Adaptive Learning in a Continuous-Time Setting: Representative Agent Exercises

Chandler Lester

June 20, 2019

Abstract

We accomplish two distinct, but closely connected, tasks in this paper. First, we look to create a connection between discrete and continuous-time models. This is done by recasting traditional discrete Ramsey models so they are dependent on the increment of time, Δt, and then taking the limit of these models as the size of this increment goes to zero. The resulting models are equivalent to continuous-time Ramsey models. Second, we examine these models in a basic adaptive learning framework. We accomplish this by applying exogenous updating rules to models with a specified stochastic process. After seeing that our misspecified models converge, we then implement a real-time updating rule where our agents update their parameter estimates, for a stochastic process, after observing output for the process itself.

1 Introduction

Macroeconomic modeling in stochastic continuous time has become increasingly popular, as solution methods for optimization problems in this setting have been introduced to economics literature. Solutions to optimization problems in this setting take the same form as fluid dynamic problems common in applied mathematics, and it has taken some time for the mathematical solution techniques to become more common in economics. The appeal for economic modeling in this framework comes from several key features of this setting, not just the availability of simple solution methods. Systems in continuous time can be summarized using sparse matrices that are simple to evaluate and use in calculations, this leads to fast algorithms that use minimal computational time. This is an attractive feature that allows for complex problems with multiple layers of heterogeneity can be easily solved. Solutions in this system also yield more detailed and easily computed probability density functions than discrete time solution methods.

These distributional advantages come from close-ties between the stochastic processes used to summarize the evolution of key variables in these models and their probability density functions. Stochastic processes are often defined according to the distribution of the random variables they represent, and optimization problems that depend on these processes inherit some of this distributional dependency (this will be described in more detail in section 2 of this paper). For instance, Gaussian processes, such as the integral of Brownian motion, have a joint normal distribution for all of the variables they define. Poisson point processes are similarly defined using a Poisson distribution. Using these processes that are defined by continuous probability density functions allows researchers to closely inspect the evolution of the distribution of variables, such as wealth, with little computational burden.

Evaluating these distributions in discrete time is more difficult since probability density functions in this setting are often point masses that truncate the tail-ends of the distribution. Going forward, these discrete methods are going to become less favorable as policy becomes more distribution-oriented. Already, the distribution of wealth and assets is becoming a popular topic when it comes to policy goals. Using the traditional discrete methods, central banks and other policymakers will be unable to properly evaluate the effects of their potential actions on distributions of wealth or assets. Since continuous-time modeling has distributional and computational advantages, this modeling framework will become more attractive and modifying modeling techniques for continuous-time models will be important. In this paper, we will take the first steps in examining adaptive learning methods in a stochastic continuous-time framework.

Many macroeconomic models in both continuous and discrete time depend on agents' expectations. Thus far, continuous-time modeling has depended solely on rational expectations. Using rational expectations limits the model by creating strong assumptions about the agents' information set; rational expectations imply that the agent knows the correct underlying model and that they will respond optimally to the actions of others. These assumptions are unlikely to hold in the real world, as agents may not correctly specify a forecast for the model and they may not understand the actions of others. Therefore, using a different form of expectations that allows for agents' to make mistakes maybe closer to reality. This motivates the use of adaptive learning, a technique that allows for agents to misspecify models and to update their misspecification once they gain more information.

Currently, adaptive learning has been widely implemented in discrete-time modeling; however, it has not been used in continuous-time models. There are two main reasons for this, most economists still use discrete-time models and learning is more difficult to intuit in continuous-time. As continuous-time modeling becomes more popular we will want to able to utilize a powerful tool, like adaptive learning, in this setting. The main goals of this paper are to make continuous-time modeling seem more intuitive and less niche to economists and to intuitively implement basic adaptive learning techniques in continuous time.

Sections 2-4 of this paper map out continuous techniques and literature to make these methods more tractable to economists that focus on discrete modeling. Section 2 gives some mathematical background so that the terminology and motivations of continuous-time literature make sense to the reader. The next section provides a literature review that spans a large portion of the economics continuous-time literature and provides more background on adaptive learning. Despite not being widely popular, continuous-time literature spans several decades, has many significant contributions, and includes a large number of papers by notable economists. The fourth section of this paper explores the mathematical relationship between a variety of discrete and continuous-time models, this section should provide a clear link between these models and make continuous-time modeling more intuitive to those who use discrete-time models.

Section 5 begins the task of implementing adaptive learning techniques in continuous-time. A key part of this section is the methodology for finding steady-state solutions in continuous-time. Although the solution methods for stochastic continuous time models currently used in economics have only recently been introduced to the literature, interest in this class of models has existed in the field for a long time. Exploration of the Ramsey model in stochastic continuous time has been previously studied, notably by Merton (1975a), Mirrlees (1966), and Mirman (1973). There are several different methods for implementing a stochastic process in this modeling framework. Some such as Merton (1975a) have introduced a stochastic process for capital accumulation. Others such as Achdou *et al.* (2014) have used stochastic processes to model productivity. In this paper, we will look at modeling the changes in technological progress and capital as stochastic processes.

This implementation is more intuitive for several different reasons. First, capital accumulation in part depends on technological progress thus if technological progress can change according to this type of process capital accumulation with inherently depends on this process as well. Additionally, technological progress is a variable that, in the real world, often seems to constantly change and improve. Therefore, it is reasonable to assume that other variables like capital stock evolve continuously as they depend on variables we may model continuously, such as technological progress. We can observe technological progress growing over time, so agents are likely to forecast a positive mean and an upward trend. In practice, though we often are unsure of what sectors or improvements will happen over time, and technological progress is almost constantly evolving. Technological progress is something that most believe is constantly improving, because of open-source software and near-constant technological improvements in modern productivity.

Before further discussing the work in this paper, it worth reiterating the benefits that come from continuous solution methods. Continuous time models have unique solutions that can be found using a simple and portable algorithm, and these models only need a few weak boundary conditions to obtain unique solutions. Additionally, these solution methods are computationally faster than discrete methods. This means solving complex economic models with heterogeneity can be done with fewer boundary conditions and in less time. A simple description of an algorithm to solve for a steady state solution in this setting is as follows. First, we discretize over the state spaces in our model. This allows us to maintain the continuous time setting while giving us discrete state spaces to use in a finite difference algorithm. We then implement a finite difference scheme until we get a stable steady state estimate of our value function. Despite the discretization, this solution method is different and faster than most discrete methods. Because in this setting, we can summarize the evolution of our system in large sparse matrices.

We can then take advantage of this discretization to implement traditional discrete learning algorithms in continuous-time. The main difference will be the agent's observation over a given time period. When altering adapting learning algorithms for continuous-time it is tempting to simply use discrete-methods, since the solution methods for continuous-time problems are discretized. However, this discretization is only over state-spaces so, we must be careful to maintain continuity in our time-dimension. This will be important in section 5 when we examine adaptive learning methods in stochastic continuous time models.

In this paper we work to accomplish this through two different methods, one method uses supposes that an agent uses a misspecified process to solve for their steady state and then at discrete time periods gain more information and resolves the continuous model. The other method supposes that an agent uses ordinary least squares to create a forecast of model parameters and then at updates this forecast, using recursive least squares, over intervals of time. The first method demonstrates that continuous-models respond in a predictable manner when presented with misspecification and an exogenous updating rule, and the second provides an intuitive method for adapting learning techniques to continuous models. As we proceed with learning in continuous-time, it will be important to picture our time periods as disjoint intervals of time. Thus, in our forecasting model, each forecasting period contains several observations from our continuous stochastic process. In future work, this could be a key feature of continuous-time learning.

In sections below, we develop two key results that serve as a primer to learning in continuous-time. First, continuous-time models and discrete-time models are fairly comparable mathematically. This can be seen in section 4, which derives discrete models with an unknown time step (Δt) and then limits these models to their continuous-time counterparts. Our second result is that basic models in this setting respond in an expected fashion to new information, through an exogenous and more discrete updating rule and a more continuous forecasting method. Together these results reveal that further studies on adaptive learning in continuous-time may be promising.

The rest of the paper precedes as follows. The next section gives a basic mathematical background for modeling in this framework. Section three discusses the literature relevant to stochastic continuous time modeling and adaptive learning techniques in economics. Section four derives the representative agent model in discrete and continuous time Section five describes the exogenous learning rule model and provides the numerical results of this exercise, and section 6 concludes.

2 Mathematical Background

Continuous time optimization problems in economics have a simple general form, and the continuous time analog of the Bellman equation, the Hamilton-Jacobi-Bellman can be intuitively derived from the discrete model (Dixit, 1992). Suppose we have a simple Ramsey model where agents maximize their expected utility per unit time over time t

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \sum_{n=0}^{\lfloor \frac{1}{\Delta t} \rfloor} e^{-\rho(t+n\Delta t)} u(c_{t+n\Delta t}) \Delta t, \tag{1}$$

where capital evolves according to the following stochastic differential equation

$$\Delta k_{t+\Delta t} = a(k_t, c_t)\Delta t + b(k_t, c_t)\Delta W_t, \tag{2}$$

where ΔW_t is the increment of the Wiener process and the maximum value of n, $\lfloor \frac{1}{\Delta t} \rfloor$, limits value of n to integer values. This floor function will be equal to one when $\Delta t = 1$. Note that as $\Delta t \to 1$ equation (1) limits to the typical discrete utility maximization problem with a constant discount factor. The Wiener process can be written as $\varepsilon\sqrt{\Delta t}$, where $\varepsilon \sim N(0, 1)$. Thus, we can calculate the expectation and variance of ΔW_t

$$\mathbb{E}[\Delta W_t] = 0 \quad \text{and} \quad \mathbb{E}[(\Delta W_t)^2] = \Delta t.$$

The Bellman equation for this system can then be written as follows,

$$V(k, t) = \max_{c} \; u(c)\Delta t + e^{-\rho \Delta t}\mathbb{E}[V(k + \Delta k, t + \Delta t)] \tag{3}$$

in this setting the value function can be thought of as: the value of capital today is equal to the gain from the utility of consumption over one interval of time (Δt) plus expected discounted value the agent receives at $t + \Delta t$. The utility function in (3) is multiplied by the length of our time period as we care about the benefits that will accrue in that first period relative to its size (Dorfman, 1969). Since our value function is defined recursively, this expectation captures all future value of capital over time. In order to get the desired continuous-time value function, we can transform this discrete version (Dixit, 1992). First, using the power series expansion of $e^{-\rho \Delta t}$

we rewrite this problem. [1]

$$\rho\Delta t V(k,t) = \max_c \; u(c)\Delta t + (1-\rho\Delta t)\mathbb{E}[V(k+\Delta k, t+\Delta t) - V(k,t)] \tag{4}$$

Next we have to use stochastic calculus to find the value of this expectation. In stochastic calculus, we need to apply Itô's lemma in order to properly take the derivative of a function that depends on a stochastic process. This is necessary because these processes are continuous everywhere, but due to their volatile nature, they are nowhere differentiable.

Suppose, for a moment, that we are in a continuous setting with the following diffusion process,

$$dX_t = \mu dt + \sigma dW_t$$

in this setting μ is a drift term, σ is a variance term, and dW_t is the increment of a Wiener process. If we have a function $f(X_t, t)$ that depends on X_t and time t, we cannot take its derivative using traditional methods since X_t is nowhere differentiable. Instead, we must use Itô's lemma, this will yield

$$df(X_t,t) = \frac{\partial f}{\partial t}dt + \frac{\partial f}{\partial x}\cdot dX_t + \frac{1}{2}\frac{\partial^2 f}{\partial x^2}\cdot (dX_t)^2 + \mathcal{O}(dt^{\frac{3}{2}}).$$

Note, the application of Itô's lemma is essentially just a Taylor expansion of the series using particular assumptions about the stochastic nature of the system. A key assumption of stochastic calculus is at work in the equation above, we assume that all terms with dt^n where $n \geq \frac{3}{2}$ are approximately zero. This will lead to the cancellation of several terms in the expansion of dX_t^2 and all of terms in $\mathcal{O}(dt^{\frac{3}{2}})$. After expanding

[1] The power series expansion of $e^{-\rho\Delta t} = 1 - \rho\Delta t + \rho\Delta t^2 + \mathcal{O}(\Delta t^3)$. One of the common assumptions of stochastic calculus is that terms with including Δt to the power of $3/2$ or higher will be approximately zero in the limit. Thus, we will approximate $e^{-\rho\Delta t}$ as $1 - \rho\Delta t$.

terms and rearranging the equation we will be left with,

$$df(X_t, t) = \left(\frac{\partial f}{\partial t} + \frac{\partial f}{\partial x} \cdot \mu + \frac{1}{2} \frac{\partial^2 f}{\partial x^2} \cdot \sigma^2 \right) dt + \frac{\partial f}{\partial x} \cdot \sigma dW_t.$$

Now, if we take the expectation of this the last term will drop out since $\mathbb{E}[dW_t] = 0$. Thus, we will be left with

$$\mathbb{E}[df(X_t, t)] = \left(\frac{\partial f}{\partial t} + \frac{\partial f}{\partial x} \cdot \mu + \frac{1}{2} \frac{\partial^2 f}{\partial x^2} \cdot \sigma^2 \right) dt.$$

Following a similar set of steps, we can look at the difference in our value function over time. Approximating dV as $V(k + \Delta k, t + \Delta t) - V(k, t)$ we can write this as

$$V(k + \Delta k, t + \Delta t) - V(k, t) = V_t(k, t)\Delta t + V_k(k, t)(\Delta k) + \frac{1}{2} V_{kk}(k, t)(\Delta k)^2,$$

here we have already dropped out most terms with t^n where $n \geq \frac{3}{2}$. Carrying through the expectation will the give us the original term from our Bellman equation on the left hand side.

$$\mathbb{E}[V(k+\Delta k, t+\Delta t) - V(k, t)] = V_t(k, t)\Delta t + V_k(k, t)a(k, c)\Delta t + \frac{1}{2} V_{kk}(k, t)b(k, c)^2 \Delta t,$$

the $a(k, c)$ and $b(k, c)$ terms come from the original equation for our capital accumulation process given by equation (2). Plugging our expectation term into our value function in (4) we get,

$$\rho \Delta t V(k, t) = \max_c \; u(c)\Delta t + (1 - \rho \Delta t)\big(V_t(k, t) + V_k(k, t)a(k, c) + \frac{1}{2} V_{kk}(k, t)b(k, c)^2\big)\Delta t.$$

Then if we divide by Δt and take the limit as $\Delta t \to 0$ we get the standard HJB

$$\rho V(k,t) = \max_c \; u(c) + V_t(k,t) + V_k(k,t)a(k,c) + \frac{1}{2}V_{kk}(k,t)b(k,c)^2.$$

This HJB equation represents a solution to the given continuous-time maximization problem,

$$\max_{c_t} \int_{t=0}^{\infty} e^{-\rho t} u(c_t) dt.$$

Often, when we are concerned with infinite-horizon problems the $V_t(k,t)$ term will be left out of the HJB. This term is assumed to be zero in infinite horizon problems because as our time dimension becomes infinitely large changes in our value function over (the already infinitely small) increments of time become negligible.

Additionally, in this setting, we might care about the distribution of our state variable k, $g(k,t)$. This distribution is particularly of interest in a setting with heterogeneous agents because heterogeneity and idiosyncratic shocks will impact the evolution of this distribution over time. We can find this distribution using the Kolmogorov Forward Equation (KF), sometimes called the Fokker-Planck Equation. Given an initial distribution $g_0(k)$ the distribution $g(k,t)$ satisfies,

$$\frac{\partial g(k,t)}{\partial t} = -\frac{\partial}{\partial k}[a(k,c)g(k,t)] + \frac{1}{2}\frac{\partial^2}{\partial k^2}[b(k,c)^2 g(k,t)].$$

If a stationary distribution for $g(k)$ exists, it satisfies the ordinary differential equation (ODE)

$$0 = -\frac{\partial}{\partial k}[a(k,c)g(k)] + \frac{1}{2}\frac{\partial^2}{\partial k^2}[b(k,c)^2 g(k)].$$

In a model with multiple agents, the KF equation is one of the key equations that describe the system. In an Aiyagari model, for instance, the KF will determine prices and market clearing, since market clearing is dependent on the distribution of the agents and their preferences. The KF equation is an important feature in stochastic continuous-time literature; however, it is not used in the representative agent setting present in the rest of this paper. For more information on the derivation and key concepts of the KF equation see the appendix.

