Consumption and Saving Problem with Stochastic Interest Rates

by

Kyung Rok Min

An honors thesis submitted in partial fulfillment

of the requirements for the degree of

Bachelor of Science

Undergraduate College

Leonard N. Stern School of Business

New York University

May 2017

Professor Marti G. Subrahmanyam — Faculty Adviser

Professor Eduardo Dávila — Thesis Adviser

1 Introduction

What followed the Great Recession was a prolonged period of low interest rates and high liquidity. After fixing the target federal funds rate at 0 ~ 0.25% for years, the U.S. Federal Reserve has signalled the transition of the U.S. economy by incrementally raising this rate. While these rate hikes are coming at a pace below the historical average, there seems to be little disagreement over the view that change is due.

A raise in interest rates affects the whole economy, albeit in different pace, direction, and magnitude. For example, suppose the Federal Open Market Committee announces a higher target rate. The financial markets tend to respond almost immediately, whereas it takes much longer for the effect to be realized in a more widespread manner. The profitability of banks increase, whereas other firms may respond negatively to the added cost of borrowing. But amidst all this, it is quite possible for the average consumer to not notice any difference after an interest rate hike, while many analysts cautiously watch the markets over a couple of basis point changes in interest rates. In the current world with the globalization of economies, such changes also induce action from foreign investors.

Quantifying these effects in the macroeconomic scale would be a profoundly complex task. With many segments of the economy working together and against each other, both the cause and the effect appear to be elusive and not very amenable to analysis. What this paper tries to achieve is a reasonable first step in the right direction in developing a model to study and quantify such effects.

In this paper, I develop and solve a simple model of an agent that faces income and interest rate fluctuations in the economy. The uncertainty in the income and interest rate processes are characterized by two separate finite state, discrete time Markov Chains. In this setting, the agent has the option to save in order to smooth consumption over his lifetime, while borrowing is not allowed. After numerically solving for the optimal policy function, I run simulations for seven different scenarios and present plots and empirical statistics for each.

There are two main features of this paper that are important to put forward from the onset. The first is the restriction of the scope of analysis to an arbitrary agent. In a more general context, this might be the individual household that provides the microeconomic foundations of the aggregate macroeconomic sector. This restriction has two advantages: one is that this allows us to focus our attention to individuals that make up one of the most crucial sectors in the economy, and the other is that it allows us to directly study behavior without relying on other equilibrium arguments in the economy. Using a more parsimonious model gives us a partial snapshot of the bigger picture that is the macroeconomy. The tradeoff here is that to gain a clearer sense of what happens to the individual agent and what drives his or her decisions, I abstract away the details of other sectors of the economy that the agent interacts with.

The second feature is that by assuming incomplete markets, I expose the agent to risks that stem from income and interest rate fluctuations. Standard economic theory tells us that a risk averse agent, given access to financial markets and the appropriate instruments, will hedge out risk factors and smooth consumption over different time periods. I deviate from this standard by only allowing the agent to save and not borrow, in order to more closely examine consumption and saving behavior of individuals under exposure to risk. This not only gives rise to more interesting behavior, but also dispenses with what is perhaps a less favorable concept of the standard complete markets assumption. What results is a "self-insurance" type of behavior of an individual, shown through a series of saving and consumption rates throughout her lifetime.

There are a couple of forces in play when considering the effect of interest rates on income and savings. Substitution effect suggests that an increase in interest rates makes savings more attractive, making people shift consumption to the future and save more today. The income effect tells us that an increase in the interest rate expands the feasible consumption set, inducing the agent to raise present consumption and save less. And lastly, the wealth effect implies that a raise in interest rate will decrease the discounted lifetime income, and thus reinforces the interest rate's substitution effect in decreasing present consumption and raising saving. These concepts will serve as important tools of analysis, especially as I move on towards comparing results in later sections of this paper.

Now, what we observe in the outcome is a mix of these effects, and thus the overall effect is somewhat ambiguous. The substitution and income effects are offsetting forces in this dynamic, and what we can say a priori about which effect dominates is limited. However, studying a simple two period case gives us an idea of where this is headed, although it is not the full model considered and solved in the paper. Consider an agent who wishes to maximize the two period utility

$$\max_{c_1,c_2} u(c_1) + \beta u(c_2),$$

with $0 < \beta < 1$ as the discount factor of future utility. The agent makes the decision subject to the lifetime budget constraint

$$c_1 + \frac{c_2}{1+r} = y_1 + \frac{y_2}{1+r},$$

with an isoelastic utility function $u(c) = \frac{c^{1-\theta}}{1-\theta}$, where θ is the risk aversion parameter. Here c_1 and c_2 denote the period 1 and period 2 consumption, respectively, whereas y_1 and y_2 are the period 1 and period 2 deterministic endowments. Interest rates are fixed at r, and the utility is loosely interpreted as the total value derived from consumption.

Calculating the first order condition and applying the utility function gives the optimal future consumption rate, $c_2 = c_1 \beta^{1/\theta} (1+r)^{1/\theta}$. Substituting this value into the lifetime budget constraint yields the optimal present consumption,

$$c_1 = \frac{1}{1+\beta^{\frac{1}{\theta}} (1+r)^{\frac{1}{\theta}-1}} \left(y_1 + \frac{y_2}{1+r} \right).$$

One of the parameters that drives the agent's decision on how much to consume and save is the risk aversion parameter, θ. In this simple two period case, it is possible to tell which effect dominates depending on the risk aversion parameter θ. The two offsetting forces of substitution effect and the income effect are characterized by the $(1+r)^{\frac{1}{\theta}-1}$ term in the above formula. If $\theta < 1$, then the substitution effect dominates. If $\theta > 1$, the income effect will drive the decision. Each of those two scenarios will be considered in the general model setting below. What is interesting is that for the $\theta = 1$ case where $u(c) = \ln(c)$, the ratio of present consumption to the lifetime income is independent of the interest rate. The ratio is

$$\frac{c_1}{y_1 + \frac{y_2}{1+r}} = \frac{c_1}{c_1 + \frac{c_2}{1+r}} = \frac{c_1}{c_1 + c_1 \frac{\beta^{1/\theta}(1+r)^{1/\theta}}{1+r}} = \frac{1}{1+\beta^{\frac{1}{\theta}} (1+r)^{\frac{1}{\theta}-1}},$$

where I applied the lifetime budget constraint in the first equality, and the optimal future consumption rate in the second. Letting $\theta = 1$, this quantity simplifies to

$$\frac{1}{1+\beta},$$

as the fraction of lifetime income consumed today. Again, this is not the actual model considered in the rest of the paper, but it gives us some insight to the mechanism at work behind the full model. It also shows some of the elements that the consumption and saving behaviors depend on - namely the risk aversion parameter θ and the discount factor β.

Related Literature

The savings problem is one of the classical problems in macroeconomics. One can trace its beginnings from Irving Fisher's and Milton Friedman's consumption Euler equation, and the permanent income hypothesis proposed by Friedman (1957). The central message of the permanent income hypothesis is that consumption levels will be consistent with the expected long term average income of the consumers. The rich analysis that followed paved the way for a vast amount of literature that formed the cornerstones of modern macroeconomic theory.

Some of the earlier papers such as Schectman (1976), Yaari (1976), and Bewley (1977) aimed to give a more formal treatment of Friedman (1957). Both Schectman (1976) and Yaari (1976) proposed a utility maximization framework for the permanent income hypothesis in a finite horzion, zero interest rate, and undiscounted ($\beta = 1$) setting. The first two papers formulate income fluctuations as independent and identically distributed random variables, whereas Bewley (1977) formulates it as a stationary stochastic process.

Such developments then motivated another line of work that analyzed agent behavior. Brock and Mirman (1972) introduces uncertainty to the model through random shocks in the production function. Aiyagari (1994) then modifies this model to include uninsured idiosyncratic risks and liquidity contraints. The Aiyagari (1994) paper was motivated by the observation that the behavior of individual's consumption rates are strongly at odds with the complete markets assumption, which had been the building blocks of the representative agent model. Krussell and Smith (1998) then extends Aiyagari (1994) by adding a technology shock (presented as an aggregate state variable) that follows a Markov process. Krusell and Smith (1998) also proposes a numerical method to resolve the complexities that arise with introducing heterogeneous agents.

The notion of "self-insurance" characterizes an important feature of this paper. When faced with uncertainty and no means to perfectly hedge risk, the agent saves in order to self-insure against possible future drops in income. Similarly, the agent may also draw from savings in periods of low realized income to avoid drastic drops in consumption rates. Both Sotomayor (1984) and Chamberlain and Wilson (2000) analyzes economic settings where this type of behavior is observed. They both discuss conditions that will induce the optimal consumption path to eventually diverge to infinity.

On the other hand, this paper also assumes incomplete markets, which describes environments with restrictions to hedging or exchanging risks. In recent years, significant progress has been made to depart from the traditional complete markets setting. For example, Marcet and Singleton (1999) does this to compute equilibrium consumption-savings plans and asset prices. İmrohoroğlu (1989) considers two versions of incomplete insurance markets to study the welfare costs of business cycles. Huggett (1993) also works in this setting to develop models for risk-free interest rates that better explains an empirical observation than those formulated in a tradional Arrow-Debreu market context.

As these models evolved with increasing complexity, it also became nearly impossible to achieve analytical solutions in many cases. Accordingly, an important line of development in the literature involved numerical methods. The papers den Haan (1997) and the previously mentioned Krusell and Smith (1998) show that models with both aggregate and idiosyncratic risk are computable. In particular, the methods described in the second paper became very popularized, and led to more papers in a similar strand of logic. For example, Young (2010) presents a similar solution algorithm with a different simulation procedure; den Haan (2010) et al. compares some of the properties of algorithms used to solve models that feature incomplete markets with aggregate uncertainty; and Maliar, Maliar, and Valli (2010) extends the analysis on the properties of numerical solutions. The more recent paper Chipeniuk et al. (2016) offers a formal treatment to the numerical findings of Krusell and Smith (1998).

Outline

Section 2 describes the economic environment of the model and the underlying assumptions. Section 3 explains the approach and the methodology for solving the model. Section 4 makes a brief detour to the solutions of simplified versions of the model. Section 5 presents the solution and simulation results through a comparison scheme across different scenarios. Section 6 concludes, followed by a technical appendix and a list of references.

2 Economic Environment

In this section of the paper, I will lay out the setup of the model and outline its assumptions. The model seeks to solve the savings problem of an arbitrary consumer who wants to maximize the discounted expected utility by choosing a sequence of consumption rates.

Main problem The central component of the consumption and saving problem for the consumer is to find the maximized lifetime utility by choosing a sequence of consumption rates, $\{c_t\}_{t=0}^{\infty}$, at the initial time period $t = 0$. As it will be explained later, instead of designating some fixed consumption rate for every future time period t, I do this by choosing a consumption "plan" associated with every possible state variable that may arise in the future. This effectively chooses the optimal saving and consumption rate as the agent goes through the random fluctuations in the economy.

Here, how much a consumer values the consumption path, or the total sequence of consumption rates, is expressed as follows:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t).$$

The expectation operator conditioned at time $t = 0$ signals that there is uncertainty in this total sum, no matter how the consumer decides. The uncertainty stems from two main sources - changes in the endowment (or income) and interest rates over the consumer's lifetime. The consumer, however, has perfect a priori information on the *probabilities* of each events. In other words, the consumer knows the distribution of possible events and acts accordingly by maximizing the *expectation*. The key insights will follow from how the agent varies the consumption in between periods, while faced with the said uncertainty. In addition, all future values are discounted by a constant β, which is called the intertemporal discount factor. This agent-specific parameter takes a real number between 0 and 1, and reflects "patience," or how much the agent values future utility. Note that each successive time period is discounted equally - by $\beta, \beta^2, \beta^3, \cdots$, so on and so forth. A β value closer to 1 describes a consumer that is more patient, and a value closer to 0 someone who is less so.

The agent saves and consumes a single good at time t, where c_t represents the amount consumed at time t. In this setting, the time is discrete, and takes values $t = 0, 1, 2, \cdots$. The real-valued utility function $u(\cdot)$ then indicates how much the agent values the consumption rate at each period. This mapping is taken to be twice continuously differentiable, increasing, and strictly concave. One function that satisfies these requirements, and is used in this paper, is the Constant Relative Risk Aversion (CRRA) utility function:

$$u(c) = \begin{cases} \frac{1}{1-\theta} c^{1-\theta} & \text{if } 0 < \theta < 1 \\ \ln(c) & \text{if } \theta = 1. \end{cases} \tag{1}$$

Its namesake feature is that the relative risk aversion does not change with the level of consumption. Relative risk aversion is defined as

$$-c\frac{u''(c)}{u'(c)},$$

which in this case is equal to θ. This has the implication that the decision-making process is independent of scale.

Then, the problem is to find the maximum possible value over all feasible consumption paths:

$$\max_{\{c_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t).$$

What I mean by feasible is explained in the following section.

Constraints and timeline In essence, what is presented here is an infinite horizon, sequential decision-making process. The tricky part of the problem is that this decision-making process is inherently an intertemporal one. The tradeoff of consuming more today is that, all else equal, there will be less to consume tomorrow. Similarly, while increasing savings today in the face of uncertainty will insure the agent against future downturns, he will be forced to lower consumption today. This tradeoff is expressed mathematically in the following form of contraints at each time period t:

$$a_{t+1} + c_t = (1 + r_{t+1}) a_t + y_t \tag{2}$$

$$a_{t+1} \geq 0. \tag{3}$$

In the above expression, a refers to the level of assets held by the agent (hence one of the state variables in this case), r the interest rate, and y the income. The initial wealth a_0 is given. The equality in the first line is interpreted as the sum of next period assets and current consumption (left hand side) is equal to the return on current assets at the rate of $(1+r_{t+1})$ plus the income that period (right hand side). The second line delivers the additional condition that level of assets at every period must be nonnegative. This may seem like a somewhat arbitrary lower bound, but it has the interpretation that the agent can only save, not borrow. It is equivalent to having access to capital markets with a single, one-period security that has a return of $(1+r_{t+1})$ the next period, but without the option to short it. This condition will prevent the agent from reacting to economic dowturns simply by borrowing to avoid low consumption rates. Instead, it will help bring focus to the "self-insurance" type of behavior, where one saves in response to the *possibility* of financial turmoil.

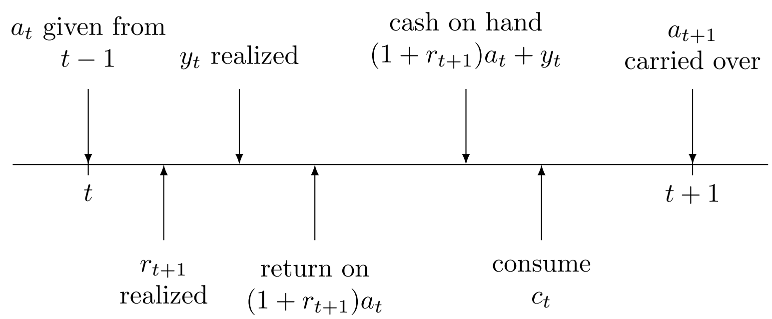

Figure 1: Timeline

The agent's actual decision-making process is best illustrated by a timeline. At the beginning of every period t, the agent is left with assets a_t from the previous period. While this is a consequence of the decision made in the previous period $t-1$, the consumer cannot change its level of assets a_t at the current period - instead, it is simply given. Then, the interest rate r_{t+1} is realized exogenously with respect to some probability, and the leftover assets a_t are saved at that rate. Overtime, the income y_t is also realized and the agent receives y_t, along with the return from the savings $(1+r_{t+1})a_t$. The total amount of wealth available at the time of decision-making is then $(1+r_{t+1})a_t + y_t$, which is called the cash on hand for that period. From a total of $(1+r_{t+1})a_t + y_t$, the agent then chooses how much to consume ("c_t") and how much to leave out for the next period ("a_{t+1}").

Notice that the agent practically only decides how much to consume; how much to leave out for the next period is determined by exclusion. This is guaranteed mathematically by our constraint (2) because it is possible to solve for a_{t+1} given the rest of the variables c_t, r_{t+1}, a_t, and y_t. Then there is a certain duality to the agent's problem - the savings problem is the consumption problem, and vice versa.

