outlook. A retiree with a mortality hazard rate that fluctuates stochastically in time faces not only a possible sudden depreciation in

the value of the annuity contract, viewed as an asset, but perhaps more importantly, can also face serious liquidity risk associated with the sudden need for cash to meet a jump in health costs.

The effect of health shocks on investment and consumption have been studied somewhat. For instance, in [6], Huang et. al. study a simple consumption problem for a retiree with an uncertain time of death using a stochastic force of mortality. However, they modeled the force of mortality using a diffusion. It is certainly not clear that a diffusion is a good paradigm for a force of mortality, as changes in health tend to be less constant but more sudden. Regime-switching provides a more natural framework for changes in health, and the augmentation of our model with a regime-switching health process permits us to study the following interesting problem:

Problem. Find the optimal annuity decision for a retiree who faces a regime-switching force of mortality, both when the annuity is not defaultable and when it is defaultable.

5.2.4 Callable Bonds. One of the ways in which corporations finance activities is through the issuance of debt in the form of corporate bonds. Typically, a corporate bond provides its the obligee with a regular, constant rate of income until a fixed expiration date, at which time the initial investment is repaid. Naturally, bonds are subject to counterparty risk for which they receive an interest premium. Some bonds, known as callable, or redeemable, bonds, carry a provision that allow for the obligor to cancel the bond before the nominal expiration in exchange for par value plus a premium. The fixed income nature of our model matches the cash flows associated with a callable bond, and the random times of our model can be interpreted as a random time of bond default and a random time of redemption. Coupling this basic set-up with regime switching allows us to follow stochastic improvements and declines in the corporate credit rating. The application of our model with regime-switching leads to an interesting problem:

Problem. Given a financial market consisting of a risk-free bond and a corporate bond,

determine the an individual or institutional investor's optimal division of wealth between the two bonds.

The financial market could also be expanded to include a stock.

APPENDIX A. PROOFS AND DERIVATIONS

A.1 DERIVATION OF (2.39)

We apply dynamic programming to by studying control policies in (3.24) over a short time interval, h ,

$$\begin{aligned}
 V(t, x) &= \sup_{c, \pi, \zeta_M \in \mathcal{U}} \mathbb{E}^{t, x} \left[\int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}_M(s | t) ds \right. \\
 &\quad + \int_{t+h}^T (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}_M(s | t) ds \\
 &\quad \left. + B(T, \zeta_M(T)) \bar{F}_M(T | t) \right] \\
 &= \sup_{c, \pi, \zeta_M \in \mathcal{U}} \mathbb{E}^{t, x} \left[\int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}_M(s | t) ds \right. \\
 &\quad + \bar{F}_M(t+h | t) \left(\int_{t+h}^T (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}(s | t+h) ds \right. \\
 &\quad \left. \left. + B(T, \zeta_M(T)) \bar{F}_M(T | t+h) \right) \right] \tag{A.1}
 \end{aligned}$$

From iterated conditioning and the Markov property for X ,

$$\begin{aligned}
 V(t, x) &= \sup_{c, \pi, \zeta_M \in \mathcal{U}} \mathbb{E}^{t, x} \left[\int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}_M(s | t) ds \right. \\
 &\quad + \bar{F}_M(t+h | t) \mathbb{E}^{t+h, X(t+h)} \left[\int_{t+h}^T (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}(s | t+h) ds \right. \\
 &\quad \left. \left. + B(T, \zeta_M(T)) \bar{F}_M(T | t+h) \right) \right] \\
 &= \sup_{c, \pi, \zeta_M \in \mathcal{U}} \mathbb{E}^{t, x} \left[\int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}_M(s | t) ds \right. \\
 &\quad \left. + \bar{F}_M(t+h | t) V(t+h, X(t+h)) \right]. \tag{A.2}
 \end{aligned}$$

