

Machine Learning Calvo's Optimal Plan

Thomas J. Sargent

New York University

Ziyue Yang

Australian National University

March 14, 2025

Abstract

A computer program calculates a pair of infinite sequences of money creation and price level inflation rates $(\vec{\theta}, \vec{\mu})$ that maximizes a benevolent time 0 government's objective function. The limit of a monotonically declining sequence of continuation values is a worst continuation value associated with a "timeless perspective". The time-invariant inflation rate associated with the worst continuation Ramsey plan is not the inflation rate associated with a restricted Ramsey plan in which a time 0 government is constrained to choose a time-invariant money creation rate. We Bellmanize the continuation Ramsey problem.

Key words: Artificial intelligence, machine learning, fake data, Ramsey plan, time inconsistency, open loop, closed loop, inflation, money.

Dedication

We are honored to have been invited to celebrate the contributions of Michel Jullard to quantitative dynamic macroeconomics by trying to study aspects of some of Michel's favorite topics with some of the important tools that Michel has taught and enabled us to use through the creation of dynare. Michel's work illustrates how to combine high human intelligence with machine learning in ways that help make good public policies. Along with [Kydland and Prescott \(1977, 1980\)](#), [Calvo \(1978\)](#) set the stage for the conceptual and computational challenges that Michel took the lead in confronting. That explains our choice of topics.

1 Introduction

Many applications of *machine learning* compute a nonlinear function $f : X \rightarrow Y$ that satisfies context-specific conditions. Popular contexts include:

- (a) f maximizes some functional or solves some functional equation.
- (b) $\{x_i, y_i\}_{i=1}^I \in X^I \times Y^I$ is a data set and f is a non-linear least squares regression function.

This paper provides interrelated examples of both types. Our first example runs regressions on 'fake data' generated by our second application. We run those regressions to uncover interpretable economic structure concealed by outcomes of our second application, a discrete-time version of an optimum problem of [Calvo \(1978\)](#) that seeks a sequence of money growth rates $\{\mu_t\}_{t=0}^{\infty}$ that maximizes a government's objective function at time 0. The optimizer takes the form of a function f that maps times $t \in X = \{0, 1, 2, \dots\}$ into \mathbb{R} . Let:

- p_t be the log of the price level,
- m_t be the log of nominal money balances,
- $\theta_t = p_{t+1} - p_t$ be the net rate of inflation between t and $t + 1$,
- $\mu_t = m_{t+1} - m_t$ be the net rate of growth of nominal balances.

The government's problem is cast in terms of these components:

- $\vec{\mu} = \{\mu_t\}_{t=0}^{\infty}$ is a time series of money growth rates,
- $\mu^t = \{\mu_{t+s}\}_{s=0}^{\infty}$ is a time t **future** or **tail** of a sequence of money growth rates,

- $\vec{\theta} = \{\theta_t\}_{t=0}^{\infty}$ is a sequence of inflation rates in the price level,
- a function g that maps the future μ^t of $\vec{\mu}$ at t into the inflation rate at t , so that $\theta_t = g(\mu^t)$,
- a social welfare criterion

$$v_0 = \sum_{t=0}^{\infty} \beta^t r(\mu^t) \quad (1)$$

where $r(\mu^t) = s(g(\mu^t), \mu_t)$, g , and $s(\cdot, \cdot)$ are known functions and $\beta \in (0, 1)$.

The function g describes the behavior of private agents and markets that determine the inflation rate π_t at t as a function of the future μ^t of money growth rates from time t forward. The government knows the functions g and r and wants an **open loop** plan $\mu_t = f(t)$, i.e., a function of time that describes a sequence of money growth rates that maximizes welfare criterion v_0 defined in (1). Calvo (1978) is our source of s and g .

The plan that optimizes criterion (1) is **time inconsistent**. The presence of the function g in the government's objective function tells it to recognize effects that μ_s for all $s \geq 0$ have on θ_0 when it chooses a time series $\vec{\mu}$. A government that at time 1 chooses a sequence $\vec{\mu}$ to maximize a welfare criterion

$$v_1 = \sum_{t=1}^{\infty} \beta^{t-1} r(\mu^t) \quad (2)$$

would not care about μ_0 or θ_0 and consequently would select a different $\vec{\mu}$ time series than the maximizer of criterion (1).

This paper extends an analysis of a linear quadratic of Calvo's model presented in Sargent and Yang (2025a,b) in which the representative household's one-period utility function is a quadratic function of inflation and the money growth rate. That special LQ setting affected both our choice of tools and some of the findings in ways that we shall highlight below. While discounted linear quadratic dynamic programming was the appropriate tool for our LQ setting, it isn't here. And while linear regressions on "fake data" were adequate machine learning torches there, they aren't enough here. Furthermore, in this paper we go further by studying a version of a classic "aggregation over time" problem by setting the decision interval specifies the timing protocol governing when decision makers act: the government when it sets a money growth rate, and the representative household when it sets an expected inflation rate and a demand for money. We study how outcomes depend on the time increment. By driving the time increment to zero we approach Calvo's continuous time specification.

Section 2 describes components of the model and poses the government's planning problem. Section 3 describes a restricted plan that by construction is time consistent. Section

4 writes the government planner’s time 0 objective as a function of a money growth rate sequence and hands it over to a **gradient ascent** optimizer, an application of the same approach applied in a linear-quadratic setting by [Sargent and Yang \(2025a\)](#). Section 5 extends the discrete-unit-time interval model of section 2 to a discrete Δ -time-interval model, allowing us to approximate the original continuous time setting of [Calvo \(1978\)](#). We apply our gradient ascent algorithm to this model and compare outcomes to the unit-time-interval specification. Our gradient ascent recovers an **open loop** representation of an optimal plan. Section 6 uses some “human intelligence” to guide specification of some nonlinear least squares regressions that we apply to “fake data” in the form of the open loop optimal plans that we computed in section 5. Those regressions detect two recursive representations of the optimal plan; the representations differ in their specification of the key state variable. Section 7 uses insights of [Chang \(1998\)](#) to decide which of the two representations is more enlightening economically and extends an analysis presented in [Sargent and Yang \(2025b\)](#) to formulate the government’s problem as a recursive decision problem. Section 8 offers concluding remarks about the roles artificial and human intelligence in our Calvo model laboratory.

2 The Model

Calvo’s model focuses on intertemporal tradeoffs between:

- utility accruing from a representative agent’s anticipations of future deflation that lower the agent’s cost of holding real money balances and thereby induces the agent to increase his stock of real money balances, and
- social costs associated with the distorting taxes that a government levies to acquire the paper money that it withdraws from circulation in order to generate prospective deflation.

The model features:

- rational expectations,
- costly government actions at all dates $t \geq 1$ that increase the representative agent’s utilities at dates before t .

The model combines a demand function for real balances formulated by [Cagan \(1956\)](#) with the perfect foresight or rational expectations assumed by [Sargent and Wallace \(1973\)](#) and [Calvo \(1978\)](#).¹

¹Work of [Olivera \(1970, 1971\)](#) about passive money influenced [Sargent and Wallace \(1973\)](#).

2.1 Components

There is no uncertainty. A representative agent's demand for real balances is governed by a perfect foresight version of a Cagan (1956) demand function:

$$m_t - p_t = -\alpha(p_{t+1} - p_t), \quad \alpha > 0, \quad (3)$$

for all $t \geq 0$.