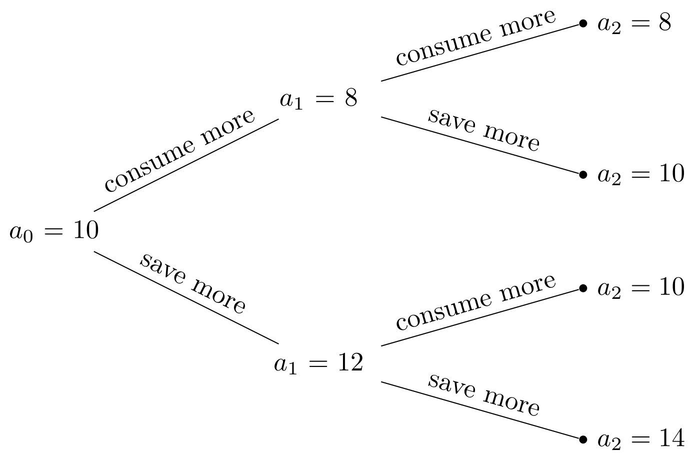

Figure 2: Sample Decision Tree

Another important observation is that every choice has consequences for all future periods. A decision made today not only affects tomorrow, but the entire future (Naturally, however, future decision-making cannot affect the present). It is also formulated recursively, which is a feature that will be exploited in obtaining the numeircal solution. The above figure shows one possible scenario of a decision making process, with the level of assets a_t describing the transition in between states.

Random components: income and interest rates Other crucial aspects of this model include its stochastic elements, the income and interest rates. Every period, the income and interest rate are realized by a random mechanism that is detached from the agent's decisions. Thus the consumer is a price taker in this setting, and does not have enough market power to affect this process. His individual savings are inconsequential as to how the interest rate is set, and his consumption (or lack thereof) will not start a recession. While incorporating firms, production, labor, among others can be done, the reason for pursuing this particular, rather parsimonious model is twofold: 1) restricting our attention to the household sector helps us focus on agent behavior under the specific circumstances of random interest rates without having to rely on any equilibrium argument; and 2) it streamlines the formulation and computation of the model.

Both income and interest rates follow random processes generated by discrete time, finite state Markov Chains. These stochastic processes exhibit the Markov Property, or "memorylessness." This means that the probability of moving to another state only depends on the present state, and not on the previous states. The process in a sense "forgets" its past states when evolving from the current state to the next.

In this environment, the agent faces income fluctuations according to transitions between three states in the macroeconomy - normal growth, mild recession, and severe recession. At each period t, the agent is assigned one value for the endowment, which we will label with a superscript, y^1, y^2, and y^3, respectively. We denote the state space of income with an upper case Y with no labels,$Y = \left\{y^1, y^2, y^3\right\}$. Note that these y^1, y^2, and y^3 are fixed values of endowment corresponding to each state. Now, the transitions in the states for income follow a Markov Chain with a transition matrix P that has the following form:

$$[P] = \begin{bmatrix} P_{11} & P_{12} & P_{13} \\ P_{21} & P_{22} & P_{23} \\ P_{31} & P_{32} & P_{33} \end{bmatrix} \text{ where } \sum_{j=1}^{3} P_{ij} = 1 \text{ for } i = 1, 2, 3$$

so that each row of the matrix adds up to one. Here each entry P_{ij} refers to the probability of transitioning from state y^i to y^j:

$$P_{ij} = \mathbb{P}\left[y_{t+1} = y^j \middle| y_t = y^i\right] \text{ for all } t = 0, 1, 2, \cdots.$$

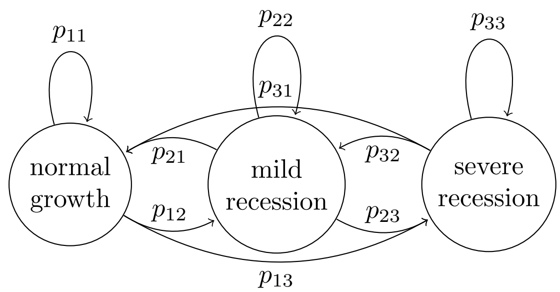

Figure 3: Markov States

On the other hand, the interest rate transitions between two states in the state space R, and is noted as $R = \{r^1, r^2\}$. This second process for interest rates is independent of the process for income, and is also realized exogenously. Here it can be interpreted that the economy transitions between two states of high and low interest rates, of low liquidity and high liquidity in the economy, or of high cost of capital and low cost of capital. It may also be interpreted as the per unit cost of present consumption in terms of future consumption. Similar to income, the interest rates in this economy evolves according to a Markov Chain with a transition matrix Q:

$$[Q] = \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix} \text{ where } \sum_{j=1}^{2} Q_{ij} = 1 \text{ for } i = 1,2,$$

and each entry Q_{ij} refers to the probability of transition from some state r^i to r^j :

$$Q_{ij} = \mathbb{P}\left[r_{t+1} = r^j \middle| r_t = r^i \right] \text{ for all } t = 0,1,2,\cdots.$$

Similarly, r^1 and r^2 are some fixed values, and the economy simply moves between the two.

Note that there are then six possible states associated with the random components, $|Y| \cdot |R| = 3 \cdot 2 = 6$. However, it may be simpler to think of two separate, independent Markov processes evolving throughout time, than to think of economy moving between six different states.

3 Methodology

Solving the main problem To reiterate, the goal is to find the objective function

$$\max_{\{c_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subjective to

$$a_{t+1} + c_t = (1 + r_{t+1}) a_t + y_t$$
$$a_{t+1} \geq 0,$$

for each time period t.

I will approach this problem using dynamic programming. More specifically, I will solve this problem by iterating the value function (also called the Bellman operator) until convergence. Here, the value function expresses the maximized value of current utility and discounted expected future utility for each period.

For ease of presentation, the subcript t denoting the period will be dropped from all variables. Next period variables will be labelled with a superscript $\prime$, such as a', y', r', so on and so forth. For example, for period t, a_t becomes a, a_{t+1} becomes a'. Similarly, r_{t+1} becomes r, and r_{t+2} becomes r'. Then the Bellman equation has the following form:

$$v(a,y,r) = \max_{c, a' \geq 0} \left\{ u(c) + \beta \mathbb{E}\left[v\left(a', y', r'\right) \middle| y, r \right] \right\}$$

subject to the contraints

$$c + a' = \left(1 + r'\right) a + y$$
$$a' \geq 0.$$

The value function takes a, y, and r (the level of assets, income, and interest rate for the period considered) as given. These are the three state variables. Then, the maximum is taken over all possible choices c and a' (the consumption and the assets carried over to the next period) so that it maximizes the quantity $u(c) + \beta \mathbb{E}\left[v\left(a', y', r'\right) | y, r\right]$. Here, c and a' are called the control variables. Notice that this formulation is recursive.

An essential feature of this problem is that it has optimal substructure, so that the principle of optimality applies. Borrowing from Bellman (1957), the principle of optimality states that "An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision." This means that the consumer's consumption-saving problem can be broken down into smaller subproblems, which is what makes this problem tractable. In other words, choice made at time 1 will put the agent in a certain state, affecting the decision problem from time 1 and onwards.

From here, I will take steps to simplify this problem and make it solvable by numerical means. Since I will be using the CRRA utility function (1), the Bellman equation can be rewritten as

$$v(a, y, r) = \max_{c, a' \geq 0} \left\{ \frac{c^{1-\theta}}{1-\theta} + \beta \mathbb{E}\left[v\left(a', y', r'\right) \middle| y, r \right] \right\}.$$

Furthermore, notice that consumption can be rewrriten as $c = (1 + r') a + y - a'$. This is in line with the duality mentioned before - the constraint effectively eliminates one of the two control variables (c and a') by substitution. Thus, the maximum is now only taken over a' instead of both a' *and* c. Also, the expectation operator can be calculated explicitly because the probability distribution for the two stochastic elements in this model - y and r - are known. The probabilities, however, depend on two of the current state variables y and r due to the Markov Property. Then, making these changes to the Bellman equation,

$$\begin{aligned} v(a, y, r) &= \max_{a'} \left\{ \frac{\left((1+r) a + y - a'\right)^{1-\theta}}{1-\theta} + \beta \mathbb{E}_{y', r'} \left[v\left(a', y', r'\right) \middle| y, r \right] \right\} \\ &= \max_{a'} \left\{ \frac{\left((1+r) a + y - a'\right)^{1-\theta}}{1-\theta} + \beta \left[\sum_{j=1}^{2} \sum_{i=1}^{3} \mathbb{P}\left(y' = y^i \middle| y\right) \mathbb{P}\left(r' = r^j | r\right) v\left(a', y' = y^i, r' = r^j\right) \right] \right\}, \end{aligned}$$

where the index i sums over the three possible states for come, and j over the two possible states for interest rates. The conditional probabilities involving income and interest rates are given by the entries in the transition matrices P and Q, respectively.

4 Sample Scenarios: Analytical Solutions

Two periods with no uncertainty and log utility Suppose we have a simplified version of the model with only two periods and deterministic income and interest rates. Take $\theta = 1$, so that the CRRA utility simplifies to log utility. Then the problem at hand becomes

$$\max_{c_1, c_2} \left\{ \ln(c_1) + \beta \ln(c_2) \right\}$$

subject to

$$
\begin{aligned}
a_2 + c_1 &= (1+r_2)\,a_1 + y_1 \\
c_2 &= (1+r_3)\,a_2 + y_2 \\
a_2 &\geq 0.
\end{aligned}
$$

The initial values are given: $a_1$the initial wealth, y_1 and r_2. The terminal wealth a_3 is implicitly assumed to be zero, so that $a_3 = 0$. Previously random values r_3 and y_2 are also known ahead of time. Then the problem becomes

$$\max_{a_2} \left\{ \ln\left((1+r_2)\,a_1 + y_1 - a_2\right) + \beta \ln\left((1+r_3)\,a_2 + y_2\right) \right\}.$$

The corresponding first order condition is

$$\beta\,(1+r_3) = \frac{c_2}{c_1}.$$

Here it can be interpreted that at the optimal consumption path, the ratio of tomorrow's consumption to today's consumption is given by $\beta\,(1+r_3)$. This value may be rewritten in the more familiar form

$$\frac{\beta u'(c_2)}{u'(c_1)} = \frac{1}{1+r_3}.$$

The left hand side gives the ratio of utility earned by increasing a unit of tomorrow's consumption to the utility earned by increasing a unit of today's consumption. The right hand side tells us that at the optimal path is when this ratio is equal to $\frac{1}{1+r_3}$. In this simplified version, it is possible to solve for a_2^* value that maximizes the above objective function:

$$a_2^* = \frac{\beta\left[(1+r_2)\,a_1 + y_1\right] - \frac{y_2}{1+r_3}}{(1+\beta)}.$$

All the values on the right hand side of the above equation are constants in this setting - $\beta, a_1, y_1, y_2, r_{2,}$ and r_3. Then, the optimal consumption path associated with this policy is

$$
\begin{aligned}
c_1^* &= \frac{1}{1+\beta}\left[(1+r_2)\,a_1 + y_1 + \frac{y_2}{1+r_3}\right] \\
c_2^* &= \frac{\beta}{1+\beta}\left[(1+r_3)\left[(1+r_2)\,a_1 + y_1\right] + y_2\right].
\end{aligned}
$$

Finite periods with no uncertainty and log utility First suppose we have a finite horizon case with terminal period T. Assume the deterministic and log utility setting as before. The objective is to choose a consumption path that will maximize

$$\sum_{t=0}^{T} \beta^t u(c_t).$$

As before, I will use the constraints to eliminate the consumption variable. This yields

$$\max_{\{a_{t+1}\}_{t=0}^{T}} \sum_{t=0}^{T} \beta^t u\left((1+r_{t+1})\,a_t + y_t - a_{t+1}\right).$$

Expanding this sum will give an expression for an arbitrary period t such that $0 \leq t < T$,

$$\cdots + \beta^t u\left((1+r_{t+1})\,a_t + y_t - a_{t+1}\right) + \beta^{t+1} u\left((1+r_{t+2})\,a_{t+1} + y_{t+1} - a_{t+2}\right) + \cdots.$$

Observe that these two terms are the only terms that include the variable a_{t+1}. Taking the derivative respect to a_{t+1} gives the first order condition:

$$-\beta^t u'\left((1+r_{t+1})\,a_t + y_t - a_{t+1}\right) + \beta^{t+1} u'\left((1+r_{t+2})\,a_{t+1} + y_{t+1} - a_{t+2}\right)(1+r_{t+2}) = 0.$$

Then, rearranging this expression,

$$u'\left(\left(1+r_{t+1}\right)a_t+y_t-a_{t+1}\right)=\beta u'\left(\left(1+r_{t+2}\right)a_{t+1}+y_{t+1}-a_{t+2}\right)\left(1+r_{t+2}\right).$$

Again, the familar form arises:

$$\frac{\beta u'\left(c_{t+1}\right)}{u'\left(c_t\right)}=\frac{1}{1+r_{t+2}}.$$

This applies to all the periods except the last one. As in the two period case, the terminal wealth is assumed to be zero, hence it is implied that

$$a_{T+1}=0.$$

Now I will apply the log-utlity form. This means the first order condition becomes

$$\frac{1}{\left(1+r_{t+1}\right)a_t+y_t-a_{t+1}}=\frac{\beta\left(1+r_{t+2}\right)}{\left(1+r_{t+2}\right)a_{t+1}+y_{t+1}-a_{t+2}}.$$

This gives the quasi-analytical solution

$$a_{t+1}=\frac{\beta}{1+\beta}\left[\left(1+r_{t+1}\right)a_t+y_t+\frac{a_{t+2}-y_{t+1}}{\beta\left(1+r_{t+2}\right)}\right].$$

Notice that this is not a real solution because of the a_{t+2} term in the expression.

5 Comparing Computational Results

A roadmap I will present the results of the calculation on a case-by-case basis. The emphasis will be on how changing the various statistical properties of the interest rate process will affect the optimal policy and other results. Different values for the risk aversion parameter will also be considered. After obtaining the optimal policies in each setting, I will then apply the policy to simulated income and interest rate processes. By obtaining the empirical variance-covariance matrix for the implied consumption rate, asset policy, cash on hand, income and interest rates in this simulated setting, I will be able to compare how differences across parameters will cause these correlations to change.

Income process The one constant in this series of computational experiments will be the formulation of the income process. I will fix a transition matrix and state vector for the income process to better compare differences across other features of the model. The values for the transition probabilities will be borrowed from Hamilton (2005). The paper proposes a three-state Markov-switching model in studying the post-war U.S. business cycles. The author then calibrates the model to the seasonally adjusted monthly unemployment data from 1948 January to 2004 March. This dataset is available from the Bureau of Labor Statistics (http://stats.bls.gov).

The values presented in Hamilton (2005) are

$$[P]=\begin{bmatrix}0.971 & 0.029 & 0\\ 0.145 & 0.778 & 0.077\\ 0 & 0.508 & 0.492\end{bmatrix}.$$

Notice that some transition probabilties, such as the one from normal growth to severe recession and vice versa, are zero. This consequently imposes the condition that the economy cannot directly switch between extremes states in neighboring periods. Now, the state vector for the income will have to be chosen somewhat arbitrarily, with only the condition that it is the highest in periods of normal growth, lower during mild recessions, and the lowest in times of severe recessions. The vector that will be used in the following sections will be

$$y=\begin{bmatrix}5\\ 3\\ 1\end{bmatrix}.$$

Benchmark interest rate process: Rouwenhorst method In order to obtain an approximate benchmark for the interest rate process, I first fitted an AR(1) process with drift to actual data. The different cases considered below are obtained by making changes to this process, and then by discretizing it using the Rouwenhorst method. I used the monthly effective federal funds rate from 1954 July to 2004 March, fitting on 597 observations (data available at https://www.federalreserve.gov/data.htm). The model has the following form:

$$r_t = \omega + \rho r_{t-1} + \varepsilon_t,$$

where ω and ρ are constants, and the ε_t's are assumed to be independent, identically normally distributed, so that $\varepsilon_t \sim N\left(0, \sigma^2\right)$. In order to estimate these parameters, I used the following formulas:

$$\mu = \frac{\omega}{1-\rho} \approx 0.004859687$$
$$\operatorname{var}(r_t) = \frac{\sigma^2}{1-\rho^2} \approx 0.002818607^2,$$

where μ is the sample average of r_t, ρ is the first order autocorrelation of r_t. It is then possible solve for the parameters of the AR(1) process mentioned above. Making these calculations I obtain

$$\widehat{\rho} = 0.9824744$$
$$\widehat{\omega} = 8.516905 \times 10^{-5}$$
$$\widehat{\sigma} = 5.253818 \times 10^{-4}.$$

Then, I used the Rouwenhorst method to derive a Markov chain from this process. I obtain the following estimates:

$$[Q] = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$$
$$r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}.$$

From here on, I will only change the values for β (discount factor), θ (risk aversion parameter), Q (transition matrix for interest rates), and r (state vector for interest rates). At each heading, I will clearly denote which values are used for these parameters.

