Machine Learning a Ramsey Plan

Thomas J. Sargent
New York University

Ziyue Yang
Australian National University

December 26, 2024

Abstract

A computer program to calculate a pair $(\vec{\theta}, \vec{\mu})$ of infinite sequences of money creation and price level inflation rates that maximizes a benevolent time 0 government's objective function for a model of Calvo (1978). The program computes an associated monotonically declining, bounded from below, sequence of continuation values whose limit is a worst continuation value that is 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.

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

1 Introduction

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

(a) $\{x_i, y_i\}_{i=1}^I \in X^I \bigtimes Y^I$ is a data set and f is a non-linear least squares regression function.

(b) f maximizes some functional or solves some functional equation.

This paper provides an instance of the second context, a classic optimum problem that seeks a time series 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 $t \in X = \{0, 1, 2, \ldots\}$ into $\mathbf{R}$. 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

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

where $r(\mu^t) = \tilde{r}(g(\mu^t), \mu_t), g,$ and $\tilde{r}(\cdot, \cdot)$ 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 describe a sequence of money growth rates that maximizes welfare criterion v_0 defined in (1). The classic paper by Calvo (1978) is our source of the functions $\tilde{r}$ and g.

The presence of the function g in the government's objective function tells it to take into account effects that μ_s for all $s \geqslant 0$ have on θ_0 when it chooses a time series $\vec{\mu}$. A government that at time 1 instead sought to choose a sequence $\vec{\mu}$ to maximize the alternative welfare criterion

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

would not care about μ_0 or θ_0 and consequently would choose a different $\vec{\mu}$ time series than the maximizer of criterion (1), so the plan that optimizes criterion (1) is **time inconsistent**.

We deploy two machine learning approaches. The first is quite lazy: it writes an algorithm that computes the government planner's objective as a function of a money growth rate sequence and hands it over to a **gradient ascent** optimizer. The appendix describes a less lazy approach that expresses the planner's objective as an affine quadratic form in $\vec{\mu}$, computes first-order conditions for an optimum, arranges them into a system of simultaneous linear equations for $\vec{\mu}$ and then $\vec{\theta}$, and solves them. The second approach uses less computer time to calculate the Ramsey plan.

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 $t \geqslant 1$ 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]

[1]Work 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:

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

for all $t \geqslant 0$.

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

$$\theta_t = \frac{\alpha}{1+\alpha}\theta_{t+1} + \frac{1}{1+\alpha}\mu_t. \tag{4}$$

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:

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

The government values a representative household's utility of real balances at time t according to the utility function:

$$U(m_t - p_t) = u_0 + u_1(m_t - p_t) - \frac{u_2}{2}(m_t - p_t)^2, \quad u_0 > 0,\ u_1 > 0,\ u_2 > 0. \tag{6}$$

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

The money demand function (3) and the utility function (6) imply that[3]:

$$U(-\alpha\theta_t) = u_0 + u_1(-\alpha\theta_t) - \frac{u_2}{2}(-\alpha\theta_t)^2. \tag{7}$$

Via equation (5), a government plan $\vec{\mu} = \{\mu_t\}_{t=0}^{\infty}$ implies a sequence of inflation rates $\vec{\theta} = \{\theta_t\}_{t=0}^{\infty}$.

The government incurs social costs $\frac{c}{2}\mu_t^2$ when it changes the stock of nominal money balances at rate μ_t at time t. Therefore, the one-period welfare function of a benevolent government is:

$$s(\theta_t, \mu_t) = U(-\alpha\theta_t) - \frac{c}{2}\mu_t^2.$$

The government chooses everything it can at time $t = 0$ and wants to maximize

$$V = \sum_{t=0}^{\infty} \beta^t s(\theta_t, \mu_t), \tag{8}$$

where $\beta \in (0, 1)$ is a discount factor.

The Ramsey planner chooses a vector of money growth rates $\vec{\mu}$ to maximize criterion (8) subject to equation (5) and the restriction:

$$\vec{\theta} \in L^2. \tag{9}$$

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

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

where

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

2.2 Basic Objects

Let's remind ourselves of the mathematical objects in play.

We have a pair of sequences of inflation rates and money rates

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

[3] A "bliss level" of real balances is $\frac{u_1}{u_2}$; the inflation rate that attains it is $-\frac{u_1}{u_2\alpha}$.

and a planner's value function:

$$V = \sum_{t=0}^{\infty} \beta^t \left(h_0 + h_1 \theta_t + h_2 \theta_t^2 - \frac{c}{2} \mu_t^2 \right), \tag{11}$$

where we set h_0, h_1, h_2 to match:

$$u_0 + u_1(-\alpha\theta_t) - \frac{u_2}{2}(-\alpha\theta_t)^2$$

with

$$h_0 + h_1\theta_t + h_2\theta_t^2.$$

To make our parameters match as desired, we set:

$$\begin{aligned} h_0 &= u_0, \\ h_1 &= -\alpha u_1, \\ h_2 &= -\frac{u_2 \alpha^2}{2}. \end{aligned} \tag{12}$$

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) asks 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** in the sense that it chooses all components of a bivariate time series $(\vec{\mu}, \vec{\theta})$ at time 0.