Equation (3) asserts that the demand for real balances is inversely related to the representative agent's expected rate of inflation. Because there is no uncertainty, the expected rate of inflation equals the actual rate of inflation.²

Subtracting equation (3) at time t from the same equation at time $t + 1$ gives:

$$\mu_t - \theta_t = -\alpha\theta_{t+1} + \alpha\theta_t,$$

or equivalently,

$$\theta_t = \frac{\alpha}{1 + \alpha}\theta_{t+1} + \frac{1}{1 + \alpha}\mu_t. \quad (4)$$

Because $\alpha > 0$, $0 < \frac{\alpha}{1 + \alpha} < 1$, so difference equation (4) in the θ sequence with sequence $\vec{\mu}$ as the "forcing sequence" is stable when "solved forward."

Definition 2.1. For scalar b_t , let L^2 be the space of sequences $\{b_t\}_{t=0}^{\infty}$ that satisfy

$$\sum_{t=0}^{\infty} b_t^2 < +\infty.$$

We say that a sequence that belongs to L^2 is *square summable*.

When we assume that $\vec{\mu} = \{\mu_t\}_{t=0}^{\infty}$ is square summable and also require that $\vec{\theta} = \{\theta_t\}_{t=0}^{\infty}$ is square summable, the linear difference equation (4) can be solved forward to get:

$$\theta_t = \frac{1}{1 + \alpha} \sum_{j=0}^{\infty} \left(\frac{\alpha}{1 + \alpha} \right)^j \mu_{t+j}, \quad t \geq 0. \quad (5)$$

The government $\vec{\mu}$ once and for all at time $t = 0$ and wants to maximize

$$V = \sum_{t=0}^{\infty} \beta^t s(\theta_t, \mu_t), \quad (6)$$

²When there is no uncertainty, an assumption of **rational expectations** becomes equivalent to **perfect foresight**.

where $\beta \in (0, 1)$ is a discount factor, and $s(\theta_t, \mu_t)$ is a one-period welfare function of a benevolent government.

To capture a tradeoff between the utility of real balances and the social costs of money creation, we assume that $s(\theta_t, \mu_t)$ is a function of θ_t and μ_t that satisfies:

$$s(\theta_t, \mu_t) = U(-\alpha\theta_t + u_0) - c(\mu_t), \quad (7)$$

where

- The government values a representative household's utility of real balances at time t according to the utility function $U : \mathbb{R} \rightarrow \mathbb{R}$.
- The government incurs social costs $c(\mu_t)$ when it changes the stock of nominal money balances at rate μ_t at time t with $c : \mathbb{R} \rightarrow \mathbb{R}$ measuring the social costs of money creation.

We assume that U is twice continuously differentiable, with $U'(0) > 0$, $U''(0) < 0$, and satisfies Inada-type conditions: $\lim_{x \rightarrow 0^+} U'(x) = \infty$ and $\lim_{x \rightarrow \infty} U'(x) = 0$. This prevents corner solutions.

We assume that cost function c is twice continuously differentiable, with $c(0) = 0$, $c'(\mu_t) > 0$ for $\mu_t > 0$, $c'(\mu_t) < 0$ for $\mu_t < 0$, and $c''(\mu_t) > 0$. These assumptions ensure the optimization problem is well-behaved and has a unique solution.

We focus on the case where the utility function U is logarithmic, i.e., $U(x) = \log(x)$, and the cost function c is polynomial. Specifically, we let

$$U(-\alpha\theta_t + u_0) = \log(-\alpha\theta_t + u_0), \quad c(\mu_t) = \frac{c_2}{2}\mu_t^2 + \frac{c_3}{3}\mu_t^3 + \frac{c_4}{4}\mu_t^4, \quad (8)$$

where $\alpha > 0$, $u_0 > 0$, $c_2 > 0$, $c_3^2 < 3c_2c_4$, and $c_4 > 0$ are parameters.

The Ramsey planner chooses a vector of money growth rates $\vec{\mu}$ to maximize criterion (6) subject to equation (5) and the restriction:

$$\vec{\theta} \in L^2. \quad (9)$$

Equations (5) and (9) imply that $\vec{\theta}$ is a function of $\vec{\mu}$. In particular, the inflation rate θ_t satisfies:

$$\theta_t = (1 - \lambda) \sum_{j=0}^{\infty} \lambda^j \mu_{t+j}, \quad t \geq 0, \quad (10)$$

with

$$\lambda = \frac{\alpha}{1 + \alpha}.$$

2.2 Basic Objects

Mathematical objects in play include a pair of sequences of inflation rates and money growth rates

$$(\vec{\theta}, \vec{\mu}) = \{\theta_t, \mu_t\}_{t=0}^{\infty},$$

and a planner's value function:

$$\begin{aligned} V &= \sum_{t=0}^{\infty} \beta^t [U(-\alpha\theta_t + u_0) - c(\mu_t)] \\ &= \sum_{t=0}^{\infty} \beta^t \left[\log(-\alpha\theta_t + u_0) - \left(\frac{c_2}{2}\mu_t^2 + \frac{c_3}{3}\mu_t^3 + \frac{c_4}{4}\mu_t^4 \right) \right]. \end{aligned} \tag{11}$$

Definition 2.2. A Ramsey planner chooses $\vec{\mu}$ to maximize the government's value function (11) subject to equation (10). A $\vec{\mu}$ that solves this problem is called a *Ramsey plan*.

2.3 Timing Protocol

Calvo (1978) instructs the government to choose the money growth sequence $\vec{\mu}$ once and for all, at or before time 0. By choosing the money growth sequence $\vec{\mu}$, the government indirectly chooses the inflation sequence $\vec{\theta}$. So the government effectively chooses a bivariate **time series** $(\vec{\mu}, \vec{\theta})$. The government's problem is **static**: it chooses all components of a bivariate time series $(\vec{\mu}, \vec{\theta})$ at time 0.

2.4 Approximation and Truncation Parameter T

We start by guessing that the sequence $\vec{\mu}$ converges under a Ramsey plan:

$$\lim_{t \rightarrow +\infty} \mu_t = \bar{\mu}.$$

Convergence of μ_t to $\bar{\mu}$ together with equation (10) imply that:

$$\lim_{t \rightarrow +\infty} \theta_t = \bar{\theta}.$$

We'll guess a time T large enough that μ_t has gotten very close to the limit $\bar{\mu}$. Then we'll approximate $\vec{\mu}$ by a truncated vector with the property:

$$\mu_t = \bar{\mu} \quad \forall t \geq T.$$

Similarly, we'll approximate $\vec{\theta}$ with a truncated vector with the property:

$$\theta_t = \bar{\theta} \quad \forall t \geq T.$$

In light of our approximation that $\mu_t = \bar{\mu}$ for all $t \geq T$, we seek a function that takes

$$\tilde{\mu} = \left[\mu_0 \quad \mu_1 \quad \cdots \quad \mu_{T-1} \quad \bar{\mu} \right]$$

as an input and gives as an output the vector

$$\tilde{\theta} = \left[\theta_0 \quad \theta_1 \quad \cdots \quad \theta_{T-1} \quad \bar{\theta} \right],$$

where $\bar{\theta} = \bar{\mu}$ and θ_t satisfies:

$$\theta_t = (1 - \lambda) \sum_{j=0}^{T-1-t} \lambda^j \mu_{t+j} + \lambda^{T-t} \bar{\mu}, \quad (12)$$

for $t = 0, 1, \dots, T - 1$.

Having defined vector $\tilde{\mu}$ and computed the vector $\tilde{\theta}$ using formula (12), we can rewrite the government's value function (11) as

$$\tilde{V}(\tilde{\theta}) = \sum_{t=0}^{T-1} \beta^t [U(-\alpha\theta_t + u_0) - c(\mu_t)] + \frac{\beta^T}{1 - \beta} [U(-\alpha\bar{\theta} + u_0) - c(\bar{\mu})], \quad (13)$$

where θ_t for $t = 0, 1, \dots, T - 1$ satisfies formula (12).

