Beeps

Jeffrey Ely

December 13, 2013

The Problem

Email Beeps

How do you maximize the time spent not checking email? (and instead presumably being productive)

Email Beeps: Model

- Email arrives at Poisson rate λ.
- Agent has beliefs at each time t whether an email has arrived.
- As soon the probability is at least p^* that an email is waiting, he checks.
- Email software is programmed to beep or not.

No Beep

The probability that an email arrives by time t is

$$1 - e^{-\lambda t}$$

When the beep is off this expression also represents the agent's beliefs at time t about whether there is an email waiting.

No Beep

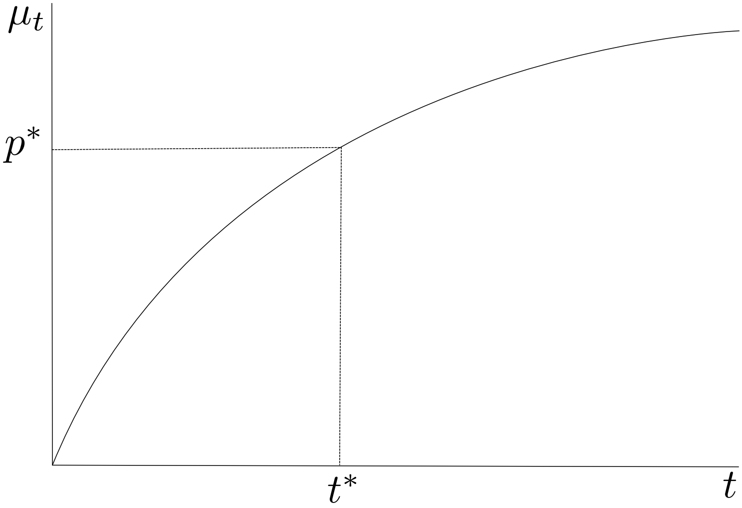

No Beep

When the beep is off, the agent will work for a length of time t^* given by

$$1 - e^{-\lambda t^*} = p^*$$

or

$$t^* = -\frac{log(1 - p^*)}{\lambda}$$

after which he will stop working and read email.

Beep On

When the beep is on the agent learns immediately when an email arrives. Let $\tilde{t}$ be the random time when an email arrives. Then the agents beliefs are a step function.

Beep On

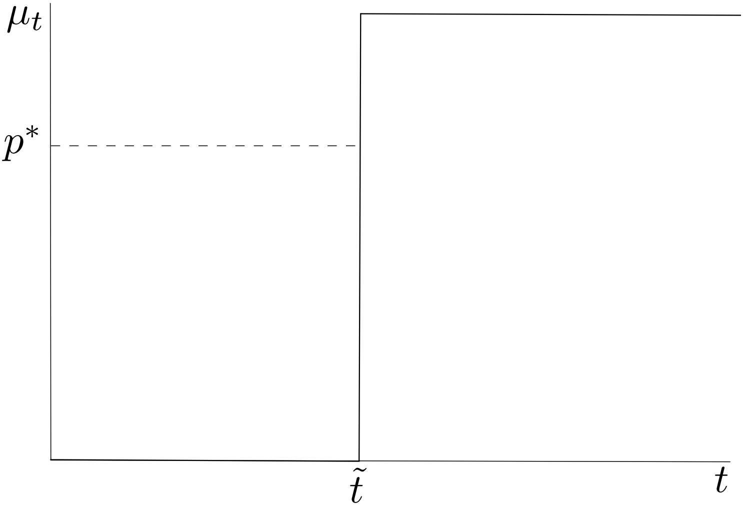

Beep On

With the beep on the agent checks precisely when an email arrives. The expected time spent working is therefore

$$\mathbf{E}\tilde{t} = 1/\lambda$$

Beep or No Beep?

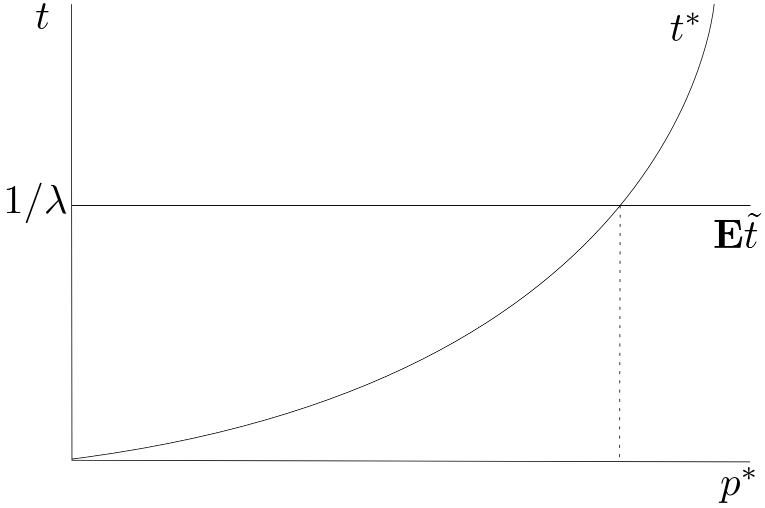

Beep or No Beep?