2.4 Approximation and Truncation Parameter T

It turns out that the sequences $\{\theta_t\}$ and $\{\mu_t\}$ both converge to stationary values under a Ramsey plan. Consequently, we impose the guess that

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

Convergence of μ_t to $\bar{\mu}$ together with formula (10) for the inflation rate then implies that:

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

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:

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

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

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

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

$$\tilde{\mu} = \begin{bmatrix} \mu_0 & \mu_1 & \cdots & \mu_{T-1} & \bar{\mu} \end{bmatrix}$$

as an input and gives as an output the vector

$$\tilde{\theta} = \begin{bmatrix} \theta_0 & \theta_1 & \cdots & \theta_{T-1} & \bar{\theta} \end{bmatrix},$$

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

$$\theta_t = (1-\lambda) \sum_{j=0}^{T-1-t} \lambda^j \mu_{t+j} + \lambda^{T-t} \bar{\mu}, \tag{13}$$

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

Having defined vector $\tilde{\mu}$ and computed the vector $\tilde{\theta}$ using formula (13), we can rewrite the government's value function (11) as

$$\tilde{V} = \sum_{t=0}^{\infty} \beta^t \left(h_0 + h_1 \theta_t + h_2 \theta_t^2 - \frac{c}{2} \mu_t^2 \right), \tag{14}$$

or more precisely as

$$\tilde{V} = \sum_{t=0}^{T-1} \beta^t \left(h_0 + h_1 \theta_t + h_2 \theta_t^2 - \frac{c}{2} \mu_t^2 \right) + \frac{\beta^T}{1-\beta} \left(h_0 + h_1 \bar{\mu} + h_2 \bar{\mu}^2 - \frac{c}{2} \bar{\mu}^2 \right), \tag{15}$$

where θ_t for $t = 0, 1, \ldots, T-1$ satisfies formula (13).

3 Gradient Ascent Algorithm

We now describe an algorithm that maximizes the criterion function (11) subject to equations (10) by choice of the truncated vector $\tilde{\mu}$. Based on the discussion in Section 2.4, we compute the gradient of the objective function (15) with respect to $\tilde{\mu}$. We can compute it using the following simple algorithm.

Algorithm 1: Compute $\tilde{V}(\tilde{\mu})$ (`Compute_V`)

Require: Parameters $\tilde{\mu}$, $\beta = 0.85$, $c = 2$, $\alpha = 1$, $u_0 = 1$, $u_1 = 0.5$, $u_2 = 3$, $T = 40$

1: Compute $\tilde{\theta}$ using (13)

2: Compute coefficients h_0, h_1, h_2 using (12)

3: Compute $\tilde{V}$ using

$$\tilde{V}(\tilde{\mu}) = \sum_{t=0}^{T-1} \beta^t \left(h_0 + h_1 \tilde{\theta}_t + h_2 \tilde{\theta}_t^2 - \frac{c}{2}\tilde{\mu}_t^2\right) + \frac{\beta^T}{1-\beta}\left(h_0 + h_1 \bar{\mu} + h_2 \bar{\mu}^2 - \frac{c}{2}\bar{\mu}^2\right),$$

4: **return** $\tilde{V}$

We use a Python function `Compute_V` to compute a value $\tilde{V}$ associated with given a vector $\tilde{\mu}$.[4] Our algorithm for maximizing the value function $\tilde{V}$ with respect to $\tilde{\mu}$ employs autodifferentiation in `JAX` (Bradbury et al., 2018) and the Adam optimizer (`AdamOptimizer`) (Kingma, 2014) from `optax` (DeepMind et al., 2020). Autodifferentiation computes the gradient directly from the function `Compute_V`. The `optax` and machine learning libraries typically implement gradient descent, so we reformulate our maximization problem as an equivalent minimization of $-\tilde{V}$ with respect to $\tilde{\mu}$ in Algorithm 2.

We initiate the gradient ascent algorithm with a money growth sequence $\mu_t = 0$ for all $t \geqslant 0$, then iteratively update $\tilde{\mu}$ until convergence. Figure 1 plots the Ramsey plan's μ_t and θ_t for $t = 0, \ldots, T$ against t computed by the algorithm. Note that while θ_t is less than μ_t for low t's, it eventually converges to a limit $\bar{\mu}$ of μ_t as $t \to +\infty$, a consequence of how formula (5) makes θ_t be a weighted average of future μ_t's.

4 Continuation Values

It is useful to compute a sequence $\{v_t\}_{t=0}^T$ of "continuation values"

$$v_t = \sum_{s=t}^{\infty} \beta^{s-t} s(\theta_{t+s}, \mu_{t+s})$$

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

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

[4] We describe and deploy a more sophisticated method to compute the value function in the appendix.