3 A Restricted Optimal Plan

Our Ramsey planner chooses $\vec{\mu}$ to maximize the government's value function (11) subject to equations (10). It is useful to consider a distinct problem in which a planner again chooses $\vec{\mu}$ to maximize the government's value function (11), but now subject to equation (10) and the additional restriction that $\mu_t = \bar{\mu}$ for all t . The solution of this problem is a time-invariant $\mu_t = \mu^{CR}$ for all $t \geq 0$ that attains a value V^{CR} of the Ramsey plan.³

³Computing μ^{CR} with the section 4 gradient ascent algorithm by setting the sequence μ_t to be constant for all t and iterating until convergence is easy.

4 Brute Force Gradient Ascent Algorithm

In this section, we compute the Ramsey plan by applying the same gradient ascent algorithm that we used in [Sargent and Yang \(2025a\)](#). We write a Python function that inputs a truncated $\vec{\mu}$ sequence and returns the time 0 objective of the Ramsey planner, then hand that function to a gradient ascent optimizer. We initiate the algorithm with a money growth sequence $\mu_t = 0$ for all $t \geq 0$, then iteratively update $\vec{\mu}$ until convergence. Figure 1 plots the Ramsey plan's μ_t and θ_t for $t = 0, \dots, T$ against t computed by the




Figure 1: Ramsey plan ($\vec{\mu}$ and $\vec{\theta}$)

algorithm. While θ_t is less than μ_t for low values of t , it eventually converges to the limiting value $\bar{\mu}$ of the sequence $\{\mu_t\}$ as $t \rightarrow +\infty$, a consequence of how formula (5) makes θ_t be a weighted average of future μ_t 's.

To compute a sequence $\{v_t\}_{t=0}^T$ of “continuation values”

$$v_t = \sum_{j=t}^{\infty} \beta^{j-t} s(\theta_j, \mu_j)$$

along a Ramsey plan, we'll start at our truncation date T and compute

$$v_T = \frac{1}{1-\beta} s(\bar{\mu}, \bar{\mu}).$$

Then starting from $t = T - 1$, we'll iterate backwards on the recursion

$$v_t = s(\theta_t, \mu_t) + \beta v_{t+1}$$

for $t = T - 1, T - 2, \dots, 0$.

The initial continuation value v_0 equals the optimized value of the Ramsey planner's criterion V defined in equation (6):

$$v_0 = \sum_{t=0}^{\infty} \beta^t s(\theta_t, \mu_t). \quad (14)$$

We verify approximate equality by inspecting Figure 2, which plots v_t against t for $t = 0, \dots, T$.

The limiting value of the continuation value v_t in Figure 2 is evidently approached from above, so it is the value of the **worst** continuation Ramsey plan. Some researchers recommend following this plan starting at time 0 because it is time invariant and thus time consistent. But as Figure 2 also shows, there is another time consistent plan that is also time consistent and attains a higher value: the Ramsey plan with μ_t constrained to be a constant.⁴

5 A Δ -Time Interval Specification

Although Calvo (1978) formulated his model in continuous time, thus far we have formulated our model in discrete time with a unit time interval. In the interests of approaching Calvo's continuous time specification, in this section we consider a version of the discrete-time Ramsey problem in which the time increment is $\Delta \in (0, 1)$, so that now we assume the $t = 0, \Delta, 2\Delta, 3\Delta, \dots$, $p_{t+\Delta} = p_t + \Delta\theta_t$, and $m_{t+\Delta} = m_t + \Delta\mu_t$. The government's value function becomes:

$$V = \sum_{t=0}^{\infty} \exp(-\rho t\Delta) s(\theta_t, \mu_t) \Delta \quad (15)$$

for $t = 0, \Delta, 2\Delta, \dots$, subject to the constraints

$$\theta_t = \exp(-\gamma\Delta)\theta_{t+\Delta} + (1 - \exp(-\gamma\Delta))\mu_t, \quad (16)$$

⁴For further analysis of this topic in the setting a linear-quadratic version of Calvo's model, please see Sargent and Yang (2025b).




Figure 2: Continuation values v_t

where ρ, γ are related to β and λ through:

$$\begin{aligned} \exp(-\rho) &= \beta \\ \exp(-\gamma) &= \lambda = \frac{\alpha}{1 + \alpha}. \end{aligned}$$

When computing the Ramsey plan using the gradient ascent algorithm, we set $T_\Delta = \frac{T}{\Delta}$. We again approximate the Ramsey plan using both the gradient ascent algorithm.

Figure 3 indicates that shapes of $(\vec{\mu}, \vec{\theta})$ resemble those for the unit-time-interval Ramsey problem.

5.1 Outcomes with different Δ time intervals

Figure 4 displays the computed $\vec{\mu}$ sequences for $\Delta = 0.001$ (designed to approximate a continuous-time setting) and $\Delta = 1$ (our initial unit-time interval specification). The time interval Δ affects outcomes because of how it influences both the objective function in equation (15) and the constraint equation (16).

Several key distinctions emerge from varying Δ :




Figure 3: Ramsey plan ($\vec{\mu}$ and $\vec{\theta}$) at $\Delta = 0.001$.




Figure 4: Comparison of money growth rates $\vec{\mu}$ between near-continuous time ($\Delta = 0.001$) and discrete time ($\Delta = 1$) Ramsey plans. Constrained-to-constant μ^{CR} is invariant to Δ specifications.

- The government in the small Δ -time specification makes decisions more often, controlling the money growth rate at each smaller time step Δ .
- The effective discount rate between adjacent decision points is $\exp(-\rho\Delta)$ rather than $\exp(-\rho)$, resulting in less aggressive discounting between adjacent decision times when Δ is small.
- Equation (16) shows that relationships between current and future inflation rates differ across different Δ 's. That affects intertemporal tradeoffs confronting the government.

Comparing the optimal money growth trajectories, we observe two notable differences: the Δ -time Ramsey plan with $\Delta = 0.001$ features a higher initial money growth rate μ_0 but converges to a lower value $\bar{\mu}$. Evidently as we approach continuous time, the optimal policy front-loads inflation more but ultimately converges to a lower rate of money growth and inflation.

These differences highlight the influence of the time increment Δ on the optimal policy. The ability to adjust the money growth rate more often shapes an optimal policy.

6 Human Intelligence

We have represented a Ramsey plan in the **open loop** form of a function

$$\mu_t = f(t) \tag{17}$$

that maps $t \in \{0, \Delta, 2\Delta, \dots, \}$ to $\mu_t \in \mathbb{R}$. Figures 1 and 4 indicate that the Ramsey planner makes both $\vec{\mu}$ and $\vec{\theta}$ vary over time.

- $\vec{\theta}$ and $\vec{\mu}$ both decline monotonically.
- $\vec{\theta}$ and $\vec{\mu}$ converge from above to a common constant $\bar{\mu}$.

While the **open loop** representation of a Ramsey plan respects the Ramsey problem's purpose to choose a **sequence** $\vec{\mu}$ once-and-for-all at time 0. Many macroeconomists and control theorists prefer a **closed loop** representation of a Ramsey plan that takes the form of a pair of functions

$$\begin{aligned} \mu_t &= m(z_t) \\ z_{t+\Delta} &= n(z_t), \end{aligned}$$

where z_t is a **state vector**. The second equation is a transition equation for $z_{t+\Delta}$, and z_0^R is a value that the Ramsey planner chooses for the initial state vector.