Another way to view the KF equation is as the continuous time analog to the multiplication of transition matrices (in a Markovian setting). The time-dependent version of this equation describes the evolution of the probability density function of the key variables under the influence of deterministic and random forces found in diffusion processes. The continuous probability distributions that come from our KF equation are one of the most attractive features of continuous time modeling. Since often with modern policies, we care most about the distribution of goods, wealth, or assets.

Now, that we have explored both the HJB and KF equations it is important to note that the HJB equation is closely related to the maximized Hamiltonian, this is easily shown. First, if we have the system defined in this section with $b(k, c) = 0$ our Hamiltonian is

$$\mathcal{H}(k_t, c_t, \lambda_t) = u(c_t) + \lambda_t a(k_t, c_t),$$

while our HJB equation is

$$\rho V(k) = \max_c \; u(c) + V'(k)a(k_t, c_t).$$

Connecting the two we see $\lambda_t = V'(k)$, i.e. the shadow price of k is equivalent to the marginal value of k. Thus we can rewrite the HJB as

$$\rho V(k) = \max_c \; H(k, c, V'(k))$$

where,

$$H(k, V'(k)) = u(c) + V'(k)a(k, c).$$

The HJB and KF equations, though compact and simple in appearance, can be used to solve complex economic and financial problems. Closed form solutions to these problems are often impossible to calculate by hand, but with new computational developments finding solutions to these systems has become more plausible and these solution methods show some advantages to long-popular discrete models.

Not only is continuous time modeling is more intuitive, but it also provides more information about the distribution of parameters with convenience. This comes from the KF equation that summarizes the distribution of parameters, using the distribution from this equation researchers can analyze the distribution of a variable over time or after a shock. The distribution that solves the KF equation can also be used for estimating model parameters and can provide a likelihood estimator for the model. Additionally, the algorithms for solving continuous time systems are fast due to the sparsity of the matrices that determine the evolution of the system.

These modern advances have made continuous time modeling more attractive to economists since solutions to these systems can now be found without a large number of assumptions. Though these continuous time problems did not have simple solution methods until more recently, many researchers have explored modeling in a stochastic continuous time setting.

3 Literature

This paper works to develop learning techniques in stochastic continuous time. Therefore we blend two distinct kinds of literature together, stochastic continuous time modeling and adaptive learning. In this section, we will first review the stochastic continuous-time literature. Research on these models in economics has been sparse but spread widely throughout time. For a deeper understanding of this setting and on why it is becoming more relevant today a historical overview of these modeling techniques is necessary. Learning literature, on the other hand, has been consistently studied for a long time. There is a wealth of knowledge on this topic, and we only examine a small part of this literature that is relevant to our work.

3.1 Stochastic Continuous-Time Literature: A Historic Overview

The stochastic continuous time setting has become increasingly popular in macroeconomic modeling. Interest in this framework first arose in the early 1970s with financial economic models, these early works include Merton (1969), Merton (1971), and Black & Scholes (1973). In financial economics casting models in continuous time is particularly intuitive as many financial variables evolve, such as stock prices, can be observed on very small intervals; making their prices essentially a continuous variable instead of a discrete one.

Some early works in continuous time financial models include Black & Scholes (1973), Eaton (1981), Merton (1971), Merton (1969), and Mirrlees (1971). These papers set up continuous time models and solve them as rigorously as possible without the aid of modern computational techniques, often by using comparative statics. This is done because the system of partial differentials that describes equilibrium in this class of models is often unsolvable unless specific forms for the value function are assumed. Due to these identification issues, most of the papers mentioned above focused on solving for the distributional steady state of their models.

Black & Scholes (1973) develops a method for determining fair prices for European call options. Unlike many economic models, Black & Scholes (1973) is able to assume several boundary conditions and functional forms that aid in solving their key partial differential equations. These boundary conditions and functional forms are such that the HJB is able to be written in the same form as a standard heat equation. Once getting the HJB problem into this format is easily solve for using Fourier transformations. Most optimization problems in this setting cannot be solved for explicitly like the Black-Scholes problem. Part of the reason why this is possible to solve the Black-Scholes model is that it is defined specifically for European call options that can only be called at the end of their lifespan.

Eaton (1981) explores the effects of fiscal policies on the composition of portfolios and the accumulation of capital. This model defines the net output, government expenditure, and tax revenue as stochastic processes. All of these processes depend on aggregate capital stock, which allows the government in this model to tax the random component of capital income at a different rate than the deterministic part and defines government expenditure to be depend differently on the deterministic and random parts of capital. After setting up this model the author uses comparative statics and some simplifying assumptions to conclude that fiscal policy changes impact the average yield and riskiness of capital relative to government debt.

Robert Merton has several papers from this period that develop models in stochastic continuous-time. Merton (1969) develops a model for optimal portfolio selection where the agents' income is generated by returns on assets. Merton (1971) further examines this problem and uses explicit forms for the utility function to derive optimal consumption and portfolio rules. This paper also uses comparative statics to examine the response of these rules to certain parameter changes, a popular technique during this time. Merton (1975b) examines common economics growth models in this setting. The model discusses in Merton (1975b) is a one-sector neoclassical growth model where the size of the labor force evolves according to a stochastic process. The paper then takes the neoclassical growth model and expands it into a stochastic Ramsey problem. Merton (1975b) is one of the first publications that use more traditional economic modeling in this stochastic continuous time setting. Another paper that implements traditional economic models is Brock & Mirman (1972).

Brock & Mirman (1972) differs from these other papers because in this model a solvable steady state exists. This growth model is special due to the linearity of the consumption function, this allows for the steady state of the stochastic model to be equal to the steady state of the non-stochastic model. Due to the tractability of this model the Brock-Mirman model is one of the most common stochastic continuous models used prior to the introduction of more advanced computational methods.

Dixit (1989) models firm entry and exit decisions where output price follows a geometric Brownian motion. The model in this paper is solved by simplifying the system of PDEs into a simpler system of ordinary differentials. This produces a solution that consists of trigger prices for firm entry and exit. Prices in between the entry trigger and the exit trigger price produce "hysteresis" which appears in the model even with small sunk costs.

During the late 1990s and early 2000s a number of books were published on continuous time models in financial economics; these include Merton (1992), Dixit (1992), Dixit & Pindyck (1994), and Stokey (2009). The publication of these works formalized the use of continuous time models, particularly in finance. Merton (1992) contains several of Merton's papers mentioned earlier in this literature review and is directed at finance graduate students. Dixit & Pindyck (1994) is also targeted at finance graduate students and contains some of the most intuitive derivations of the HJB equation out of all economics and finance literature. Dixit (1992) contains intuitive mathematical introductions and focuses on how to implement boundary conditions in the stochastic continuous time setting. Stokey (2009), differs from the other books on stochastic continuous time modeling. This book focuses more on the mathematical background and measure theory that is necessary for a deeper understanding of this material. The main contribution of this work is the understanding that continuous time modeling better captures the dynamics of inaction and boundaries that are rarely binding. This setting's ability to capture inaction and boundary conditions is the reason why stochastic continuous time modeling has become so popular in financial economics.

With the availability of better computational methods more econometric papers have been written on stochastic continuous time models. Hansen & Scheinkman (1995) derive moment conditions for estimating and testing continuous-time Markov models using discrete time data. Aït-Sahalia has a number of economics and finance papers published throughout the 1990s and early 2000s on econometric tests for diffusion processes. Aït-Sahalia (2002) constructs a maximum-likelihood approach to estimating parameters in discretely sampled diffusion models. Aït-Sahalia (2004) furthers the methods from Aït-Sahalia (2002) and constructs an approach to estimating parameters in these models when the sampling intervals are not uniform. Posch (2009) solves continuous time dynamic stochastic general equilibrium models with jumps and shows how the continuous time setting can make it simpler to estimate the likelihood function. This paper solves the model by introducing several simplifying assumptions and confirming the results with Monte Carlo estimates.

Most stochastic continuous time modeling in the early 2000s was done using assumptions about the form of the value function or by imposing multiple boundary conditions. Financial economists such as Sannikov extended stochastic continuous time modeling to a microeconomic setting. In Sannikov (2008) and DeMarzo & Sannikov (2006) a principal-agent setting is developed in continuous-time. Solutions to these principal agents are found by implementing several boundary constraints, which at the time of their publication was an innovative technique. This technique opened up the door for more publications in the stochastic continuous time setting.

Hansen *et al.* (2006) takes a more theoretical approach to stochastic continuous time modeling and explores model misspecification in this setting. Duffie & Epstein (1992) develops a stochastic differential formulation of recursive utility. Gabaix (2009) has a section on continuous time approaches to power laws. In this paper, the size of an economic unit (cities or firms) is modeled as a stochastic process that can hit reflective boundaries at some points. Using this process one can use the KF equation to describe the evolution of this distribution using power laws a unique solution to this system can be found.

Prior to 2015 economists were not widely implementing computational methods to find solutions to this class of optimization problems. However, Forsyth & Labahn (2007), a computational finance paper, studies numerical methods for solving HJB equations in finance. This paper finds that discretizing the HJB and solving it numerically will converge to the viscosity solution. The viscosity solution is the same solution that economists began focusing on around 2015. Viscosity solutions are continuous and differentiable solutions to the HJB that are in most cases unique. Forsyth & Labahn (2007) also analyzes Newton-type iterations schemes and finds that these also solve the HJB equation, another result economists realized later.

With the rise of heterogeneity in macroeconomics, economic models have developed new more complexity. Discrete time models can capture rich heterogeneity; however, these methods are time-consuming and cannot provide the same level information about the distribution of key variables as continuous time models. To solve these new richer models economists have developed more advanced solution techniques many recent papers have been focused on developing and implementing these algorithms.

Achdou *et al.* (2014) uses tools from applied mathematics to solve the HJB equation. The algorithm outlined in the paper uses finite difference methods to solve for an approximate solution to the HJB. This approximate solution, called the viscosity solution, assumes that the value function is differentiable on its entire domain. Viscosity solutions are unique given that several weak conditions hold. In Achdou *et al.* (2014) this method allows the authors to find both steady-state and time-dependent solutions for their models. Other papers such as Kaplan *et al.* (2018), Achdou *et al.* (Working), and Parra-Alvarez *et al.* (Working) implement the same techniques. This paper uses the steady state solution methods presented in Achdou *et al.* (2014) in the exogenous learning rule model.

A key issue with the solution methods presented in Achdou *et al.* (2014) is that the time-dependent solutions cannot be used in conjunction with random aggregate shocks. Ahn *et al.* (2018) uses the foundation developed in Achdou *et al.* (Working) to create a more complicated algorithm for analyzing models with heterogeneous agents that are subject to shocks. This algorithm calculates the steady-state versions of the HJB and KF equations and then linearizes the system around that steady state without aggregate shocks. Linearization around this steady state involves using a first order Taylor expansion since this system has a large number of variables the derivatives needed for this Taylor expansion cannot be taken by hand and must be calculated using automatic differentiation.

After the system is linearized it can be easily solved and the Schur decomposition of the coefficient matrix can be used to check for stable roots. Using this algorithm one can look at impulse response functions and the effects of shocks on a continuous model. The algorithm in this paper is an important innovation as previous solution methods prevented researchers from analyzing random macroeconomic shocks. Being unable to analyze these types of shocks was a drawback of stochastic continuous time modeling in macroeconomics. Now that a simple portable algorithm for analyzing these types of models exists the stochastic continuous time setting is likely to become increasingly popular among researchers in theoretical macroeconomics.

The representative agent model outlined in this paper will use the same approach as Ahn *et al.* (2018) to solve the model and to implement learning in this framework.

3.2 Learning Literature

The motivation of this paper is to develop adaptive learning techniques in the stochastic continuous-time setting. Adaptive learning is a statistical approach that overcomes the strict model assumptions implied by traditional rational expectations. In learning models, agents use statistical techniques to estimate model parameters and update their expectations of parameters and other values over time. Most learning papers involve direct feedback from the agents' estimates through a special mapping called a T-map. This paper relies on exogenous learning rules that appear similar to simple econometric learning as described in Evans & Honkapohja (2001); however, the algorithms in section five do not have this feedback rule. Instead, the learning in this paper comes from simple information drops and all new information is used to directly update parameter estimates.

This method of learning is more similar to the early works in this literature. For instance, Bray (1982) looks at a more simple version of updating estimates via OLS. Some of the models explored in this paper do not look directly at feedback rules and instead focus on seeing if agents can get rational expectations estimates of parameters when presented with additional information. The agents do this by implementing OLS each period with updated information. This paper found that under some assumptions the OLS learning converged to the rational expectations equilibrium's values. Also, learning in this paper focuses on learning parameters in a steady-state setting. Similar environments have explored previously work, notable steady learning as mentioned in Evans & Honkapohja (2009).

There does exist some literature similar to stochastic continuous-time adaptive learning in asset pricing literature. Veronesi (2019) examines a Bayesian learning rule in an asset pricing model with heterogeneous risk preferences. Some other papers such as Bhamra & Uppal (2014), also discuss implementing a similar learning rule. The work in these papers is distinctly different than what we will proceed with, since the focus of these works is finding parameters based on distributions.

In this paper, one of the main focuses in our learning section is adapting misspecification. There has been some work on this within asset pricing literature, notably Hansen & Sargent (2019b) and Hansen & Sargent (2019a). These papers look at misspecification within models and also look at an agent's choice between several well-defined models. All of these asset pricing models are cast in stochastic continuous-time. This is done in order to exploit the convenient properties of Brownian motion and continuous likelihood functions.

4 A Representative Agent Model

In this section, we develop discrete and continuous models in deterministic and stochastic settings to better understand the connects between discrete and continuous models. This is done with a few simple Ramsey models. We first develop the discrete and continuous models separately before comparing them closely. One important feature of the discrete methods is the inclusion of time increments Δt, which allows us to compare our discrete and continuous models. The use of Δt in the following sections is based on previous work by Dorfman (1969). Doing this allows us to better understand the similarities of discrete and continuous time systems, and creates a discrete setting to develop a benchmark for how learning should impact a system with infinitely small time intervals. This section of the paper proceeds by first developing deterministic and stochastic versions of a model in discrete and continuous time. Next, we will compare these models and show how they are related as increments of time get infinitely small. In both of the cases outlined below the discrete model limits to the continuous version.

4.1 A Deterministic Model

Before worrying about systems with stochasticity, we first outline a simple Ramsey model in a deterministic setting. First, we will describe the discrete case and the continuous case separately. Then, we will compare the two models.

4.1.1 The Discrete-Time Deterministic Model

Discrete time models in economic often assume that $\Delta t = 1$, this assumption makes models less notationally bulky. However, in doing so information is lost, the utility functions used in economics are utility *per unit time* and our discrete discount factor is dependent on units of time as well. The model outlined below takes these units of time into consideration, and carefully examines the optimality conditions with this Δt term.

Before describing the model it is important to understand the discount factor's dependence on time. The discount factor β is defined as the discount rate per unit of time and can be written as a function of the increment of time $\beta(\Delta t)$. Furthermore, $\lim_{\Delta t \to 0} [\beta(\Delta t)]^t = e^{-\rho t}$.