Simulation and computed statistics Every simulation was done using a random number generator in Python. The length of the series is 500 periods for all cases. I began by obtaining a simulated array for income and interest rate processes, as those are the purely random processes. The random seed for the income process was fixed throughout this paper, so that the income process is identical for all cases. This was done to better compare differences across other components. Then, I applied the previously calculated policy function to derive the other simulated processes.

As for the computed statistics, I presented below the correlation matrix and the standard deviations. Unlike the plots, these values are obtained by a much longer simulated process. For each sample time series of length $T = 5000$, I computed the correlation matrix and the standard deviations. Then, I averaged these statistics across 1000 samples, and presented the outcome.

5.1 Case 1: Low mean

For the first case, I would like to refer back to the Markov process and the AR(1) process for the interest rates. To make comparisons between one process with a lower mean and higher mean, suppose the underlying data for the interest rate process had a 20% lower mean. Then I can fit an AR(1) process with drift, and again derive the transition matrix and the state vectors. Then the values for $(\omega, \rho, \sigma) = (0.00006813524, 0.9824744, 0.0005253818)$. Using the Rouwenhorst method, I can then find the associated transition matrix and state vector. To summarize, the following values are used for the first case:

$\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00106915 \\ 0.00670636 \end{bmatrix}$.

Here are the plots of the solved model:

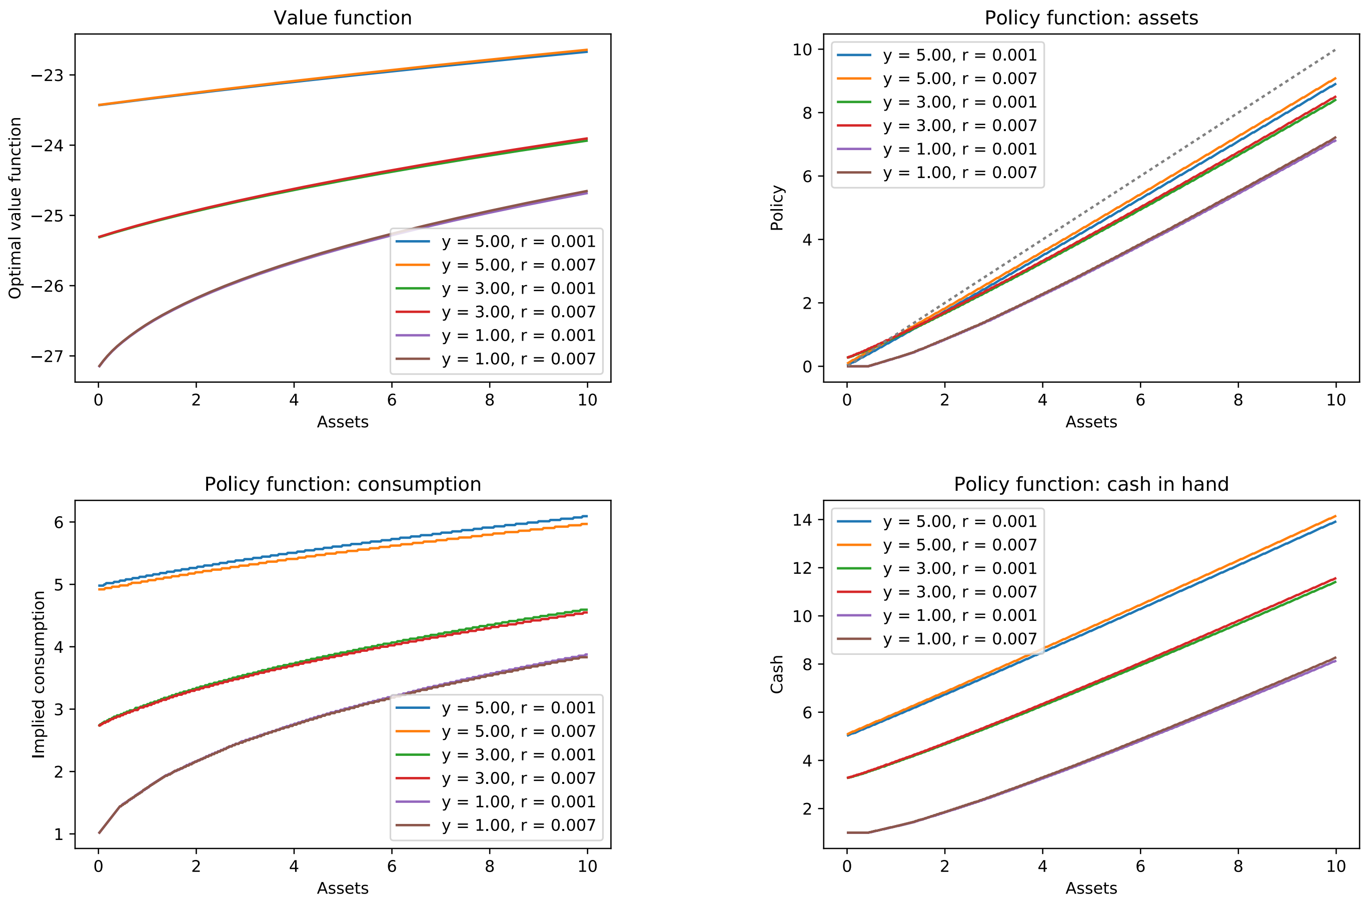

*Case 1: Low mean $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00106915 \\ 0.00670636 \end{bmatrix}$.

Figure 4: Solution results

Figure 4 shows the four "solution" plots for the first case. Beginning from the top left hand side, I have the value function, followed by the three policy functions associated with assets, consumption, and cash in hand.

The first plot is the direct outcome of the calculation, as it is the optimal value function that results after iterating the Bellman operator until convergence. What is plotted is the value function associated with various state variables (a, y, r). The first state variable a is continuous and is shown on the x-axis. Six different curves correspond to six different finite state values of y and r $(3 \times 2 = 6)$ for each asset value. Notice that the y-axis values are negative - this is a result of setting $\theta > 1$, which pushes the utility to negative values. This should not be interpreted as these functions having "negative" value, per se, as utility is simply a way to present a preference relationship. What really matters is the ordered relationship between those values. In addition, observe that the value function curves inherit its shape from the utility function - it is also increasing and concave.

The second plot on the top right hand side shows the assets policy function, which is derived from the value function. Again, the state variables are represented similarly. The current assets are on the x-axis, and the different curves correspond to different y and r values. The asset policy function tells the agent that given these state variables a_t, y_t, and r_{t+1}, the optimal response in that state is to set the a_{t+1} (assets to carry over to the next period) equal to the value of the asset policy function. Because choosing a_{t+1} effectively chooses the consumption rate and vice versa, this fully characterizes what the optimal policy is in any given state.

Notice that these calculations fully incorporate the uncertainty associated with the income and interest rate fluctuations.

The bottom two plots are the consumption policy function and the cash in hand policy function. Both are derived from the asset policy function, and both represent optimal policies. For example, the consumption policy function "instructs" the agent that given the state variables (a, y, r), the optimal policy is to consume the amount shown on the y-axis. Note that here a lower interest rate value is associated with higher consumption, highlighting the substitution effect. A higher interest rate makes savings more attractive, inducing the agent to shift from consumption to saving. However, as shown in the plots, there are much larger differences in consumption rates across income states than across interest rate states. This hints that the decision is driven mostly by the income fluctuations.

In Figure 5 on the next page, I have the five different simulation plots for Case 1. The first three plots, titled "Consumption," "Assets," and "Cash in hand" are results of applying the previously mentioned three policy functions to simulated random processes for y and r. Since the policy functions instructs the agent what to do in every possible state, I test it to simulated processes, and plot the agent's optimal response in every period. The last two plots, titled "Income" and "Interest rates" are purely random simulated Markov Chains. Remember that the random seed was fixed for the income process.

Beginning with the plot on the left hand side of the first row, I have the simulated consumption paths. Notice that the consumption path is not smooth at all - instead it fluctuates very frequently, and somewhat resembles the income process. This is a result of the agent's exposure to the risks associated with income and interest rate processes. Without the means to hedge out risks, the agent has to face the drops in consumption levels during downturns in the economy. In the other extreme, if the agent had full access to the appropriate financial instruments, then the consumption rate of a risk averse agent would be flat.

On the right hand side of the first row is the simulated assets policy function. Whereas the consumption, cash in hand, and income plots seem to more closely resemble one another, the plot for assets look a little different. But the sudden drops in the asset levels are almost traceable by comparing them to the simulated income process. Large drops in income correspond to large drops in assets. As do the consumption and cash in hand, the asset levels seem to fluctuate a lot throughout time.

As for the income process, one can almost see the three different states in the Markov Chain - normal growth, mild recession, and severe recession. Since income only takes three different values, the plot seems like a step function with three different values. The same is true for the interest rate plot, but more noticeably so. The reason is that the transition matrix for the interest rate process is very persistent, meaning transitions between states are rare. In this case there are only a handful of interest rate changes. In a later case, I will examine the results for a less persistent interest rate process.

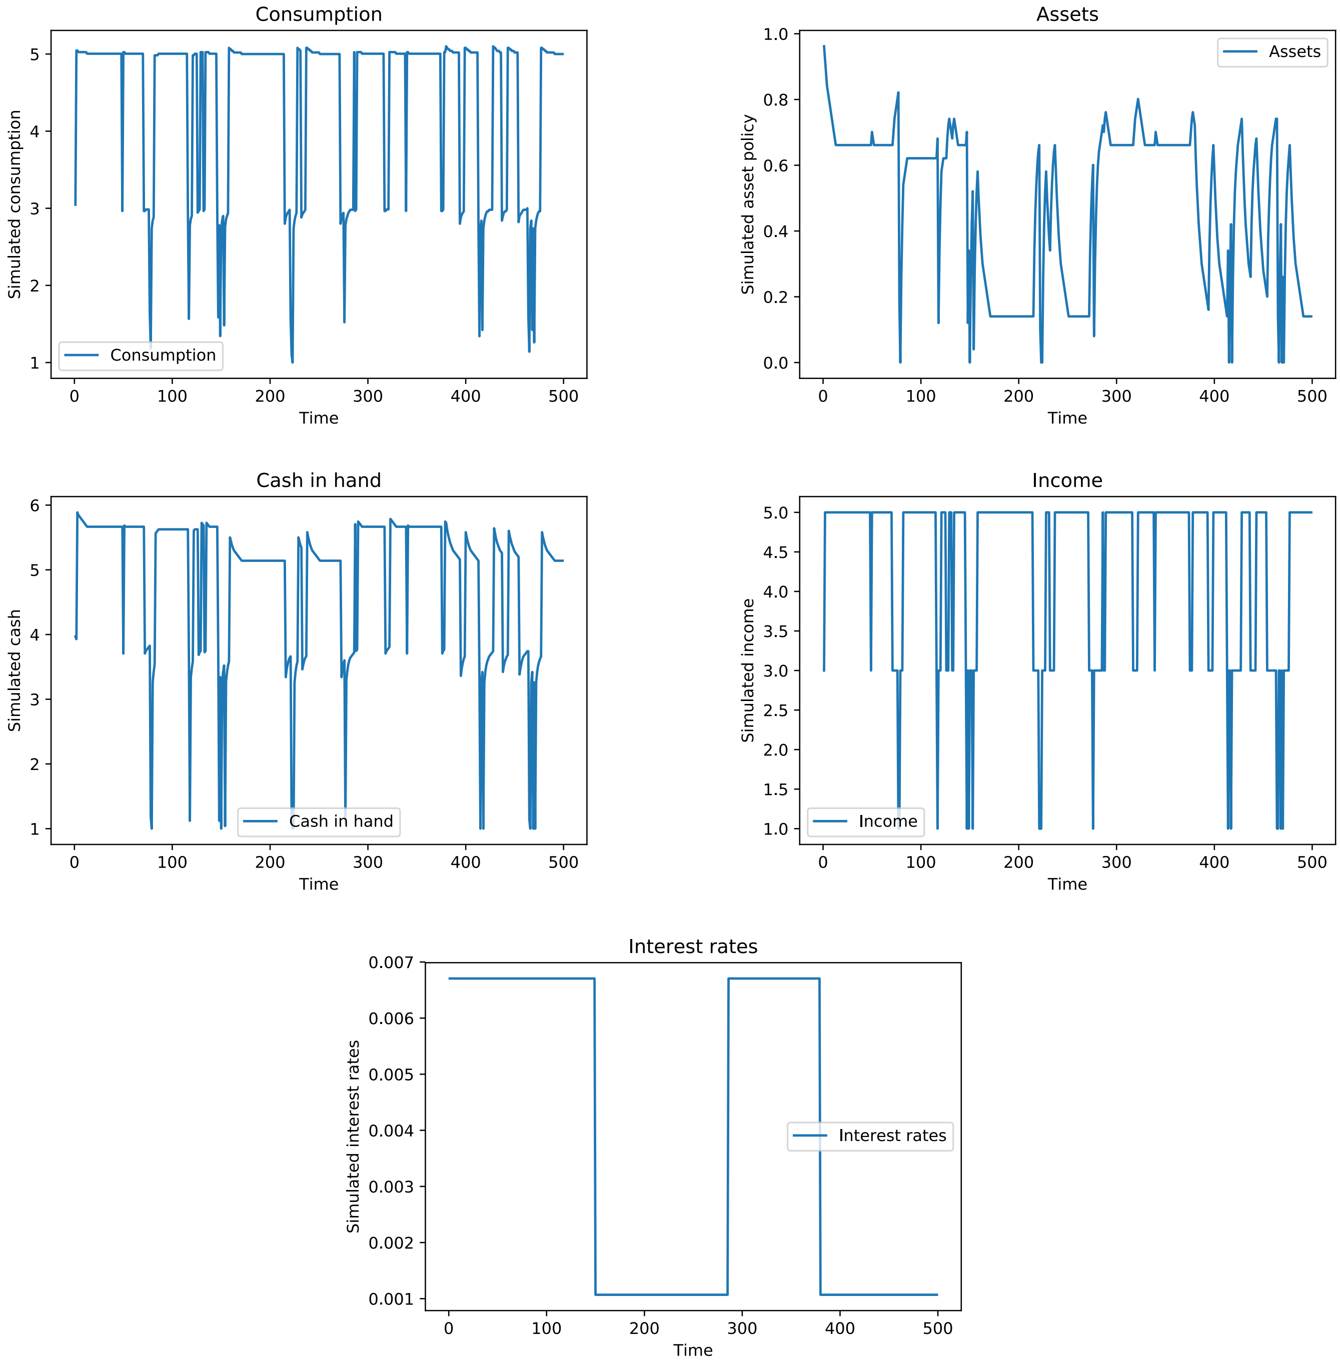

*Case 1: Low mean $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00106915 \\ 0.00670636 \end{bmatrix}$.

Figure 5: Simulation results

On the other hand, here are the computed statistics. The table on the left hand side shows the correlation matrix, followed by the standard deviations on the right. Notice that interest rate and the income processes have a correlation coefficient of -0.0001 in practice, but theoretically, those are independent processes. As expected, consumption and income are strongly correlated, at 0.9962, whereas assets and income are negatively correlated with a coefficient of -0.1120. As for the standard deviations, cash in hand is the most volatile, followed by income, consumption, assets, then interest rate.

	Consumption	Assets	Cash	Interest rate	Income
Consumption	1.0000	-0.0957	0.7882	-0.0014	0.9962
Assets	-0.0957	1.0000	0.1667	0.7590	-0.1120
Cash	0.7882	0.1667	1.0000	0.1955	0.7715
Interest rate	-0.0014	0.7590	0.1955	1.0000	-0.0001
Income	0.9962	-0.1120	0.7715	-0.0001	1.0000

	Standard deviation
Consumption	0.9106
Assets	0.2373
Cash	0.9246
Interest rate	0.0028
Income	0.9123

*Case 1: Low mean $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00106915 \\ 0.00670636 \end{bmatrix}$.

5.2 Case 2: High mean

In the second case, I approach similarly as the first.

The AR(1) parameters are $(\omega, \rho, \sigma) = (0.0001022029, 0.9824743761, 0.0005253818)$. The parameters changed are then $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00301303 \\ 0.00865024 \end{bmatrix}$.