Algorithm 2: Optimization Algorithm for Computing $\tilde{V}$

Require:

Functions: `Compute_V`, `AdamOptimizer`

Parameters: $\eta = 0.1$ (learning rate)

$\varepsilon = 10^{-7}$ (convergence tolerance)

$N = 10,000$ (max number of iterations)

Step 1: Initialization

Set initial guess $\tilde{\mu}_0 \leftarrow \vec{0}$

Compute gradient function $\nabla_{\tilde{\mu}} V = \left[\frac{\partial \tilde{V}}{\partial \mu_1}, \frac{\partial \tilde{V}}{\partial \mu_2}, \ldots, \frac{\partial \tilde{V}}{\partial \mu_T}\right]$ using automatic differentiation `jax.grad`.

Step 2: Optimization

Initialize `AdamOptimizer` with learning rate η, and exponential decay rates.

Set iteration counter $i \leftarrow 0$

repeat

  Compute gradients: $g_i \leftarrow -\nabla_{\tilde{\mu}} V(\tilde{\mu}_i)$ *# For maximization*

  Update parameters: $\tilde{\mu}_{i+1} \leftarrow$ `AdamOptimizer`$(\tilde{\mu}_i, g_i)$

  if $\|g_i\| < \varepsilon$ **then**

    Convergence achieved: $\tilde{\mu}^* \leftarrow \tilde{\mu}_i$

    break

  end if

  $i \leftarrow i + 1$

until $i \geqslant N$

$V^* \leftarrow$ `Compute_V`$(\tilde{\mu}^*)$

return $\tilde{\mu}^*, \tilde{V}^*$

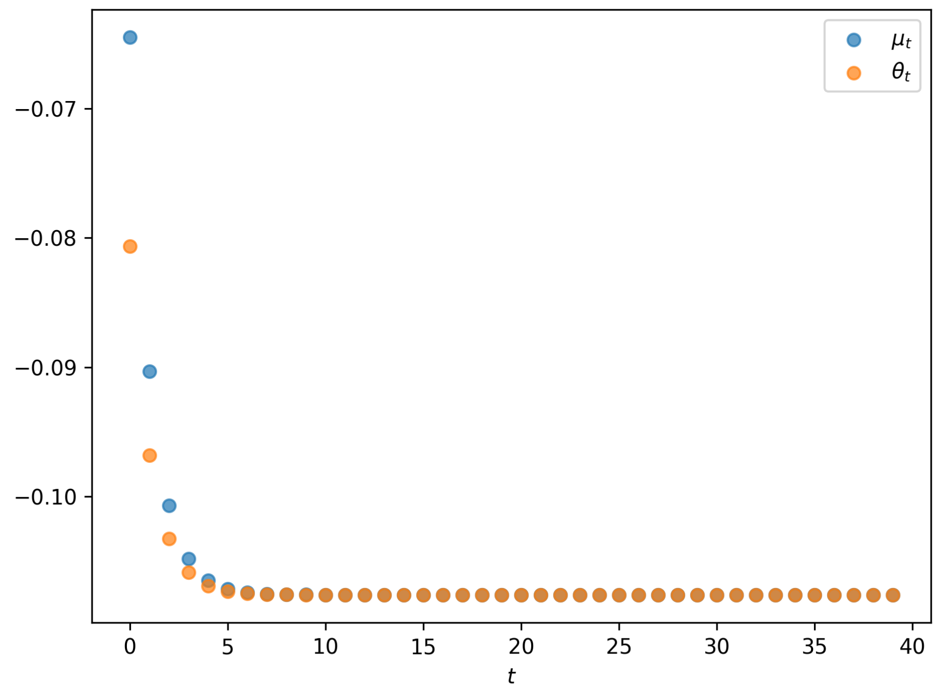

Figure 1: Ramsey plan ($\vec{\mu}$ and $\vec{\theta}$)

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 $t = T - 1, T - 2, \ldots, 0$.

The initial continuation value v_0 should equal the optimized value of the Ramsey planner's criterion V defined in equation (8). We verify approximate equality by inspecting Figure 2, which plots v_t against t for $t = 0, \ldots, T$.

Before studying Figure 2 in detail, we take a brief detour. Recall that 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 $t \geqslant 0$. Computing μ^{CR} with a gradient ascent algorithm is easy.