Using this discount factor we can proceed with our model. A representative agent in this setting will maximize utility per unit time according to

$$\max_{c_t} \mathbb{E}_0 \sum_{t=0}^{\infty} \sum_{n=0}^{\lfloor \frac{1}{\Delta t} \rfloor} \beta^{t+n\Delta t} u(c_{t+n\Delta t})\Delta t, \tag{5}$$

here $\lfloor \frac{1}{\Delta t} \rfloor$ limits n to integer values, since $\Delta t \leq 1$. In this setting, capital evolves according to the following process[2]

$$k_{t+\Delta t} = (e^{z_t} f(k_t) - \delta k_t - c_t)\Delta t + k_t \tag{6}$$

where $f(k_t) = k_t^{\alpha}$.[3] In this deterministic setting we will have the following process for the evolution of productivity z_t, [4]

$$z_{t+\Delta t} = (1 - \eta \Delta t) z_t \tag{8}$$

This model is closely related to the stochastic continuous model outlined later in this section. We can note that Δt becomes dt in the limit, using this the equations (5)-(8) will be equivalent to the ones used for the continuous deterministic model.

Optimization problems in this setting can take several different forms. First, we

[2]Setting

$$\dot{k} = \frac{k_{t+\Delta t} - k_t}{\Delta t} = e^{z_t} f(k_t) - \delta k_t - c_t$$

(6) is the typical equation for the evolution of capital in a discrete Ramsey model

[3]In this discrete model if we normalize $\Delta t = 1$, (6) is the standard equation for capital accumulation.

$$k_{t+1} = e^{z_t} f(k_t) + (1 - \delta) k_t - c_t$$

[4]In a stochastic setting productivity z_t will evolve according to the following AR(1) process. This process was derived from the standard Ornstein-Uhlenbeck process in (30) using the Euler-Maruyama method

$$z_{t+\Delta t} = (1 - \eta \Delta t) z_t + \sigma \epsilon_t \sqrt{\Delta t} \tag{7}$$

Here $\epsilon_t \sim N(0, 1)$.

can write out the Lagrangian.

$$\mathscr{L}(z_0, c_0, \lambda_0) = \mathbb{E}_0 \sum_{t=0}^{\infty} \sum_{n=0}^{\lfloor \frac{1}{\Delta t} \rfloor} \beta^{t+n\Delta t} \{ u(c_{t+n\Delta t})\Delta t$$
$$+ \lambda_{t+n\Delta t}[k_{t+n\Delta t} + (e^{z_t} f(k_{t+n\Delta t}) - \delta k_{t+n\Delta t} - c_{t+n\Delta t})\Delta t - k_{t+(n+1)\Delta t}]\}$$

In this setting our first order conditions will be the following,

$$\frac{\partial \mathscr{L}}{\partial c_t} = \frac{\partial u}{\partial c}\Delta t - \lambda_t \Delta t = 0 \tag{9}$$

$$\frac{\partial \mathscr{L}}{\partial k_{t+\Delta t}} = \beta^{t+\Delta t}\mathbb{E}_t \lambda_{t+\Delta t}\big[\big(e^{z_t} f'(k_{t+\Delta t}) - \delta\big)\Delta t + 1\big] - \beta^t \lambda_t = 0 \tag{10}$$

$$\frac{\partial \mathscr{L}}{\partial \lambda_t} = k_t + (e^{z_t} f(k_t) - \delta k_t - c_t)\Delta t - k_{t+\Delta t} = 0 \tag{11}$$

where we have supposed that $\lambda_{t+\Delta t} = \lambda_t + \dot{\lambda}\Delta t$, where $\dot{\lambda}$ is that rate at which it will change over our interval of time. Since this setting is deterministic we can now drop the expectation term. Then we can rewrite (10).

$$\frac{\partial \mathscr{L}}{\partial k_{t+\Delta t}} = \lambda_t[e^{z_t} f'(k_{t+\Delta t}) - \delta] = -\lambda_t \ln \beta - \dot{\lambda} \tag{12}$$

For a full derivation of (12) see the appendix.

We can look at this problem in from a Hamiltonian framework. In this setting the current value Hamiltonian is,

$$J(k_t, \mu_{t+\Delta t}, c_t, t, t+\Delta t) = u(c_t) + \mu_{t+\Delta t}(e^{z_t} f(k_t) - \delta k_t - c_t) + \gamma_{t+\Delta t}\big(- \eta z_t \Delta t\big) \tag{13}$$

with the transversality condition

$$\lim_{t\to\infty} \beta^t \mu_t k_t \leq 0. \tag{14}$$

The first order conditions for this system are given by the following equations,

$$\frac{\partial J}{\partial c_t} = u'(c_t) - \mu_{t+\Delta t} = 0 \tag{15}$$

$$\frac{\partial J}{\partial k_t} = \mu_{t+\Delta t}(e^{z_t} f'(k_t) - \delta) = -\frac{\mu_{t+\Delta t} - \mu_t}{\Delta t} - \ln(\beta)\mu_{t+\Delta t} \tag{16}$$

in this setting $\frac{\mu_{t+\Delta t}-\mu_t}{\Delta t} = \dot{\mu}_t$. At a glance (16), looks very similar to (12). Using (15) and (16) we can get the typical first order conditions for a Hamiltonian system,

$$\left(e^{z_t} f'(k_t) - \delta + \ln(\beta)\right) = -\frac{u''(c_t)}{u'(c_t)}\dot{c}_t.$$

This is similar to the continuous time version; however, the multiplier in this case is incremented forward one unit of time, and our discrete discount rate causes our first order conditions to include a $\ln(\beta)$ term. As the increment of time approaches zero, the discrete Hamiltonian outlined here will be equivalent to the continuous Hamiltonian outlined in the following section.

We can also write a discrete Bellman equation for this system

$$V(z_t, k_t) = \max_{c_t}\ u(c_t)\Delta t + \beta^{\Delta t}[V(z_{t+\Delta t}, k_{t+\Delta t})]. \tag{17}$$

This setting will have similar first conditions. First we can take the first order condition of this system with respect to c_t

$$\frac{\partial u}{\partial c_t}\Delta t + \beta^{\Delta t}\frac{\partial}{\partial c_t}V(z_{t+\Delta t}, k_{t+\Delta t}, t + \Delta t) = 0. \tag{18}$$

In this case $\frac{\partial}{\partial c_t}V(k_{t+\Delta t}, z_{t+\Delta t}, t + \Delta t) = \frac{\partial V(\cdot)}{\partial k_{t+\Delta t}}\frac{\partial k_{t+\Delta t}}{\partial c_t}$. Thus, we will have

$$\frac{\partial u}{\partial c_t}\Delta t = \beta^{\Delta t}\frac{\partial V(\cdot)}{\partial k_{t+\Delta t}}\frac{\partial k_{t+\Delta t}}{\partial c_t}\Delta t.$$

Simplifying and denoting the marginal value of capital at time t as $\mu_t = \frac{\partial}{\partial k}V(k,t)$ we will have the following equation

$$\frac{\partial u}{\partial c_t} = \beta^{\Delta t}\mu_{t+\Delta t}\frac{\partial k_{t+\Delta t}}{\partial c_t},$$

this is equivalent to (15). Taking the first order condition with respect to k will then yield,

$$\mu_t = [1 + \big(e^{z_t}f'(k_t) - \delta\big)\Delta t](\mu_{t+\Delta t})\beta^{\Delta t}$$

Simplifying this will give us (16) from our discrete Hamiltonian. Examining this we can see that the value function and Hamiltonian are closely related as in Dorfman (1969).

4.1.2 The Continuous-Time Deterministic Model

The continuous time version of this model can be described according to the following equations. Our agent will maximize expected utility according to the following equation, here $e^{-\rho t}$ will be the continuous time equivalent of the discrete discount factor β^t,

$$\max_{c_t} \mathbb{E}_0 \int_{t=0}^{\infty} e^{-\rho t}u(c_t)dt. \tag{19}$$

This is setting capital will evolve according to the following process

$$dk_t = (e^{z_t}f(k_t) - \delta k_t - c_t)dt, \tag{20}$$

where the production function is the same as before. Productivity will evolve according to

$$dz_t = -\eta z_t dt, \tag{21}$$

the continuous time analog to the discrete process in the previous section.

In this setting, the current value Hamiltonian can be rewritten as follows,

$$H(k_t, z_t, c_t, \gamma_t, \mu_t, t) = u(c_t) + \mu_t(e^{z_t} f(k_t) - \delta k_t - c_t) - \gamma_t(\eta z_t).$$

It is clear that $H(\cdot) = \lim_{\Delta t \to 0} J(\cdot)$, thus this directly related to our discrete time problem. The first order conditions for this system will be given by the following equations.

$$\frac{\partial H}{\partial k_t} = \mu_t(e^{z_t} f'(k_t) - \delta) = -\frac{d\mu_t}{dt} + \rho\mu_t \tag{22}$$

$$\frac{\partial H}{\partial c_t} = u'(c_t) - \mu_t = 0 \tag{23}$$

The transversality condition in continuous time can be written as follows.

$$\lim_{t \to \infty} e^{-\rho t} \mu_t k_t \leq 0$$

Together the first order conditions (22) and (23) imply,

$$u'(c_t)(e^{z_t} f'(k_t) - \delta - \rho) = -\frac{d\mu_t}{dt}$$

We can also write a HJB for this system, since we are in a continuous time setting.

$$\rho V(k, z) = \max_c \; u(c) + \partial_k V(k, z)(e^z f(k) - \delta k - c) - \partial_z V(k, z)(\eta z) \tag{24}$$

Setting μ_t in the current value Hamiltonian equal to $\partial_k V(k, z)$, and γ_t equal to $\partial_z V(k, z)$ we can rewrite this again.

$$\rho V(k, z) = \max_c \; H\big(k, z, c, \partial_k V(k, z), \partial_z V(k, z)\big) \tag{25}$$

4.1.3 Comparing the Deterministic Models

For a clear comparison of the discrete and continuous time models outline in this section, we can examine the discrete Bellman equation (17) as $\Delta t \to 0$. First, we can take an approximation of $V(k_{t+\Delta t}, z_{t+\Delta t})$, in a method similar to Dorfman (1969).

$$V(k_{t+\Delta t}, z_{t+\Delta t}) = V(k_t, z_t) + \partial_k V(k_t, z_t)(k_{t+\Delta t} - k_t) + \partial_z V(k_t, z_t)(z_{t+\Delta t} - z_t) + \mathcal{O}\big(\Delta t^{\frac{3}{2}}\big)$$

All other partials and cross partial derivatives will be in the $\mathcal{O}$ term. These terms will all be approximately zero in the limit as $\Delta t \to 0$. Next, we will approximate $\beta^{\Delta t} \approx e^{-\rho \Delta t} \approx (1 - \rho \Delta t)$. Using these two approximations we can rewrite (17) as follows.

$$V(k_t, z_t) = \max_{c_t} \; u(c_t) + (1-\rho\Delta t)[V(k_t, z_t) + \partial_k V(k_t, z_t)(k_{t+\Delta t} - k_t) + \partial_z V(k_t, z_t)(z_{t+\Delta t} - z_t)]$$

Simplifying and substituting in for the changes in k and z, this will yield

$$\rho V(k, z) = \max_{c} \; u(c) + \partial_k V(k, z)(e^z f(k) - \delta k - c) - \partial_z V(k, z)(\eta z). \qquad (26)$$

This is the same equation as the HJB derived earlier in this section (24). In the deterministic system comparing the HJB and the Bellman equation is more simple, since we do not need to worry about expectation terms. This is because the deterministic version of this model does not have uncertainty, adding in a continuous time version of our process in (7) will give us a more complicated optimization problem.

4.2 A Stochastic Model

Now, we build a stochastic model in discrete and continuous time. Adding in stochasticity will yield more complex models and additional terms in the HJB equation. These stochastic models are more common in literature and are closely related to the Ramsey models used later in this paper.

4.2.1 The Discrete-Time Stochastic Model

This model will be a stochastic version of the discrete time model defined previously. In this setting agents will maximize utility according to (5) and capital will evolve according to (6) with the same Cobb-Douglas production function. The main difference between this model and the previous deterministic model is that z_t evolves according to the following AR(1) process,

$$z_{t+\Delta t} = (1 - \eta\Delta t)z_t + \sigma\epsilon_t\sqrt{\Delta t}, \tag{27}$$

where $\epsilon_t \sim N(0,1)$. This model is closely related to the stochastic continuous model outlined later in this section.

Optimization problems in this setting can take several different forms. The current value Hamilton for this problem is,

$$J(k_t, \mu_{t+\Delta t}, c_t, t, t+\Delta t) = u(c_t) + \mu_{t+\Delta t}(e^{z_t} f(k_t) - \delta k_t - c_t) + \gamma_{t+\Delta t}\big(-\eta z_t + \sigma\epsilon_t\sqrt{\Delta t}\big). \tag{28}$$

The transversality condition will be the same as in the discrete deterministic model (14). Despite the presence of an additional term, the first order conditions for this system will be the same as the ones from the discrete deterministic model. Also as the increment of time approaches zero, the discrete Hamiltonian outlined here will be equivalent to the continuous Hamiltonian outlined in the following section.

We can also write the discrete Bellman equation for this system

$$V(k_t, z_t) = \max_{c_t}\ u(c_t)\Delta t + \beta^{\Delta t}\mathbb{E}[V(z_{t+\Delta t}, k_{t+\Delta t})]. \tag{29}$$

This setting will have similar first conditions to the discrete model previously studied.

4.2.2 The Continuous-Time Stochastic Model

One of the key differences between the continuous-time model in this section and the one previously outlined is the process for productivity. Productivity in the continuous time setting will evolve according to the following Ornstein-Uhlenbeck process, the continuous-time analog of (27).

$$dz_t = -\eta z_t dt + \sigma dW_t \tag{30}$$

Where dW_t is the increment of the Wiener process.

Equilibrium in the continuous time setting is given by the following equations. First, equilibrium will depend on the HJB equation (31), the continuous time analog of the Bellman equation. We can first write this equation in a form similar to (17).

$$V(k, z) = \max_{c_t}\ u(c_t) + e^{-\rho t}\mathbb{E}[V(k', z')]$$

The expectation term in this model will differ from the expectations in (17). This is because the Wiener process in (30) is continuous, but is nowhere differentiable making it impossible to treat this expectation like a normal Riemann integral. Using stochastic calculus to solve for this expectation will yield the following HJB equation.

$$\rho V(k,z) = \max_c \ u(c) + \partial_k V(k,z)(e^{z_t} f(k) - \delta k - c_t) + \partial_z V(k,z)(-\eta z_t) + \frac{1}{2}\partial_{zz} V(k,z)\sigma^2 \tag{31}$$

Taking the first order condition with respect to consumption for the HJB equation will give us (32).

$$u'(c_t) = \partial_k V(k,z)$$

This is analogous to (9) in the discrete model or (23) in the deterministic continuous model. The term setting the μ_t from the continuous time current value Hamiltonian (13) equal to $\partial_k V(k,z)$ we can rewrite the HJB.

$$\rho V(k,z) = \max_c \ H\big(k, c, z, \partial_k V(k,z), \partial_z V(k,z)\big) + \frac{1}{2}\partial_{zz} V(k,z)\sigma^2 \tag{32}$$

This equation links the HJBs of our stochastic and non-stochastic continuous time models.