By Taylor's theorem,

$$\bar{F}_M(t+h | t) = \exp \left(- \int_t^{t+h} \lambda_M(u) du \right) = 1 - \lambda_M(t)h + o(h^2), \tag{A.3}$$

and by Ito's lemma,

$$\begin{aligned} V(t+h, X(t+h)) &= V(x) + \int_t^{t+h} (V_t(s, X(s)) + b(s, X(s))V_x(s, X(s))) ds \\ &\quad + \int_t^{t+h} \sigma\pi(s)V_x(s, X(s))dW(s) + \int_t^{t+h} \frac{1}{2}\sigma^2\pi^2V_{xx}(s, X(s)) ds, \end{aligned} \quad (\text{A.4})$$

where

$$b(t, x) = -c(t) - \eta_M(t)\zeta_M(t) + (r + \eta_M(t))x + (\mu - r)\pi(t).$$

Substituting (A.3) and (A.4) into (A.2), we obtain

$$\begin{aligned} V(t, x) &= \sup_{c, \pi, \zeta_M \in \mathcal{U}} \mathbb{E}^{t, x} \left[\int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))\bar{F}_M(s | t)) ds \right. \\ &\quad + (1 - \lambda_M(t)h + o(h^2)) \left(V(x) + \int_t^{t+h} (V_t(s, X(s)) + b(s, X(s))V_x(s, X(s))) ds \right. \\ &\quad \left. \left. + \int_t^{t+h} \sigma\pi(s)V_x(s, X(s)) dW(s) + \int_t^{t+h} \frac{1}{2}\sigma^2\pi^2V_{xx}(s, X(s)) ds \right) \right]. \end{aligned} \quad (\text{A.5})$$

Multiplying each side of (A.5) by $1/h$, taking the limit as $h \downarrow 0$, and rearranging, we arrive at the HJB equation for V ,

$$\begin{aligned} 0 &= \sup_{c, \pi, \zeta_M \in \mathcal{U}} \{ -\lambda_M(t)V(t, x) + V_t(t, x) + b(t, x)V_x(t, x) \\ &\quad + \pi^2(t)\sigma^2V_{xx}(t, x) + U(t, c) + \lambda_M(t)B(t, x) \}, \quad 0 \leq t < T. \end{aligned} \quad (\text{A.6})$$

A.2 PROOF OF LEMMA 3.28

We assume the strong Markov property for X at the \mathbb{F} -stopping time, τ_D , i.e.,

$$\mathbb{E}^{t, x} [f(X(\tau_D + s)) | \mathcal{F}_{\tau_D}] = \mathbb{E}^{\tau_D, X(\tau_D)} [f(X(s))], \text{ for all Borel-measurable } f. \quad (\text{A.7})$$

Sketch of Proof. Since the expectation in (3.24) is conditioned on $\tau_D > t$, iterated condition-

ing gives,

$$\begin{aligned} & \mathbb{E}^{t,x} \left[\left(\int_{\tau_D}^{\tau_M \wedge T} U(s, c(s)) + B(\tau_M, \zeta_M(\tau_M)) \right) \mathbf{1}_{\tau_D > \tau_M \wedge T} \right] \\ &= \mathbb{E}^{t,x} \left[\mathbb{E} \left[\left(\int_{\tau_D}^{\tau_M \wedge T} U(s, c(s)) + B(\tau_M, \zeta_M(\tau_M)) \right) \mathbf{1}_{\tau_D > \tau_M \wedge T} \mid \mathcal{F}_{\tau_D} \right] \right]. \end{aligned} \quad (\text{A.8})$$

By definition,

$$\mathcal{F}_{\tau_D} = \{A \in \mathbb{F} : A \cap \{\tau_D < t\} \in \mathcal{F}_t\},$$

and for each t , $\{\tau_D < \tau_M \wedge T\} \cap \{\tau_D < t\} = \cup_{s \leq t} (\{\tau_M \wedge T \geq s\} \cap \{\tau_D < s\})$. Moreover, $\{\tau_M \wedge T \geq s\} \cap \{\tau_D < s\} \in \mathcal{F}_s \subseteq \mathcal{F}_t$ by construction, so $\{\tau_D < \tau_M \wedge T\}$ is \mathcal{F}_{τ_D} -measurable.