The beep should be turned on if and only if

$$\begin{aligned} t^* &\leq \mathbf{E}\tilde{t} \\ -\frac{log(1-p^*)}{\lambda} &\leq 1/\lambda \\ p^* &\leq 1 - 1/e \end{aligned}$$

- This is approximately 0.63
- If you are easily tempted by email, turn your beeper on!
- Not because you want to know when email arrives, but because you want to know when it *hasn't*.
- Doesn't depend on arrival rate of email.
- Note: assuming no discounting.

Random Beeps

- When an email arrives, beep with probability z.
- A hybrid of the previous two.
- Conditional on a beep he checks.
- Conditional on no beep his posterior increases, but more slowly than no beeps.
- If there is no beep for a sufficiently long time he will check.
- Gives initial value

$$\frac{1-(1-p^*)^{-\frac{z}{z-1}}}{\lambda z}$$

- The payoff is not monotonic in z and typically an interior beep probability is optimal.
- In fact beep with probability 1 is never optimal.

Other Ways To Beep

1. Beep with a delay
2. Beep with a random delay
3. Beep with variable volume
4. Random, delayed, variable volume beeps...

Optimal Beeps: Model Overview

- The principal continuously observes a stochastic process.
- The agent knows the law of the process but doesn't directly observe its realizations.
- The agent forms beliefs about the state of the process an continuously takes actions.
- The principal sends messages to the agent about the state of the process.
- In order to influence the agent's action.
- The principal has commitment power.
- Dynamic extension of Kamenica and Gentzkow (2011)

Examples

- Fending off audits by the board.
- Moral Hazard/performance reviews.
- Selling a good of fluctuating/depreciating quality...

The Stochastic Process

- Finitely many states $s \in S$.
- Conditional on being in state s_t,
  - State transitions occur at Poisson rate $\lambda \geq 0$,
  - When a transition occurs the new state is drawn from a distribution

$$M_s \in \Delta S.$$

- If $\mu_t \in \Delta S$ is the probability over states at time t, then
  - In continuous time,

$$\dot{\mu}_t = \lambda \left(M_\mu - \mu \right)$$

    where $M_\mu = \sum_s \mu(s) M_s$.
  - In discrete time, $\mu_{t+1} = f(\mu_t) := \mathbf{P}_{t+1}(\cdot \mid \mu_t)$

Timing

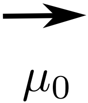

The agent has prior belief μ_0 about the initial state s_0.

Timing

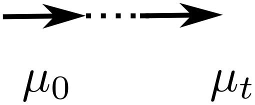

In each period t, the agent enters with a belief $\mu_t \in \Delta S$.

Timing

The principal observes the current state s_t.

Timing

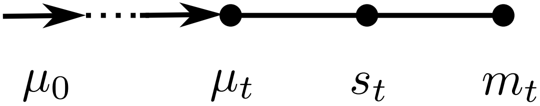

The principal sends a message $m_t \in M_t$.

Timing

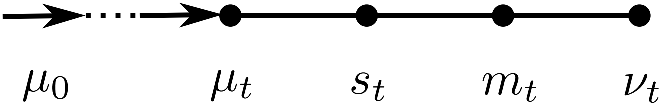

The agent updates to an interim belief $\nu_t \in \Delta S$ and takes an action.

Timing

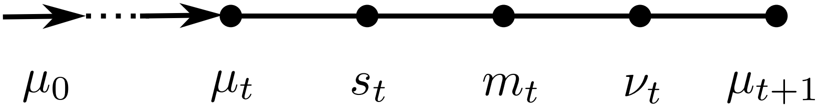

Time passes and the agent updates his belief to

$$\mu_{t+1} = f(\nu_t) := \mathbf{P}_{t+1}(\cdot \mid \nu_t)$$

Timing

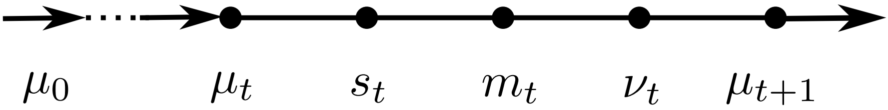

The principal knows μ_{t+1}.

Policy, Commitment

The Principal's message m_t can be an arbitrary function of the complete prior history.

$$\sigma(h_t) \in \Delta M_t$$

The Principal commits to a policy $\sigma(\cdot)$.

Indirect Utility

The agent's action is chosen to maximize his payoff

$$a_t \in \mathrm{argmax}_a \mathbf{E}_{\nu_t} v(a, s)$$

And the Principal obtains (flow) indirect utility

$$u(a_t(\nu_t))$$

Indeed we will take $u(\nu_t)$ as the primitive assuming only that it is bounded and upper-semicontinuous.

Principal's Objective

The principal chooses a policy to maximize the expectation of

$$(1-\delta)\sum_{t=0}^{\infty}\delta^t u(\nu_t)$$

Beeps

- $S = \{0, 1\}$. (no email, email waiting.)
- μ_t, ν_t are probabilities that email has arrived.
- $$u(\nu) = \begin{cases} 1 & \text{if } \nu \leq p^* \\ 0 & \text{otherwise} \end{cases}$$
- Policies
  - Beep on: $M_t = \{0, 1\}$, $\sigma_t(h_t) = s_t$.
  - Random beeps: $\sigma_t(s_t) \in \Delta\{0, 1\}$.
  - Variable volume beeps: $M_t = [0, 1]$.
  - etc.