If present at all, a recursive structure in the $\vec{\mu}, \vec{\theta}$ chosen by our machine-learning Ramsey planner lies hidden from view. Let's try to bring it out by again using machine learning. We'll proceed by viewing Ramsey outcomes $\vec{\mu}^R, \vec{\theta}^R, \vec{v}^R$ as “fake data” on which we'll run some exploratory, possibly nonlinear, least squares regressions.⁵

We add some **human intelligence** to the **artificial intelligence** embodied in our Python programs by formulating specifications of regressions to run on our “fake data”. We begin by computing least squares linear regressions of some components of $\vec{\theta}^R, \vec{\mu}^R$, and \vec{v}^R on components of $\vec{\theta}^R$ or $\vec{\mu}^R$, hoping that these regressions will reveal structure hidden within the $\vec{\mu}^R, \vec{\theta}^R$ sequences associated with a Ramsey plan.

Let's pause to think about roles being played here by **human** and **artificial** intelligence. Artificial intelligence in the form of a computer program runs the regressions. But one can regress anything on anything else. Human intelligence, such as it is, must tell us what regressions to run. Human intelligence will be required fully to appreciate what those regressions reveal about the structure of a Ramsey plan.

Our machine-learned Ramsey plan $\vec{\mu}^R, \vec{\theta}^R$ constitutes the “fake” data set that we use to run regressions in Table 1 and Table 2. Table 1 reports several regressions with μ_t as the independent variable, while Table 2 reports several regressions with θ_t as the independent variable.

We begin by focusing on the first entry in Table 1 that reports outcomes from regressing θ_t on a constant and μ_t . This seems natural because equation (5) asserts that inflation at time t is determined by the money growth sequence $\{\mu_s\}_{s=t}^{\infty}$. After all, since a Ramsey planner chooses a money growth sequence, shouldn't money growth be the “exogenous variable” in our regressions? We'll return to this question soon.

The first entry of Table 1 reports the least squares affine regression $\theta_t = \tilde{b}_0 + \tilde{b}_1\mu_t + \varepsilon_t$, where ε_t is a least squares residual that is by construction orthogonal to μ_t .

That the R^2 is 0.989 indicates that there is a non-trivial residual ε_t that is orthogonal to μ_t . To improve the fit, the second entry in Table 1 shows a cubic regression model. The regression uncovers the following representation of θ_t as a function of μ_t

$$\theta_t = -0.0701 + 1.0759\mu_t + 1.4626\mu_t^2 + 0.6133\mu_t^3.$$

We plot the regression line in the left panel of Figure 5. The dots indicate (μ_t, θ_t) pairs for $t = 0, \Delta, 2\Delta, \dots$ that converge from above to a limiting pair $(\bar{\mu}, \bar{\theta})$.

In hopes of discovering a law of motion for $\vec{\mu}$ under the Ramsey plan, the third entry of Table 1 reports the least squares affine regression $\mu_{t+\Delta} = \tilde{d}_0 + \tilde{d}_1\mu_t + \varepsilon_t$, where, recycling

⁵Thus, our “fake data” set is just the Ramsey plan generated by our open loop formula for μ_t as a function of t and formula (5) that takes the future μ^t and maps it into θ_t .

Table 1: Regression results with μ_t as independent variable

Model	Variable	Coefficient	Std. Error	t-statistic
$\theta_t = \tilde{b}_0 + \tilde{b}_1\mu_t + \varepsilon_t$	Constant (\tilde{b}_0)	-0.0914	0.000	-786.548
	μ_t (\tilde{b}_1)	0.6588	0.000	1346.446
	$R^2 = 0.989$			
$\theta_t = \tilde{b}_0 + \tilde{b}_1\mu_t + \tilde{b}_2\mu_t^2 + \tilde{b}_3\mu_t^3 + \varepsilon_t$	Constant (\tilde{b}_0)	-0.0701	3.68×10^{-5}	-1905.508
	μ_t (\tilde{b}_1)	1.0759	0.001	1002.001
	μ_t^2 (\tilde{b}_2)	1.4626	0.008	174.001
	μ_t^3 (\tilde{b}_3)	0.6133	0.019	32.817
	$R^2 = 1.000$			
$\mu_{t+\Delta} = \tilde{d}_0 + \tilde{d}_1\mu_t + \varepsilon_t$	Constant (\tilde{d}_0)	-7.627×10^{-5}	2.03×10^{-7}	-375.168
	μ_t (\tilde{d}_1)	0.9997	8.56×10^{-7}	1.17×10^6
	$R^2 = 1.000$			
$v_t = \tilde{g}_0 + \tilde{g}_1\mu_t + \tilde{g}_2\mu_t^2 + \varepsilon_t$	Constant (\tilde{g}_0)	4.4947	4.76×10^{-5}	9.44×10^4
	μ_t (\tilde{g}_1)	0.3896	0.001	547.298
	μ_t^2 (\tilde{g}_2)	-0.0792	0.002	-36.494
	$R^2 = 0.998$			
$v_t = \tilde{g}_0 + \tilde{g}_1\mu_t + \tilde{g}_2\mu_t^2 + \tilde{g}_3\mu_t^3 + \varepsilon_t$	Constant (\tilde{g}_0)	4.4897	3.88×10^{-5}	1.16×10^5
	μ_t (\tilde{g}_1)	0.1906	0.001	168.432
	μ_t^2 (\tilde{g}_2)	-1.7427	0.009	-196.743
	μ_t^3 (\tilde{g}_3)	-3.7386	0.020	-189.847
	$R^2 = 0.999$			
$v_t = \tilde{g}_0 + \tilde{g}_1\mu_t + \tilde{g}_2\mu_t^2 + \tilde{g}_3\mu_t^3 + \tilde{g}_4\mu_t^4 + \varepsilon_t$	Constant (\tilde{g}_0)	4.4855	9.54×10^{-6}	4.7×10^5
	μ_t (\tilde{g}_1)	-0.0963	0.000	-210.737
	μ_t^2 (\tilde{g}_2)	-6.1852	0.006	-960.114
	μ_t^3 (\tilde{g}_3)	-27.5885	0.034	-822.267
	μ_t^4 (\tilde{g}_4)	-41.3858	0.058	-715.488
	$R^2 = 1.000$			




Figure 5: Regression of θ_t on a constant, μ_t , μ_t^2 , and μ_t^3 (left), regression of $\mu_{t+\Delta}$ on a constant and μ_t (center), and regression of ν_t on a constant, μ_t , μ_t^2 , μ_t^3 , and μ_t^4 (right). The orange line denotes the value of V^{CR} (right).

notation, ε_t is now a least squares residual that is by construction orthogonal to μ_t . We obtained a nearly perfect fit ($R^2 = 1.000$) and have discovered the following approximate Ramsey planner's law of motion for $\vec{\mu}^R$

$$\mu_{t+\Delta} = -7.627 \times 10^{-5} + 0.9997\mu_t.$$

We plot the regression line in the middle panel of Figure 5. Here the dots indicate $(\mu_t, \mu_{t+\Delta})$ pairs for $t = 0, \Delta, 2\Delta, \dots$ that converge from above to a limiting pair $(\bar{\mu}, \bar{\mu})$.

The fourth entry of Table 1 reports a least squares regression of ν_t on a constant, μ_t , and μ_t^2 . The R^2 is 0.998, indicating that there is a small residual ε_t that is orthogonal to both μ_t and μ_t^2 . In search for a perfect fit, we increase the order of the polynomial in the fifth and sixth entries of Table 1 to third and fourth degree, respectively, and obtain a perfect fit ($R^2 = 1.000$) with the fourth-degree polynomial. The regression uncovers the following representation of ν_t as a function of μ_t

$$\nu_t = 4.4855 - .0963\mu_t - 6.1852\mu_t^2 - 27.5885\mu_t^3 - 41.3858\mu_t^4.$$

with the regression curve plotted in the right panel of Figure 5. The results show that the relationship between ν_t and μ_t is well-approximated by a quadratic function, though the small residuals indicate some additional complexity in the relationship that is only fully captured by quartic terms.