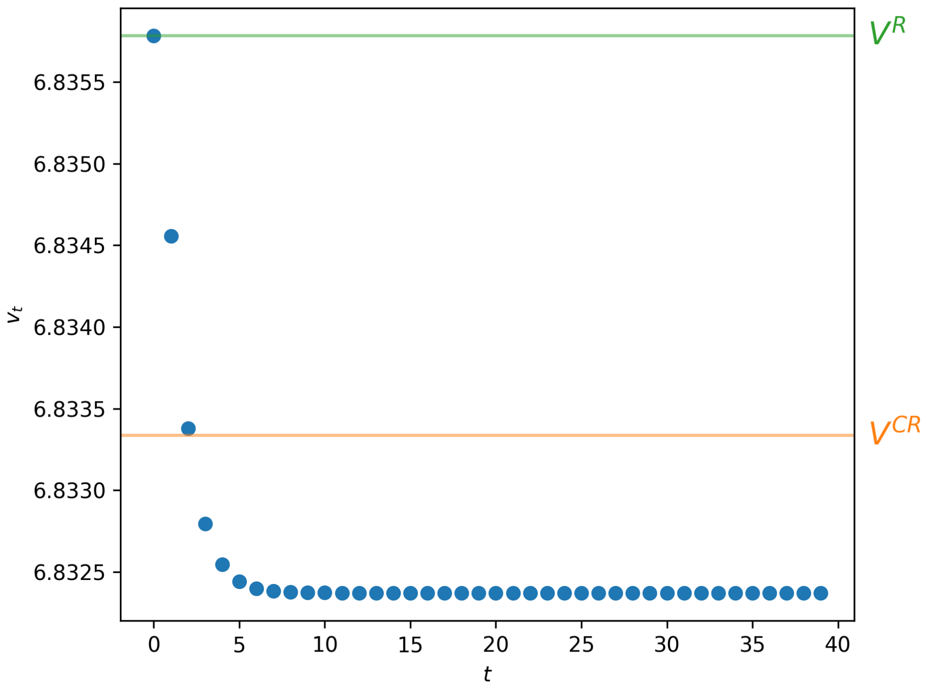

Figure 2: Continuation values

Now turn to figure Figure 2 and observe (a) that the sequence of continuation values $\{v_t\}_{t=0}^T$ is monotonically decreasing; (b) that $v_0 > V^{CR} > v_T$ so that (c) the value v_0 of the ordinary Ramsey plan exceeds the value V^{CR} of the special Ramsey plan in which the planner is constrained to set $\mu_t = \mu^{CR}$ for all t; (d) the continuation value v_T of the ordinary Ramsey plan for $t \geqslant T$ is constant and is less than the value V^{CR} of the special Ramsey plan in which the planner is constrained to set $\mu_t = \mu^{CR}$ for all t, and (e) the worst continuation value v_T is what some macroeconomists call the "value of a Ramsey plan under a time-less perspective." (We'll have more to say about this concept soon.)

5 Applying Human Intelligence

So far, we have represented a Ramsey plan in the **open loop** form of a function

$$\mu_t = f(t) \tag{16}$$

that maps $t \in \{0, 1, 2, \ldots, \}$ to $\mu_t \in \mathbf{R}$.

As indicated in Figure 1, the Ramsey planner makes $\vec{\mu}$ and $\vec{\theta}$ both vary over time.

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

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

$$\begin{aligned} \mu_t &= m(z_t) \\ z_{t+1} &= n(z_t), \end{aligned}$$

where z_t is a **state vector**, the second equation is a transition equation for z_{t+1}, and z_0^R is a value that the Ramsey planner chooses for the initial state vector.

Such 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 the Ramsey pair $\vec{\mu}^R, \vec{\theta}^R$ as "fake data" on which we'll run some exploratory least squares regressions.[5] In what follows, we use $\vec{\mu}^R, \vec{\theta}^R$ to denote the "fake data".

We add some **human intelligence** to the **artificial intelligence** embodied in our least squares 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 $\vec{\theta}^R$ and $\vec{\mu}^R$ on other components and hoping that these regressions will reveal structure hidden within the $\vec{\mu}^R, \vec{\theta}^R$ sequences associated with a Ramsey plan.

It is worth pausing to think about roles being played here by **human** and **artificial** intelligence. Artificial intelligence in the form of a computer program runs the regressions.

[5]Thus, 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.

But one is always free to regress anything on anything else. Human intelligence, such as it is, must tell us which regressions to run.

Additional inputs of 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 on the right sides. Table 2 reports several regressions with θ_t on the right sides. 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.

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 ($\tilde{b}_0$)	-0.0403	1.59×10^{-8}	-2.53×10^{6}
	μ_t ($\tilde{b}_1$)	0.6252	1.5×10^{-7}	4.16×10^{6}
	$R^2 = 1.000$			
$\mu_{t+1} = \tilde{d}_0 + \tilde{d}_1 \mu_t + \varepsilon_t$	Constant ($\tilde{d}_0$)	-0.0645	3.61×10^{-8}	-1.79×10^{6}
	μ_t ($\tilde{d}_1$)	0.4005	3.4×10^{-7}	1.18×10^{6}
	$R^2 = 1.000$			
$v_t = \tilde{g}_0 + \tilde{g}_1 \mu_t + \varepsilon_t$	Constant ($\tilde{g}_0$)	6.8417	0.000	2.09×10^{4}
	μ_t ($\tilde{g}_1$)	0.0864	0.003	27.927
	$R^2 = 0.954$			
$v_t = \tilde{g}_0 + \tilde{g}_1 \mu_t + \tilde{g}_2 \mu_t^2 + \varepsilon_t$	Constant ($\tilde{g}_0$)	6.8281	1.92×10^{-6}	3.55×10^{6}
	μ_t ($\tilde{g}_1$)	-0.2370	4.55×10^{-5}	-5213.119
	μ_t^2 ($\tilde{g}_2$)	-1.8369	0.000	-7125.667
	$R^2 = 1.000$			

Notice that the R^2 statistic is 1, so the error term is nearly zero and we have discovered that

$$\theta_t = -.0403 + .6252 \mu_t$$

We plot the regression line in the left panel of Figure 3. The dots indicate μ_t, θ_t pairs for $t = 0, 1, 2, \ldots$ that converge from above to a limiting pair $\underline{\mu}, \underline{\theta}$.