The Obfuscation Principle

Lemma

Any policy σ induces a stochastic process (μ_t, ν_t) satisfying

1. $\mathbf{E}(\nu_t \mid \mu_t) = \mu_t$,
2. $\mu_{t+1} = f(\nu_t)$.

Any stochastic process with initial belief μ_0 satisfying these properties can be generated by a policy σ which depends only on the current belief μ_t and the current state s_t.

Reformulate The Problem

The principal chooses the stochastic process (μ_t, ν_t) directly.

- Given μ_t, the natural state variable, choose a random variable ν_t,
- $\mathbf{E}\nu_t = \mu_t$.
- Then $\mu_{t+1} = f(\nu_t)$.

The principal is just telling the agent what his beliefs should be and the principal's payoff is the expected discounted value of $u(\nu_t)$.

Discrete Time

Normalize the length of a "period" to 1 and let $\delta = e^{-r}$. Here is the value function

$$V(\mu_t) = \max_{\substack{p \in \Delta(\Delta S) \\ \mathbf{E}p = \mu_t}} \mathbf{E}_p \left[(1-\delta)u(\nu_t) + \delta V(f(\nu_t))\right]$$

Given the current state μ_t, the principal chooses a lottery p over his current payoff $u(\nu_t)$ resulting in a transition to the new state $f(\nu_t)$ and continuation value $V(f(\nu_t))$.