Assembling our regressions, we have discovered that along a single Ramsey outcome path

$\vec{\mu}^R, \vec{\theta}^R$ the following relationships prevail:

$$\begin{aligned}\mu_0 &= \mu_0^R \\ \theta_t &= \tilde{b}_0 + \tilde{b}_1\mu_t + \tilde{b}_2\mu_t^2 + \tilde{b}_3\mu_t^3 \\ \mu_{t+\Delta} &= \tilde{d}_0 + \tilde{d}_1\mu_t,\end{aligned}\tag{18}$$

where $\tilde{b}_0, \tilde{b}_1, \tilde{b}_2, \tilde{b}_3, \tilde{d}_0, \tilde{d}_1$ are parameters whose values we estimated with our regressions; we unearthed initial value μ_0^R along with other components of $\vec{\mu}^R, \vec{\theta}^R$ when we computed the Ramsey plan. In addition, we learned that along our Ramsey plan, continuation values are described by the quartic equation

$$v_t = \tilde{g}_0 + \tilde{g}_1\mu_t + \tilde{g}_2\mu_t^2 + \tilde{g}_3\mu_t^3 + \tilde{g}_4\mu_t^4.$$


We discovered these relationships by running regressions and noticing that the R^2 's of approximately unity tell us that the fits are nearly perfect.

6.1 Direction of fit?

Instead of taking μ_t as the “independent” (i.e., right hand side) variable, let’s temporarily put θ_t on the right hand side. A plausible case for putting θ_t and not μ_t on the right hand side could be that the Ramsey planner is “inflation targeting,” just as many governments today tell their central banks to do. The four entries of Table 2 report results.

Table 2: Regression results with θ_t as independent variable

Model	Variable	Coefficient	Std. Error	t-statistic
$\mu_t = b_0 + b_1\theta_t + \varepsilon_t$	Constant (b_0)	0.1347	0.000	492.179
	θ_t (b_1)	1.5012	0.001	1346.446
	$R^2 = 0.989$			
$\mu_t = b_0 + b_1\theta_t + b_2\theta_t^2 + b_3\theta_t^3 + \varepsilon_t$	Constant (b_0)	0.0916	0.000	278.792
	θ_t (b_1)	1.5197	0.006	247.017
	θ_t^2 (b_2)	4.5951	0.036	128.601
	θ_t^3 (b_3)	14.9799	0.065	229.423
	$R^2 = 1.000$			
$\theta_{t+\Delta} = d_0 + d_1\theta_t + \varepsilon_t$	Constant (d_0)	-9.34×10^{-5}	1.9×10^{-7}	-492.204
	θ_t (d_1)	0.9997	7.73×10^{-7}	1.29×10^6
	$R^2 = 1.000$			
$v_t = g_0 + g_1\theta_t + g_2\theta_t^2 + g_3\theta_t^3 + \varepsilon_t$	Constant (g_0)	4.4713	4.94×10^{-5}	9.05×10^4
	θ_t (g_1)	-0.5225	0.001	-564.994
	θ_t^2 (g_2)	-4.5673	0.005	-850.320
	θ_t^3 (g_3)	-5.1838	0.010	-528.132
	$R^2 = 1.000$			

**Figure 6:** Regression of μ_t on a constant, θ_t , θ_t^2 , and θ_t^3 (left), regression of $\theta_{t+\Delta}$ on a constant and θ_t (center), and regression of v_t on a constant, θ_t , θ_t^2 , and θ_t^3 (right). The orange line depicts V^{CR} (right).

Taking stock, our regression with θ_t on the right side tells us that along the Ramsey

outcome $\vec{\mu}^R, \vec{\theta}^R$, the affine function $\mu_t = 0.1347 + 1.5012\theta_t$ provides a good initial fit for the relationship between μ_t and θ_t , with $R^2 = 0.989$. However, the cubic model improves the fit to be nearly perfect ($R^2 = 1.000$):

$$\mu_t = 0.0916 + 1.5197\theta_t + 4.5951\theta_t^2 + 14.9799\theta_t^3.$$

Similarly, the following linear regression of $\theta_{t+\Delta}$ on θ_t fits perfectly:

$$\theta_{t+\Delta} = -9.34 \times 10^{-5} + 0.9997\theta_t.$$

For the continuation values, the cubic regression has $R^2 = 1.000$, indicating a perfect fit:

$$\nu_t = 4.4713 - 0.5225\theta_t - 4.5673\theta_t^2 - 5.1838\theta_t^3.$$

Thus, we have discovered that along a single Ramsey outcome path $\vec{\mu}^R, \vec{\theta}^R$ the following relationships prevail:

$$\begin{aligned} \theta_0 &= \theta_0^R \\ \mu_t &= b_0 + b_1\theta_t + b_2\theta_t^2 + b_3\theta_t^3 \\ \theta_{t+\Delta} &= d_0 + d_1\theta_t, \end{aligned} \tag{19}$$

where $b_0, b_1, b_2, b_3, d_0, d_1$ are parameters whose values we estimated with our regressions; we unearthed the initial value θ_0^R along with other components of $\vec{\mu}^R, \vec{\theta}^R$ when we computed the Ramsey plan. In addition, we learned that along our Ramsey plan, continuation values are perfectly described by the cubic function

$$\nu_t = g_0 + g_1\theta_t + g_2\theta_t^2 + g_3\theta_t^3.$$

As with our earlier regressions with μ_t on the right side, we discovered these relationships by running regressions, examining the results, and noting the perfect R^2 values that indicate excellent fits.

The right panel of Figure 6 shows that the highest continuation value ν_0 at $t = 0$ appears near the peak of the cubic function $g_0 + g_1\theta_t + g_2\theta_t^2 + g_3\theta_t^3$. Subsequent values of ν_t for $t \geq 1$ appear to the lower left of the pair (θ_0, ν_0) and converge monotonically from above to the limiting value at time T . The value V^{CR} attained by the Ramsey plan that is restricted to use a constant $\mu_t = \mu^{CR}$ sequence appears as a horizontal line.

6.2 What machine learning taught us

We have discovered that the Ramsey plan for $\vec{\mu}$ seems to have a recursive structure. But by using the methods and ideas that we have deployed here, it is challenging to say more.

We have discovered **two** closed-loop representations of a Ramsey plan and the associated continuation value sequence, one with μ_t as the right-hand side “independent variable”, the other with θ_t as the right-hand side variable. Both are valid representations. Which representation is better in terms of understanding forces shaping the plan?

To answer that question, we deploy economic theory presented by [Chang \(1998\)](#), who showed that (19) is actually a better way to represent a Ramsey plan.

[Chang](#) noted that equation (5) indicates that an equivalence class of continuation money growth sequences $\{\mu_{t+j}\}_{j=0}^{\infty}$ deliver the same θ_t . Consequently, equations (3) and (5) describe how θ_t intermediates how the government’s choices of μ_{t+j} , $j = 0, \Delta, \dots$ impinge on time t real balances $m_t - p_t = -\alpha\theta_t$ and thereby on time t welfare.

We can appreciate Chang’s reasoning by thinking about the following “machine learning” procedure for computing continuation values from time 0 that start from an arbitrary initial inflation rate θ_0 . For each $\theta_0 \in \mathbb{R}$, define a set

$$\Omega(\theta_0) = \left\{ \{\theta_{t+\Delta}, \mu_t\}_{t=0}^{\infty} : \theta_{t+\Delta} = \frac{1}{\tilde{\lambda}}\theta_t - \frac{1-\tilde{\lambda}}{\tilde{\lambda}}\mu_t, \quad \forall t \geq 0 \right\}, \quad (20)$$

where $\tilde{\lambda} = \exp(-\gamma\Delta)$.