In hopes of discovering a law of motion of $\vec{\mu}$ under the Ramsey plan, the second entry of Table 1 reports the least squares affine regression $\mu_{t+1} = \tilde{d}_0 + \tilde{d}_1 \mu_t + \varepsilon_t$, where, recycling notation, ε_t is now a least squares residual that is by construction orthogonal to μ_t.

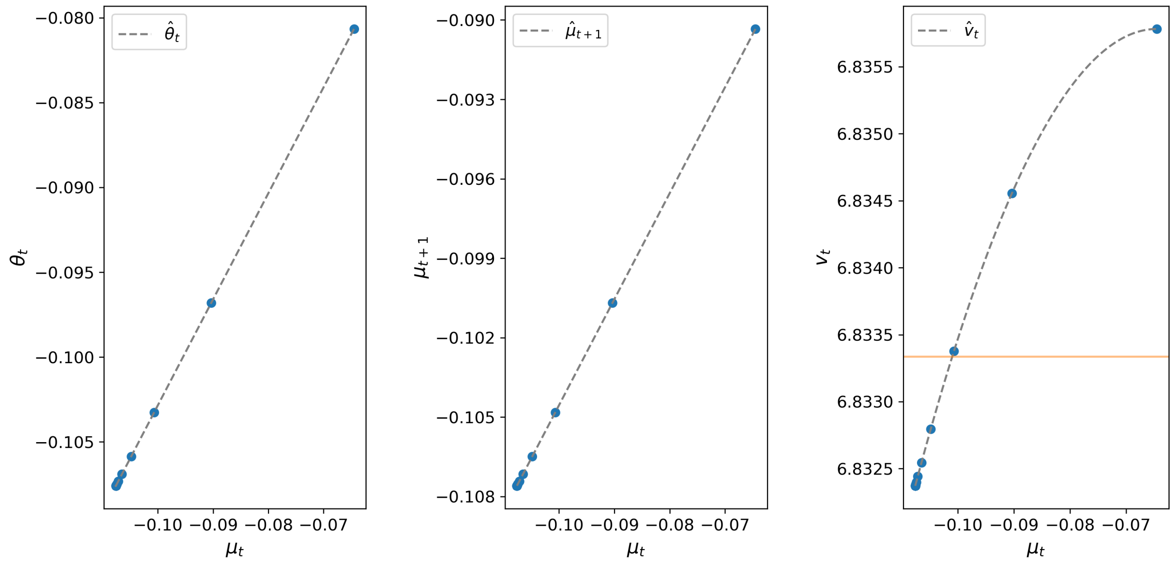

Figure 3: Regression of θ_t on a constant and μ_t (left), regression of μ_{t+1} on a constant and μ_t (center), and regression of v_t on a constant, μ_t, and μ_t^2. The orange line denotes the value of V^{CR} (right).

We again got a perfect fit and have now discovered the following Ramsey planner's law of motion for $\vec{\mu}^R$:

$$\mu_{t+1} = -.0645 + .4005\mu_t$$

We plot the regression line in the middle panel of Figure 3. Here the dots indicate μ_t, μ_{t+1} pairs for $t = 0, 1, 2, \ldots$ that converge from above to a limiting pair $\underline{\mu}, \underline{\mu}$.

The third entry of Table 1 reports the least squares affine regression $v_t = \tilde{g}_0 + \tilde{g}_1\mu_t + \varepsilon_t$, where, again recycling notation, ε_t is now a least squares residual that is by construction orthogonal to μ_t. The R^2 indicates the affine regression line explains only 95.4% of the variation in v_t. Since the R^2 is less than 1, the non-zero least squares residual ε_t that is orthogonal to μ_t remains in the equation. We find this displeasing because neither the government nor the representative agent faces any uncertainty and μ_t seems to be the only thing that can affect v_t. In hopes of reducing the error terms, the fourth entry of Table 1 reports a least squares regression of v_t on a constant, μ_t, and μ_t^2. Now the R^2 is 1, so ε_t disappears and we have succeeded in unearthing the following representation the time t continuation valuation v_t as a function of the time t money growth rate μ_t:

$$v_t = 6.8281 - .2370\mu_t - 1.8369\mu_t^2,$$

with regression line plotted in the right panel of Figure 3.