4.2.3 Comparing the Stochastic Models

Furthermore, we can compare the discrete and continuous stochastic models we have outlined thus far. If we take the discrete Bellman in (29), we can recast it and make it more similar to (32). First, we can take an approximation of $V(k_{t+\Delta t}, z_{t+\Delta t})$, in a method similar to Dorfman (1969).

$$V(k_{t+\Delta t}, z_{t+\Delta t}) = V(k_t, z_t) + \partial_k V(k_t, z_t)dk + \partial_z V(k_t, z_t)dz + \frac{1}{2}\partial_{zz} V(k_t, z_t)dz^2 + \mathcal{O}(\Delta t^{\frac{3}{2}})$$

All other partials and cross partial derivatives will be in the $\mathcal{O}$ term. These terms will all be approximately zero in the limit as $\Delta t \to 0$. Next, we will approximate

$\beta^{\Delta t} \approx e^{-\rho\Delta t} \approx (1 - \rho\Delta t)$. Using these two approximations we can rewrite (17) as follows.

$$\begin{aligned} V(k_t, z_t) &= \max_{c_t}\ u(c_t) + (1 - \rho\Delta t)[V(k_{t+\Delta t}, z_{t+\Delta t})] \\ &= V(k_t, z_t) + \partial_k V(k_t, z_t)dk + \partial_z V(k_t, z_t)dz + \frac{1}{2}\partial_{zz} V(k_t, z_t)dz^2 \end{aligned}$$

Simplifying and taking the limit as $\Delta t \to 0$ we wil be left with the following equation.

$$\rho V(k, z) = \max_{c}\ u(c) + \partial_k V(k, z)(e^{z_t} f(k) - \delta k - c) - \partial_z V(k, z)(\eta z_t) + \frac{1}{2}\partial_{zz} V(k, z)\sigma^2$$

Using this derivation we have gotten the stochastic HJB in equation (31). Thus, we have connected our discrete and continuous models in both deterministic and continuous settings.

4.3 Results

We have built four closely related models in this section and shown how discrete time models limit to their continuous time counterparts. With the correct setup, discrete time models will be the same in the limit as the continuous time models. The model comparisons in this section have demonstrated clear connections between discrete and continuous models. These connections are especially clear in the deterministic version of our models; however, with the use of stochastic calculus, they are easily seen.

Furthermore, in this section, we have recast discrete time models so that they contain the increment of time, Δt. This alone is a contribution to current literature as few economists examine models where $\Delta t = 1$. Within this class of models where Δt is built into the model, one could explore and compare many models with different values for Δt.

5 Adaptive Learning Rules

Now that we have developed our modeling framework for this paper, we will move on to examining representative agent exercises in learning. The first group of exercises will focus on an “stylized” learning rule. In this setting we build models where our agents have a misperception of the true underlying parameters, then our agents receive information dumps where they get some insight into the true model parameters. Here agents are trying to update their parameters to make optimal steady state decisions. Thus, our system is not time dependent. Our agents do recalculate the model for a number of periods, but these periods do not correspond to time periods in our model.

Three models will be explored in the following section. The first examines the stylized learning rule when the unknown model parameter is part of the exogenous stochastic process. Next, the stylized learning rule is applied to a model with misspecification in an endogenous stochastic process for the evolution of capital stock. Lastly, we modify the model with a stochastic processes for productivity and implement a real time updating rule that utilizes recursive least squares, a more meaningful and realistic approach.

5.1 Learning the Process for Productivity

There is a representative agent that makes consumption choices c and has capital stock k. The state of the economy depends on the flow of capital stock. The agent has standard preferences over utility flows based on capital discounted at rate $\rho \geq 0$. This can be written as the following equation:

$$\mathbb{E}_0 \int_{t=0}^{\infty} e^{-\rho t} u(c_t) dt \tag{33}$$

Here consumption, $c_t \geq 0$ for all periods. The agent's capital stock will evolve according to the following stochastic process used in Achdou *et al.* (2014).

$$dk_t = \big(z_t k_t^{\alpha} - \delta k_t - c_t\big)dt \tag{34}$$

This is the continuous time analog of the typical equation for the evolution of capital stock. The production function used in this section is Cobb-Douglas, $f(k_t) = k_t^{\alpha}$. Technological progress z_t will evolve according to the following equation

$$d\log(z_t) = -\theta \log(z_t)dt + \sigma dW_t. \tag{35}$$

This is a logged version of an Ornstein-Uhlenbeck process, this means that z_t will follow a stationary continuous process that is analogous to an AR(1) process. This can be rewritten in terms of z_t,

$$dz_t = \Bigg(-\theta \log(z_t) + \frac{\sigma^2}{2}\Bigg) z_t dt + \sigma z_t dW_t. \tag{36}$$

In this form we can more clearly see the drift for this process will be, $\big(-\theta \log(z_t) + \frac{\sigma^2}{2}\big) z_t$, and the variance term will be, σz_t.

The utility function used throughout this project will have constant relative risk aversion (CRRA),

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

here the CRRA parameter will be γ and $\gamma > 0$.

5.1.1 Stationary Equilibrium

A stationary equilibrium in this setting is given by the following equations. Our HJB for this problem is

$$\rho V(k,z) = \max_{c} \; u(c) + \partial_k V(k,z) \cdot \big(z_t f(k) - \delta k - c\big) + \\ + \partial_z V(k,z) \cdot \bigg(- \theta \log(z) + \frac{\sigma^2}{2} \bigg) z + \partial_{zz} V(k,z) \cdot \frac{1}{2} \sigma^2 z^2.$$

The derivation of this HJB can be found in the appendix along with a description of the algorithm used to solve this value function problem.

The agents in this simple model hold an incorrect belief about the diffusion process for technological progress. In this setting with exogenous learning, they predict that the diffusion process is given by the equation below,

$$d \log(z_t) = -\theta_g \log(z_t) dt + \sigma_g dW_t$$

There are two parameters that the agent misspecifies in this setting, σ and θ. These misspecifications could be modeled a number of different ways, but in this section, we have selected misspecified values of θ and σ that move the drift of the z_t in the same direction. The results from other specifications are shown in the appendix. In the results presented in this section, the agent initially believes that θ is larger than the true value and that σ is smaller than the true value. Specifically, in period one, $\theta_g = 0.25$ while $\theta = 0.105$ and $\sigma_g^2 = 0.008$ when $\sigma^2 = 0.015$.

To test how a learning process could evolve in this environment we first introduce an exogenous learning process. Since the process is exogenous, the agents will repeatedly solve the steady state of the HJB with different amounts of information at each period. In each one of these periods, there is a chance that the agents will have a chance to gain more information in the form of noisy observations of the true parameter values. In this model these noisy observations will be of the form,

$$\tilde{\theta}_i = \theta + \epsilon_{i,\theta}, \quad \epsilon_\theta \sim N(0, 0.1) \tag{37}$$

$$\tilde{\sigma}_i^2 = \sigma^2 + \epsilon_{i,\sigma}, \quad \epsilon_\sigma \sim N(0, 0.01) \tag{38}$$

The information will be given to an agent based on a draw from a standard Bernoulli distribution, and the agents will update their estimate of both parameters using the following equations

$$\theta_{g,i+1} = \theta_{g,i} + 0.01(\tilde{\theta}_i - \theta_{g,i}),$$

$$\sigma^2_{g,i+1} = \sigma^2_{g,i} + 0.01(\tilde{\sigma}^2_i - \sigma^2_{g,i}).$$

In this problem θ and σ are the true values of the parameters, and i is an index for the updating period. These parameters are updated according to the algorithm above and then used to calculate the steady state of our system, this steady state algorithm is described in the appendix (Achdou *et al.* , 2014).

5.1.2 Productivity Process Results

Below are the convergence results for the stylized learning rule in this setting. The following figure displays the value function over z and k. Looking at the convergence in the value function over z for a median value of k we can see clear convergence, here our value function starts off flat and develops the correct slope and curvature as our updating procedure continues. However, after $1,000$ periods we are still some distance from the true value function. Convergence over k for a median value of z is less interesting. In this case there is appropriate convergence; however, the difference between the misspecification and the true value is much smaller than in the z dimension.

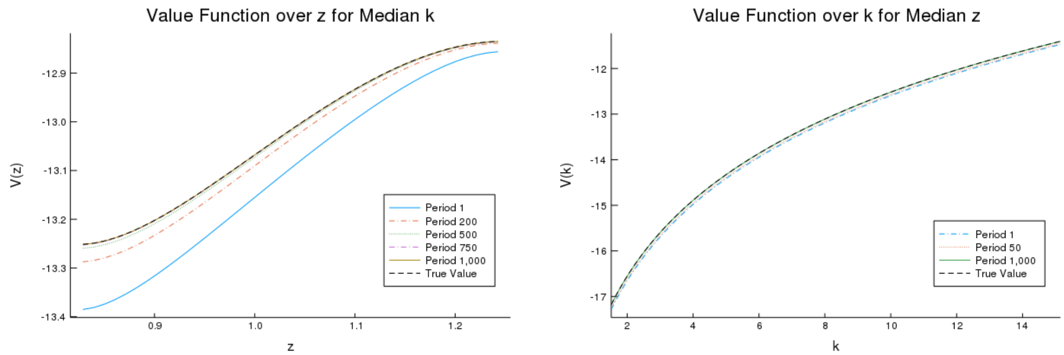

(a) (b)

The misspecified parameters, θ and σ, converge as we would thought. Below is a graph of the values of σ and θ at each period, including those in which the system does not update.

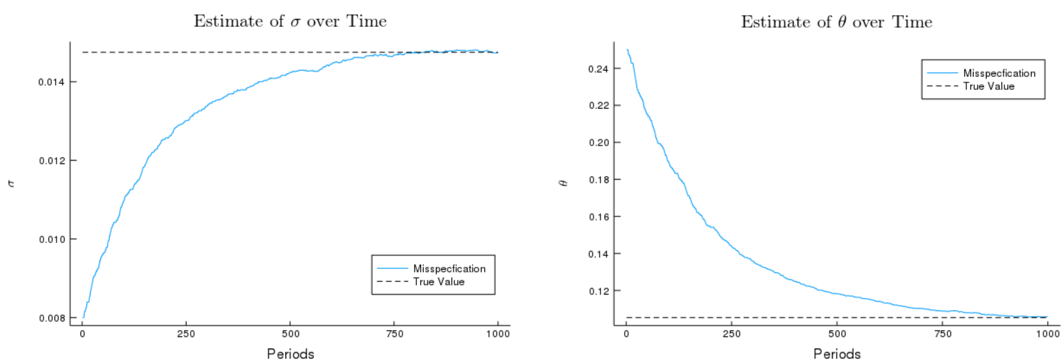

(a) (b)

This exercise displays the type of convergence we would have predicted.. Thus, we expect that learning rules would perform in a predictable manner in a stochastic continuous time setting.

5.2 Learning the Process for Capital

After examining the stylized learning rule's impacts on a model with a misspecified exogenous process, we investigate a model with a misspecified endogenous process. In this model, we have a diffusion process that summarizes the evolution of capital stock. Misspecification in this diffusion process impacts optimal savings and therefore the optimal consumption choice in the model. Thus, an incorrect specification of this process directly impacts our equilibrium choices. Furthermore, a poor consumption choice directly impacts the drift term in our diffusion process.

In our endogenous process model, there is a representative agent that makes consumption choices c and has capital stock k. The state of the economy depends on the flow of capital stock. The agent has standard preferences over utility flows based on capital discounted at rate $\rho \geq 0$. This can be written as the following equation:

$$\mathbb{E}_0 \int_{t=0}^{\infty} e^{-\rho t} u(c_t) dt$$

Here consumption, $c_t \geq 0$ for all periods. The agent's capital stock will evolve according to the following stochastic process used in Merton (1975a). This change has been made so that we can model learning with stochastic process capital. The earlier specification where our stochasticity came from z_t is more common in the literature. In this setting, capital will follow the stochastic process

$$dk_t = \big(f(k_t) - (\delta + n - \sigma^2)k_t - c_t\big)dt + \sigma k_t dW_t.$$

Here n measures the growth of the work force and dW_t is the increment of a Wiener process. In this setting, $f(k_t) - (\delta + \sigma^2)k_t - c_t$ summarizes the drift of capital and σk_t describes the variance.

5.2.1 Stationary Equilibrium

Stationary equilibrium in this setting will be given by several equations. The HJB for this problem will be

$$\rho V(k) = \max_{c} \ u(c) + V'(k) \cdot \big(f(k) - (\delta + n - \sigma^2)k - c\big) + \frac{1}{2}V''(k) \cdot (\sigma k).$$

The derivation of the HJB can be found in the appendix. This system will be defined on $(\bar{k}, \infty)$ where $\bar{k}$ is the value of capital at which the agent would consume nothing.

The agents in this simple model hold an incorrect belief about the diffusion process for capital stock. In this setting with exogenous learning they predict that the diffusion process is given by equation (39).

$$dk_t = \big(f(k_t) - (\delta + n - \sigma_g^2)k_t - c_t\big)dt + \sigma_g k_t dW_t$$

In this model the agent believes that the parameter σ is smaller than it should be, $\sigma_g < \sigma$. Specifically, $\sigma_g = 0.02$ when the true value $\sigma = 0.5$. With this misspecification, the agent believes the drift is larger than it should be *and* the variance is smaller than the true variance of the process. Other misspecifications for this process were examined, these results are in the appendix.

To test how a learning process could evolve in this environment we first introduce a stylized learning process. Since the information gain is exogenous the agents will repeatedly solve the steady state of the HJB with different amounts of information at each period. In each one of these periods, there is a chance that the agents will have a chance to gain more information in the form of a noisy observation of the true parameter estimate. The noisy parameter estimate will take the form,

$$\tilde{\sigma}_i = \sigma + \epsilon_i, \quad \epsilon_i \sim N(0, 0.1). \tag{39}$$

The information will be given to the agent based on draw from a standard Bernoulli distribution and the agents will update their estimate of σ_g according to

$$\sigma_{g,i+1} = \sigma_{g,i} + 0.01(\tilde{\sigma}_i - \sigma_{g,i}).$$

Here i is the index for the updating period and this updating process will continue for 1,000 periods.

5.2.2 Capital Process Results

Below are several results, the figures on the left show all the output from all 1,000 iterations of the endogenous learning algorithm. Figures to the right display select output from different periods of the iteration.

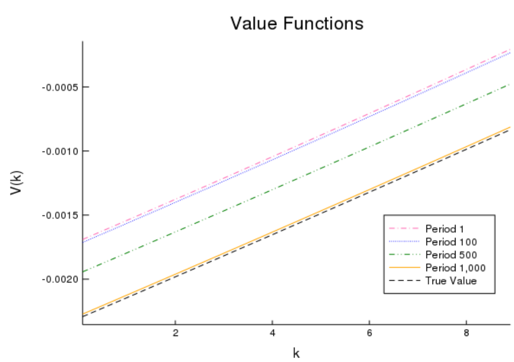

(a)

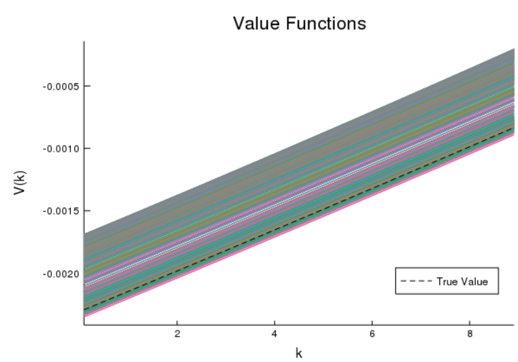

(b)

Using the learning rule, the value function converges to the true estimate over time. In this setting convergence is slow and even after 1,000 periods, the value function is still a small distance from the true value. Convergence is equally slow for some measures such as savings.