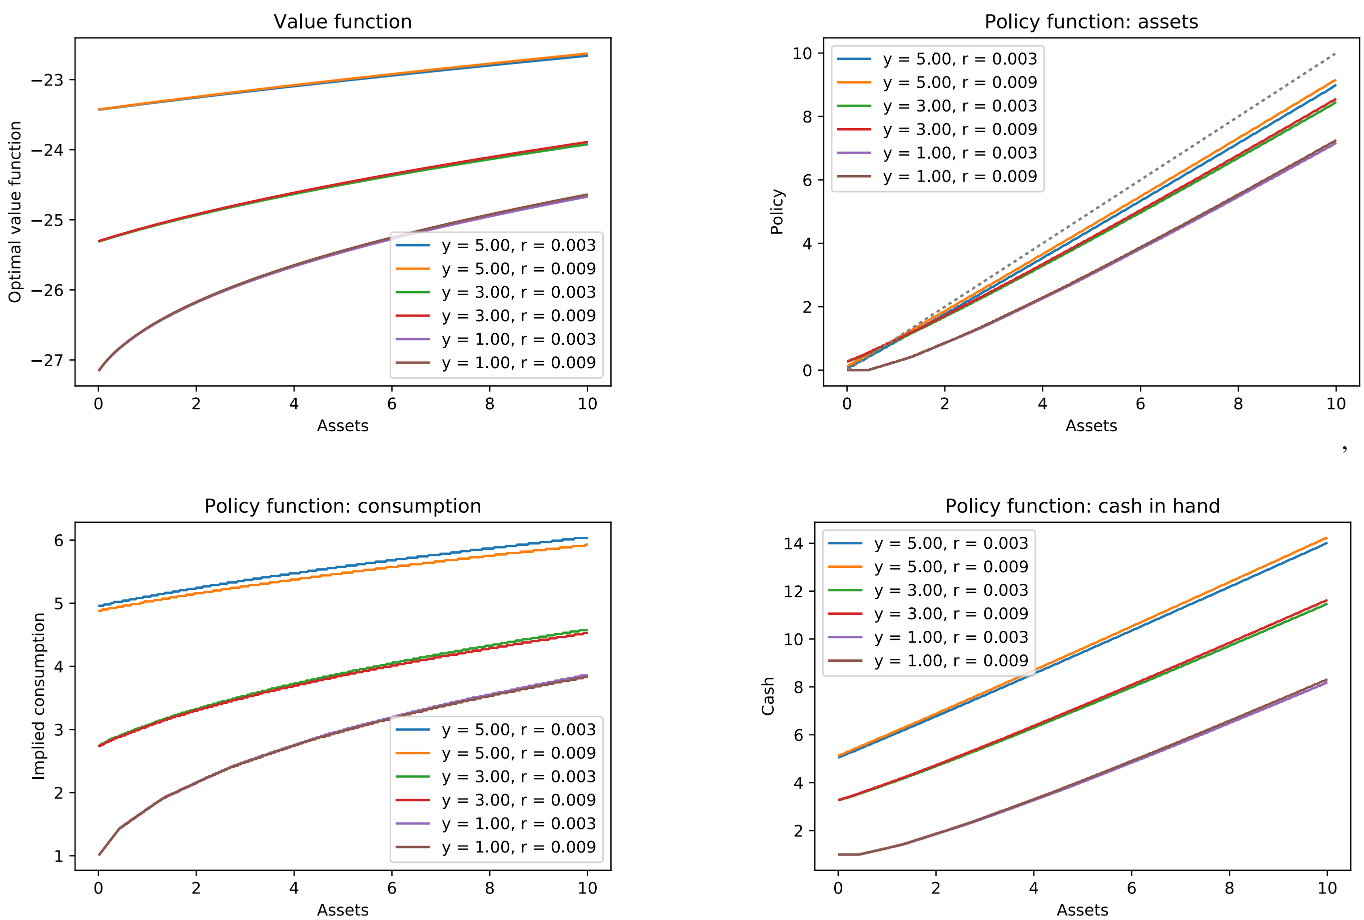

*Case 2: High mean $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00301303 \\ 0.00865024 \end{bmatrix}$.

Figure 6: Solution results

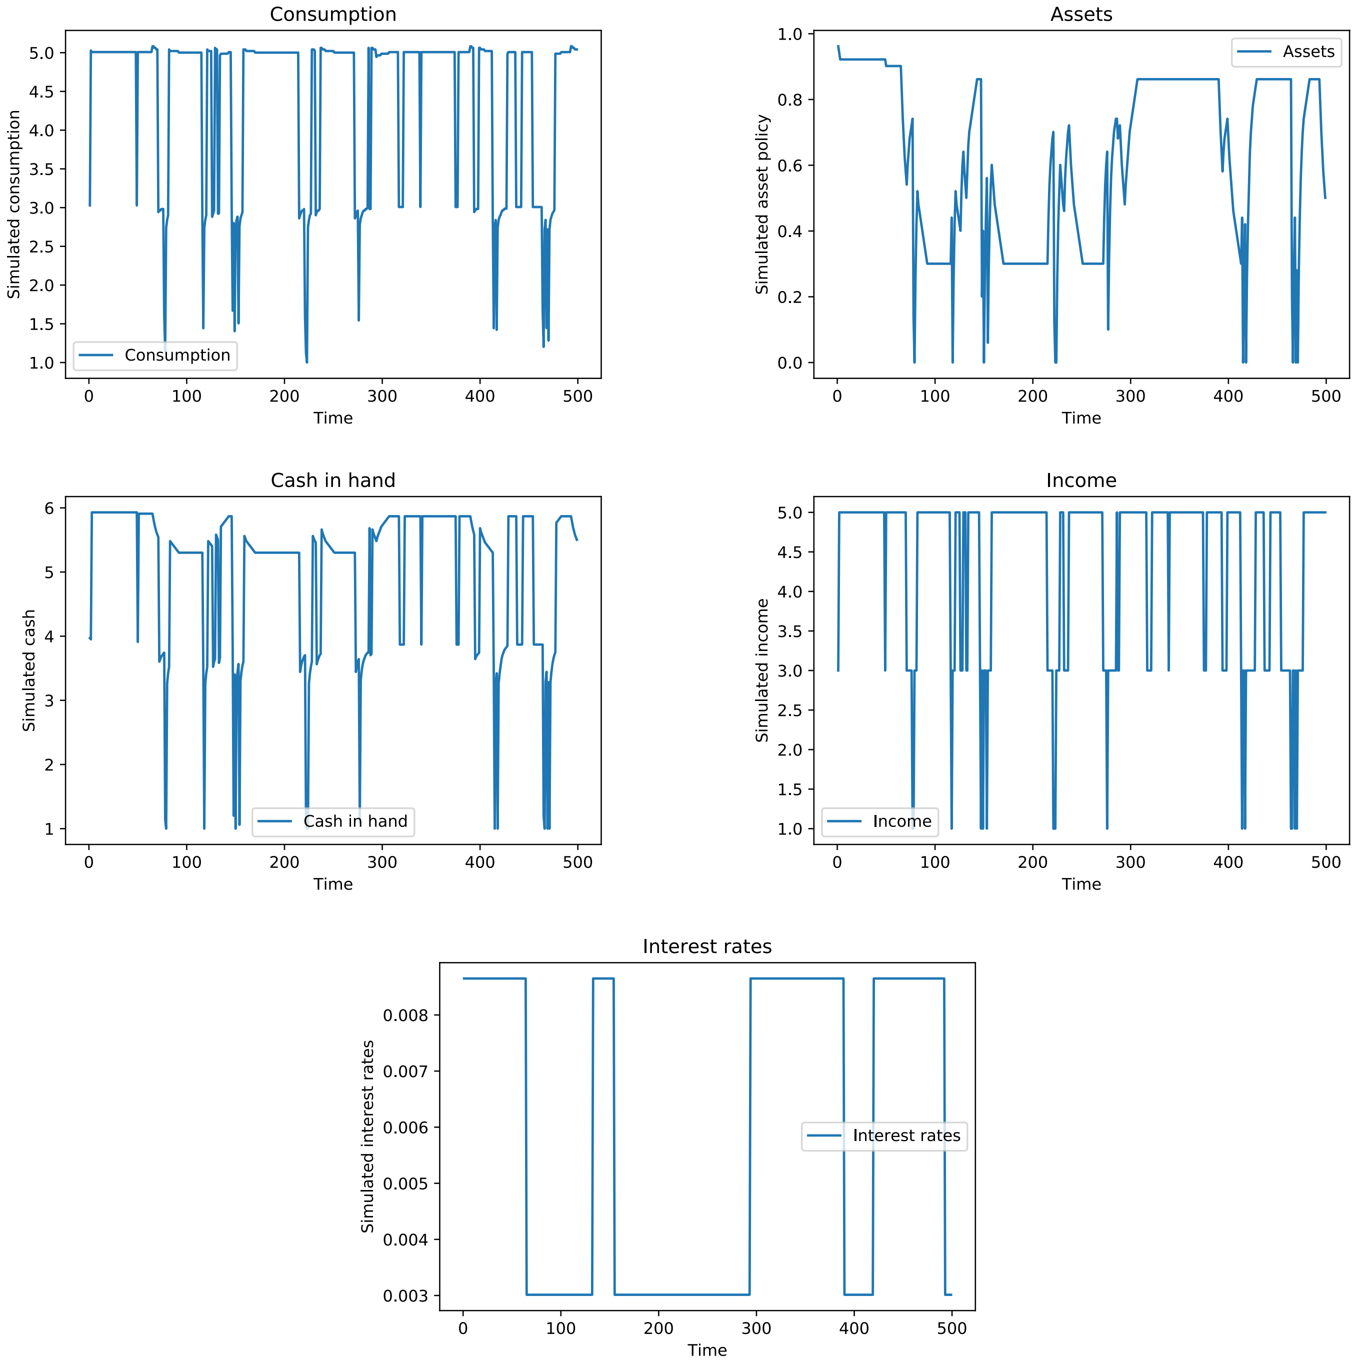

*Case 2: High mean $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00301303 \\ 0.00865024 \end{bmatrix}$.

Figure 7: Simulation results

The preceding figures show the solution and simulation plots for the second case with high mean. Notice that while the income process is identical to Case 1 with low mean, the interest rate process is different. This is because the random seed is fixed for the income process, whereas the interest rates are newly drawn each time. Furthermore, in Case 2 with high mean, the interest rate state values both shift upwards.

Here are the computed statistics:

	Consumption	Assets	Cash	Interest rate	Income
Consumption	1.0000	0.0697	0.7928	0.0017	0.9961
Assets	0.0697	1.0000	0.3579	0.7979	0.0520
Cash	0.7928	0.3579	1.0000	0.2114	0.7724
Interest rate	0.0017	0.7979	0.2114	1.0000	0.0026
Income	0.9961	0.0520	0.7724	0.0026	1.0000

	Standard deviation
Consumption	0.9029
Assets	0.2538
Cash	0.9754
Interest rate	0.0028
Income	0.9123

*Case 2: High mean $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00301303 \\ 0.00865024 \end{bmatrix}$.

5.3 Case 3: Low volatility

In the case of comparing high and low volatilities, I will change the standard deviation of the ε_t's (that is, the σ) associated with the AR(1) process. As before, the changes will be 20% lower and 20% above. Using $0.8 \times \sigma$ for the standard deviation, I obtain the transition matrix and state vectors. As a result, the parameters used in this case are $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00260481 \\ 0.00711458 \end{bmatrix}$.

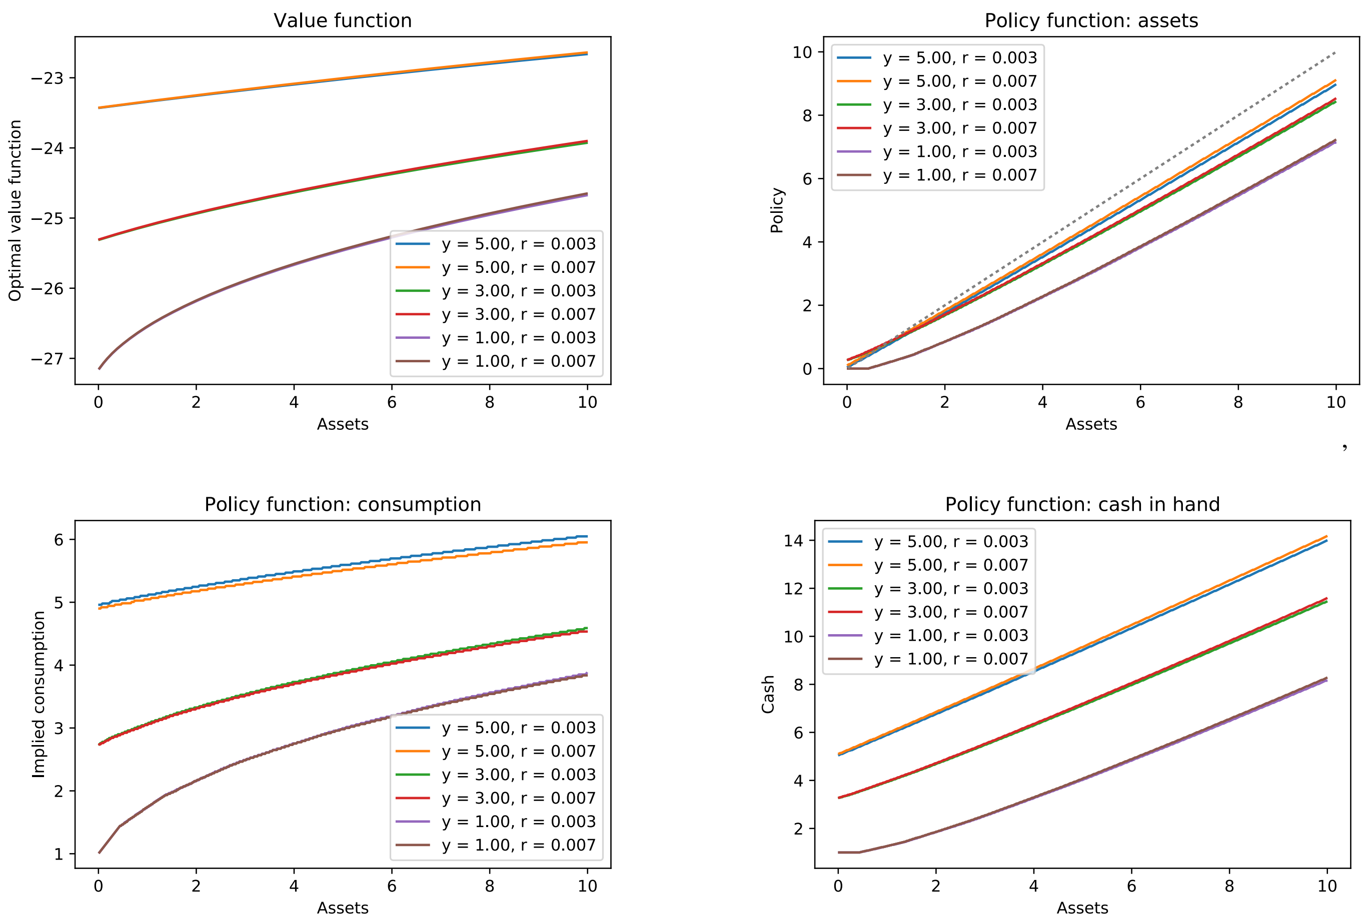

*Case 3: Low volatility $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00260481 \\ 0.00711458 \end{bmatrix}$.