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

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

$$\begin{aligned} \mu_0 &= \mu_0^R \\ \theta_t &= \tilde{b}_0 + \tilde{b}_1 \mu_t \\ \mu_{t+1} &= \tilde{d}_0 + \tilde{d}_1 \mu_t, \end{aligned} \tag{17}$$

where $\tilde{b}_0, \tilde{b}_1, \tilde{d}_0, \tilde{d}_1$ 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 quadratic function

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

5.1 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 there central banks to do. The three 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)	0.0645	4.42×10^{-8}	1.46×10^6
	θ_t (b_1)	1.5995	4.14×10^{-7}	3.86×10^6
	$R^2 = 1.000$			
$\theta_{t+1} = d_0 + d_1\theta_t + \varepsilon_t$	Constant (d_0)	-0.0645	4.84×10^{-8}	-1.33×10^6
	θ_t (d_1)	0.4005	4.54×10^{-7}	8.82×10^5
	$R^2 = 1.000$			
$v_t = g_0 + g_1\theta_t + g_2\theta_t^2 + \varepsilon_t$	Constant (g_0)	6.8052	5.91×10^{-6}	1.15×10^6
	θ_t (g_1)	-0.7581	0.000	-6028.976
	θ_t^2 (g_2)	-4.6996	0.001	-7131.888
	$R^2 = 1.000$			

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

$$\mu_t = .0645 + 1.5995\theta_t$$

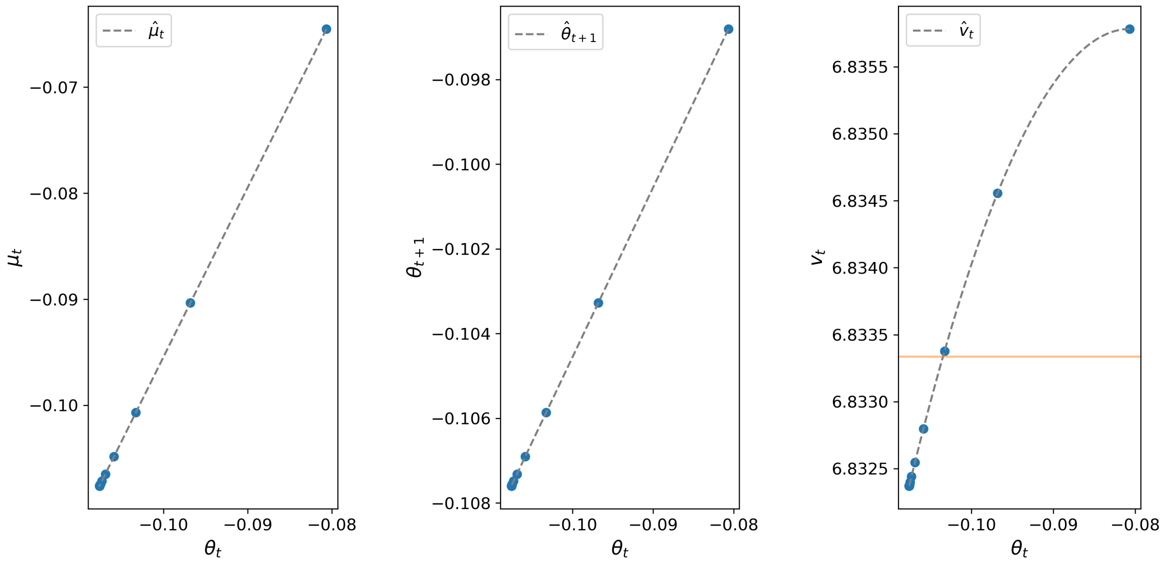

Figure 4: Regression of μ_t on a constant and θ_t (left), regression of θ_{t+1} on a constant and θ_t (center), and regression of v_t on a constant and θ_t and θ_t^2. The orange line depicts V^{CR} (right).

fits perfectly and so do the regression lines

$$\theta_{t+1} = -.0645 + .4005\theta_t$$

$$v_t = 6.8052 - .7580\theta_t - 4.6991\theta_t^2.$$

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

$$\begin{aligned} \theta_0 &= \theta_0^R \\ \mu_t &= b_0 + b_1\theta_t \\ \theta_{t+1} &= d_0 + d_1\theta_t \end{aligned} \tag{18}$$

where b_0, b_1, d_0, d_1 are parameters whose values we estimated with our regressions; we unearthed 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 described by the quadratic function

$$v_t = g_0 + g_1\theta_t + g_2\theta_t^2$$

As with our earlier regressions with μ_t on the right side, we discovered these relationships by running some regressions, staring at the results, and noticing that R^2's of unity tell us that the fits are perfect.

The right panel of Figure 4 indicates that the highest continuation value v_0 at $t = 0$

appears at the peak of the quadratic function $g_0 + g_1\theta_t + g_2\theta_t^2$. Subsequent values of v_t for $t \geqslant 1$ appear to the lower left of the pair (θ_0, v_0) and converge monotonically from above to v_T at time T.

The value V^{CR} attained by the Ramsey plan that is restricted to be a constant $\mu_t = \mu^{CR}$ sequence appears as a horizontal line. Evidently, continuation values $v_t > V^{CR}$ for $t = 0, 1, 2$ while $v_t < V^{CR}$ for $t \geqslant 3$.