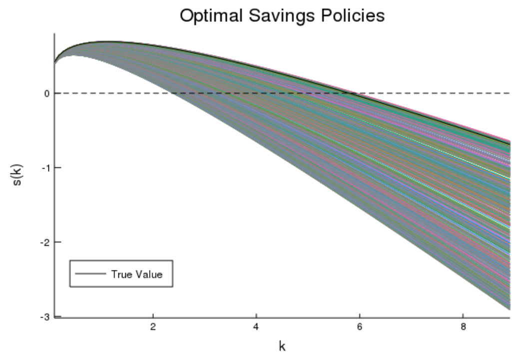

(a)

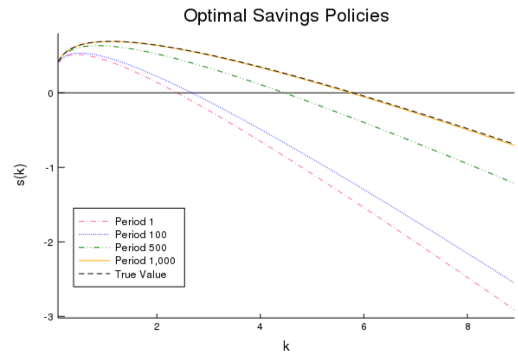

(b)

From these figures, we can see that the savings policies appear to converge more quickly to the true policy than the value functions converge to the true steady state estimates. This is likely due to the fact that optimal savings policies don't depend as strongly on the parameter σ. While σ does impact the calculations of the savings policies it is only one part of savings decision. This parameter impacts the value function more directly since it will affect the evolution of the system and the algorithm's choice of implementing a forward difference or backward difference for calculating the derivative of the value function.

Our prediction of σ converges in an expected way. We can see this in the graph below, which verifies that our updating rule works as expected. After 1,000 iterations the guess for σ is 0.005 away from the true parameter value, this is why our value functions and optimal savings policies have not completely converged to their true values.

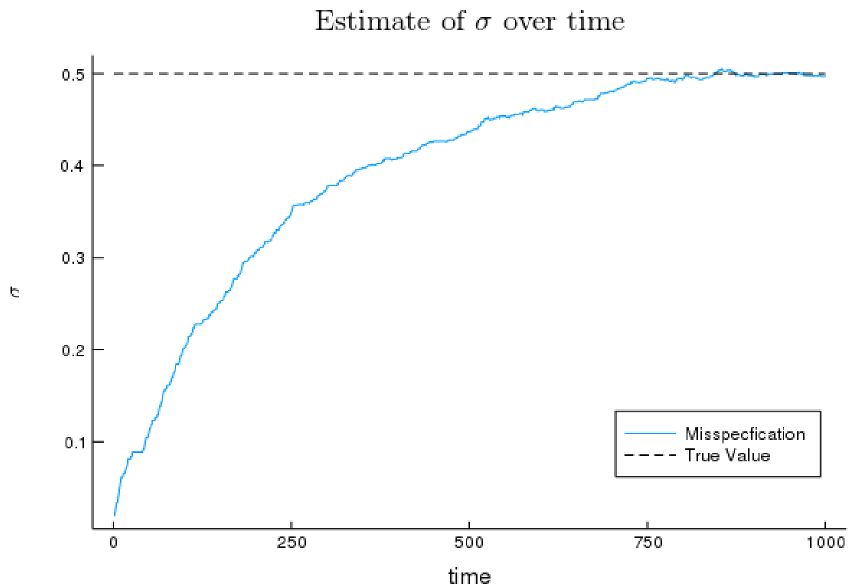

Figure 5

5.3 Learning Using Real-Time Updating

In this next section, we will explore a modified model with a stochastic process for productivity. In this model, agents will observe the process over time and update their parameter estimates based on these observations. Agents will maximize utility according to

$$\mathbb{E}_0 \int_{t=0}^{\infty} e^{-\rho t} u(c_t) dt.$$

Here productivity will evolve according to the same process as before, (35). Production will still be a standard Cobb-Douglas function used in previous sections. This means that $\log(z_t)$ is evolving according to an Ornstein-Uhlenbeck process, the continuous-time analog of an AR(1) process. Defining the process for z_t this way avoids negative values for z_t. Looking more closely at the $\log(z_t)$ process we have figure 6.

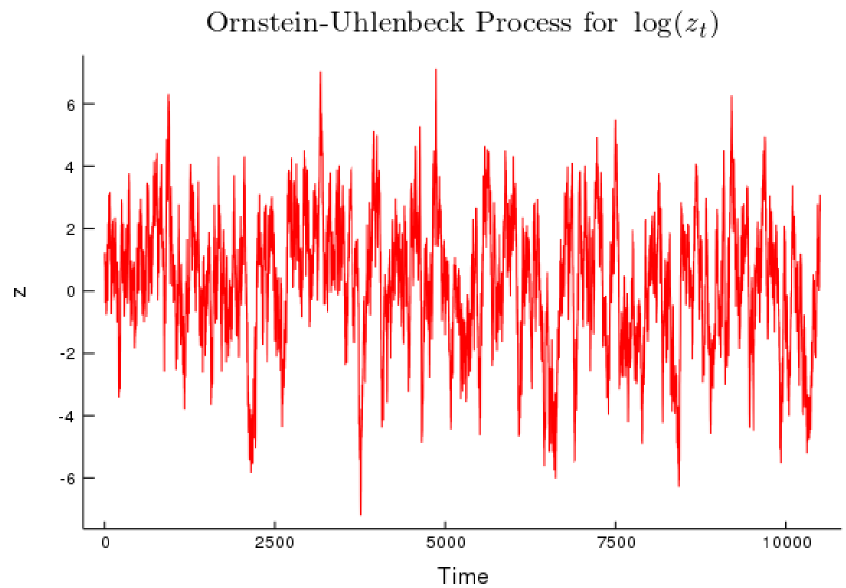

Figure 6

This process does have negative values, but the process for z_t will not.

5.3.1 Real Time Updating of Parameter Estimates

The HJB for this stochastic Ramsey model will be

$$\rho V(k,z) = \max_c \; u(c) + \partial_k V(k,z)\cdot\big(zf(k) - \delta k - c\big) + \\ \partial_z V(k,z)\cdot\left(-\theta\log(z) + \frac{\sigma^2}{2z_t}\right) + \partial_{zz}V(k,z)\cdot\frac{1}{2}\sigma^2 z^2,$$

in this setting our parameter for σ will be set equal to one. Setting $\sigma = 1$ will not only simplify our updating problem, but it will also allow for a more intuitive connection between our Ornstein-Uhlenbeck process and an AR(1) process.

In this model, agents believe that the stochastic process for productivity evolves according to

$$d\log(z_t) = -\theta_g \log(z_t)dt + dW_t.$$

Where θ_g is the agent's forecast for the process's parameter θ. Before the agents in this model begin trying to solve their value function problem, they look at the first 100 observations of the process and using ordinary least squares (OLS) predict a value for θ and a possible constant.

In this setting, the agent can use OLS to predict an initial value for θ_g, since the process for $\log(z_t)$ can be rewritten as a discrete AR(1) process using the Euler-Maruyama method. Applying this method the AR(1) process for $\log(z_t)$ will be,

$$\log(z_{t+\Delta t}) = (1 - \theta_g \Delta t)\log(z_t) + \varepsilon_t \sqrt{\Delta t}$$

for simplicity we will assume that the agent estimates these parameters as if Δt is observable.

Next, they use the finite difference algorithm described in the appendix. They implement this algorithm 10,000 times, each time they observe several additional values of the productivity process. Therefore, in this setting, we should think of the updating periods as independent intervals of time that each contains several observations. Next, using recursive least squares (RLS), the agent updates their parameter estimates. This RLS formula is given by,

$$R_{g,t+1} = R_{g,t} + \gamma_t (xx'\Delta t - R_{g,t})$$
$$\phi_{g,t+1} = \phi_{g,t} + \gamma_t R^{-1}_{g,t+1} \cdot x(y - x'\phi_{g,t})\Delta t$$

here all variables with a g subscripts represent the agent's forecast x and y are matrices that contain value of x_t and y_t for all points between $t-1$ to t and t to $t+1$ respectively. The number of points in each of these intervals will depend on dt. In the results below the agent observes 5 points of the process in each updating period, this means that after 100 periods the agent has 500 new points on which to base their estimates. This has been done in order to maintain continuity in the time dimension. Additionally, x_t and $\phi_{g,t}$ are defined as

$$x_t = \begin{bmatrix} 1 \\ \log(z_t) \end{bmatrix}, \quad \phi_{g,t} = \begin{bmatrix} c_{g,t} \\ 1 - \theta_{g,t}\Delta t \end{bmatrix},$$

where $c_{g,t}$ is our estimate for a constant in the model. The agent uses this formula to update parameter estimates and then reruns the finite difference algorithm, this is done $10,000$ times.

5.3.2 Real Time Updating Results

Some of the results from the forecasting model resemble the results from previous sections. In this setting, value functions converge quickly in the k dimension and more slowly in the z dimension. This is in line with the results from before and makes sense as the misspecification is for the process that governs z.

First, we will look at results for an algorithm where the gain $\gamma_t = \frac{1}{t}$, this means that the agent discounts the information in each updating period by $\frac{1}{t}$. Here t represents the updating period that the agent is in.

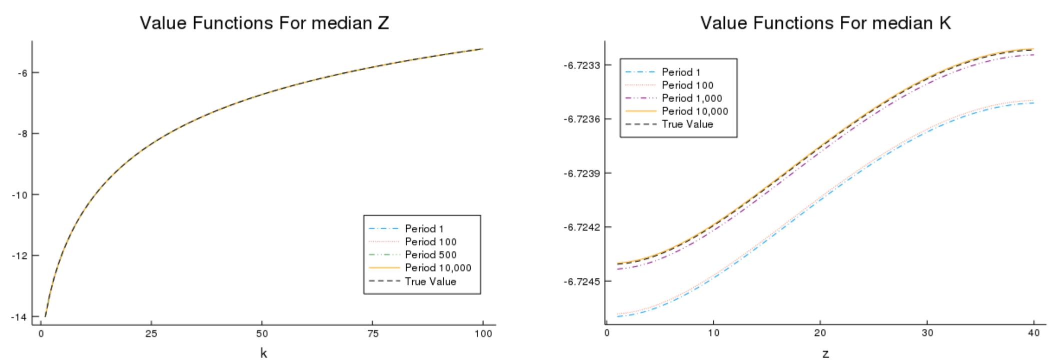

(a) (b)

Figure 7

We can take a closer look at convergence in this setting by examining our parameter estimates over time.

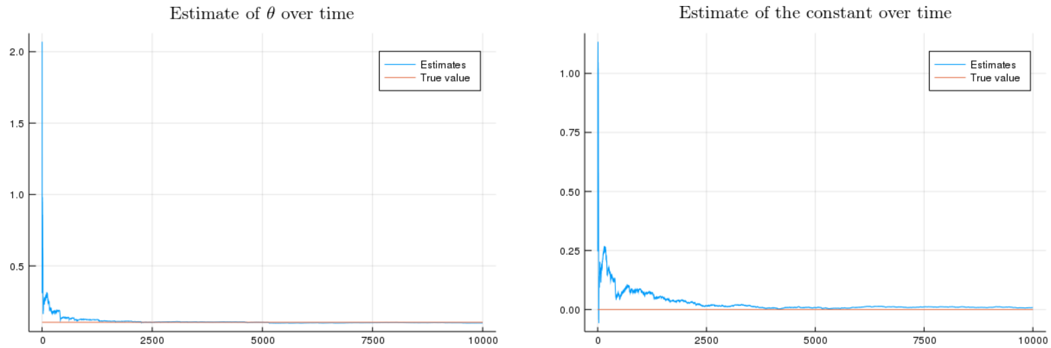

(a) (b)

Figure 8

Looking at the results above we can see that convergence in this setting is fast. Despite starting from incorrect parameter values, θ and the constant are close to their true parameter values after 200 periods.

We can also examine this real-time updating rule with a constant gain. Here we set the gain $\gamma_t = 0.01$ for all time periods. The value functions converge similarly to the decreasing gain case as seen below.

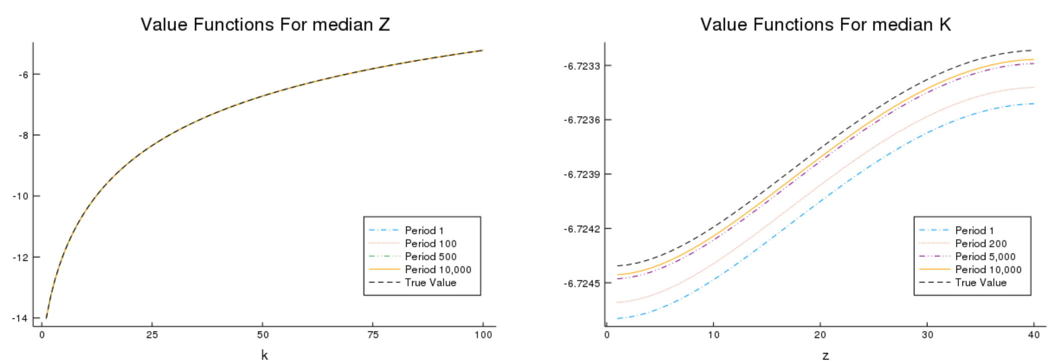

(a) (b)

Figure 9

We can again examine the convergence of θ_g and the estimate for the constant over time.

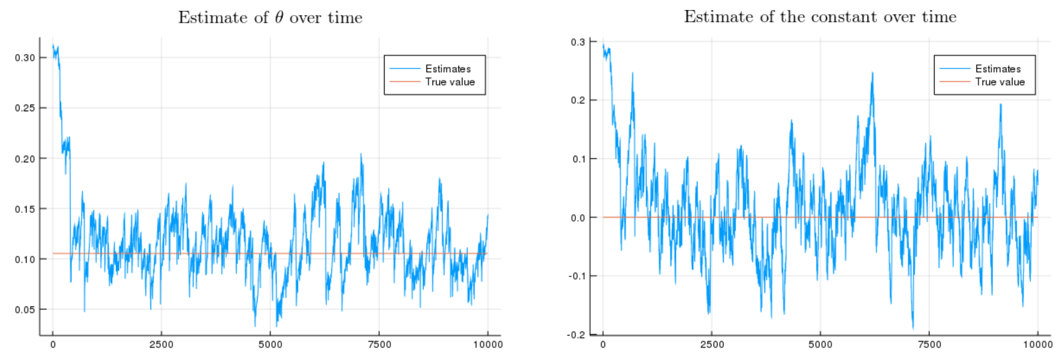

(a) (b)

Figure 10

Since, we are using a constant gain algorithm there is noise in our parameter estimates even after many periods. Constant gain algorithms place equal emphasis on all observed points from the Ornstein-Uhlenbeck process, since this is a noisy process we will see our estimates trend about the correct parameter value instead of directly to the correct value. Due to this, it is helpful to examine the mean estimates of θ and the constant over time.

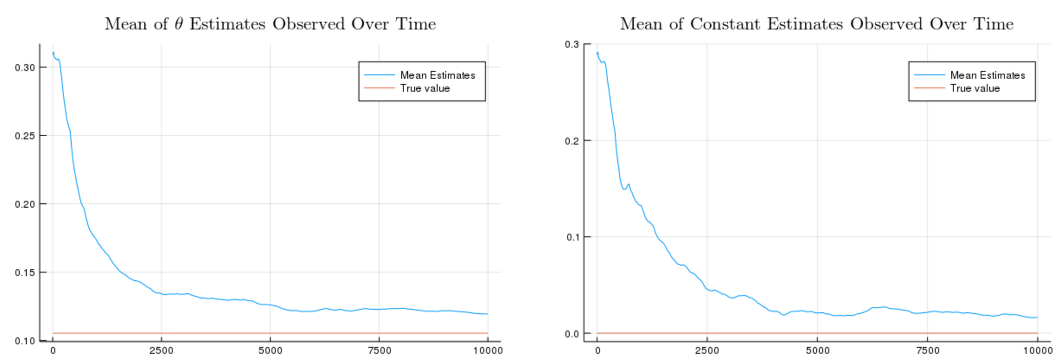

(a) (b)

Figure 11

Here the mean estimates of θ and the constant are approaching the true parameter values.