Figure 8: Solution results

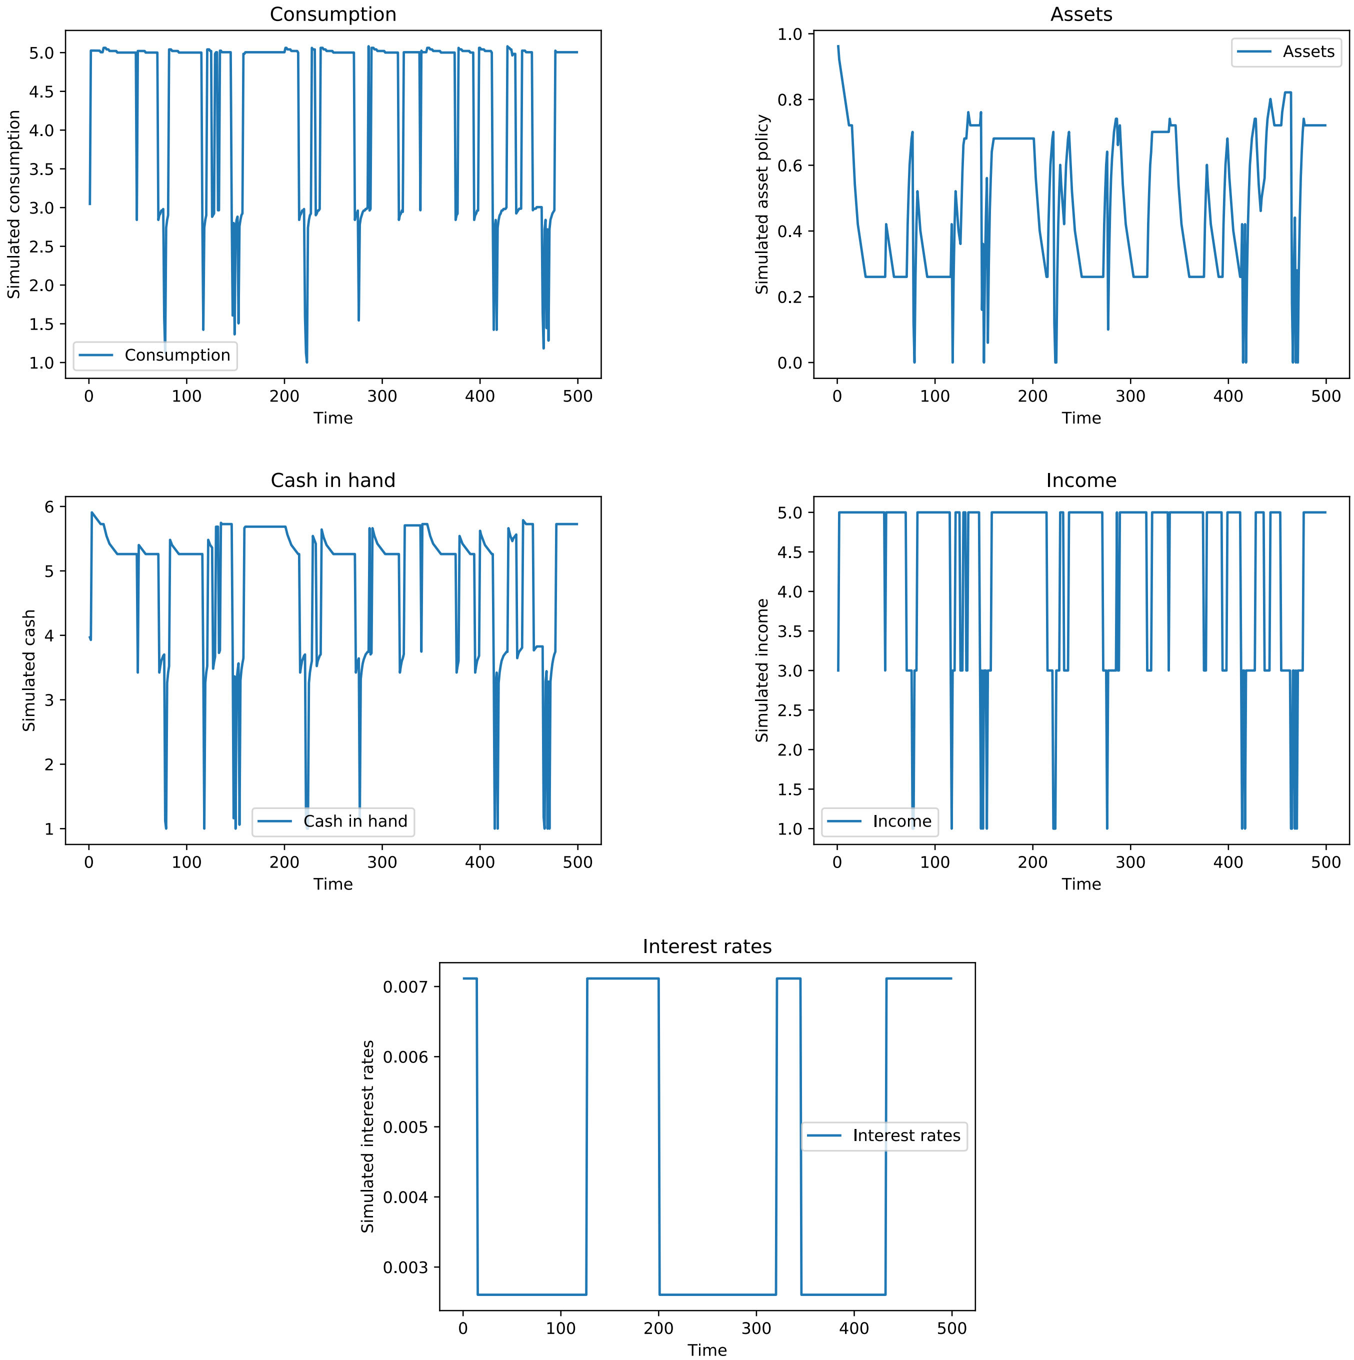

*Case 3: Low volatility $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00260481 \\ 0.00711458 \end{bmatrix}$.

Figure 9: Simulation results

For Case 3 and Case 4, the volatilities for the interest rates are changed. Applying the Rouwenhorst method only changes the state vector for the interest rate, causing the discrepancy between the high and low values to become smaller, then larger. In this case, the level of assets is slightly negatively correlated with income as before. The correlation between consumption and income seems to have become even stronger at 0.9963. This is a very large number, considering the coefficient is between -1 and 1. Noticeably, the standard deviation for interest rate has become smaller, now at 0.0022.

	Consumption	Assets	Cash	Interest rate	Income
Consumption	1.0000	-0.0200	0.7967	0.0018	0.9963
Assets	-0.0200	1.0000	0.2282	0.7487	-0.0381
Cash	0.7967	0.2282	1.0000	0.1721	0.7784
Interest rate	0.0018	0.7487	0.1721	1.0000	0.0028
Income	0.9963	-0.0381	0.7784	0.0028	1.0000

	Standard deviation
Consumption	0.9069
Assets	0.2118
Cash	0.9403
Interest rate	0.0022
Income	0.9123

*Case 3: Low volatility $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00260481 \\ 0.00711458 \end{bmatrix}$.

5.4 Case 4: High volatility

In continuation of the previous case, I will increase σ by 20%. Then, I obtain $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00147736 \\ 0.00824202 \end{bmatrix}$.

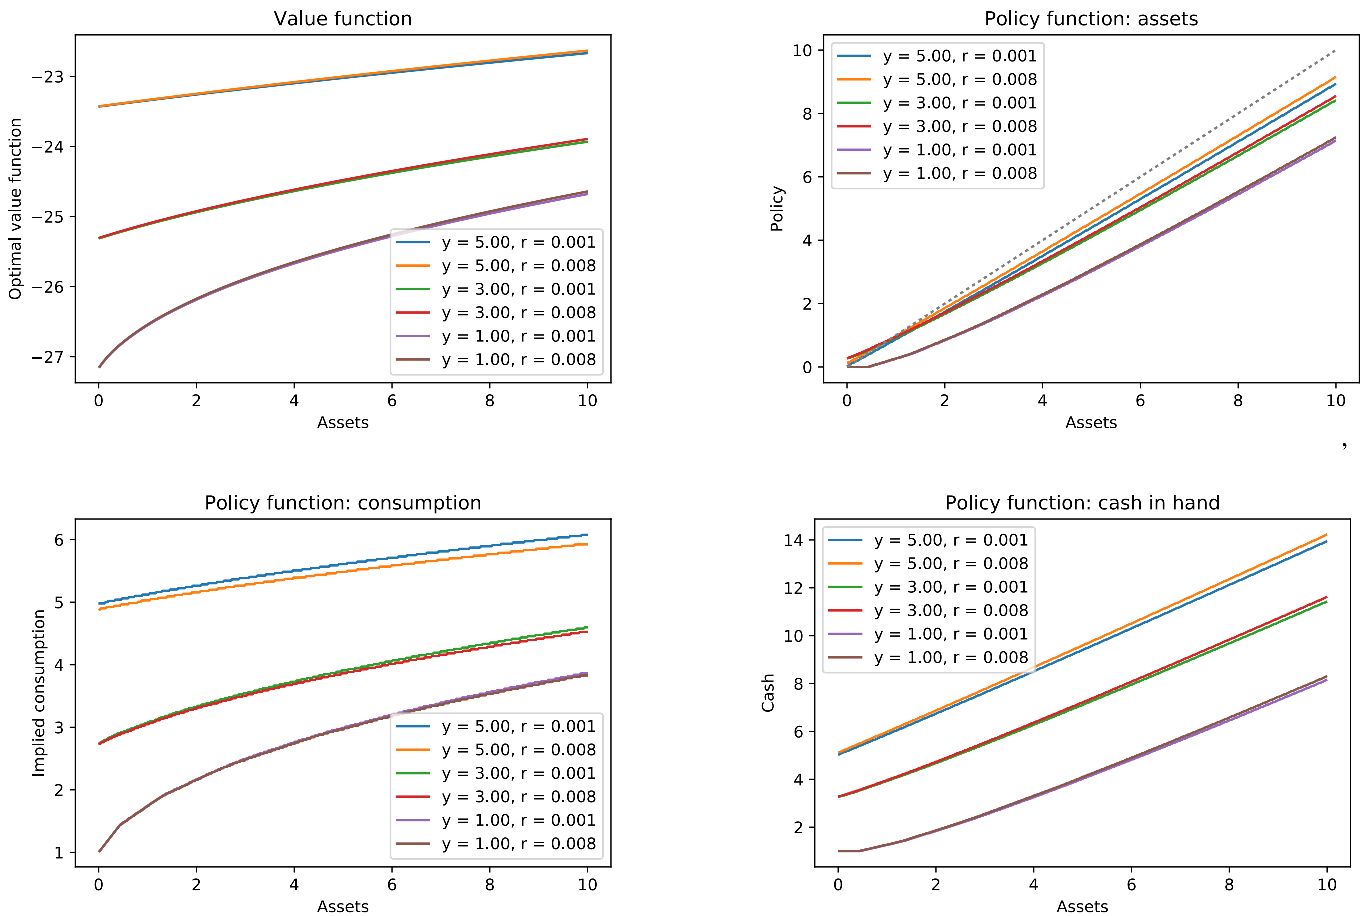

*Case 4: High volatility $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00147736 \\ 0.00824202 \end{bmatrix}$.

Figure 10: Solution results

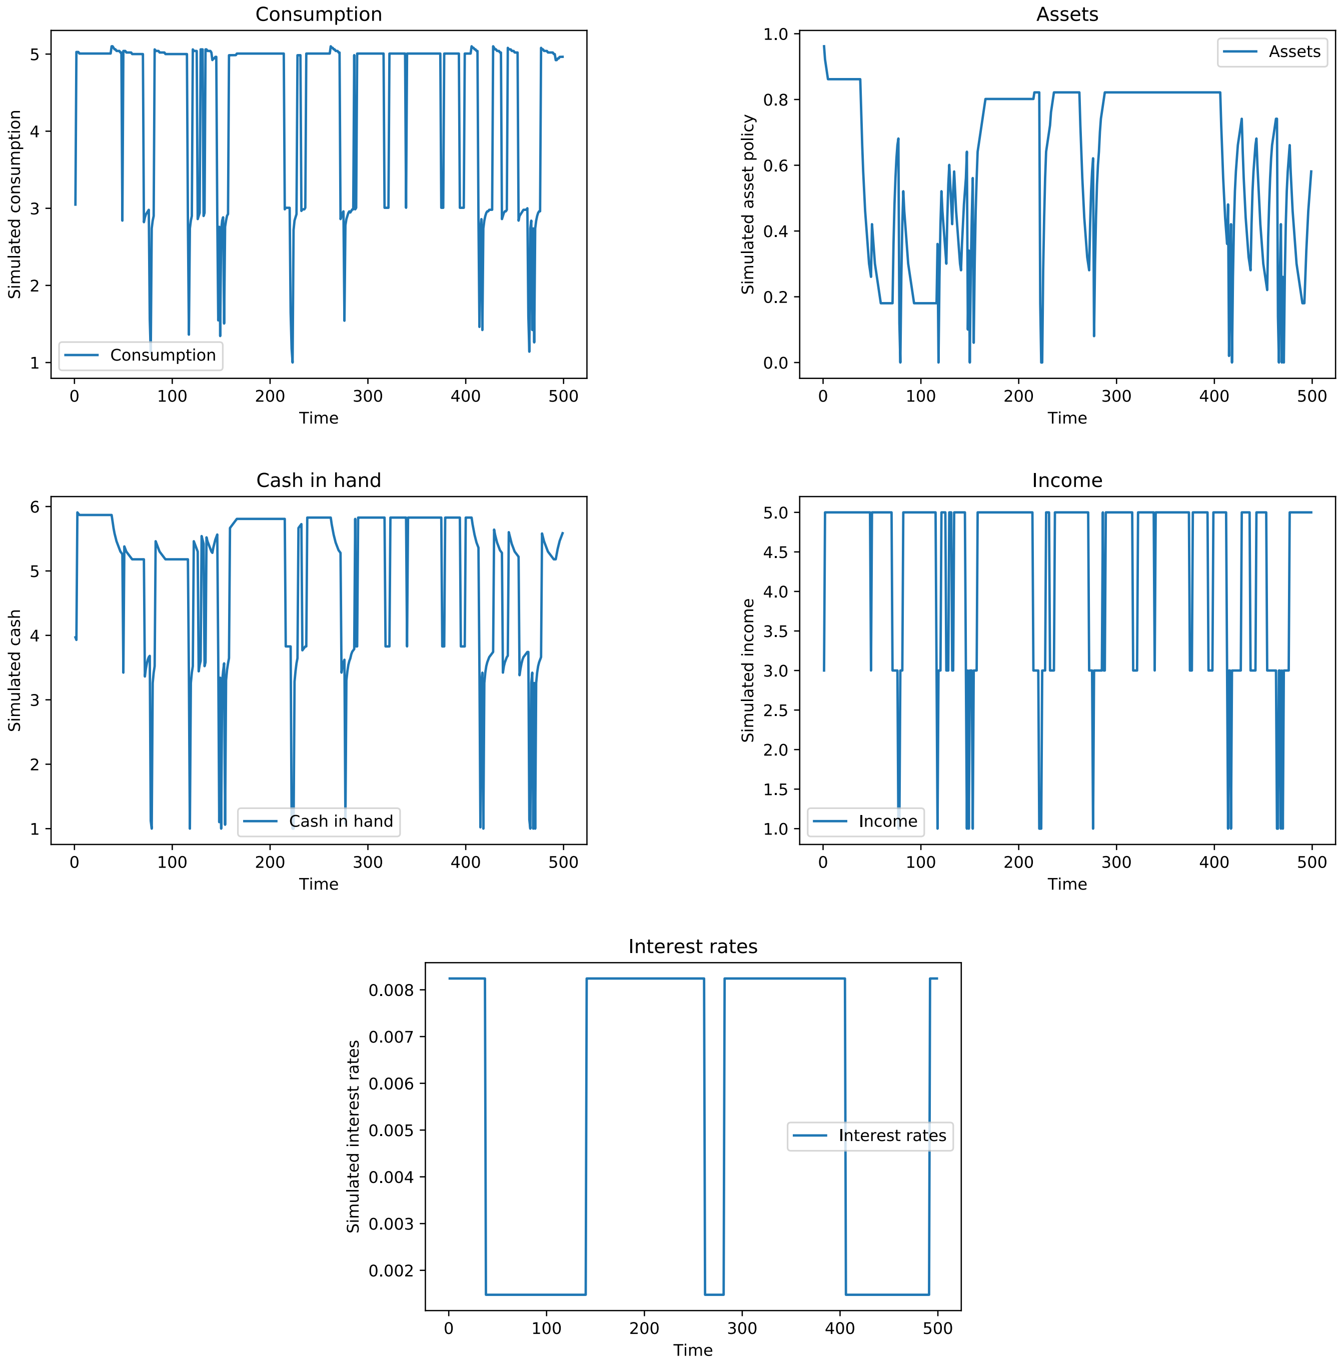

*Case 4: High volatility $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00147736 \\ 0.00824202 \end{bmatrix}$.

Figure 11: Simulation results

Similar plots follow as before. Simulated asset policy plot shows that asset levels still fluctuate throughout time. However, the correlation between assets and consumption changes from -0.02 in the low volatility case to 0.0034. The correlation between interest rate and consumption changes from 0.0018 to -0.0006. The standard deviation for assets rise from 0.2118 in the low volatility case to 0.2796 in the high volatility case. The standard deviation for cash in hand and interest rate also rises. Interestingly, while the interest rate standard deviation increases from 0.0022 to 0.0033, the standard deviation for asset increases by a larger amount.

	Consumption	Assets	Cash	Interest rate	Income
Consumption	1.0000	0.0034	0.7832	-0.0006	0.9961
Assets	0.0034	1.0000	0.3121	0.8103	-0.0125
Cash	0.7832	0.3121	1.0000	0.2375	0.7645
Interest rate	-0.0006	0.8103	0.2375	1.0000	0.0008
Income	0.9961	-0.0125	0.7645	0.0008	1.0000

	Standard deviation
Consumption	0.9063
Assets	0.2796
Cash	0.9635
Interest rate	0.0033
Income	0.9123

*Case 4: High volatility $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.00147736 \\ 0.00824202 \end{bmatrix}$.