It is reasonable to suppose that qualitatively similar relationships would hold along the Ramsey plans that our machine learning algorithms would find for other sets of parameter values $\beta, \alpha, c, u_0, u_1, u_2$, but that the parameters of the regression functions would change. These least squares regression coefficients are themselves complicated non-linear functions of the parameters $\beta, \alpha, c, u_0, u_1, u_2$ that shape the government's criterion function; but we could still expect to find that $\mathbf{R}^2 = 1$ for the corresponding regressions.

6 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 would have to deploy more economic theory in order to discover that (18) is actually a better way to represent a Ramsey plan, as Chang (1998) showed.

We close this paper by providing a preview of an insight of Chang (1998) who noticed 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 $\mu_{t+j},\ j = 0, 1, \ldots$ 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 $\theta_0 \in \mathbf{R}$, define a set

$$\Omega(\theta_0) = \{\theta_{t+1}, \mu_t\}_{t=0}^{\infty} : \theta_{t+1} = \lambda^{-1}\theta_t + (1 - \lambda^{-1})\mu_t \quad \forall t \geqslant 0, \vec{\theta} \in L^2$$

For a given θ_0, use machine learning to compute a closed loop policy

$$\theta_t = f(t; \theta_0), \quad t \geqslant 1$$

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

$$J(\theta_0) = \max_{\{\theta_{t+1},\mu_t\}_{t=0}^{\infty} \in \Omega(\theta_0)} \sum_{t=0}^{\infty} \beta^t s(\theta_t, \mu_t).$$

If we were to do this for a set of different possible θ_0's and then study how $J(\theta_0)$ varies with θ_0, we would discover that

$$J(\theta_0) = g_0 + g_1\theta_0 + g_2\theta_0^2,$$

where the right side is the same quadratic value function that we constructed earlier. 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

$$\theta_0 = \theta_0^R = \arg\max J(\theta) = -\frac{g_1}{2g_2}.$$

Finally, we could compute the value of the Ramsey plan as

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

We have come to the threshold of the formulation of the analysis of Chang (1998). He noticed that a continuation Ramsey planner's value function satisfies the Bellman equation

$$J(\theta) = \max_{\mu,\theta'} \{s(\theta,\mu) + \beta J(\theta')\}, \tag{19}$$

where maximization is subject to

$$\theta' = \lambda^{-1}\theta + (1-\lambda^{-1})\mu.$$

In a sequel to this paper, we shall describe Chang's use of **dynamic programming squared** in which state variable θ that appears in Bellman equation (19) satisfies (5). We can regard this equation as another Bellman equation, one that expresses a 'value' θ_t as a function of next period's 'value' θ_{t+1}. The argument θ in Bellman equation (19) is thus a value governed by another Bellman equation, leading us to call this an instance of a dynamic programming squared problem. [6]

We'll discuss these interpretations of Chang's state variable in the sequel.

[6] In Chang's model, θ_t simultaneously plays multiple roles as inflation target, actual inflation, promised inflation, and expected inflation.

A A Faster Machine Learning Algorithm

By thinking about the mathematical structure of the Ramsey problem and using some linear algebra, we can simplify the problem that we hand over to a **machine learning** algorithm.

We start by recalling that the Ramsey problem that chooses $\vec{\mu}$ to maximize the government's value function (11) subject to equation (10).

This turns out to be an optimization problem with a quadratic objective function and linear constraints. First-order conditions for this problem are a set of simultaneous linear equations in $\vec{\mu}$. If we trust that the second-order conditions for a maximum are also satisfied (they are in our problem), we can compute the Ramsey plan by solving these equations for $\vec{\mu}$.

To remind us of the setting, remember that we have assumed that

$$\mu_t = \mu_T \; \forall t \geqslant T$$

and that

$$\theta_t = \theta_T = \mu_T \; \forall t \geqslant T$$

Again, define

$$\vec{\theta} = \begin{bmatrix} \theta_0 \\ \theta_1 \\ \vdots \\ \theta_{T-1} \\ \theta_T \end{bmatrix}, \quad \vec{\mu} = \begin{bmatrix} \mu_0 \\ \mu_1 \\ \vdots \\ \mu_{T-1} \\ \mu_T \end{bmatrix}$$

Write the system of $T+1$ equations (13) that relate $\vec{\theta}$ to a choice of $\vec{\mu}$ as the single matrix equation

$$\frac{1}{(1-\lambda)} \begin{bmatrix} 1 & -\lambda & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & -\lambda & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & -\lambda & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & -\lambda & 0 \\ 0 & 0 & 0 & 0 & \cdots & 1 & -\lambda \\ 0 & 0 & 0 & 0 & \cdots & 0 & 1-\lambda \end{bmatrix} \begin{bmatrix} \theta_0 \\ \theta_1 \\ \theta_2 \\ \vdots \\ \theta_{T-1} \\ \theta_T \end{bmatrix} = \begin{bmatrix} \mu_0 \\ \mu_1 \\ \mu_2 \\ \vdots \\ \mu_{T-1} \\ \mu_T \end{bmatrix}$$

or

$$A\vec{\theta} = \vec{\mu}.$$

Let $B := A^{-1}$, and we can write

$$\vec{\theta} = B\vec{\mu}$$

As before, the Ramsey planner's criterion is

$$V = \sum_{t=0}^{\infty} \beta^t (h_0 + h_1 \theta_t + h_2 \theta_t^2 - \frac{c}{2} \mu_t^2)$$