For a given θ_0 , think about using machine learning to compute a closed loop policy

$$\theta_t = f(t; \theta_0), \quad t \geq \Delta$$

that solves the maximization problem on the right side of the following equation that defines a **continuation value function** $J(\theta_0)$:

$$J(\theta_0) = \max_{\{\theta_{t+\Delta}, \mu_t\}_{t=0}^{\infty} \in \Omega(\theta_0)} \sum_{t=0}^{\infty} \exp(-\rho t \Delta) s(\theta_t, \mu_t) \Delta.$$

If we were to do this for a set possible θ_0 ’s, we could then hand the function $J(\theta_0)$ over to our Ramsey planner and compute the Ramsey planner’s choice of θ_0 according to

$$\theta_0^R = \arg \max_{\theta} J(\theta)$$

and the value of the Ramsey plan as

$$v_0^R = \max_{\theta} J(\theta).$$

This takes us to a formulation of [Chang \(1998\)](#) that we present in the next section.

7 Dynamic Programming

We present [Chang \(1998\)](#)'s recursive formulation of the Ramsey problem and solve it using optimistic policy iteration (OPI). We show that both discrete time and near-continuous time models can be solved using the same dynamic programming (DP) formulation. We first treat the discrete time model with $\Delta = 1$ and then show how to extend the analysis to the near-continuous time model with $\Delta = 0.001$ in section 7.1.1 by replacing the discount factor and scaling the welfare function.

We approximate the feasible inflation rates via discretization. Let $\Theta \subset \mathbb{R}$ denote a finite set of feasible inflation rates. In our setting, Θ is a discretized points in the interval $[\underline{\theta}, \bar{\theta}]$. The following version of equation (4) links the current state θ , the control variable μ , and the future state θ'

$$\theta = \frac{\alpha}{1+\alpha}\theta' + \frac{1}{1+\alpha}\mu. \quad (21)$$

Define a correspondence $\Gamma_\mu : \Theta \rightarrow 2^{\mathbb{R}}$ that describes the feasible set of choices of μ given the current state θ

$$\Gamma_\mu(\theta) = \left\{ \mu \in \mathbb{R} : \frac{(1+\alpha)\theta - \mu}{\alpha} \in \Theta \right\}.$$

We formulate the Ramsey problem sequentially by posing two subproblems.

7.1 Subproblem 1

We seek a function $J : \Theta \rightarrow \mathbb{R}$ that satisfies the Bellman equation

$$J(\theta) = \sup_{(\theta', \mu) \in \Theta \times \Gamma_\mu(\theta)} \{s(\theta, \mu) + \beta J(\theta')\}. \quad (22)$$

By solving equation (21) for μ , we get

$$\mu = (1+\alpha)\theta - \alpha\theta'. \quad (23)$$

Substitute this into (22) to get

$$J(\theta) = \sup_{\theta' \in \Theta} \{s(\theta, (1+\alpha)\theta - \alpha\theta') + \beta J(\theta')\}. \quad (24)$$

Bellman equation (24) is an instance of a Recursive Decision Process (RDP) $\mathcal{R} = (\Gamma, \mathcal{V}, B)$ (see [Sargent and Stachurski \(2025, ch. 8\)](#)) in which

- $\Gamma(\theta) := \{\theta' \in \Theta : \mu = (1+\alpha)\theta - \alpha\theta', \mu \in \Gamma_\mu(\theta)\}$ is the feasible correspondence;
- $\mathcal{V} := \mathbb{R}^\Theta$ is the value space;

- $B : \mathsf{G} \times \mathcal{V} \rightarrow \mathbb{R}$ is the value aggregator defined by:

$$B(\theta, \theta', J) := s(\theta, (1 + \alpha)\theta - \alpha\theta') + \beta J(\theta'), \quad (25)$$

where $\mathsf{G} := \{(\theta, \theta') \in \Theta \times \Theta : \theta' \in \Gamma(\theta)\}$.

The right side of (24) describes the decision problem of a **continuation Ramsey planner** has been told to deliver what Chang calls **promised inflation** θ today by choosing money creation today and promised inflation tomorrow. Here the understanding is that promised inflation must always equal actual inflation, a ramification of rational expectations.

Let $\Sigma := \{\sigma(\theta) \in \Gamma(\theta) \text{ for all } \theta \in \Theta\}$ denote the set of feasible policies where $\sigma(\theta)$ is the policy function that maps θ to θ' .

Proposition 7.1. *The triple $\mathcal{R} = (\Gamma, \mathcal{V}, B)$ forms a well-defined RDP.*

Proof. We must verify required consistency and monotonicity conditions. Consistency condition clearly holds. So we only need to verify the monotonicity condition.

For any $Q, K \in \mathcal{V}$ with $Q \leq K$ and any feasible pair $(\theta, \theta') \in \mathsf{G}$:

$$\begin{aligned} B(\theta, \theta', K) - B(\theta, \theta', Q) &= s(\theta, (1 + \alpha)\theta - \alpha\theta') + \beta K(\theta') \\ &\quad - [s(\theta, (1 + \alpha)\theta - \alpha\theta') + \beta Q(\theta')] \\ &= \beta [K(\theta') - Q(\theta')] \geq 0 \end{aligned} \quad (26)$$

since $\beta > 0$ and $K(\theta') \geq Q(\theta')$ for all $\theta' \in \Theta$. □

Define the Bellman operator $T : \mathbb{R}^\Theta \rightarrow \mathbb{R}^\Theta$ by:

$$(TJ)(\theta) = \sup_{\theta' \in \Theta} \{s(\theta, (1 + \alpha)\theta - \alpha\theta') + \beta J(\theta')\},$$

where \mathbb{R}^Θ is the space of real-valued functions on Θ equipped with the sup norm $\|J\| = \sup_{\theta \in \Theta} |J(\theta)|$.

For any policy $\sigma \in \Sigma$, define the policy operator $T_\sigma : \mathbb{R}^\Theta \rightarrow \mathbb{R}^\Theta$ by:

$$(T_\sigma J)(\theta) = B(\theta, \sigma(\theta), J) = s(\theta, (1 + \alpha)\theta - \alpha\sigma(\theta)) + \beta J(\sigma(\theta)).$$

For a policy σ , we define $s_\sigma(\theta) := s(\theta, (1 + \alpha)\theta - \alpha\sigma(\theta))$. Using this notation, we can write:

$$(T_\sigma J)(\theta) = s_\sigma(\theta) + \beta J(\sigma(\theta)).$$

The lifetime value of policy σ , denoted by J_σ , can be expressed as the infinite sum:

$$J_\sigma = \sum_{t=0}^{\infty} \beta^t s_\sigma(\theta_t).$$

Theorem 7.2. *For each policy $\sigma \in \Sigma$, the policy operator T_σ is a contraction mapping with modulus β .*

Proof. For any $Q, K \in \mathcal{V}$ and $\theta \in \Theta$:

$$\begin{aligned} |(T_\sigma Q)(\theta) - (T_\sigma K)(\theta)| &= |s(\theta, (1 + \alpha)\theta - \alpha\sigma(\theta)) + \beta Q(\sigma(\theta)) \\ &\quad - s(\theta, (1 + \alpha)\theta - \alpha\sigma(\theta)) - \beta K(\sigma(\theta))| \\ &= |\beta Q(\sigma(\theta)) - \beta K(\sigma(\theta))| \\ &= \beta |Q(\sigma(\theta)) - K(\sigma(\theta))| \\ &\leq \beta \sup_{\theta' \in \Theta} |Q(\theta') - K(\theta')| \\ &= \beta \|Q - K\|. \end{aligned} \tag{27}$$

$$\tag{28}$$

Taking the supremum over $\theta \in \Theta$:

$$\|T_\sigma Q - T_\sigma K\| = \sup_{\theta \in \Theta} |(T_\sigma Q)(\theta) - (T_\sigma K)(\theta)| \leq \beta \|Q - K\|. \tag{29}$$

Since $\beta \in (0, 1)$, T_σ is a contraction mapping with modulus β . □

By Theorem 7.2, the policy operator T_σ is a contraction mapping for any policy σ , hence it is globally stable. Therefore, \mathcal{R} is globally stable.