In particular, it is “known” by τ_D , so it can be factored through the inner expectation in (A.8),

$$\begin{aligned} & \mathbb{E}^{t,x} \left[\mathbb{E} \left[\left(\int_{\tau_D}^{\tau_M \wedge T} U(s, c(s)) + B(\tau_M, \zeta_M(\tau_M)) \right) \mathbf{1}_{\tau_D > \tau_M \wedge T} \mid \mathcal{F}_{\tau_D} \right] \right] \\ &= \mathbb{E}^{t,x} \left[\mathbb{E} \left[\left(\int_{\tau_D}^{\tau_M \wedge T} U(s, c(s)) + B(\tau_M, \zeta_M(\tau_M)) \right) \mid \mathcal{F}_{\tau_D} \right] \mathbf{1}_{\tau_D > \tau_M \wedge T} \right]. \end{aligned} \quad (\text{A.9})$$

Next, we apply the strong Markov property (A.7),

$$\begin{aligned} & \mathbb{E}^{t,x} \left[\mathbb{E} \left[\left(\int_{\tau_D}^{\tau_M \wedge T} U(s, c(s)) + B(\tau_M, \zeta_M(\tau_M)) \right) \mid \mathcal{F}_{\tau_D} \right] \mathbf{1}_{\tau_D > \tau_M \wedge T} \right] \\ &= \mathbb{E}^{t,x} \left[\mathbb{E}^{\tau_D, X(\tau_D)} \left[\left(\int_{\tau_D}^{\tau_M \wedge T} U(s, c(s)) + B(\tau_M, \zeta_M(\tau_M)) \right) \right] \mathbf{1}_{\tau_D > \tau_M \wedge T} \right] \end{aligned} \quad (\text{A.10})$$

Substituting (A.10) into (A.8) produces

$$\begin{aligned} V(t, x) &= \sup_{c, \pi, \zeta_M, \zeta_D} \mathbb{E}^{t,x} \left[\int_t^{\tau_M \wedge T \wedge \tau_D} U(s, c(s)) ds + B(\tau_M, \zeta_M(\tau_M)) \mathbf{1}_{\tau_D \geq \tau_M \wedge T} \right. \\ &\quad \left. + \mathbb{E}^{\tau_D, X(\tau_D)} \left[\left(\int_{\tau_D}^{\tau_M \wedge T} U(s, c(s)) + B(\tau_M, \zeta_M(\tau_M)) \right) \right] \mathbf{1}_{\tau_D > \tau_M \wedge T} \right]. \end{aligned} \quad (\text{A.11})$$

Since

$$\mathbb{E}^{\tau_D, X(\tau_D)} \left[\left(\int_{\tau_D}^{\tau_M \wedge T} U(s, c(s)) + B(\tau_M, \zeta_M(\tau_M)) \right) \right] \leq D(\tau_D, X(\tau_D)), \quad (\text{A.12})$$

$$\begin{aligned} V(t, x) \leq \sup_{c, \pi, \zeta_M, \zeta_D} \mathbb{E}^{t, x} & \left[\int_t^{\tau_M \wedge T \wedge \tau_D} U(s, c(s)) ds + B(\tau_M, \zeta_M(\tau_M)) \mathbf{1}_{\tau_D \geq \tau_M \wedge T} \right. \\ & \left. + D(\tau_D, X(\tau_D)) \mathbf{1}_{\tau_D < \tau_M \wedge T} \right]. \end{aligned} \quad (\text{A.13})$$

On the other hand, let $(c, \pi, \zeta_M, \zeta_D)$ be an arbitrary collection of decisions. By definition, for every $\epsilon > 0$ and for each (t, x) , there exists a collection $(c^\epsilon, \pi^\epsilon, \zeta_M^\epsilon)$ such that under $(c^\epsilon, \pi^\epsilon, \zeta_M^\epsilon)$,

$$D(t, x) - \epsilon \leq \mathbb{E}^{t, x} \left[\int_t^{\wedge T \wedge \tau_D} U(s, c(s)) + B(\tau_M, \zeta_M(\tau_M)) \right]. \quad (\text{A.14})$$

Let

$$(c', \pi', \zeta'_M, \zeta'_D) = \begin{cases} (c, \pi, \zeta_M, \zeta_D) & \text{if } t \leq \tau_D \\ (c^\epsilon, \pi^\epsilon, \zeta_M^\epsilon, 0) & \text{if } t > \tau_D. \end{cases}$$