Geometric Version

$$V = \mathrm{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$$

Beep On

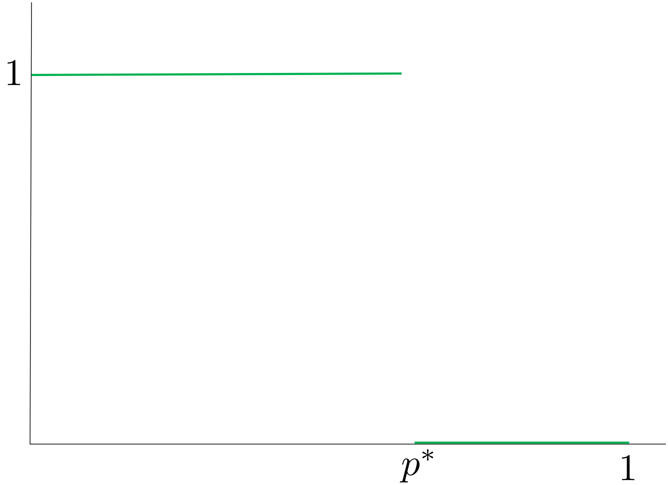

$$\text{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$$

Beep On

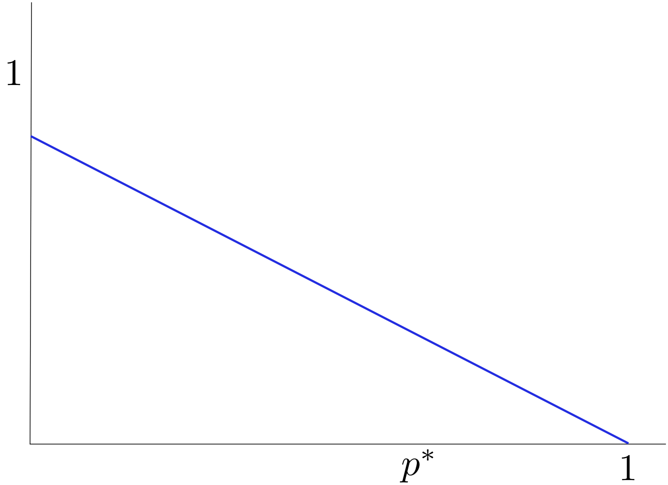

$$\text{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$$

Beep On

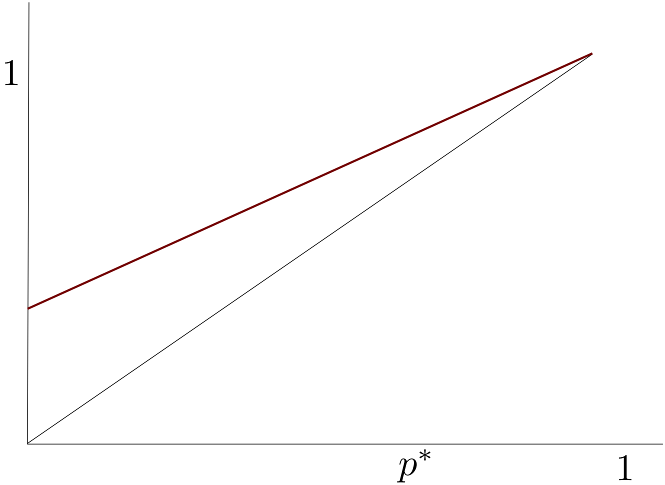

$$\text{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$$

Beep On

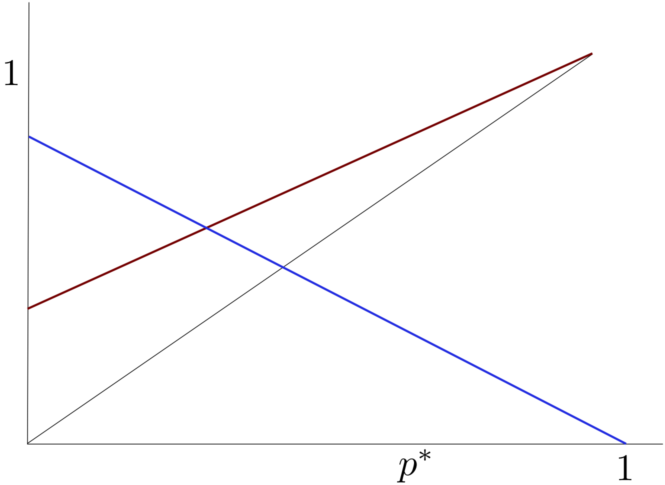

$$\text{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$$

Beep On

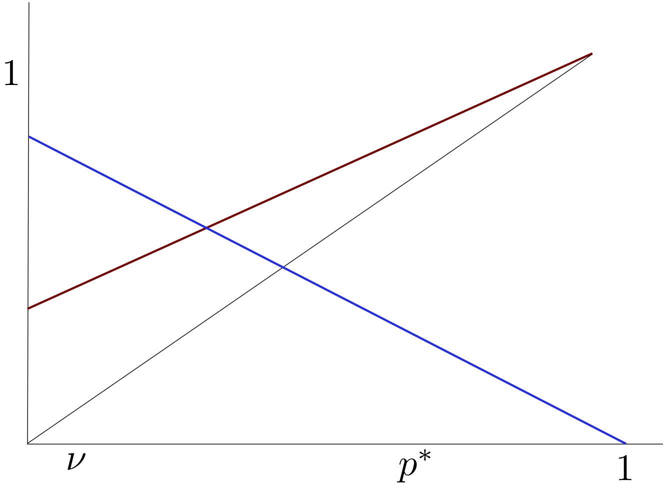

$$\text{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$$

Beep On

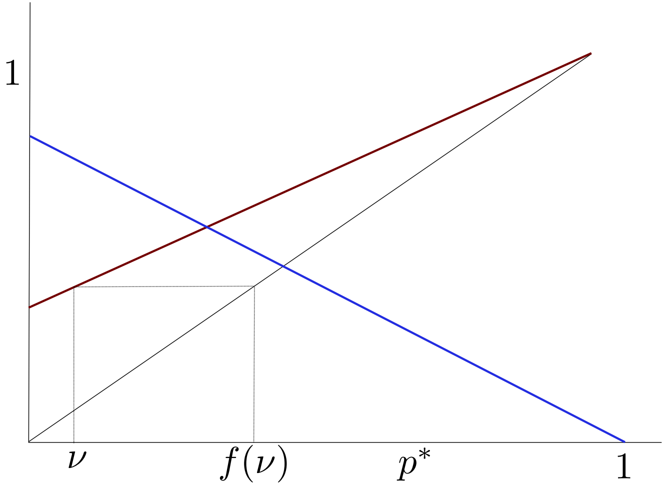

$$\text{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$$

Beep On

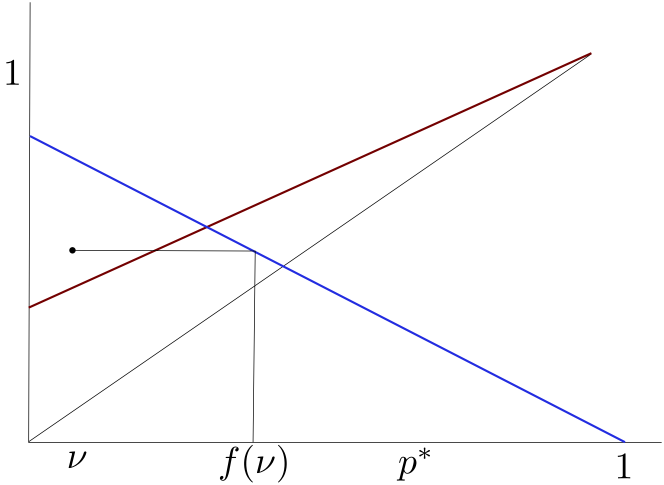