We present the OPI algorithm for solving the discrete-time Ramsey problem in Algorithm 1. Since \mathcal{R} is globally stable, Theorem 8.1.1 of Sargent and Stachurski (2025) guarantees convergence of the OPI sequence of value functions J_k to the unique solution to the Bellman equation in \mathcal{V} . Moreover, the theorem ensures the existence of a $K \in \mathbb{N}$ beyond which the policy σ_k remains optimal for all $k \geq K$. The upper panel of Figure 7 shows the value function J computed using the OPI algorithm for the discrete-time Ramsey problem.

Algorithm 1: Optimistic Policy Iteration for Discrete-time Ramsey Problem

Require:

Parameters: $m \in \mathbb{N}$ (number of policy operator iterations)
 $\tau \geq 0$ (convergence tolerance)
 $N \in \mathbb{N}$ (maximum number of iterations)
 Initial policy $\sigma_0 : \Theta \rightarrow \Theta$
 Model parameters: $\alpha = 1.0, u_0 = 2.0, c_2 = 6.0, c_3 = 25.0, c_4 = 40.0, \beta = 0.85$

- 1: Define a set of grid points $\Theta \subset \mathbb{R}$
- 2: Initialize value function $J_0(\theta)$ for all $\theta \in \Theta$ using σ_0
- 3: Set iteration counter: $k \leftarrow 0$
- 4: **repeat**
- 5: **for** each $\theta \in \Theta$ **do**
- 6: $\sigma_k(\theta) \in \arg \max_{\theta' \in \Theta} \left\{ s\left(\theta, (1+\alpha)\theta - \alpha\theta'\right) + \beta J_k(\theta') \right\}$; // Compute J_k -greedy σ_k
- 7: **end for**
- 8: $J_{k+1} \leftarrow T_{\sigma_k}^m J_k$; // Apply policy operator T_{σ_k} m times
- 9: **if** $\|J_{k+1} - J_k\| \leq \tau$ **then**
- 10: **break**
- 11: **end if**
- 12: $k \leftarrow k + 1$
- 13: **until** $k \geq N$
- 14: **return** σ_k, J_k

7.1.1 Δ -time Formulation

We can formulate the Ramsey problem in near-continuous time by rewriting the constraint on the right hand side of (20) as

$$\mu = \frac{1}{1 - \tilde{\lambda}}\theta - \frac{\tilde{\lambda}}{1 - \tilde{\lambda}}\theta'. \quad (30)$$

We now use θ and θ' to denote the current and next state with $\theta' = \theta_{t+\Delta}$ in near-continuous time.

Substituting this into the near-continuous time version of the Bellman equation (22) gives us

$$J(\theta_t) = \sup_{\theta' \in \Theta} \left\{ s\left(\theta, \frac{1}{1 - \tilde{\lambda}}\theta - \frac{\tilde{\lambda}}{1 - \tilde{\lambda}}\theta'\right) \Delta + \exp(-\rho\Delta)J(\theta') \right\}. \quad (31)$$

The corresponding policy operator is given by

$$(T_\sigma J)(\theta_t) = s_\sigma(\theta)\Delta + \exp(-\rho\Delta)J(\sigma(\theta)),$$

where $s_\sigma(\theta_t) = s\left(\theta_t, \frac{1}{1 - \tilde{\lambda}}\theta_t - \frac{\tilde{\lambda}}{1 - \tilde{\lambda}}\sigma(\theta_t)\right)$.

Since $\exp(-\rho\Delta) = \beta^\Delta < 1$ for $\Delta > 0$, the policy operator T_σ is a contraction mapping with

modulus β^Δ . Thus, results in Theorem 7.2 hold for the near-continuous time version of the Ramsey problem as well. The OPI algorithm for solving the near-continuous time Ramsey problem is similar to the discrete-time version, with the only difference being that we replace the discount factor β with $\exp(-\rho\Delta)$ and scale the welfare function by Δ . The OPI algorithm for the near-continuous time Ramsey problem is given in Algorithm 2. The middle panel of Figure 7 shows the value function J computed using the OPI algorithm for the near-continuous time Ramsey problem.

Algorithm 2: Optimistic Policy Iteration for Δ -time Ramsey Problem

Require:

- Parameters: $m \in \mathbb{N}$ (number of policy operator iterations)
 $\tau \geq 0$ (convergence tolerance)
 $N \in \mathbb{N}$ (maximum number of iterations)
Initial policy $\sigma_0 : \Theta \rightarrow \Theta$
- Model parameters: $\alpha = 1.0, u_0 = 2.0, c_2 = 6.0, c_3 = 25.0, c_4 = 40.0,$
 $\gamma = -\ln(\lambda) = 0.6931, \rho = -\ln(\beta) = 0.1625, \Delta = 0.001$
- 1: Define a set of grid points $\Theta \subset \mathbb{R}$
 - 2: Initialize value function $J_0(\theta)$ for all $\theta \in \Theta$ using σ_0
 - 3: Set iteration counter: $k \leftarrow 0$
 - 4: **repeat**
 - 5: **for** each $\theta \in \Theta$ **do**
 - 6: $\sigma_k(\theta) \in \arg \max_{\theta' \in \Theta} \left\{ s \left(\theta, \frac{1}{1-\tilde{\lambda}}\theta - \frac{\tilde{\lambda}}{1-\tilde{\lambda}}\theta' \right) \Delta + \exp(-\rho\Delta)J(\theta') \right\}$
 - 7: **end for**
 - 8: $J_{k+1} \leftarrow T_{\sigma_k}^m J_k,$
 - 9: **if** $\|J_{k+1} - J_k\| \leq \tau$ **then**
 - 10: **break**
 - 11: **end if**
 - 12: $k \leftarrow k + 1$
 - 13: **until** $k \geq N$
 - 14: **return** σ_k, J_k
-

7.2 Subproblem 2

The value of the Ramsey problem is

$$V^R = \max_{\theta_0} J(\theta_0).$$

Evidently, V^R is the maximum value of v_0 defined in equation (14). Figure 7 shows the value function $J(\theta)$ computed using the OPI algorithm, and $\theta_0^R = \arg \max_{\theta_0} J(\theta_0)$ is the optimal initial promised value.

By comparing the value function $J(\theta)$ at $\Delta = 1$ and $\Delta = 0.001$ in bottom panel of Figure 7, we find that the two value functions have almost identical shapes characterized by the




Figure 7: Upper Panel: Value function $J(\theta)$ and optimal initial state θ_0^R under discrete time ($\Delta = 1$). Middle Panel: Value function $J(\theta)$ with near-continuous time ($\Delta = 0.001$). Bottom Panel: Comparison of Value functions $J(\theta)$ with $\Delta = 1$ and $\Delta = 0.001$.

functional form of $s(\theta, \mu)$. However, the value function at $\Delta = 0.001$ is lower than the value function at $\Delta = 1$ for larger θ values signaling that the two optimization problems are fundamentally different. Note that the constrained-to-constant-money-growth-rate Ramsey plan θ^{CR} remains constant across different time increments Δ .