That $(c', \pi', \zeta'_M, \zeta'_D)$ is progressively measurable is non-trivial but can be shown by use of the measurable selection theorem. Under the controls, $(c', \pi', \zeta'_M, \zeta'_D)$, (A.13) and (A.14) combine to give

$$\begin{aligned} V(t, x) \leq \sup_{c, \pi, \zeta_M, \zeta_D} \mathbb{E}^{t, x} & \left[\int_t^{\tau_M \wedge T \wedge \tau_D} U(s, c(s)) ds + B(\tau_M, \zeta_M(\tau_M)) \mathbf{1}_{\tau_D \geq \tau_M \wedge T} \right. \\ & \left. + D(\tau_D, X(\tau_D)) \mathbf{1}_{\tau_D < \tau_M \wedge T} \right] \leq V(t, x) + \epsilon. \end{aligned} \quad (\text{A.15})$$

Since ϵ is arbitrary, (3.28) follows. □

A.3 DERIVATION OF (3.38)

We apply dynamic programming to the post-default value function by studying optimal control over a short time interval, $[t, t + h]$. From iterated conditioning and the Markov property for X_2 ,

$$\begin{aligned}
D(t, x) &= \sup_{c, \pi, \zeta_M \in \mathcal{U}_2} \mathbb{E}^{t, x} \left[\frac{1}{h} \int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}_M(s | t) ds \right. \\
&\quad \left. + \int_{t+h}^T (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}_M(s | t) ds + B(T, X(T)) \bar{F}_M(T | t) \right] \\
&= \sup_{c, \pi, \zeta_M \in \mathcal{U}_2} \mathbb{E}^{t, x} \left[\frac{1}{h} \int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}_M(s | t) ds \right. \\
&\quad \left. + \bar{F}_M(t + h | t) \left(\int_{t+h}^T (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}_M ds \right. \right. \\
&\quad \left. \left. + B(T, X(T)) \bar{F}_M(T | t + h) \right) \right] \\
&= \sup_{c, \pi, \zeta_M \in \mathcal{U}_2} \mathbb{E}^{t, x} \left[\frac{1}{h} \int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))) \bar{F}_M(s | t) ds \right. \\
&\quad \left. + \bar{F}_M(t + h | t) D(t, X_2(t + h)) \right]. \tag{A.16}
\end{aligned}$$

The survival probability $\bar{F}_M(t + h | t)$ can be expanded using Taylor's theorem,

$$\bar{F}_M(t + h | t) = \exp \left(- \int_t^{t+h} \lambda_M(u) du \right) = 1 - \lambda_M(t)h + o(h^2). \tag{A.17}$$

Assuming D is sufficiently smooth, Ito's lemma applied to D gives

$$\begin{aligned}
D(t + h, X_2(t + h)) &= D(x) + \int_t^{t+h} [D_t(s, X_2(s)) + b_2(s, X(s))D_x(s, X_2(s))] ds \\
&\quad + \int_t^{t+h} \sigma \pi(s) D_x(s, X_2(s)) dW(s) + \int_t^{t+h} \frac{1}{2} \sigma^2 \pi^2 D_{xx}(s, X_2(s)) ds, \tag{A.18}
\end{aligned}$$

where

$$b_2(t, x) = -c(t) - \eta_M(t)\zeta_M(t) + (r + \eta_M(t))x + (\mu - r)\pi(t)$$

Substituting (A.17) and (A.18) into (A.16), we obtain

$$\begin{aligned}
D(t, x) = & \sup_{c, \pi, \zeta_M \in \mathcal{U}_\epsilon} \mathbb{E}^{t, x} \left[\int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s))\bar{F}_M(s | t)) ds \right. \\
& + (1 - \lambda_M(t)h + o(h^2)) \left(V(t, x) + \int_t^{t+h} [D_t(s, X_2(s)) + b_2(s, X(s))D_x(s, X_2(s))] ds \right. \\
& \left. \left. + \int_t^{t+h} \sigma \pi(s)D_x(s, X_2(s))dW(s) + \int_t^{t+h} \frac{1}{2} \sigma^2 \pi^2 D_{xx}(s, X_2(s)) ds \right) \right]. \quad (\text{A.19})
\end{aligned}$$