5.4 Summary

Our exogenous learning rules perform well in the stochastic continuous-time steady state calculations. This is encouraging because it means that we can expect some of the familiar results from discrete time learning to carry over in our continuous setting. Although the results in this section are not particularly stunning there are several extensions to this simple learning rule that may yield more interesting results. Looking at this exogenous learning rule in a heterogeneous agent setting may allow for more feedback through the system KF equation, thus yielding less predictable results. A heterogeneous agent model creates this additional feedback through internal pricing frictions that do not exist in our representative agent model.

The performance of the forecasting rule demonstrates that using adaptive learning techniques over intervals of time works well. This method may be beneficial for future work, as it provides a clear link between discrete RLS methods and the continuous-time framework. Despite using different methodology it appears that the forecasting rule in section 5.3 and the exogenous learning rule in section 5.1 have similar convergence results, this is an interesting result that may be due to the model similarities in these sections.

6 Conclusion

This paper serves a primer on continuous-time modeling and adapting discrete adaptive learning methods to continuous-time. The mathematical results in section 4 link discrete models to their continuous-time counterparts. Section 5 contains some basic results for simple learning method applied to continuous-time models. Using the results of this paper we can conclude that the continuous-time framework is comparable to discrete-time and that learning algorithms can be adapted and form well in this setting. Future extensions to work could include implementing a continuous-time version of recursive least squares to simple continuous-models and creating a learning algorithm with more feedback in a representative agent model. There remains much to do in order to properly modify adaptive learning techniques to a continuous-time setting.

A The Kolmogorov Forward Equation

The derivation of the KF equation is not always intuitive. Dixit (1992) gives one of the clearest derivations of the KF equation targeted at economists. In this next section, we will present this derivation and compare the KF equation to discrete distributional methods. The Kolmogorov forward and backward equations govern the more general dynamics of stochastic processes. Suppose we are in a discrete system at a point $(x_1, t_1 + \Delta t)$ there two ways we could have gotten to this point. First, we could have previously been at $(x_1 - \Delta h, t_1)$ before moving forward in the x direction. Alternatively, we may have been located at $(x_1 + \Delta h, t_1)$ and then moved back in the x direction. Using this information we can write the probability of of being at (x_1, t_1) as,

$$\Pi(x_1, t_1 + \Delta t) = p\Pi(x_1 - \Delta h, t_1) + q\Pi(x_1 + \Delta h, t_1)$$

in this equation p is the probability of moving forward in the x direction and $q = 1-p$ is the probability of moving backward. For a Brownian motion, $dx = \mu dt + \sigma dW_t$, $p = \frac{1}{2}[1 + \frac{\mu}{\sigma^2}\Delta h]$ and $q = \frac{1}{2}[1 - \frac{\mu}{\sigma^2}\Delta h]$. Using a Taylor expansion our previous expression will become,

$$\Pi(x_1,t_1)+\Pi_t(x_1,t_1)\Delta t+\mathcal{O}(\Delta t)=\frac{1}{2}[1+\frac{\mu}{\sigma^2}\Delta h]\big(\Pi(x_1,t_1)-\Pi_x(x_1,t_1)\Delta h$$
$$+\frac{1}{2}\Pi_{xx}(x_1,t_1)(\Delta h)^2+\mathcal{O}(\Delta h)^2\big)+\frac{1}{2}[1-\frac{\mu}{\sigma^2}\Delta h]\big(\Pi(x_1,t_1)$$
$$-\Pi_x(x_1,t_1)\Delta h+\frac{1}{2}\Pi_{xx}(x_1,t_1)(\Delta h)^2+\mathcal{O}(\Delta h)^2\big).$$

As $\Delta t \to 0$ this equation will become

$$\Pi_t(x_1,t_1)=\frac{1}{2}\sigma^2\Pi_{xx}(x_1,t_1)-\mu\Pi_x(x_1,t_1), \tag{40}$$

our standard KF equation. This derivation is less intuitive and not as straight forward for other stochastic processes.

B Deriving Equation (12)

First simplifying the original equation (10) we get,

$$\beta^{\Delta t}\lambda_{t+\Delta t}[\big(f'(k_{t+\Delta t})-\delta\big)\Delta t+1]=\lambda_t \tag{41}$$

Then we can set $\lambda_{t+\Delta t}=\lambda_t+\dot{\lambda}\Delta t$ and use the expansion $\beta^{\Delta t}=1+\Delta t\ln\beta$

$$[1+\Delta t\ln\beta][\lambda_t+\dot{\lambda}\Delta t][\big(f'(k_{t+\Delta t})-\delta\big)\Delta t+1]=\lambda_t \tag{42}$$

Foiling this out yields the following.

$$[\lambda_t\Delta t\ln(\beta)+\lambda_t+\dot{\lambda}\Delta t+\dot{\lambda}(\Delta t)^2\ln(\beta)][f'(k_{t+\Delta t})-\delta]\Delta t=\lambda_t-\lambda_t-\lambda_t\Delta t\ln(\beta)-\dot{\lambda}\Delta t-\dot{\lambda}(\Delta t)^2\ln(\beta) \tag{43}$$

Diving through by Δt and then assuming any remaining terms with Δt are negligible we will get equation (12).

$$\lambda_t[f'(k_{t+\Delta t}) - \delta] = -\lambda_t \ln \beta - \dot{\lambda} \tag{44}$$

C Steady State Algorithm for solving the HJB

The steady-state algorithm used in section 5 of this paper comes from Achdou *et al.* (2014) and is one of the more simple solution methods in this setting.

For a simple Ramsey model as described in section 5, the algorithm proceeds as follows,

1. Compute $\partial_k V(\cdot)$ for all k
2. Compute the value of consumption from $c_i = (u')^{-1}[\partial_k V(\cdot)]$
3. Implement an upwind scheme to find "correct" $\partial_k V(\cdot)$
4. Using the coefficients determined by the upwind scheme create a transition matrix for this system
5. Solve the following system of non-linear equations

$$\rho V^{n+1} + \frac{V^{n+1} - V^n}{\Delta} = u(V) + A^n V^{n+1}$$

6. Iterate until $V^{n+1} - V^n \approx 0$

For the most part, the algorithm described above is a typical finite difference scheme. The main difference between this algorithm and what is often used for value function iteration is the upwind scheme. The upwind scheme described in this paper selects a forward difference when we experience positive drift, i.e. positive savings, in our variable of interest, a backward difference if this drift term is negative, and selects a steady state value if we see no drift. In this scheme, we continue our difference algorithms for $(n+1)$ iterations until we are no longer significantly updating our value functions.

Now, we will describe the upwind scheme in more detail. For the algorithm described above, we need to approximate three different derivatives, the backward and forward difference of the first derivative of the value function and the second derivative for the value function.

The forward difference will be given by,

$$\frac{V_{i+1} - V_i}{\Delta k}$$

and the backward difference will be

$$\frac{V_i - V_{i-1}}{\Delta k}.$$

The second derivative will be approximated by

$$\frac{V_{i+1} - 2V_i + V_{i-1}}{(\Delta k)^2},$$

where i represents the point in the k grid-space. When the drift of the state variable is positive the upwind scheme will choose the forward difference and when it is negative the upwind scheme will select the backward differences. If neither of these conditions holds, then the upwind scheme will select a steady state value.

There are several different ways to explain the upwind scheme. We can think of it as a method for consistent estimation in this setting. In this setting, we need our finite difference scheme to take the dynamics of our system into consideration.

Suppose we have the following HJB,

$$\rho V(k,z) = \max_{c}\ u(c) + \partial_k V(k,z)(f(k) - \delta k - c) - \partial_z V(k,z)(\eta z_t) + \frac{1}{2}\partial_{zz} V(k,z)\sigma^2$$

in order to approximate the derivatives of our values functions, we need to consider the flow of k and z. For z this is simple since the sign of $-\eta z_t$ will be the same for all positive values of z_t, we can use the backward difference at all points. This works as long as our z-grid contains only positive points. (The $\log(z_t)$ processes from earlier in this paper was used to help ensure we could use only one differencing method).

Since, our values of k cannot be similarly limited, especially since they rely on c, we need to use an upwind scheme in order to approximate the derivative in this dimension. Suppose we at one specific point in the k-dimension and we are unsure about the shape and differentiability of our value function. However, we do know that the drift of the k process has a positive drift at that value of k_i or in other terms, the savings function at k_i is positive. As discussed in Achdou *et al.* (2014), we can then what matters most is how our value function changes when capital increases by a small amount. Conversely, if savings are negative we want to measure how the value function changes when capital decreases by a small amount. This is our motivation for using the upwind scheme, this numerical approximation technique will take the forward difference when savings is positive and the backward difference when savings is negative.

It is worth noting that in fluid dynamics literature the upwind scheme is defined differently. In these works, the upwind scheme takes the forward difference when drift is negative and the backward difference when the drift is positive. This difference emerges because these systems of partial differentials are solved forward in time, whereas in this setting we in effect solving our system of equation backward in time. In the problems outlined in this paper, we are solving for the steady state of our system. Hence, we are effectively at $t = \infty$ meaning that our solution techniques can be thought of as working backward in time.

D Alternative Specifications

This section of the appendix outlines several different initial specifications that could have been used for the models in section 5 of this paper.

D.1 Learning the Process for Productivity

In section 5, the agents specified that θ was larger than its true value, 0.105 and σ was smaller than its true value, 0.015. Now, in the following sections, we will look at various misspecifications of these parameters and the convergence results. Below, is a table of the various initial values we examine in sections D.1.1-D.1.7.

Specification	θ_g	σ_g^2
Section 5	0.25	0.008
D.1.1	0.08	0.008
D.1.2	-0.11	0.008
D.1.3	2.0	0.008
D.1.4	0.25	0.8
D.1.5	0.25	1.5
D.1.6	-0.11	1.5
D.1.7	-0.11	0.8

Table 1: Initial values for σ and θ

D.1.1

We first examine what would happen to this model if θ was set to be smaller and the correct sign and if σ was also a smaller value. In this section the initial value for θ_g is 0.08 and the initial value for σ_g^2 is 0.008.

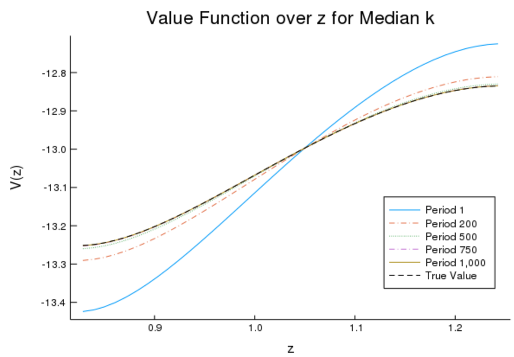

(a)

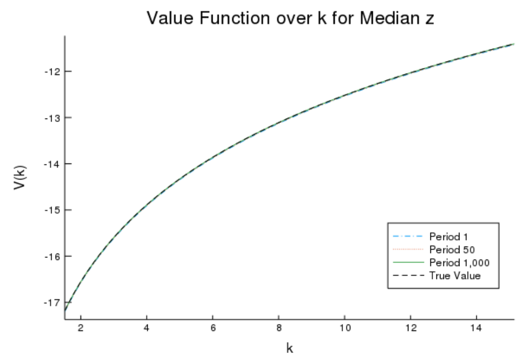

(b)

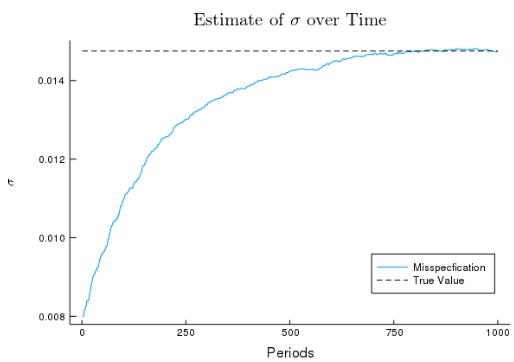

(a)

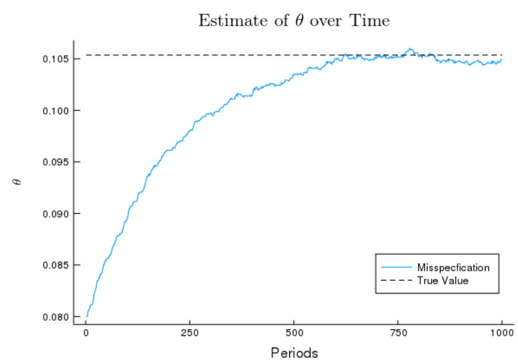

(b)

The key differences in this specification are in the value function convergence. In this setting the slope of the value function in the z dimension changes significantly as the parameters update over time.

D.1.2

Next, we examined convergence when the initial θ value was set to a negative value and left the value for σ smaller than the true value. In this section the initial value for θ_g is -0.11 and the initial value for σ_g^2 is 0.008.

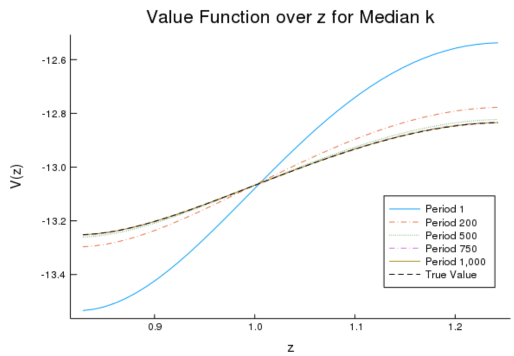

(a)

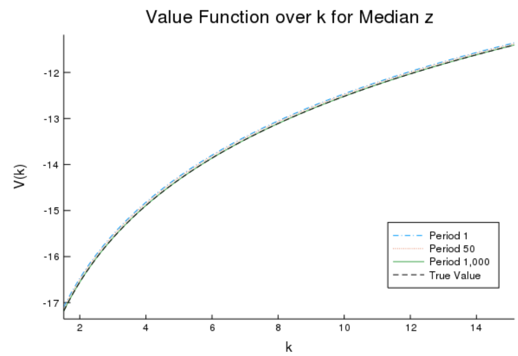

(b)

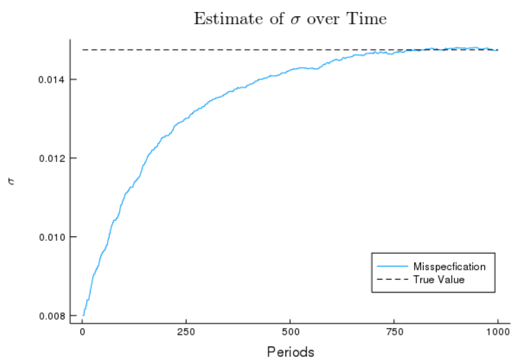

(a)

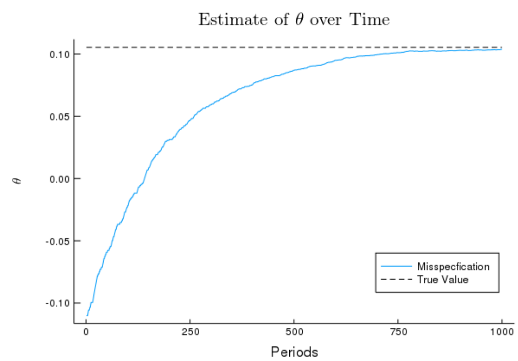

(b)

These results were similar to the previous specification's graphs.

D.1.3

The last value tested for θ was a much larger positive value, again σ was initialized with a value smaller than the true parameter value. In this section the initial value for θ_g is 2.0 and the initial value for σ_g^2 is 0.008.