With our assumption above, criterion V can be rewritten as

$$\begin{aligned} V = &\sum_{t=0}^{T-1} \beta^t (h_0 + h_1 \theta_t + h_2 \theta_t^2 - \frac{c}{2} \mu_t^2) \\ &+ \frac{\beta^T}{1-\beta} (h_0 + h_1 \theta_T + h_2 \theta_T^2 - \frac{c}{2} \mu_T^2) \end{aligned}$$

To help us write V as a quadratic plus affine form, define

$$\vec{\beta} = \begin{bmatrix} 1 \\ \beta \\ \vdots \\ \beta^{T-1} \\ \frac{\beta^T}{1-\beta} \end{bmatrix}$$

Then we have:

$$h_1 \sum_{t=0}^{\infty} \beta^t \theta_t = h_1 \cdot \vec{\beta}^T \vec{\theta} = (h_1 \cdot B^T \vec{\beta})^T \vec{\mu} = g^T \vec{\mu}$$

where $g = h_1 \cdot B^T \vec{\beta}$ is a $(T+1) \times 1$ vector,

$$h_2 \sum_{t=0}^{\infty} \beta^t \theta_t^2 = \vec{\mu}^T (B^T (h_2 \cdot \vec{\beta} \cdot \mathbf{I}) B) \vec{\mu} = \vec{\mu}^T M \vec{\mu}$$

where $M = B^T (h_2 \cdot \vec{\beta} \cdot \mathbf{I}) B$ is a $(T+1) \times (T+1)$ matrix,

$$\frac{c}{2} \sum_{t=0}^{\infty} \beta^t \mu_t^2 = \vec{\mu}^T (\frac{c}{2} \cdot \vec{\beta} \cdot \mathbf{I}) \vec{\mu} = \vec{\mu}^T F \vec{\mu}$$

where $F = \frac{c}{2} \cdot \vec{\beta} \cdot \mathbf{I}$ is a $(T+1) \times (T+1)$ matrix

It follows that

$$\begin{aligned} J = V - h_0 &= \sum_{t=0}^{\infty} \beta^t (h_1 \theta_t + h_2 \theta_t^2 - \frac{c}{2} \mu_t^2) \\ &= g^T \vec{\mu} + \vec{\mu}^T M \vec{\mu} - \vec{\mu}^T F \vec{\mu} \\ &= g^T \vec{\mu} + \vec{\mu}^T (M - F) \vec{\mu} \\ &= g^T \vec{\mu} + \vec{\mu}^T G \vec{\mu} \end{aligned}$$

where $G = M - F$.

To compute the optimal government plan we want to maximize J with respect to $\vec{\mu}$.

We use linear algebra formulas for differentiating linear and quadratic forms to compute the gradient of J with respect to $\vec{\mu}$

$$\nabla_{\vec{\mu}} J = g + 2G\vec{\mu}.$$

Setting $\nabla_{\vec{\mu}} J = 0$, the maximizing μ is

$$\vec{\mu}^R = -\frac{1}{2} G^{-1} g$$

The associated optimal inflation sequence is

$$\vec{\theta}^R = B\vec{\mu}^R$$

To implement this, we can update our gradient ascent exercise in Algorithm 2 with J and its gradient. This allows us to vectorize the operations. We find that by exploiting more knowledge about the structure of the problem, we can accelerate computation. [7]

[7]For the detailed camparison in the computation time, please see the companion QuantEcon lecture

References

Bradbury, James, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. "JAX: composable transformations of Python+NumPy programs." URL http://github.com/jax-ml/jax.

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.

DeepMind, Igor Babuschkin, Kate Baumli, Alison Bell, Surya Bhupatiraju, Jake Bruce, Peter Buchlovsky, David Budden, Trevor Cai, Aidan Clark, Ivo Danihelka, Antoine Dedieu, Claudio Fantacci, Jonathan Godwin, Chris Jones, Ross Hemsley, Tom Hennigan, Matteo Hessel, Shaobo Hou, Steven Kapturowski, Thomas Keck, Iurii Kemaev, Michael King, Markus Kunesch, Lena Martens, Hamza Merzic, Vladimir Mikulik, Tamara Norman, George Papamakarios, John Quan, Roman Ring, Francisco Ruiz, Alvaro Sanchez, Laurent Sartran, Rosalia Schneider, Eren Sezener, Stephen Spencer, Srivatsan Srinivasan, Miloš Stanojević, Wojciech Stokowiec, Luyu Wang, Guangyao Zhou, and Fabio Viola. 2020. "The DeepMind JAX Ecosystem." URL http://github.com/google-deepmind.

Kingma, Diederik P. 2014. "Adam: A method for stochastic optimization." *arXiv preprint arXiv:1412.6980* .

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 and Neil Wallace. 1973. "The stability of models of money and growth with perfect foresight." *Econometrica: Journal of the Econometric Society* :1043–1048.