5.5 Case 5: Low autocorrelation / Less persistence

For Case 5, I consider an interest rate process with less persistence. Instead of using the Rouwenhorst method on the AR(1) process as before, I do this by directly changing the autocorrelation of the transition matrix. I use a slightly different approach for coming up with parameter values here because changing the autocorrelation of the AR(1) process even by a very small amount results in negative state values for the interest rates.

Here, persistence should be interpreted as how long the process tends to stay in one state. A persistent process, such as those considered before in Cases from 1 to 4, tends to have less frequent transitions. On the other hand, a less persistent one will have much more frequent changes. This point will be best illustrated by considering the simulated interest rate process below.

Note that for any transition matrix of the form

$$\begin{bmatrix} p & 1-p \\ 1-q & q \end{bmatrix},$$

the first order autocorrelation is $p+q-1$. Because in our case $p = q$, I can directly change the autocorrelation in this case. Decreasing the value by 20%, I have

$$0.8\,(2p-1) = 0.8\,(0.9824744) = 0.78597952.$$

Then the new parameter values for Case 5 becomes

$$[Q] = \begin{bmatrix} 0.78597952 & 0.21402048 \\ 0.21402048 & 0.78597952 \end{bmatrix}$$

$$r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}.$$

Note that this method only works in the 2×2 matrix case. As before, $\beta = 0.96, \theta = 1.5$.

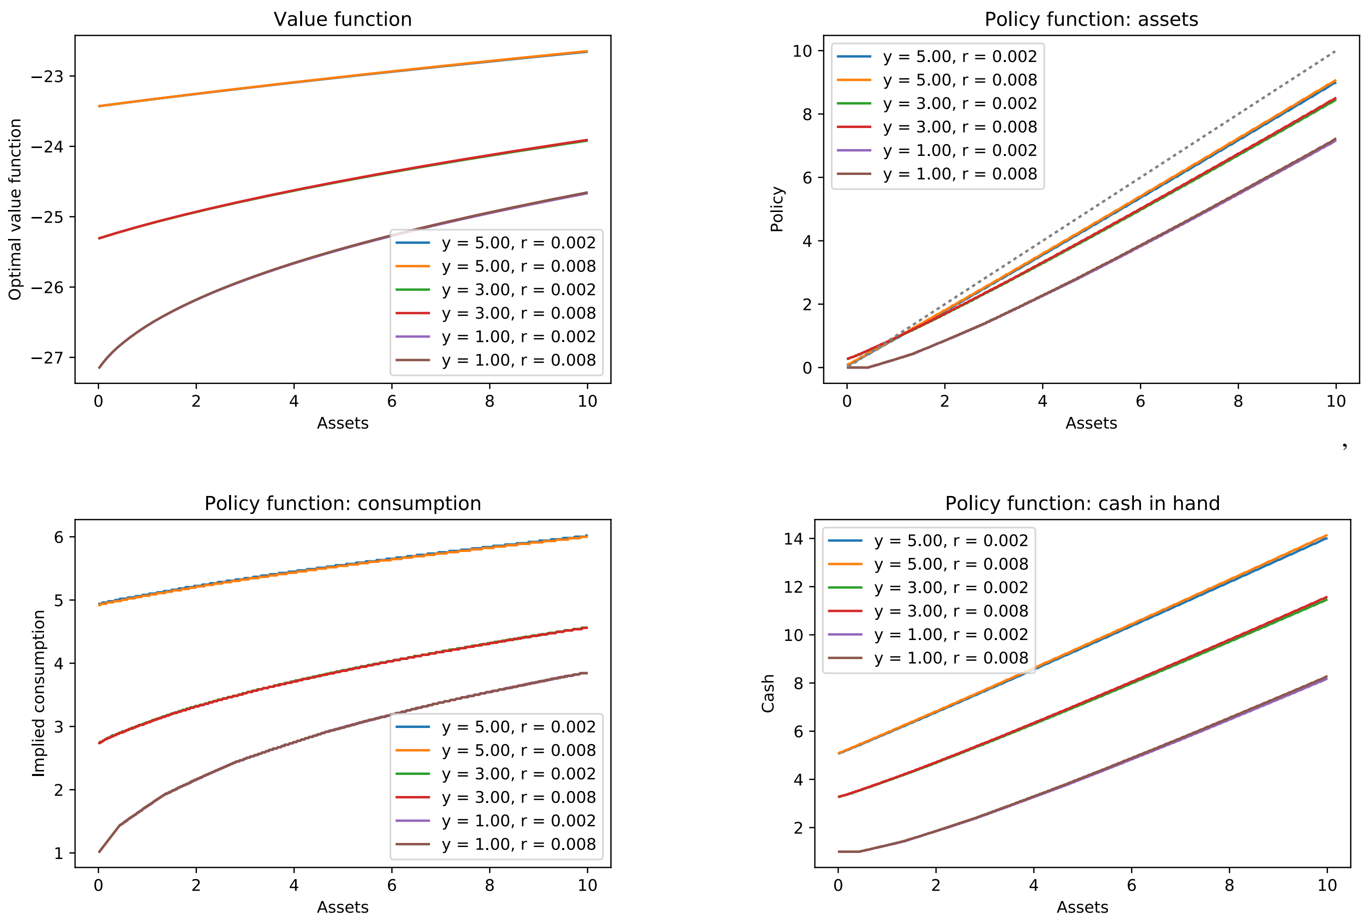

*Case 5: Low autocorrelation $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.78597952 & 0.21402048 \\ 0.21402048 & 0.78597952 \end{bmatrix}$, and $r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}$.

Figure 12: Solution results

As for the plots, the most noticeable change in this case is the simulated interest rate process. Before, there had only been a handful of interest rate changes in the time series. Here the changes are immensely more frequent. The plot for the simulated asset policy also has changed in that there are now much more frequent smaller perturbations. In previous cases, the plot seemed to be more or less charaterized by large swings - here there seems to be more frequent changes in smaller scale.

As for the statistics, the correlation between interest rate and consumption is negative at -0.0082 compared to 0.0016 for Case 6, the risk averse case. The correlation between cash in hand and assets fall to 0.0974 from 0.2756 in Case 6. Compared to Case 6, the correlation between consumption and income is very slightly stronger. The standard deviation for consumption barely changes at 0.9072 from 0.9064 for Case 6. However, the standard deviation for assets goes through a notable change, from 0.2462 in Case 6 to 0.1181 in this case.

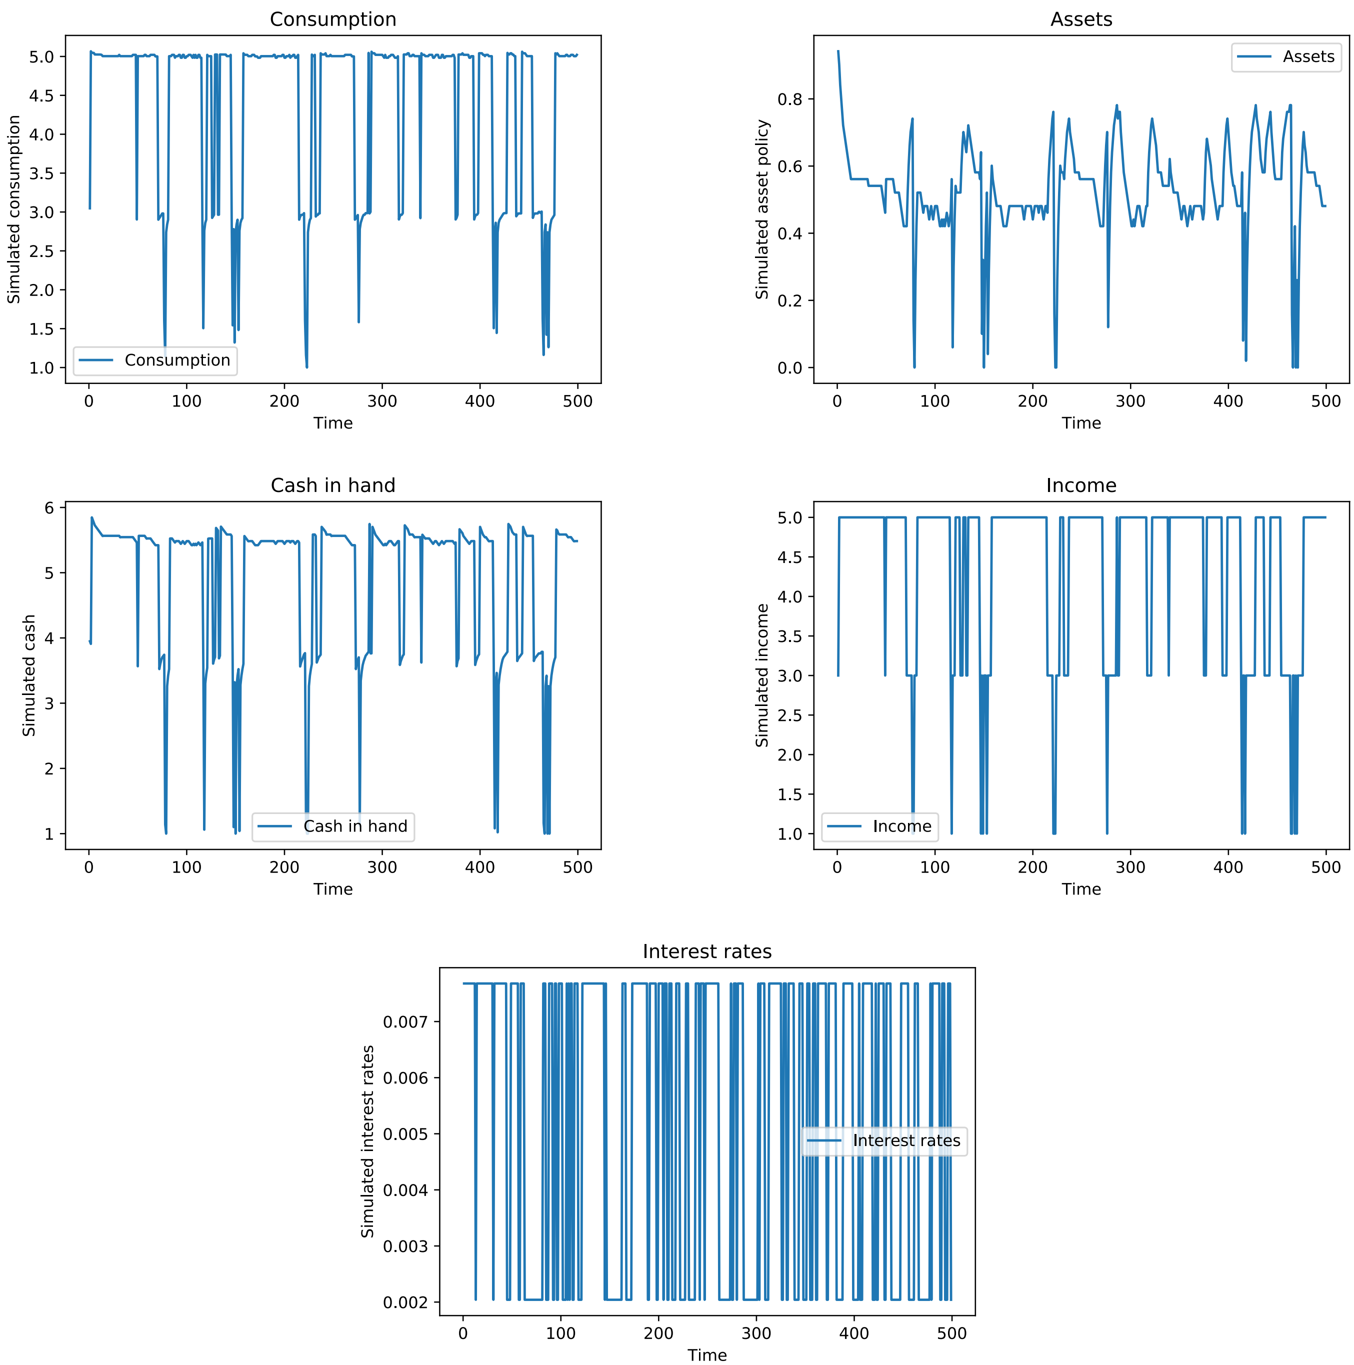

*Case 5: Low autocorrelation $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.78597952 & 0.21402048 \\ 0.21402048 & 0.78597952 \end{bmatrix}$, and $r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}$.

Figure 13: Simulation results

	Consumption	Assets	Cash	Interest rate	Income	Standard deviation
Consumption	1.0000	-0.0693	0.8107	-0.0082	0.9964	0.9072
Assets	-0.0693	1.0000	0.0974	0.0891	-0.0973	0.1181
Cash	0.8107	0.0974	1.0000	0.0107	0.7924	0.9193
Interest rate	-0.0082	0.0891	0.0107	1.0000	-0.0012	0.0028
Income	0.9964	-0.0973	0.7924	-0.0012	1.0000	0.9123

*Case 5: Low autocorrelation $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.78597952 & 0.21402048 \\ 0.21402048 & 0.78597952 \end{bmatrix}$, and $r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}$.

5.6 Case 6: Risk-averse

In the sixth case, I present the results in the same setting as the benchmark interest rate process, with the risk aversion parameter set to $\theta = 1.5$.

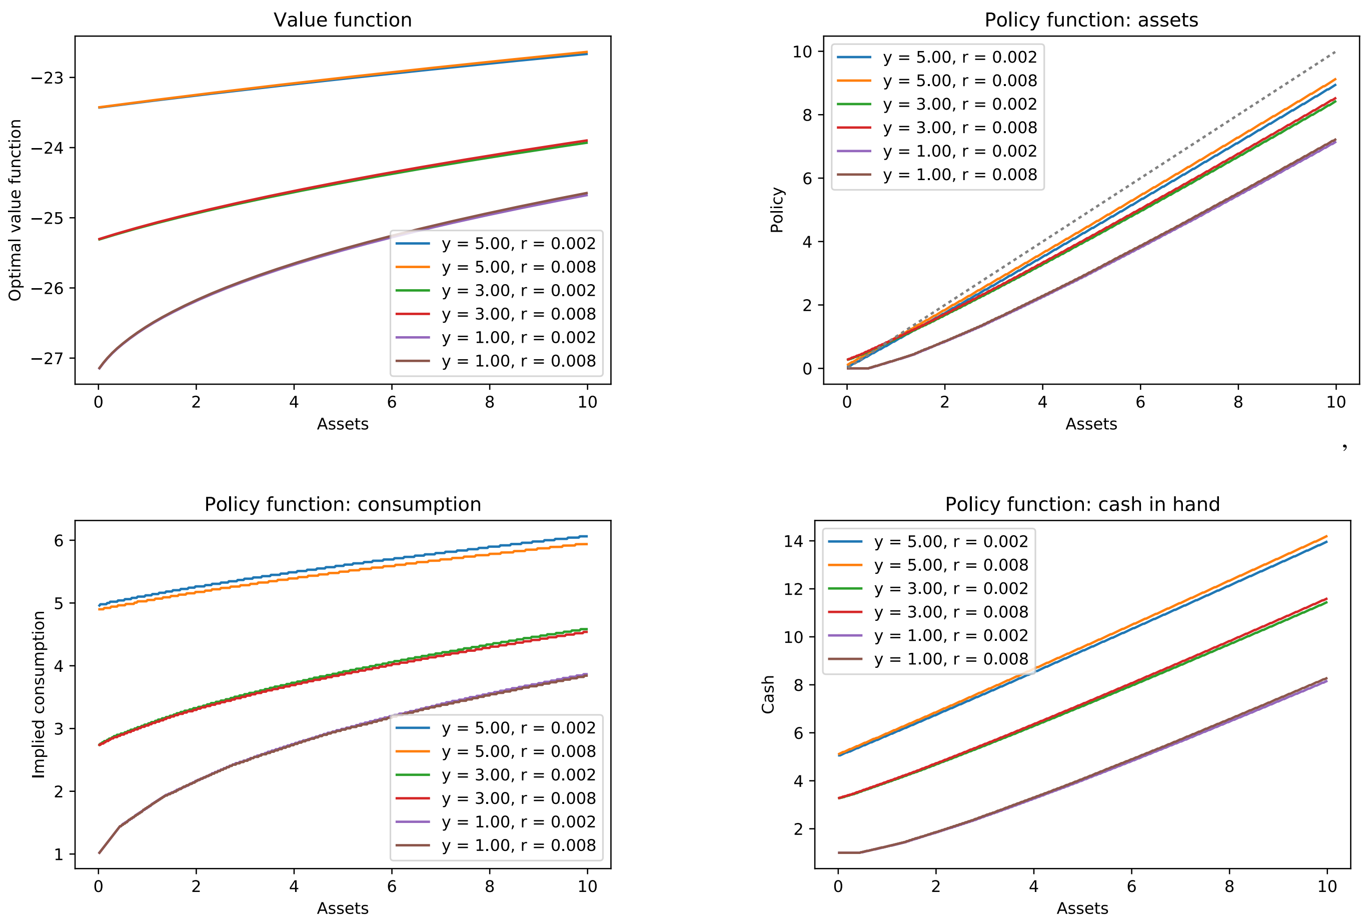

*Case 6: Risk-averse $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}$.