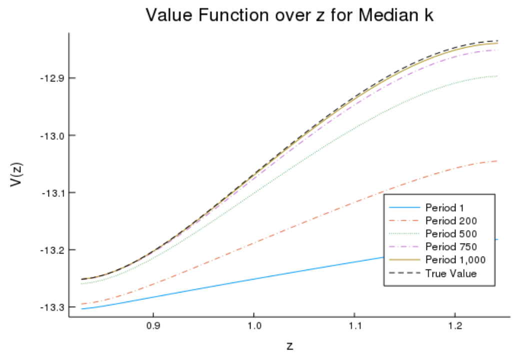

(a)

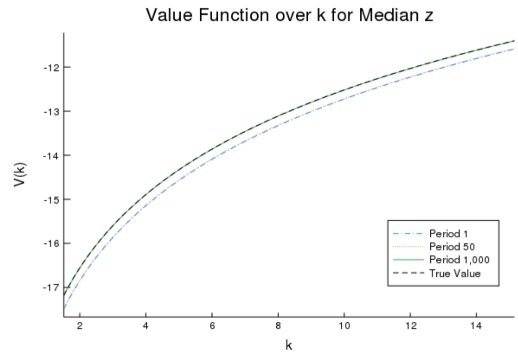

(b)

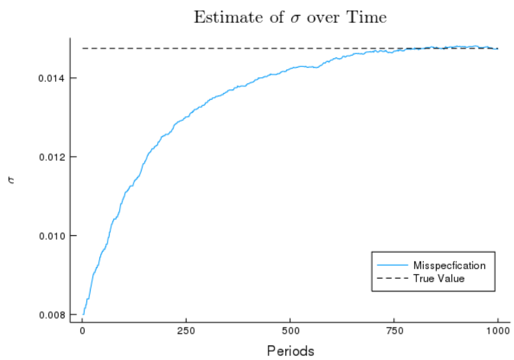

(a)

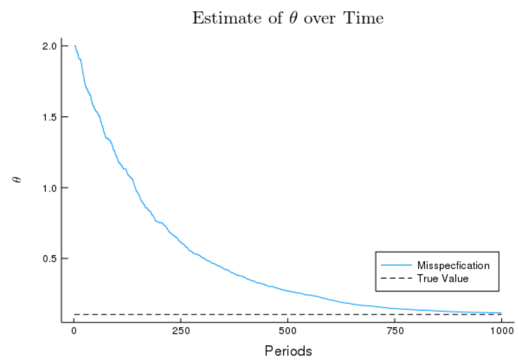

(b)

D.1.4

Next, different values for σ were explored. In the results below σ was set to be much higher than the original value but still less than one and θ was set to a larger value. Here the initial value for θ_g is 0.25 and the initial value for σ_g^2 is 0.8.

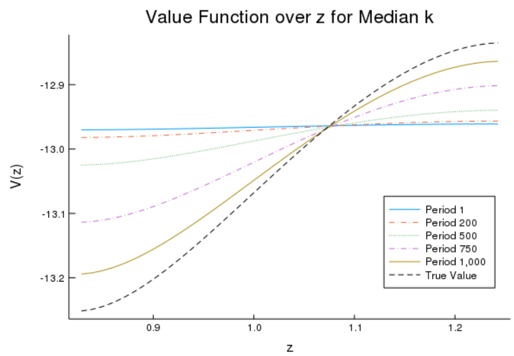

(a)

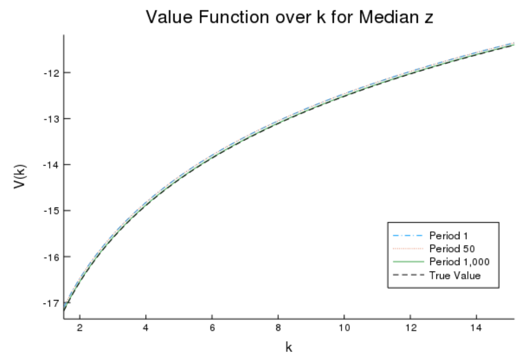

(b)

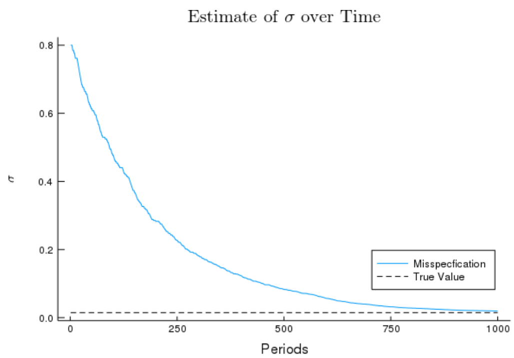

(a)

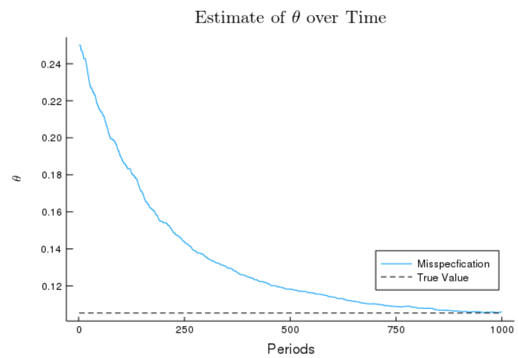

(b)

D.1.5

The same algorithm was run with a θ value that was much larger than the true value. In this section the initial value for θ_g is 0.25 and the initial value for σ_g^2 is 1.5.

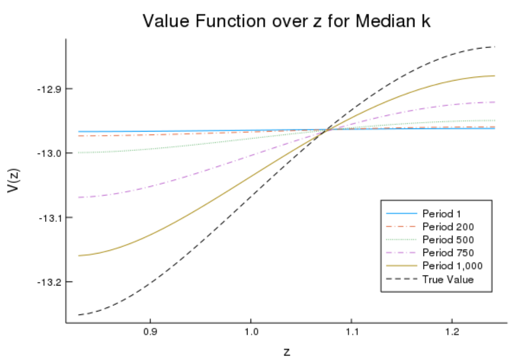

(a)

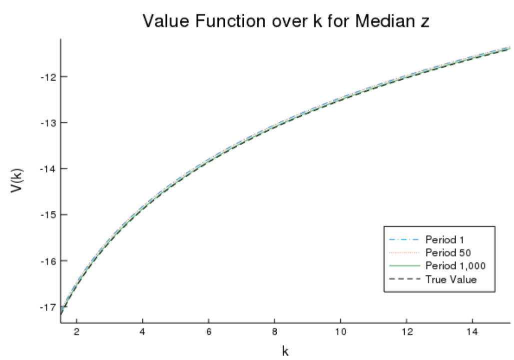

(b)

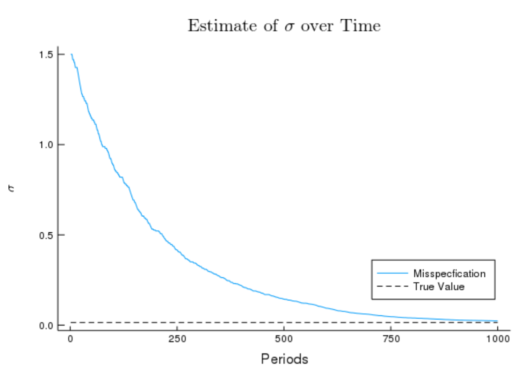

(a)

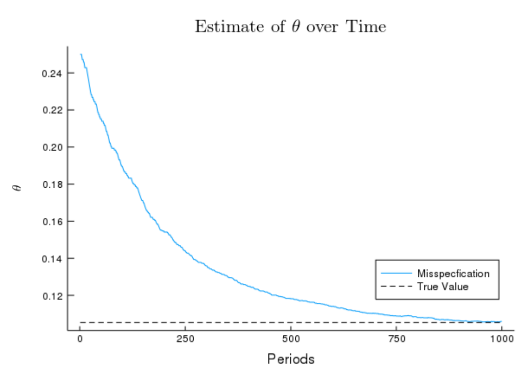

(b)

D.1.6

These initial values for σ were then run again with a small negative value for θ. In this section the initial value for θ_g is -0.11 and the initial value for σ_g^2 is 1.5.

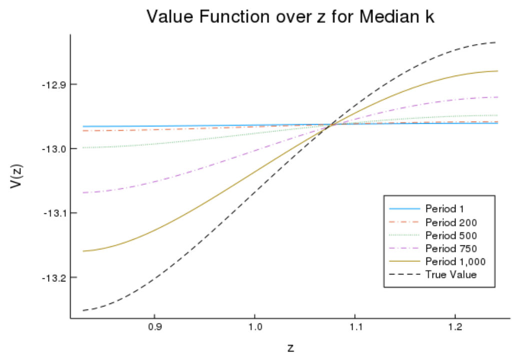

(a)

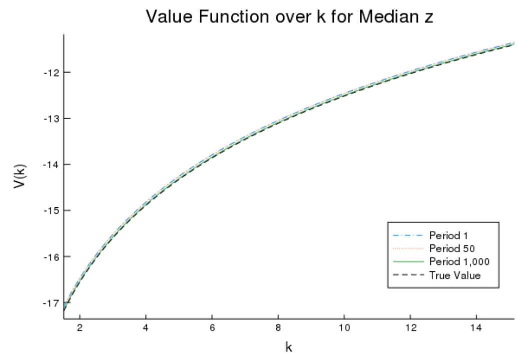

(b)

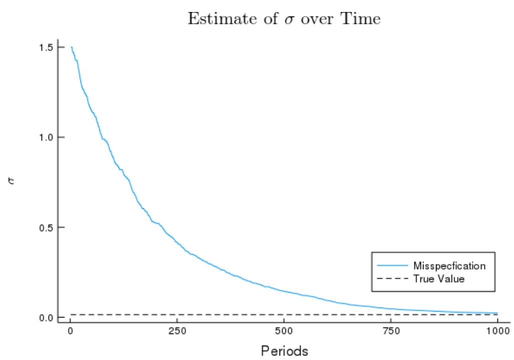

(a)

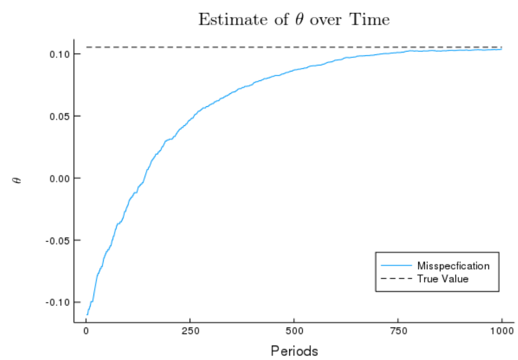

(b)

D.1.7

Last, we examine what would happen to this model if θ was set to be small and negative and if σ was a large value. In this section the initial value for θ_g is -0.11 and the initial value for σ_g^2 is 0.8.

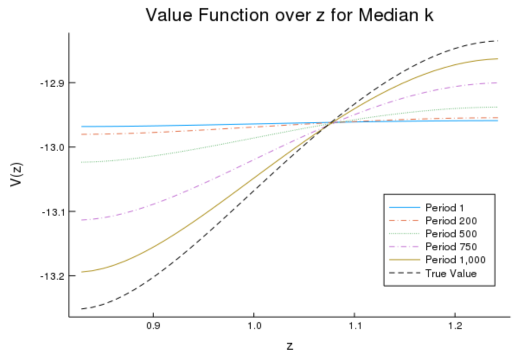

(a)

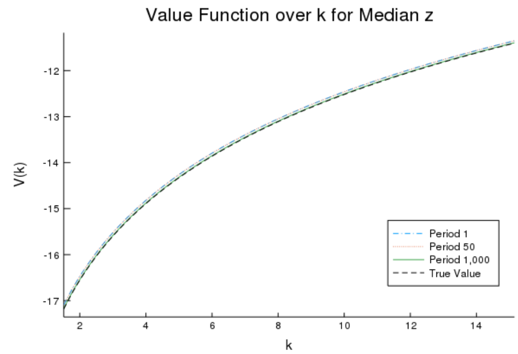

(b)

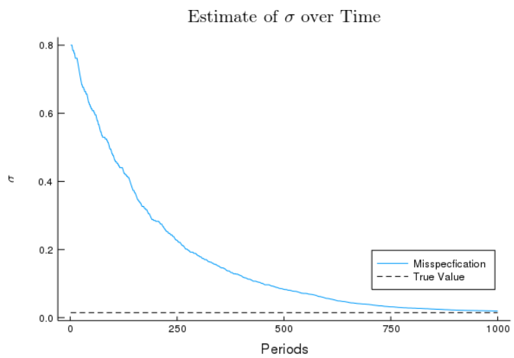

(a)

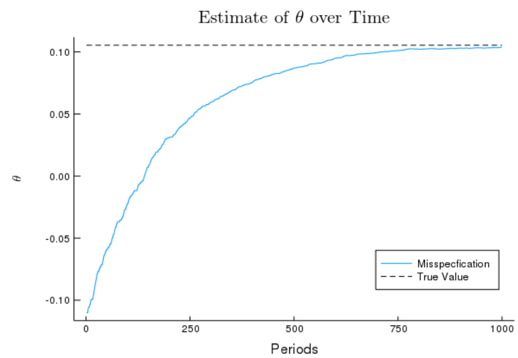

(b)

D.2 Learning the Process for Capital

Section five examined converge when σ_g was set to a lower initial value than the true parameter value, 0.5. In the following section we will explore different initial values for σ_g with varying signs and magnitudes. Below is a table of the initial values we will examine.

Specification	σ_g
Section 5	0.02
D.2.1	-0.02
D.2.2	8.0
D.2.3	-4.0

Table 2: Initial values for σ

D.2.1

In our first alternative misspecification we look at an initial value of σ_g that is the same magnitude as the correct value, but the incorrect sign. In this section $\sigma_g = -0.02$.

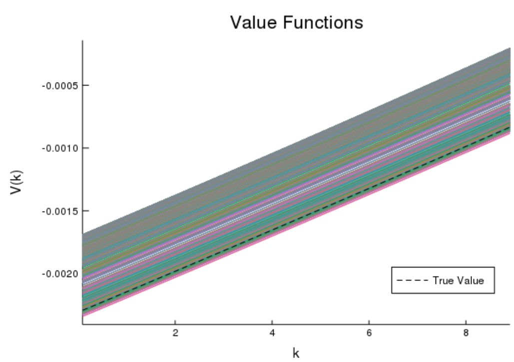

(a)

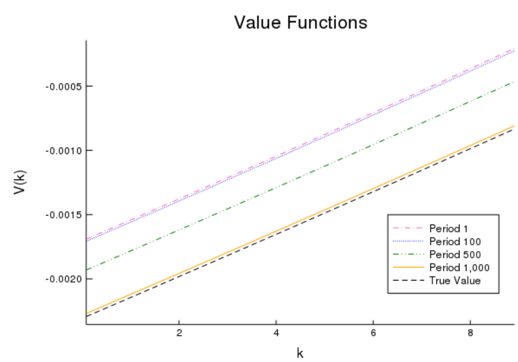

(b)

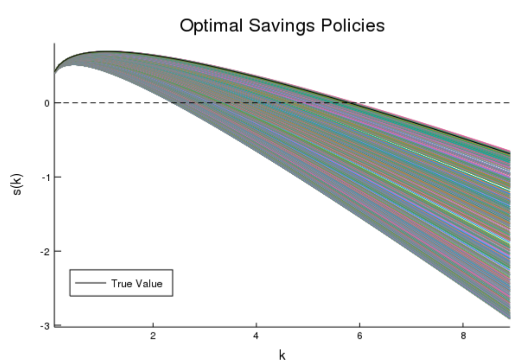

(a)

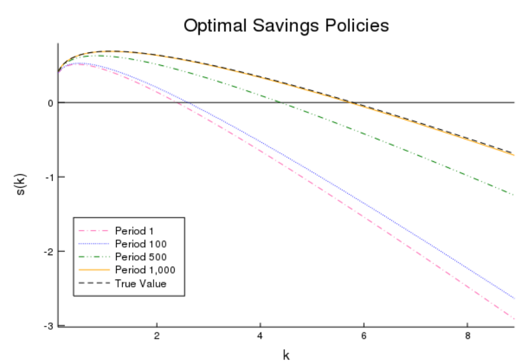

(b)

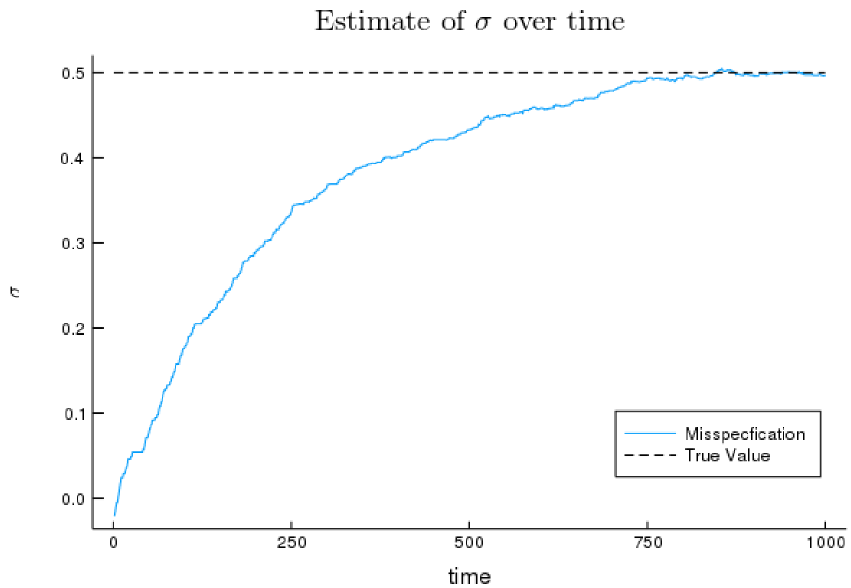

Figure 28

D.2.2

Next, we examine what would happen if the agents initial specification were much larger than the true value. Here the initial σ_g is 8.0.