Multiplying each side of (A.19) by $1/h$ and taking the limit as $h \downarrow 0$, we arrive at the HJB equation for D ,

$$\begin{aligned}
0 = & \sup_{c, \pi, \zeta_M \in \mathcal{U}_\epsilon} \left\{ -\lambda_M(t)D(t, x) + D_t(t, x) + b_2(t, x)D_x(t, x) + \pi^2(t)\sigma^2 D_{xx}(t, x) \right. \\
& \left. + U(t, c(t)) + \lambda_M(t)B(t, \zeta_M(t)) \right\}. \quad (\text{A.20})
\end{aligned}$$

A.4 DERIVATION OF (3.41)

The derivation of the HJB equation for the post-default value function, V , is analogous to the derivation for D given in the preceding section.

$$\begin{aligned}
V(t, x) &= \sup_{c, \pi, \zeta_M, \zeta_D \in \mathcal{U}_1} \mathbb{E}^{t, x} \left[\int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s)) + \lambda_D(s)D(s, \zeta_D(s))) \bar{F}(s | t) ds \right. \\
&\quad + \int_{t+h}^T (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s)) + \lambda_D(s)D(s, \zeta_D(s))) \bar{F}(s | t) ds \\
&\quad \left. + B(T, X(T)) \bar{F}_M(T | t) \right] \\
&= \sup_{c, \pi, \zeta_M, \zeta_D \in \mathcal{U}_1} \mathbb{E}^{t, x} \left[\int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s)) + \lambda_D(s)D(s, \zeta_D(s))) \bar{F}(s | t) ds \right. \\
&\quad + \bar{F}(t+h | t) \left(\int_{t+h}^T (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s)) + \lambda_D(s)D(s, \zeta_D(s))) \bar{F}(s | t) ds \right. \\
&\quad \left. \left. + B(T, X(T)) \bar{F}_M(T | t+h) \right) \right] \\
&= \sup_{c, \pi, \zeta_M, \zeta_D \in \mathcal{U}_1} \mathbb{E}^{t, x} \left[\int_t^{t+h} (U(s, c(s)) + \lambda_M(s)B(s, \zeta_M(s)) + \lambda_D(s)D(s, \zeta_D(s))) \bar{F}(s | t) ds \right. \\
&\quad \left. + \bar{F}(t+h | t) V(t, X_1(t+h)) \right]. \tag{A.21}
\end{aligned}$$

The survival probability $\bar{F}(t+h | t)$ can be expanded using Taylor's theorem,

$$\bar{F}(t+h | t) = \exp \left(- \int_t^{t+h} (\lambda_M(u) + \lambda_D(u)) du \right) = 1 - (\lambda_M(t) + \lambda_D(t))h + o(h^2), \tag{A.22}$$

and assuming V is sufficiently smooth, Ito's lemma applied to V gives

$$\begin{aligned}
V(t+h, X_1(t+h)) &= V(x) + \int_t^{t+h} (V_t(s, X_1(s)) + b_1(s, X(s))D_x(s, X_2(s))) ds \\
&\quad + \int_t^{t+h} \sigma \pi(s) D_x(s, X_2(s)) dW(s) + \int_t^{t+h} \frac{1}{2} \sigma^2 \pi^2 D_{xx}(s, X_2(s)) ds, \tag{A.23}
\end{aligned}$$

where

$$b_1(t, x) = -c(t) - \eta_M(t)\zeta_M(t) - \eta_D(t)\zeta_D(t) + (r + \eta_M(t) + \eta_D(t))x + (\mu - r)\pi(t)$$

Substituting (A.22) and (A.23) into (A.21), multiplying by $1/h$, and letting $h \downarrow 0$, we arrive

at the HJB equation for V ,

$$\begin{aligned} 0 = & -(\lambda_M(t) + \lambda_D(t))V(t, x) + V_t(t, x) + b_1(t, x)D_x(t, x) + \pi^2(t)\sigma^2 D_{xx}(t, x) \\ & + U(t, c(t)) + \lambda_M(t)B(t, \zeta_M(t)) + \lambda_D(t)D(t, \zeta_D(t)). \end{aligned} \quad (\text{A.24})$$