$$\text{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$$

Beep On

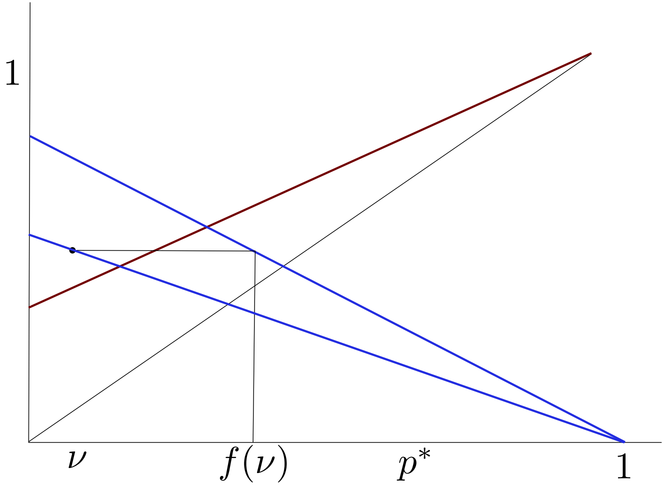

$$\text{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$$

Beep On

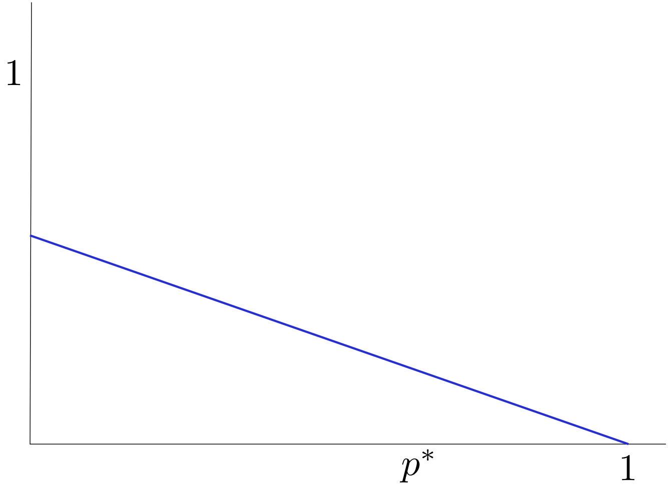

$$\text{cav}\left[(1-\delta)u + \delta(V \circ f)\right]$$

Beep On

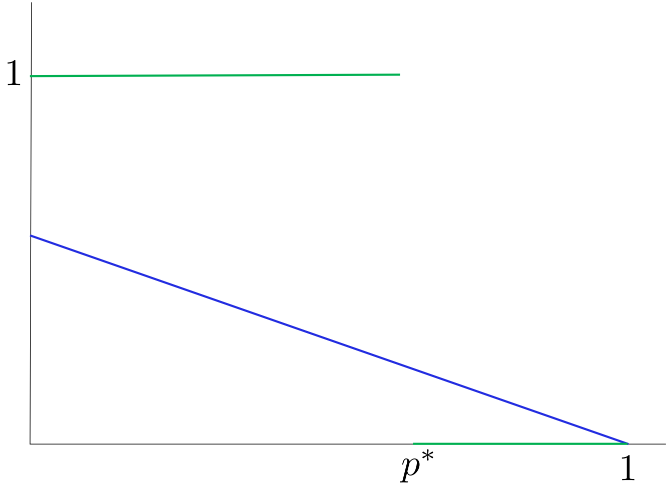

$$\text{cav}\left[(1-\delta)u + \delta(V \circ f)\right]$$

Beep On

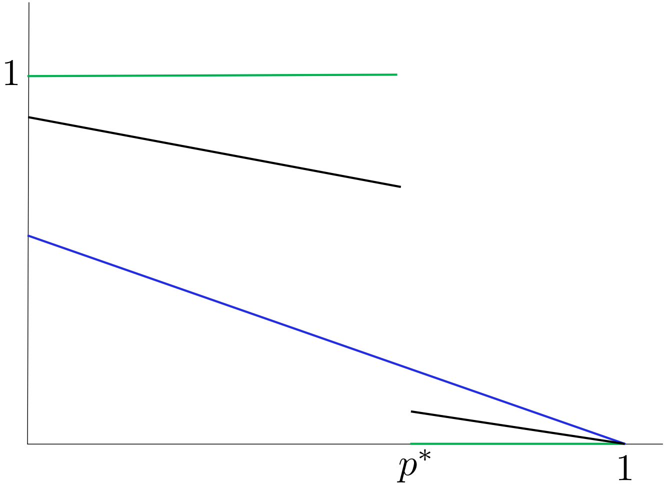

$$\text{cav}\left[(1-\delta)u + \delta(V \circ f)\right]$$

Beep On

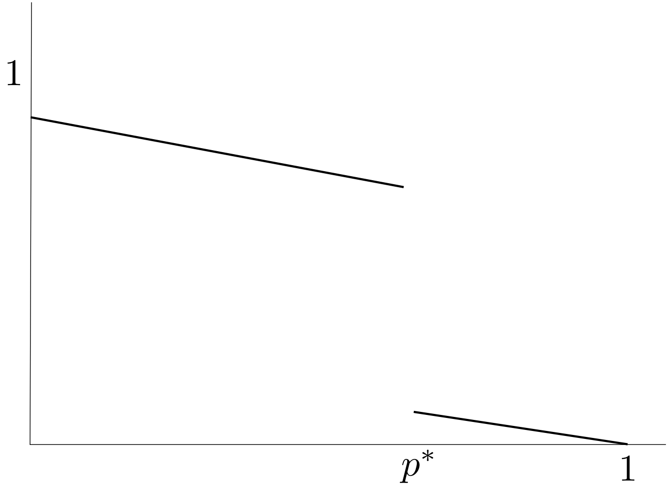

$$\text{cav}\left[(1-\delta)u+\delta\left(V\circ f\right)\right]$$

Beep On

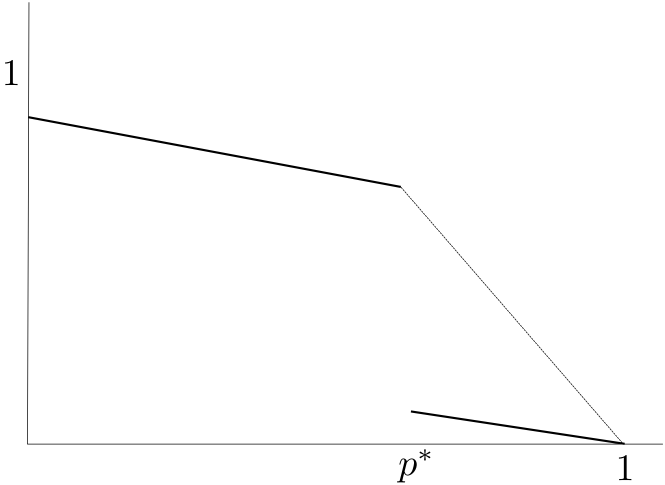

$$\text{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$$

Beep On

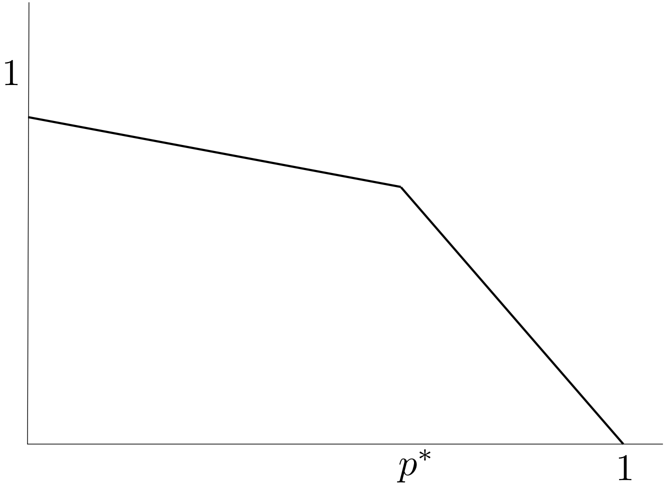

$$\text{cav}\left[(1-\delta)u+\delta\left(V\circ f\right)\right]$$