Figure 14: Solution results

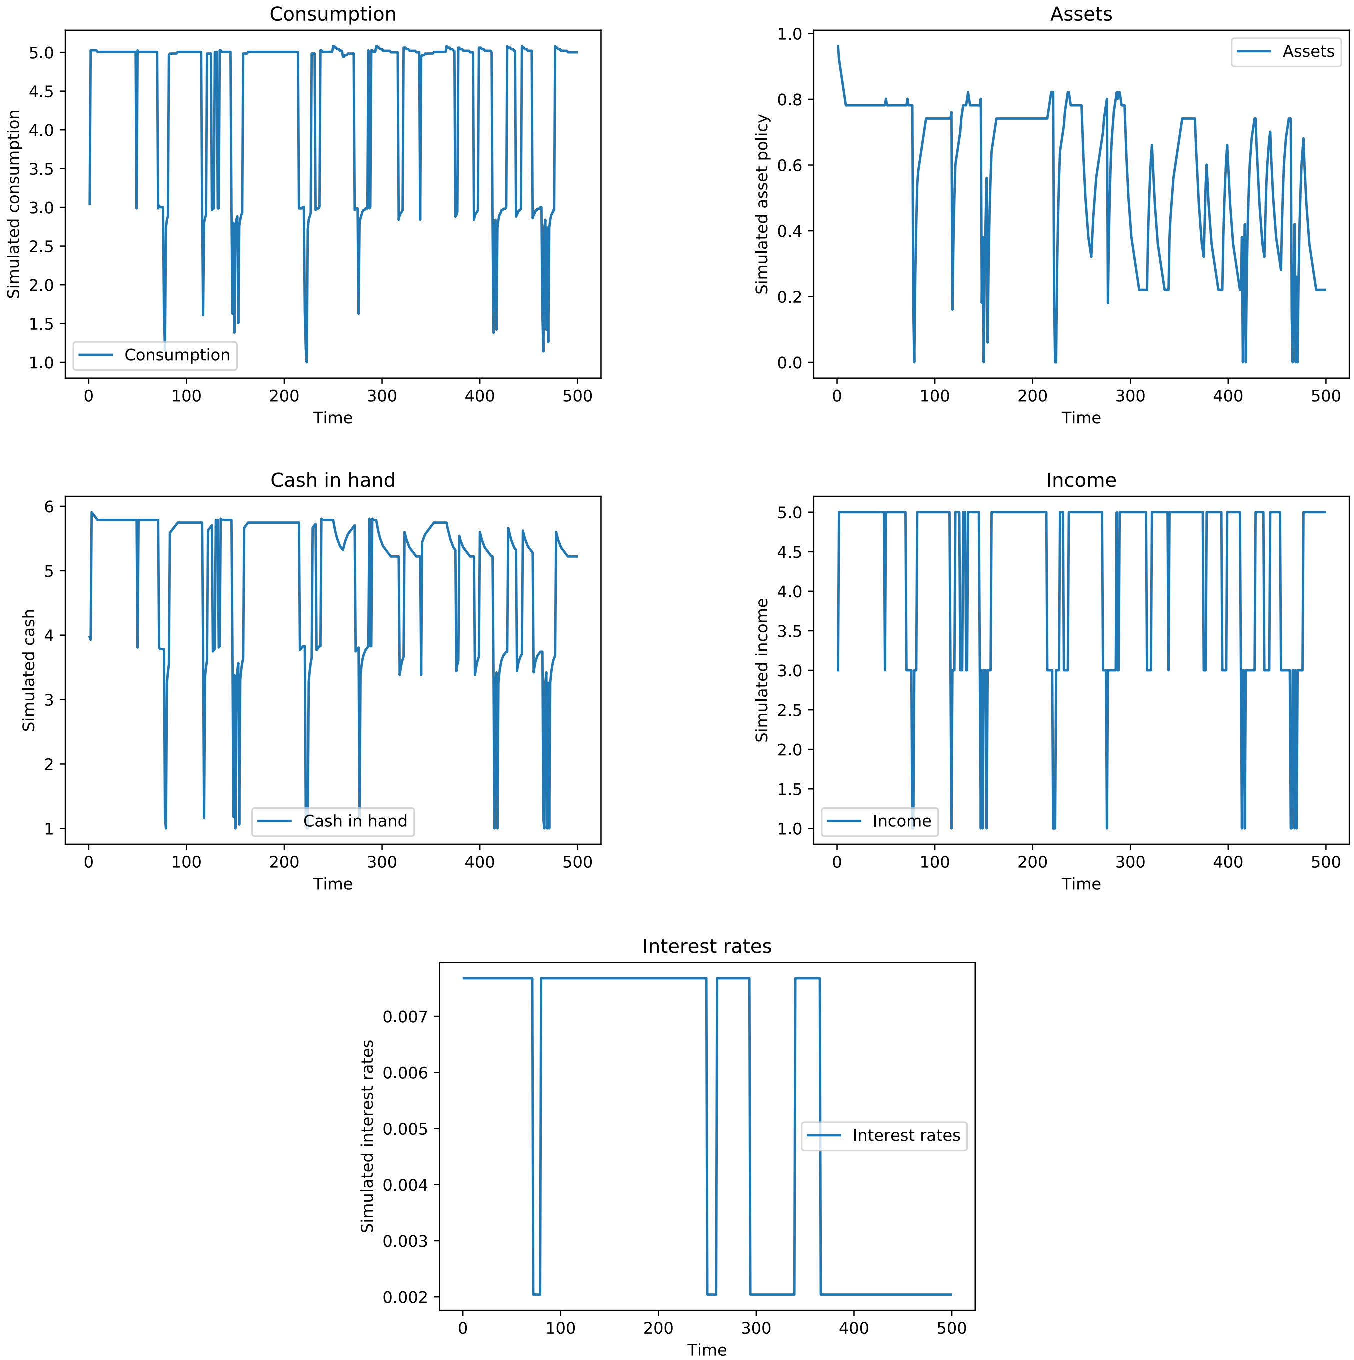

*Case 6: Risk-averse $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}$.

Figure 15: Simulation results

	Consumption	Assets	Cash	Interest rate	Income
Consumption	1.0000	-0.0042	0.7904	0.0016	0.9962
Assets	-0.0042	1.0000	0.2756	0.7883	-0.0207
Cash	0.7904	0.2756	1.0000	0.2079	0.7717
Interest rate	0.0016	0.7883	0.2079	1.0000	0.0028
Income	0.9962	-0.0207	0.7717	0.0028	1.0000

	Standard deviation
Consumption	0.9064
Assets	0.2464
Cash	0.9523
Interest rate	0.0028
Income	0.9123

*Case 6: Risk-averse $\beta = 0.96$, $\theta = 1.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}$.

5.7 Case 7: Risk-seeking

In the final case, I present the results for an individual with a relatively low risk aversion parameter, or for an individual who is a "risk-seeker." I will use $\beta = 0.96$, $\theta = 0.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}$.

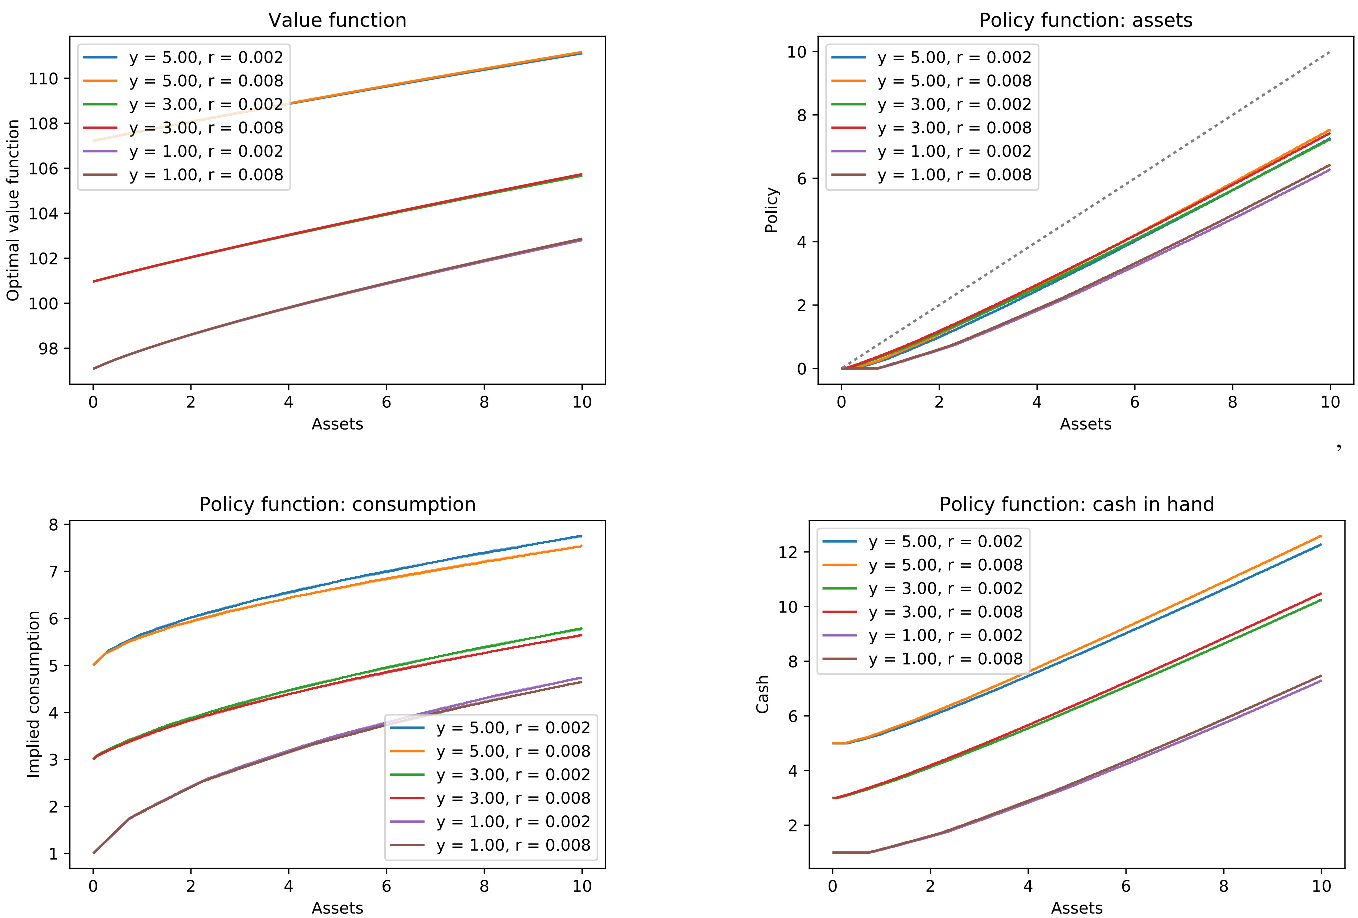

*Case 7: Risk-seeking $\beta = 0.96$, $\theta = 0.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}$.

Figure 16: Solution results

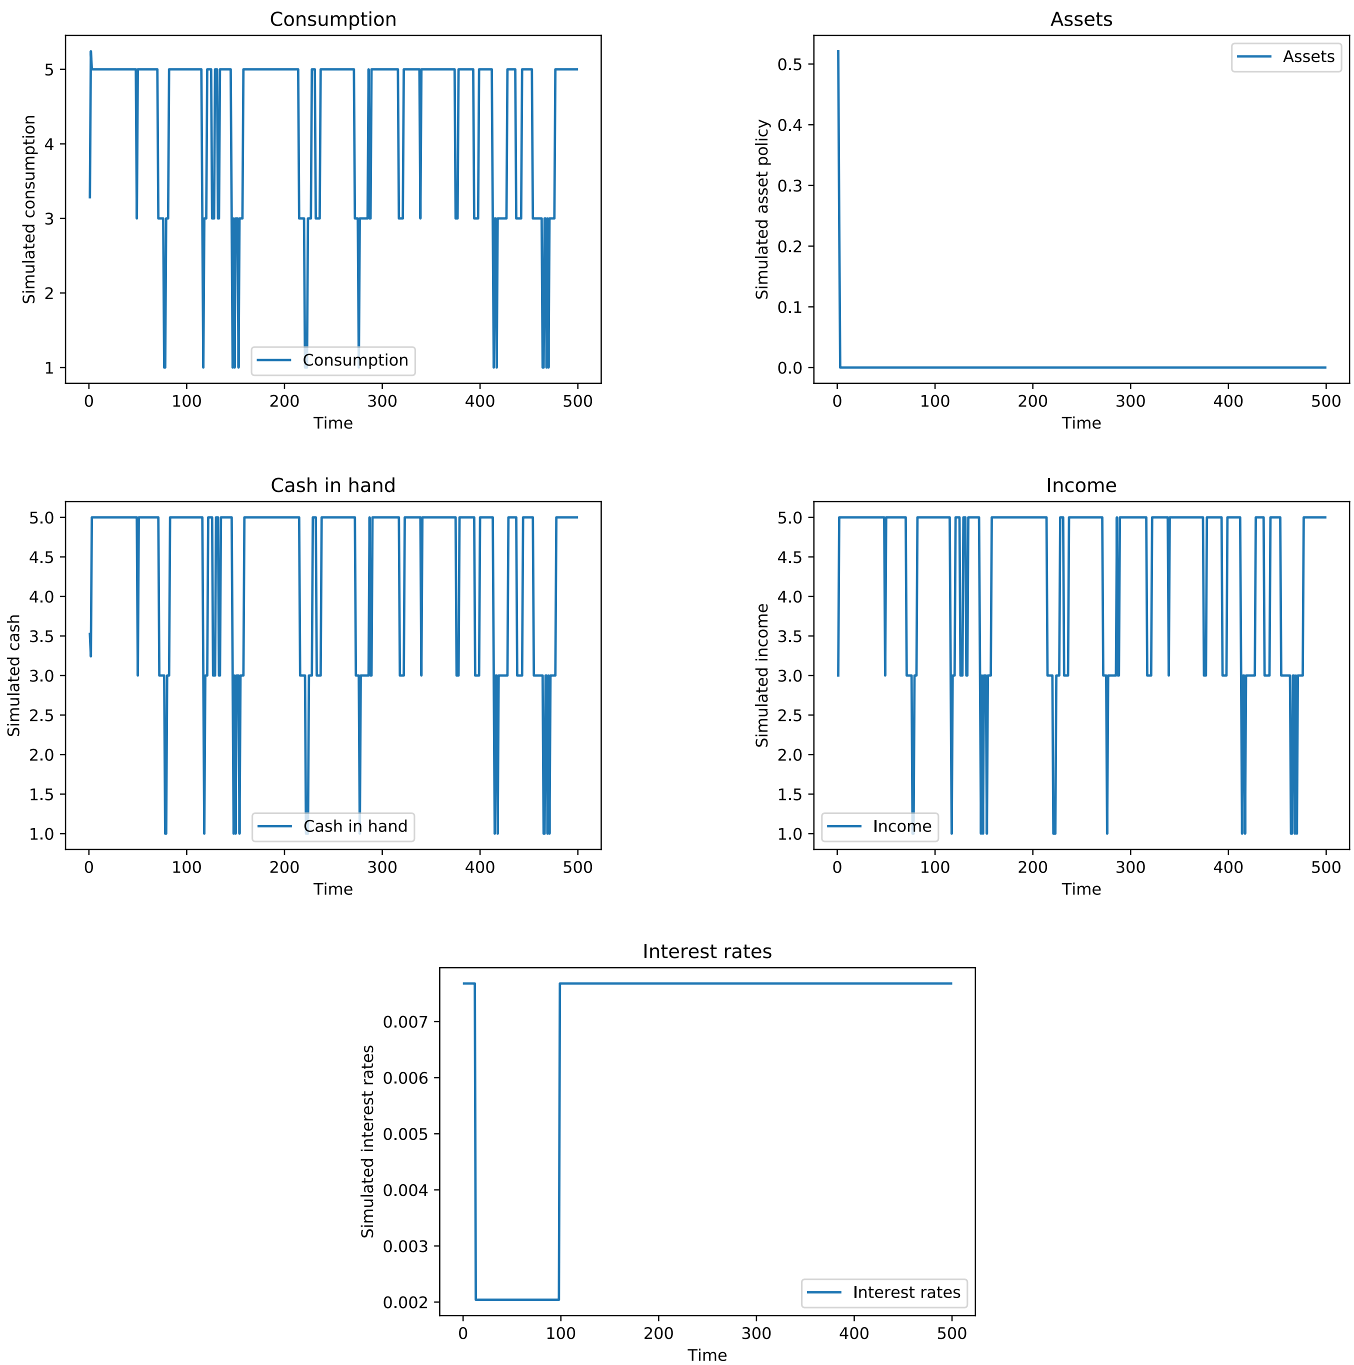

*Case 7: Risk-seeking $\beta = 0.96$, $\theta = 0.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}$.