7.3 Comparison of OPI and gradient ascent in discrete and near-continuous time

It is worth noting that our two approaches to computing a Ramsey plan seem vulnerable to different sorts of computational errors. Our “lazy” gradient ascent algorithm that simply constructs an algorithm to compute the government planner’s objective as a function of a sequence of money growth rates, then hands it over to a gradient-ascent optimizer directly optimizes over continuous controls. By not discretizing the state space, it avoids possible numerical discretization error and computational costs associated with wanting to make a fine-grained grid on the space of admissible states and controls. In contrast, the DP-based OPI solver explicitly relies on discretization of both states and controls to compute value functions and policies, leading to discretization errors and increased computational difficulty as $\Delta \rightarrow 0$. Nevertheless, Figure 8 and 9 shows a comparison of Ramsey plan $\vec{\mu}$ computed using the OPI algorithm and the gradient ascent algorithm in discrete and near-continuous time. Evidently, the two methods produce virtually the same approximation to the Ramsey plan $\vec{\mu}$.

8 Concluding remarks

The state variable θ_t plays multiple roles in our section 7 closed-loop recursive representation of the Ramsey plan: it is an inflation rate that the monetary authority yesterday **promised** to deliver today, the **forecast** of this period’s inflation rate that the representative household made yesterday, and the **actual** rate of inflation. That the same variable plays these multiple roles is a manifestation of the principle that in a rational expectations model, a representative private agent’s rule for forecasting the government’s decision equals the government’s rule for making those decisions. The two roles played by that decision rule is yet another symptom of the communism of beliefs that pervades rational expectations models. In a settings that expands the state space to include histories of government decisions, Chang (1998) studies equilibria in which a government always chooses to confirm the representative household’s forecast of its decisions.⁶ Chang’s analysis makes vivid contact with a discussion by Blinder (1999, lecture

⁶Sargent (2024) presents a simplified version of Chang’s model.

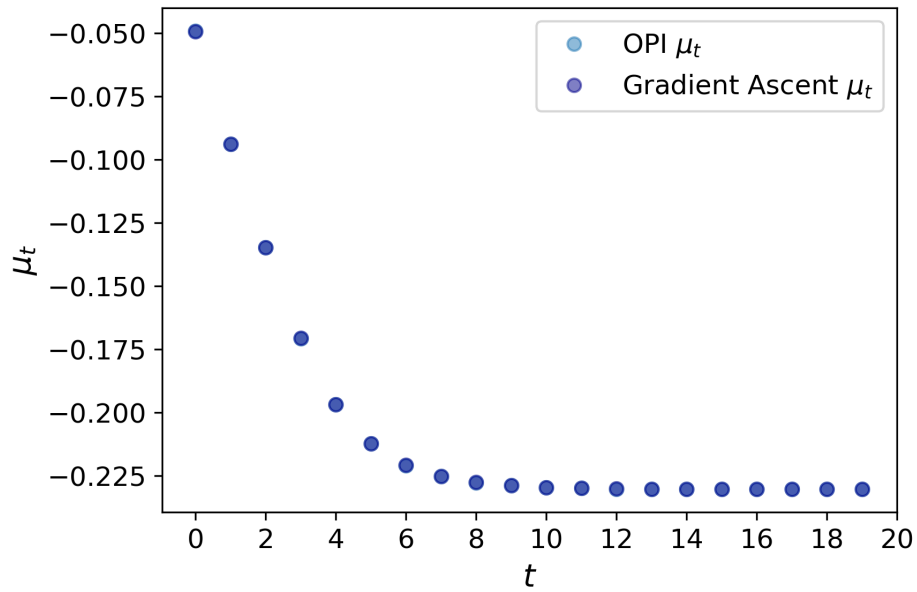


Figure 8: Comparison of Ramsey plan $\vec{\mu}$ under discrete time ($\Delta = 1$)

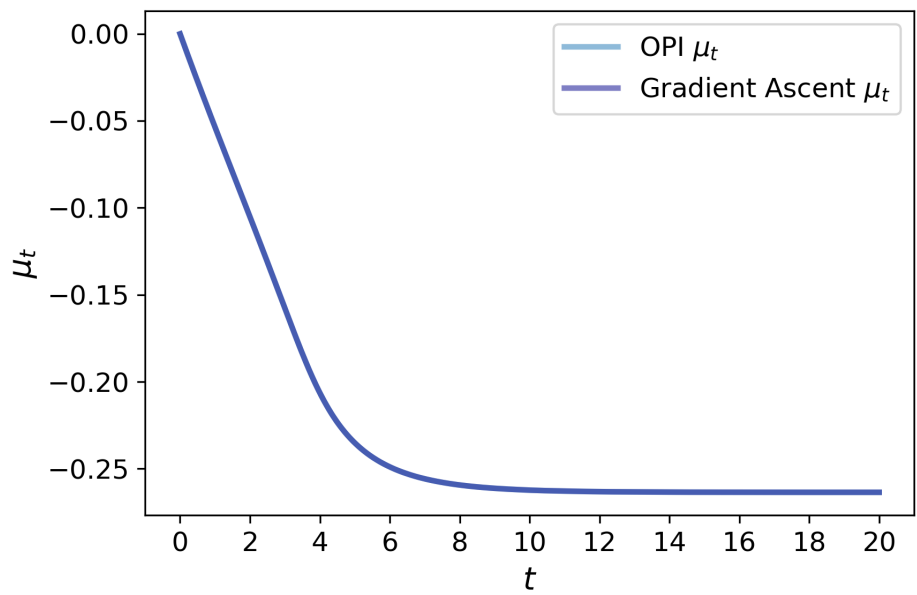


Figure 9: Comparison of Ramsey plan $\vec{\mu}$ under near-continuous time ($\Delta = 0.001$)

3, part 3) that wrestles with whether the central bank should “follow the market” by always confirming the market’s expectations of its actions. In Chang’s model, the central bank always wants to do that.

References

- Blinder, Alan S. 1999. *Central banking in theory and practice*. MIT Press.
- Cagan, Philip. 1956. “The Monetary Dynamics of Hyperinflation.” In *Studies in the Quantity Theory of Money*, edited by Milton Friedman. Chicago: University of Chicago Press, 25–117.
- Calvo, Guillermo A. 1978. “On the Time Consistency of Optimal Policy in a Monetary Economy.” *Econometrica* 46 (6):1411–1428.
- Chang, Roberto. 1998. “Credible monetary policy in an infinite horizon model: Recursive approaches.” *Journal of Economic Theory* 81 (2):431–461.
- Kydland, Finn E. and Edward C. Prescott. 1977. “Rules Rather than Discretion: The Inconsistency of Optimal Plans.” *Journal of Political Economy* 106 (5):867–896.
- . 1980. “Time to Build and Aggregate Fluctuations.” *Econometrica* 50 (6):1345–2370.
- Olivera, Julio HG. 1970. “On passive money.” *Journal of Political Economy* 78 (4, Part 2):805–814.
- . 1971. “A note on passive money, inflation, and economic growth.” *Journal of Money, Credit and Banking* 3 (1):137–144.
- Sargent, Thomas J. 2024. “Critique and consequence.” *Journal of Monetary Economics* 141:2–13.
- Sargent, Thomas J. and John Stachurski. 2025. *Dynamic Programming: Finite States*. Cambridge University Press.
- Sargent, Thomas J and Neil Wallace. 1973. “The stability of models of money and growth with perfect foresight.” *Econometrica* 41 (6):1043–1048.
- Sargent, Thomas J and Ziyue Yang. 2025a. “Machine Learning a Ramsey Plan.” *Quarterly Review* forthcoming.
- . 2025b. “Ramsey Plan for Calvo’s Model.” *Quarterly Review* forthcoming.