Beep On

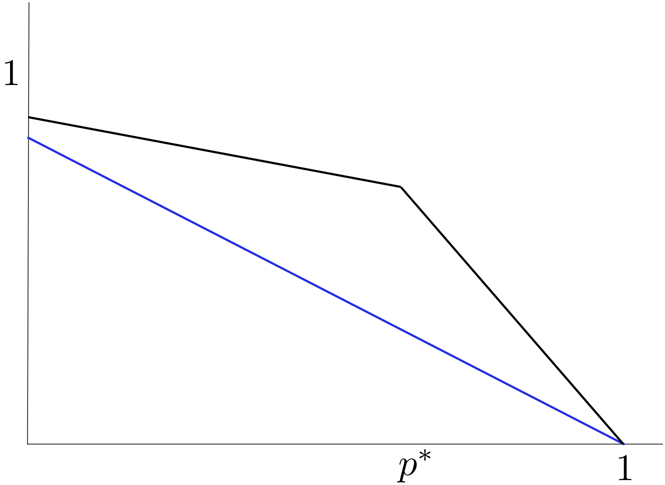

$$V \neq \text{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$$

Fixed Point

The value function is the unique fixed point of a contraction and therefore can be found by iteration.

Iteration

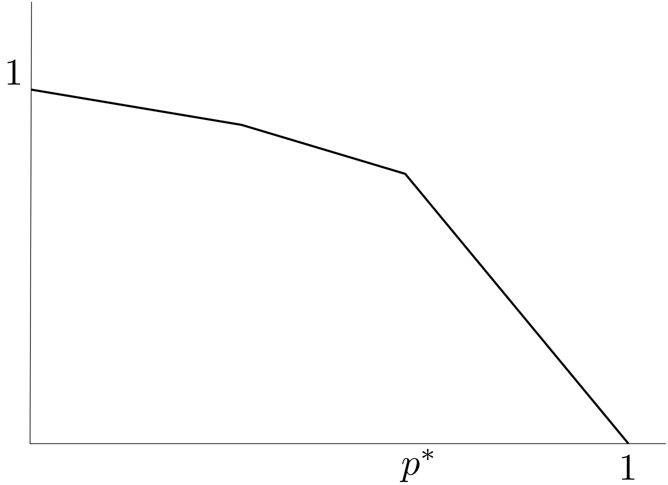

Iteration

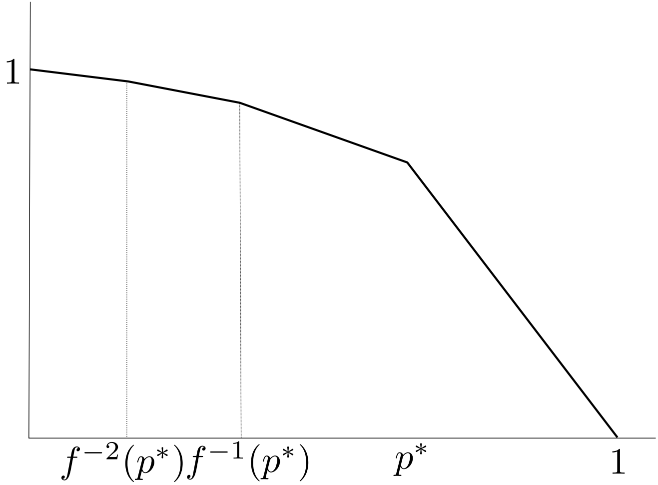

Optimal Policy

Once we have a fixed point, the optimal policy can be easily inferred:

- Where V and $\operatorname{cav}\left[(1-\delta)u + \delta\left(V \circ f\right)\right]$ coincide:
  - Any randomization along a linear segment is optimal. (Random variable volume beeps.)
  - At a kink point the null signal is uniquely optimal.
- Where they don't coincide,
  - The value function must be linear there,
  - The optimal policy jumps to the endpoints of the linear segment.