BIBLIOGRAPHY

- [1] D. F. Babbel, J. Humpherys, C. B. Merrill, and J. B. Taylor. Rational decumulation. *Working Paper*, 2013.
- [2] Zvi Bodie. Thoughts on the future: life-cycle investing in theory and practice. *Financ. Anal. J.*, 59(1):24–29, Jan.-Feb. 2003.
- [3] T. Davidoff, J. R. Brown, and P. A. Diamond. Annuities and individual welfare. *Am. Econ. Review*, 95(5):1573–1590, Dec. 2005.
- [4] H. He and H. F. Pagès. Labor income, borrowing constraints, and equilibrium asset prices. *Econ. Theory*, 3(4):663–696, Oct. 1993.
- [5] A. Hindy, C. Huang, and S. H. Zhu. Numerical analysis of a free-boundary singular control problem in financial economics. *J. Econ. Dyn. Control*.
- [6] H. Huang, M.A. Milevsky, and T. S. Salisbury. Optimal retirement consumption with a stochastic force of mortality. *Insur. Math. Econ.*, 51(2):282–291, Sept. 2012.
- [7] J. Hugonnier and D. Kramkov. Optimal investments with random endowments in incomplete markets. *Ann. Appl. Probab.*, 14(2):845–864, 2004.
- [8] I. Karatzas and S. E. Shreve. *Brownian Motion and Stochastic Calculus*. Number 113 in Graduate Texts in Mathematics. Springer, New York, 2 edition, 1991.
- [9] H. Kushner and P. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer, 1992.
- [10] L. M. Lockwood. Bequest motives and the annuity puzzle. *Preprint*, 2012.
- [11] R. C. Merton. Lifetime portfolio selection under uncertainty: the continuous-time case. *Rev. Econ. Stat.*, 51(3):247–257, Aug. 1969.
- [12] U. S. Committee of Financial Services. House report 110-835, Sep. 2008.
- [13] S. R. Pliska and J. Ye. Optimal life insurance purchase and consumption/investment under uncertain lifetime. *J. Bank. Financ.*, 31(5):1307–1319, May 2007.
- [14] S. Richard. Optimal consumption, portfolio and life insurance: rules for an uncertain lived individual in a continuous time model. *J. Financ. Econ.*, 2:187–203, Jun. 1975.
- [15] L. C. G. Rogers and D. Williams. *Diffusions, Markov Processes, and Martingales*, volume 1. Cambridge Press, Cambridge, 2 edition, 2000.
- [16] L. C. G. Rogers and D. Williams. *Diffusions, Markov Processes, and Martingales*, volume 2. Cambridge Press, Cambridge, 2 edition, 2000.
- [17] Q.S. Song, G. Yin, and Z. Zhang. Numerical methods for controlled regime-switching diffusions and regime-switching jump diffusions. *Automatica*, 42(7):1147–1157, July 2006.

- [18] L.R. Sotomayor and A. Cadenillas. Explicit solutions of consumption-investment problems in financial markets with regime switching. *Math. Financ.*, 19(2):251–279, April 2009.
- [19] M. Vellekoop and M. Davis. An optimal investment problem with randomly terminating income. In *Joint 48th IEEE Conference on Decision and Control*, December 2009.
- [20] Mao. X and C. Yuan. *Stochastic Differential Equations with Markovian Switching*. Imperial College Press, London, 2006.
- [21] M. Yaari. Uncertain lifetime, life insurance, and the theory of the consumer. *Rev. Econ. Stud.*, 32:137–150, 1965.
- [22] J. Ye. A numerical method for a continuous-time insurance-consumption-investment model. In *2010 American Control Conference*, June-July 2010.
- [23] J. Ye. A numerical method for a consumption-portfolio problems. In *2011 American Control Conference*, June-July 2011.
- [24] J. Yong and X. Y. Zhou. *Stochastic Controls: Hamiltonian Systems and HJB Equations*. Number 43 in Applications of Mathematics. Springer-Verlag, New York, 1999.