Figure 17: Simulation results

A few comments are in order. A quick examination shows that the consumption pattern seems to closely track the income process. Again, this is a consequence of the incomplete markets assumption. Left without the means to hedge her risks, the consumer is forced to essentially "eat what you have" each period. With access to capital markets that provide the appropriate financial instruments, the agent would be able to smooth her consumption over time much better. In addition, the accumulation of assets falls to the zero level very noticeably. Looking at the plot on the right hand side of the first row for Figure 17, the level of assets fall drastically to zero almost immediately, where it stays low for the remainder of the period. In the risk-seeking case, it seems like the agent does not accumulate assets at all.

Looking at the statistics, the standard deviation for assets is very low. This is expected from the plots, as it does not change much at all. Perhaps as a result of this, the consumption rates show a higher standard deviation of 0.9121. A risk seeking agent would not mind the changes in consumption levels. In this case, the correlation between income and consumption grew even higher, shown as 1.000 for the approximation to four digits. Compared to other cases considered so far, this seems to be a rather extreme scenario.

	Consumption	Assets	Cash	Interest rate	Income
Consumption	1.0000	-0.0217	0.8085	0.0022	1.0000
Assets	-0.0217	1.0000	-0.0732	0.0215	-0.0310
Cash	0.8085	-0.0732	1.0000	0.0010	0.8090
Interest rate	0.0022	0.0215	0.0010	1.0000	0.0019
Income	1.0000	-0.0310	0.8090	0.0019	1.0000

	Standard deviation
Consumption	0.9121
Assets	0.0163
Cash	0.9141
Interest rate	0.0028
Income	0.9123

*Case 7: Risk-seeking $\beta = 0.96$, $\theta = 0.5$, $Q = \begin{bmatrix} 0.9912372 & 0.0087628 \\ 0.0087628 & 0.9912372 \end{bmatrix}$, and $r = \begin{bmatrix} 0.0020411 \\ 0.0076783 \end{bmatrix}$.

6 Conclusion

Consumption and saving are perhaps one of the most extensively studied topics in economics. In this paper, I changed the economic environment in a way that allows for a closer examination of the effect of interest changes on those two components. It is, however, crucial how the uncertainty is being incorporated into the model, and how risk-averse the agent is.

It seems to be the case that in the model, most of the changes in the asset levels are driven by income. Large changes swings in endowment translate to fluctuations in asset levels due to the incomplete markets assumption. However, there are noticeable changes that follow from interest rate hikes or drops.

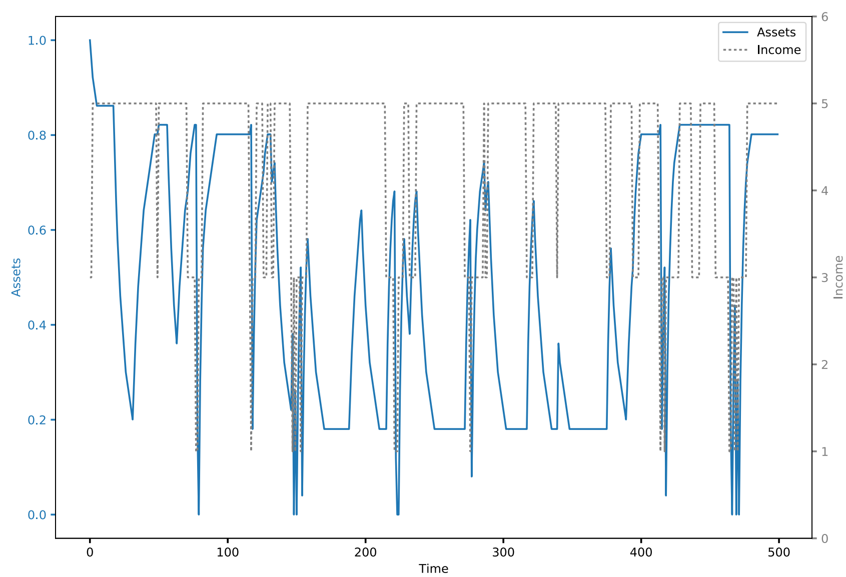

Figure 18: Assets and income

To illustrate, here is a combined plot for income and asset levels. This is the simulated asset policy plot and the simulated income plot based on the values of Case 4. Smaller decreases in income correspond to small deterioration in asset levels, and the larger decreases correspond to large drops in assets. However, even in periods of constant income levels, there are fluctuations. To get the full picture, consider the next plot.

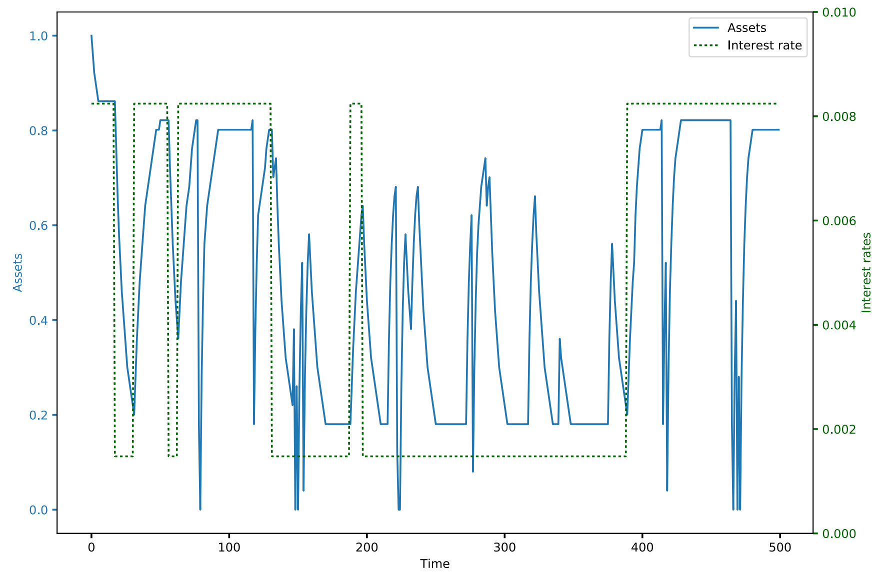

Figure 19: Assets and interest rates

Similar to the previous figure, this plot shows both the simulated asset policy and the simulated interest rate process. While there are much less fluctuations in interest rates in the persistent case, it is possible to trace some changes in asset levels back to changes in interest levels. For example, consider the time period a little before $t = 200$ in both plots. For the income process, there is a relatively sustained period of normal growth, with a constant income level. However, the asset level goes through a sharp increase and then a fall. Looking at the second plot shows that the sharp increase was induced by the hike and drop of the interest rate for that period. This is consistent with our intuition that a greater return on savings lead to an increase (albeit temporary) in asset levels.

7 Technical Appendix

In this section, I will explain some of the mathematical ideas behind this paper, along with the methodology used in obtaining the numerical solutions of the model.

7.1 Mathematical concepts

Definition 1: Markov Property *A sequence of random variables $X_1, X_2, \cdots$ has the Markov Property if*

$$\mathbb{P}\left(X_{n+1} = x | X_1 = x_1, X_2 = x_2, \cdots\right) = \mathbb{P}\left(X_{n+1} = x | X_n = x_n\right),$$

where the lower case x's denote some value in the finite state space of the chain.

Definition 2: Contraction mapping *Let (X, d) be a metric space and let $T : X \to X$. Then T is a contraction mapping if there exists a real number k, where $0 \leq \lambda < 1$, such that*

$$d\left[Tx, Ty\right] \leq \lambda d\left(x, y\right) \qquad \textit{for all } x, y \in X.$$

The full strength of a contraction mapping lies in the next theorem, where the existence and uniqueness of a fixed point is established under certain conditions.

Theorem 3: Contraction mapping principle *If X is a complete metric space, and if T is a contraction mapping of X into X, then there exists one and only one $x \in X$ such that $T(x) = x$.*

The difficulty of applying this principle may often lie in verifying whether a certain function T is actually a contraction. The following theorem provides a way of doing this.

Theorem 4: Blackwell's Sufficient Conditions *Let T be an operator on a metric space (X, d_∞), where X is a space of functions and $d_\infty(x, y) = \sup_t |x(t) - y(t)|$. Suppose that T has the following properties:*

1. *Monotonicity: For any $x, y \in X$, $x \leq y$ implies $T(x) \leq T(y)$.*
2. *Discounting: Let c denote the constant function that takes the real value c over its domain X. For any positive real c and every $x \in X$, $T(x + c) \leq T(x) + \beta c$ for some β satisfying $0 \leq \beta < 1$.*

Then T is a contraction mapping with modulus β.

These ideas provide the framework for the methodology used in this paper. To apply them, consider the set $C[a,b]$ of continuous functions on some closed interval $[a, b]$. Since the functions are continuous and are defined on a closed bounded interval, they are also bounded. Pairing this set with the $d_\infty(\cdot,\cdot)$ metric yields a complete metric space. It can be verified that the Bellman operator used in the paper is also a contraction because it satisfies Blackwell's Sufficient Conditions.

7.2 Numerical Solution

Here I explain the numerical methods used in order to obtain the solution to the models. The entire model was solved in Python, using packages such as NumPy, Matplotlib, and QuantEcon.py.

CRRA(c, theta): This function takes an array of consumption values and returns the array of CRRA utility values with the parameter theta. The formula is as follows: CRRA = $\frac{1}{1-\theta} c^{1-\theta}$.

exp_value(V0, P, Q, j, k): This function was created ad hoc to compute the particular expected value for iterating the value function. V0 is a $(N \times S \times L)$ 3-dimensional array that constitutes the value function at that iteration. P and Q are the transition matrices for income and interest rates, respectively. The last two scalars j and k are indices of the iteration step. The function handles the following computation with the double summation:

$$\mathbb{E}_{y',r'}\left[v\left(a',y',r'\right)\middle|\, y=y^j, r=r^k\right] = \sum_{m=1}^{2}\sum_{n=1}^{3}\mathbb{P}\left(y^{'}=y^n\middle|\, y=y^j\right)\mathbb{P}\left(r'=r^m\middle|\, r=r^k\right)v\left(a^{'},y^{'}=y^n,r^{'}=r^m\right).$$

The probabilities are given by the entries in the transition matrices P and Q.

Main algorithm: Given P, Q, and array a, this part of the code returns the converged value function. It does this by iterating the value function until convergence.

The maximization operation, which is noticeably a part of the Bellman operator, is done element-wise. What I mean by this is that, after computing many possible values of V1 for each a' on a grid, I use np.amax to return the maximal element on that grid. Then that maximal element is stored in V1. The corresponding a' is also stored in order to later yield the policy function.

This is done exactly $N \times S \times L$ times for each iteration, where N is the number of elements on the grid, S is the number of states for income, and L is the number of states for interest rates.

After each iteration, the code computes the error, which is defined as the maximum element of the element-wise absolute difference of the V0 and V1 arrays. This was written as err = np.amax(np.abs(V0-V1)). If the error falls below a very small constant (say, 0.001), then convergence is complete, and the code stores the V1 and a1 arrays. Otherwise, the code repeats the whole process.

Simulation. In the second part of the code, I use the computed value and policy functions to apply them to a simulated economy. The random values for income and interest rates are created using a Python package for Markov Chain simulation. Grids are created as arrays for consumption, assets, and cash on hand. For each element on the time grid t, I store a value for consumption, assets, and cash on hand by searching for indices on previously computed arrays for implied consumption, policy function, and cash on hand.

qe.MarkovChain(P, state_values=y[:,0]): This function from the QuantEcon.py library takes a transition matrix P and an array of state values and creates a class for a finite-state discrete-time Markov Chain. I only use it to create an array of Markov Chain simulations. This is done by calling .simulate on the defined object with a specified length for the array.

7.3 Table of parameters

To sum up, here is a table of all the symbols used.

variable or parameter	meaning
$V(\cdot)$	value function
$u(\cdot)$	utility
β	discount factor
θ	risk aversion
a_t	assets
c_t	consumption
y_t	income
r_t	interest rate
P	transition matrix for income
Q	transition matrix for interest rate

8 References

Aiyagari, S Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics*, 109(3): 659–684.

Bellman, Richard E. 1957. *Dynamic Programming.* Princeton University Press.

Bewley, Truman. 1977. "The Permanent Income Hypothesis: A Theoretical Formulation." *Journal of Economic Theory*, 16(2): 252–292.

Brock, William A., and Leonard J. Mirman. 1972. "Optimal Economic Growth and Uncertainty: The Discounted Case." *Journal of Economic Theory*, 4(3): 479–513.

Chamberlain, Gary, and Charles A Wilson. 2000. "Optimal Intertemporal Consumption Under Uncertainty." *Review of Economic Dynamics*, 3(3): 365–395.

Chipeniuk, Karsten O, Todd B Walker, and Nets Hawk Katz. 2016. "Approximate Aggregation in Dynamic Economies."

den Haan, Wouter J. 1997. "Solving Dynamic Models with Aggregate Shocks and Heterogeneous Agents." *Macroeconomic Dynamics*, 1(2): 355–386.

den Haan, Wouter J, Kenneth L Judd, and Michel Juillard. 2010. "Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty." *Journal of Economic Dynamics and Control*, 34(1): 1–3.

Díaz-Giménez, Javier, and Edward C Prescott. 1997. "Real returns on government debt: A general equilibrium quantitative exploration." *European Economic Review*, 41(1): 115–137.

Friedman, Milton. n.d.. *A Theory of the Consumption Function.*

Hamilton, James D. 2005. "What's Real About the Business Cycle?" *Federal Reserve Bank of St.Louis Review*, 435–452.

Huggett, Mark. 1993. "The risk-free rate in heterogeneous-agent incomplete-insurance economies." *Journal of Economic Dynamics and Control*, 17(5-6): 953–969.

Imrohoroglu, Ayse. 1989. "Cost of Business Cycles with Indivisibilities and Liquidity Constraints." *Journal of Political Economy*, 97(6).

Krusell, Per, and Anthony A Smith. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, 106(5): 867–896.

Ljungqvist, Lars, and Thomas John Sargent. 2004. *Recursive Macroeconomic Theory.* . 2 ed., MIT Press.

Maliar, Lilia, Serguei Maliar, and Fernando Valli. 2010. "Solving the incomplete markets model with aggregate uncertainty using the Krusell-Smith algorithm." *Journal of Economic Dynamics and Control*, 34(1): 42–49.

Marcet, Albert, and Kenneth J Singleton. 1999. "Equilibrium Asset Prices and Savings of Heterogeneous Agents in the Presence of Incomplete Markets and Portfolio Constraints." *Macroeconomic Dynamics*, 3: 243–277.

Marcet, Albert, Francesc Obiols-Homs, and Philippe Weil. 2007. "Incomplete markets, labor supply and capital accumulation." *Journal of Monetary Economics*, 54: 2621–2635.

Obstfeld, Maurice, and Kenneth Rogoff. 1996. *Foundations of International Macroeconomics.* MIT Press.

Schechtman, Jack. 1976. "An Income Fluctuation Problem." *Journal of Economic Theory*, 12: 218–241.

Sotomayor, Marlida A. de Oliveira. 1984. "On Income Fluctuations and Capital Gains." *Journal of Economic Theory*, 32(1): 14–35.

Yaari, Menahem E. 1976. "A Law of Large Numbers in the Theory of Consumer's Choice under Uncertainty." *Journal of Economic Theory*, 12: 202–217.

Young, Eric R. 2010. "Solving the incomplete markets model with aggregate uncertainty using the Krusell Smith algorithm and nonstochastic simulations." *Journal of Economic Dynamics and Control*, 34(1): 36–41.

Zhu, Shenghao. 2009. "Existence of Equilibrium in an Incomplete Market Model with Endogenous Labor Supply."