Continuous Time Limit

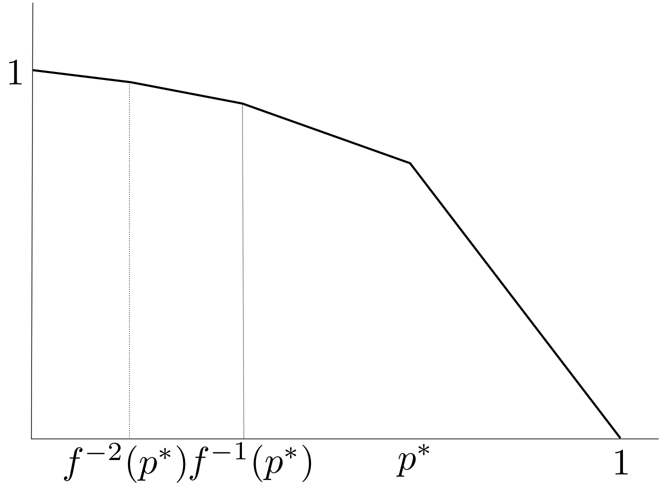

As the period length shrinks.

Continuous Time Limit

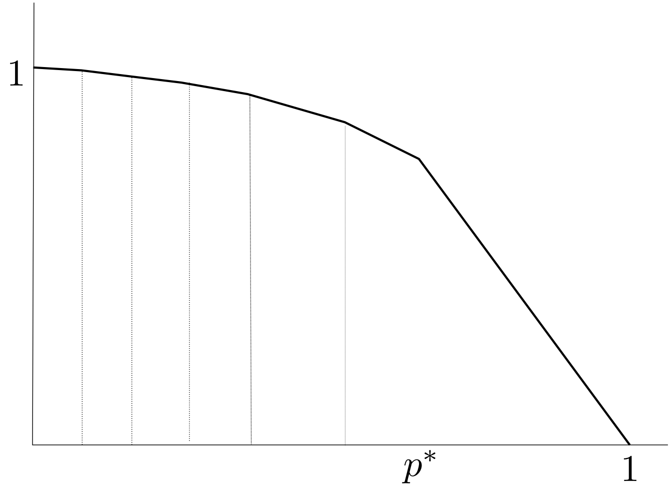

As the period length shrinks.

Continuous Time Limit

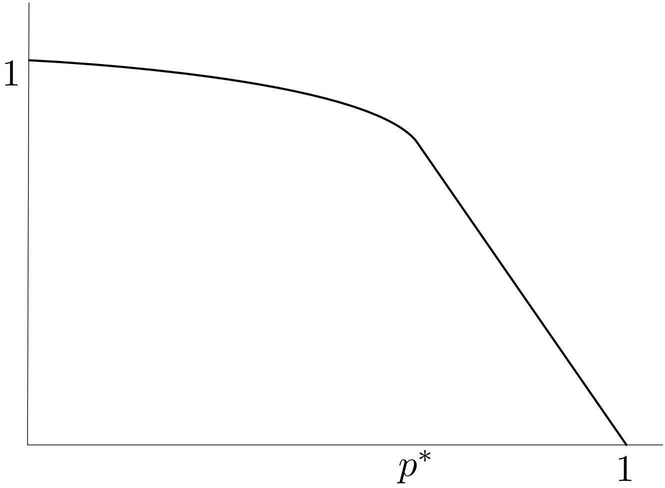

The continuous time limit.

Implementation

Inspecting the limit value function we see that:

- Below p^* it is uniquely optimal to send no message.
- Above p^* it is uniquely optimal to randomize between p^* and 1.
- At p^*?

Differentiable

The value function is differentiable.

- Obviously below and above p^*.
- At p^*:
  - The value function is left- and right-differentiable
  - The left-hand slope must be no steeper than the right-hand (concavity)
  - It cannot be strictly flatter.

$V^*(p^*)$

Differentiability allows us to calculate $V^*(p^*)$:

- To the left of p^* the agent is earning 1 and waiting for $V^*(p^*)$. This gives value

$$V^*(\mu) = \int_0^{t(\mu)} e^{-rt} dt + e^{-rt(\mu)} V^*(p^*)$$

  where $\mu + (1-\mu)(1 - e^{-\lambda t(\mu)}) = p^*$

- We can calculate the derivative with respect to μ.
- To the right of p^*, the value function is linear with slope

$$\frac{-V^*(p^*)}{1-p^*}$$

- Differentiability says these must be equal at p^*.
- Yields $V^*(p^*) = r/(r+\lambda)$.