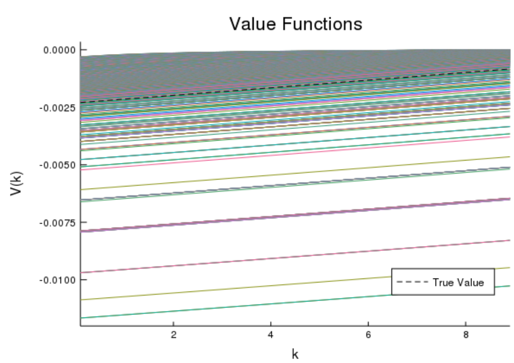

(a)

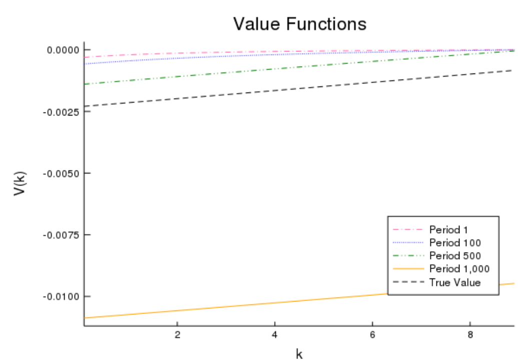

(b)

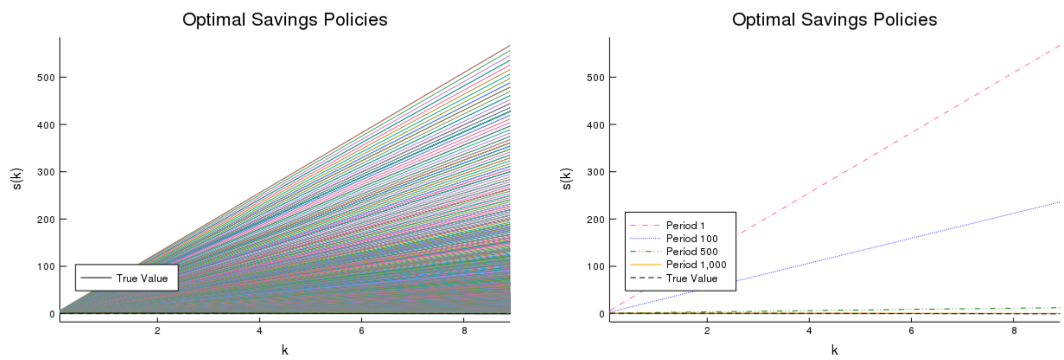

(a) (b)

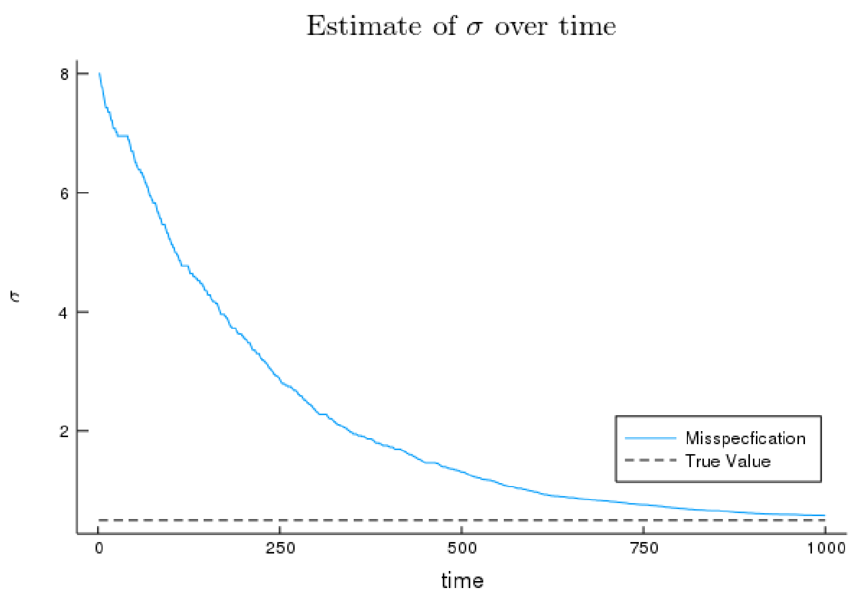

Figure 31

D.2.3

Last, we set the initial value for σ so that it is negative and has a large magnitude, $\sigma_g = -4.0$.

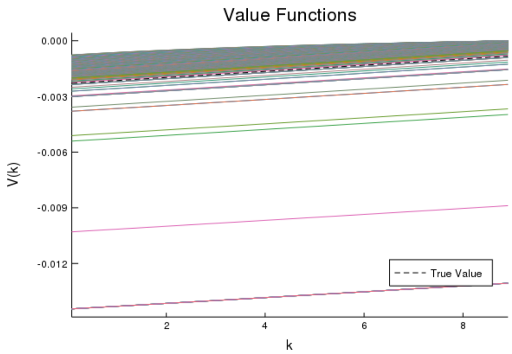

(a)

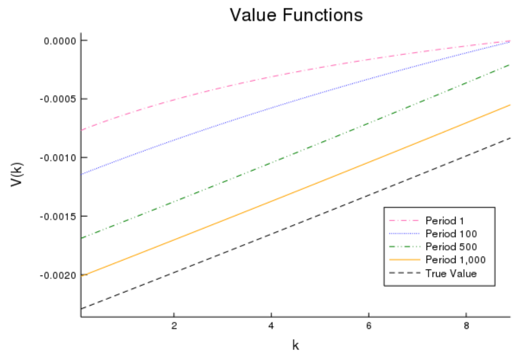

(b)

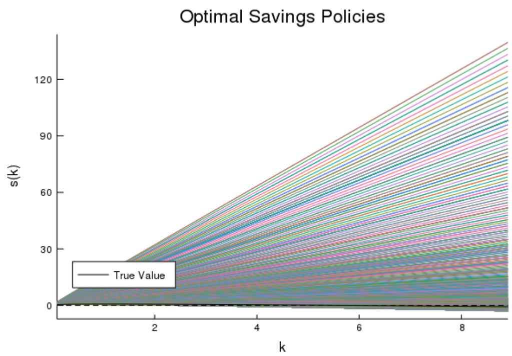

(a)

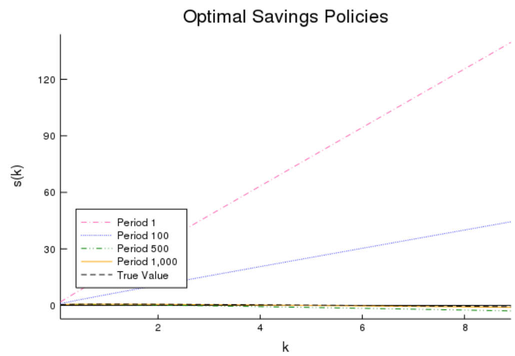

(b)

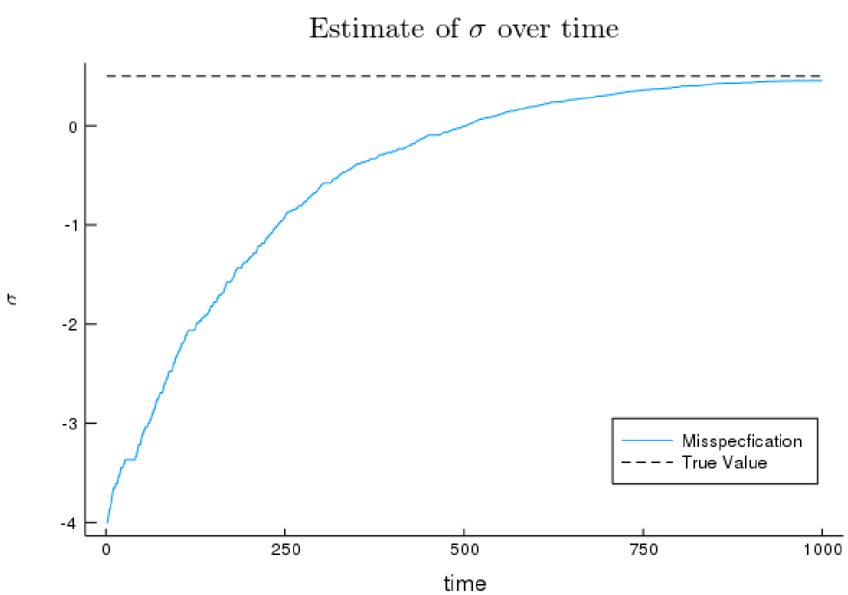

Figure 34

References

Achdou, Y, Buera, F, Lasry, J, Lions, P, & Moll, B. 2014. Partial Differential Equation Models in Macroeconomics. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, **372**(2028).

Achdou, Y, Han, J, Lasry, J, Lions, P, & Moll, B. Working. Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *Forthcoming Review of Economic Studies*, ().

Ahn, S, Kaplan, G, Moll, B, Winberry, T, & Wolf, C. 2018. When Inequality Matters for Macro and Macro Matters for Inequality. *NBER Macroeconomics Annual*, **32**(1).

Aït-Sahalia, Y. 2002. Maximum Likelihood Estimation of Discretely Sampled Diffusions: A Closed-Form Approximation Approach. *Econometrica*, **70**, 202–262.

Aït-Sahalia, Y. 2004. Disentangling Diffusion from Jumps. *Journal of Financial Economics*, **74**, 487–528.

Bhamra, Horjat, & Uppal, Raman. 2014. Asset Prices with Heterogeneity in Prices. *The Review of Financial Studies*.

Black, F, & Scholes, M. 1973. Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, **81**(3), 637–654.

Bray, M. 1982. Learning, Estimation, and the Stability of Rational Expectations. *Journal of Economic Theory*, **26**, 318–339.

Brock, W, & Mirman, L. 1972. Optimal Economic Growth and Uncertainty: The Discounted Case. *Journal of Economic Theory*, **4**(3), 479–513.

DeMarzo, PM, & Sannikov, Y. 2006. Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance*, **61**(6), 2681–2724.

Dixit, AK. 1989. Entry and Exit Decisions under Uncertainty. *Journal of Political Economy*, **97**(3), 620–638.

Dixit, AK. 1992. *The Art of Smooth Pasting*. STICERD - Theoretical Economics Paper Series. Suntory and Toyota International Centres for Economics and Related Disciplines, LSE.

Dixit, AK, & Pindyck, RS. 1994. *Investment under Uncertainty*. Princeton University Press.

Dorfman, R. 1969. An Economic Interpretation of Optimal Control Theory. *The American Economic Review*, **59**(7), 817 – 831.

Duffie, D, & Epstein, LG. 1992. Stochastic Differential Utility. *Econometrica*, **60**(2), 353–394.

Eaton, J. 1981. Fiscal Policy, Inflation and the Accumulation of Risky Capital. *Review of Economic Studies*, **48**(3), 435–445.

Evans, G, & Honkapohja, S. 2009. Learning and Macroeconomics. *Annual Review of Economics.*

Evans, George W., & Honkapohja, Seppo. 2001. *Learning and Expectations in Macroeconomics.* Princeton University Press.

Forsyth, PA, & Labahn, G. 2007. Numerical Methods of Controlled Hamilton-Jacobi-Bellman PDEs in Finance. *Journal of Computational Finance*, **11**(2), 1–44.

Gabaix, X. 2009. Power Laws in Economics and Finance. *Annual Review of Economics*, **1**, 255–293.

Hansen, LP, & Sargent, TJ. 2019a. Macroeconomic Uncertainty Prices when Beliefs are Tenuous. *NBER working paper.*

Hansen, LP, & Sargent, TJ. 2019b. Structured Uncertainty and Model Misspecification. *University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2018-77.*

Hansen, LP, & Scheinkman, JA. 1995. Back to the Future - Generating Moment Implications for Continuous-Time Markov Processes. *Econometrica*, **63**, 767–804.

Hansen, LP, Sargent, TJ, Turmuhambetova, G, & Williams, N. 2006. Robust Control and Model Misspecification. *Journal of Economic Theory*, **128**(1), 45–90.

Kaplan, G, Moll, B, & Violante, GL. 2018. Monetary Policy According to HANK. *American Economic Review*, **108**(3), 697–743.

Merton, RC. 1969. Lifetime Portfolio Selection Under Uncertainty - Continuous-Time Case. *Review of Economics and Statistics*, **51**(3), 247–257.

Merton, RC. 1971. Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory*, **3**(4), 373–413.

Merton, RC. 1975a. An Asymptotic Theory of Growth Under Uncertainty. *The Review of Economic Studies*, **42**(3), 375–393.

Merton, RC. 1975b. An Asymptotic Theory of Growth Under Uncertainty. *Review of Economic Studies*, **42**(3), 375–393.

Merton, RC. 1992. *Continuous-Time Finance*. Macroeconomics and Finance Series. Wiley.

Mirman, LJ. 1973. Steady State Behavior of a Class of One Sector Growth Models with Uncertain Technology. *Journal of Economic Theory*, **6**(3), 219–242.

Mirrlees, J.A. 1966. Optimum Accumulation Under Uncertainty. *Econometrica*, **34**(5S), 26–&.

Mirrlees, JA. 1971. *Optimum Accumulation Under Uncertainty: the Case of Stationary Returns to Investment*. London: Palgrave Macmillan UK. Pages 36–50.

Parra-Alvarez, JC, Posch, O, & Wang, M. Working. Identification and Estimation of Heterogeneous Agent Models: A Likelihood Approach . *CREATES Research Paper*.

Posch, O. 2009. Structural Estimation of Jump-Diffusion Processes in Macroeconomics. *Journal of Econometrics*, **153**(2), 196 – 210.

Sannikov, Y. 2008. A Continuous-Time Version of the Principal-Agent Problem. *Review of Economic Studies*, **75**(Feburary), 957–984.

Stokey, NL. 2009. *The Economics of Inaction: Stochastic Control Models with Fixed Costs*. Princeton University Press.

Veronesi, Pietro. 2019 (January). *Heterogeneous Households under Uncertainty*. Working Paper 25448. National Bureau of Economic Research.