Implementation

- $r/(r+\lambda)$ is the expected discounted value of a flow payoff of 1 that terminates with Poisson rate λ.
- Therefore it is also the *initial value* (i.e. starting at $\mu = 0$) of the Beep On policy. A hint!
- Consider a beep with a delay of length t^* given by

$$1 - e^{-\lambda t^*} = p^*$$

(Recall that this is the time it takes to reach belief p^* when the beep is off.)

Beep With Delay t^* Is Optimal.

When the policy is to beep with delay t^*,

- Before date t^* there will be no beep for sure.
  - Thus the policy is equivalent to sending no message, as is optimal.
- At date t^* (and after),
  - His belief will be p^*.
  - There will be a beep if and only if an email arrived exactly t^* earlier.
  - These beeps arrive with Poisson rate λ.
  - After hearing a beep the belief moves to $\mu = 1$.
  - Consider the agent's updated belief after no beep.
    - He learns that no email has arrived more than t^* periods ago.
    - He learns nothing about whether an email has arrived at some point later than that.
    - It is therefore as if he has been in a Beep Off policy, starting with belief $\mu = 0$ for the past t^* periods.
    - His belief therefore remains pegged at p^* and he does not check email.
- Thus, at belief p^* the principal earns a flow payoff of 1 until a beep which arrives with Poisson rate λ. This implements the optimal value function. (Beliefs between p^* and 1 are never reached.)

Coincidentally (?)

The optimal policy in continuous time coincides with the (a) static optimum from Kamenica-Gentzkow.

$$\operatorname{cav} u$$

In Fact

Proposition

If u is

- *Concave,*
- *Convex,*
- *Or a step function with 2 steps*

then the static optimal policy is also optimal in continuous time.

Concave u Iteration

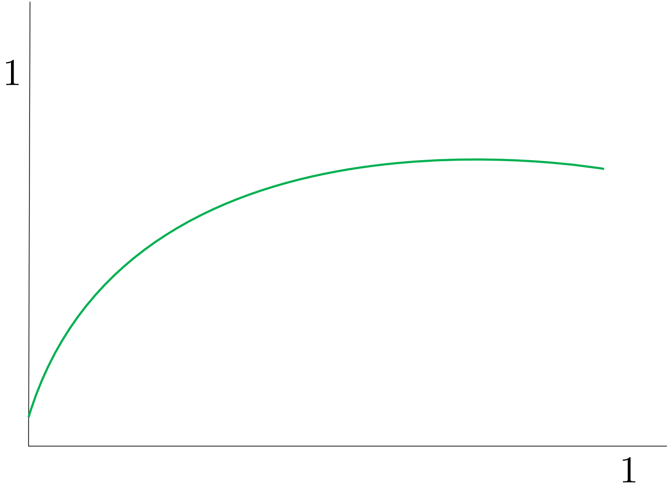

Concave u.

Concave u Iteration

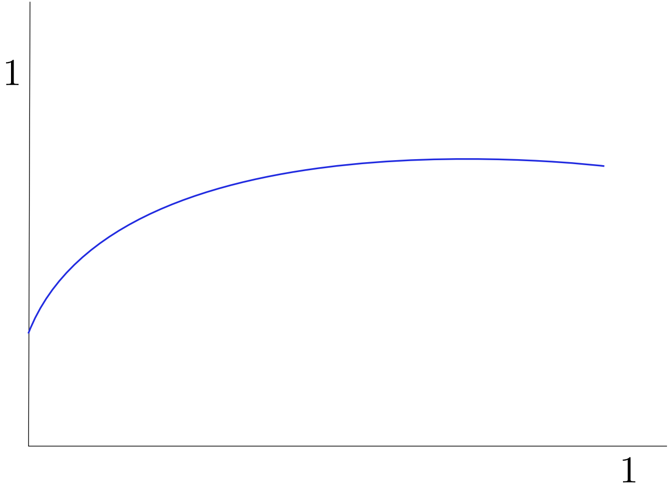

$u \circ f$ is still concave.

Concave u Iteration

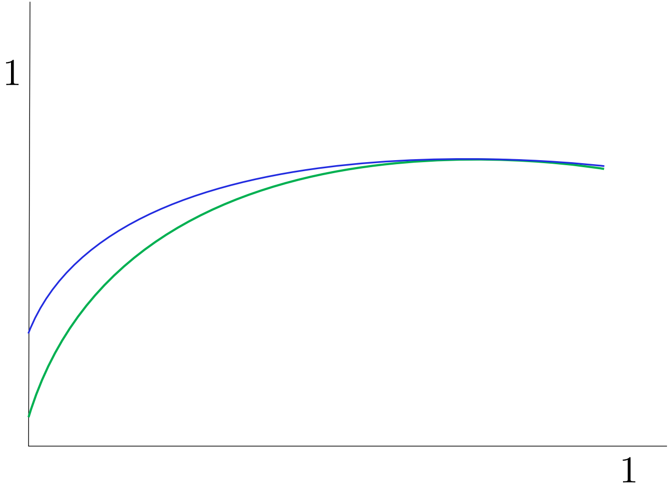

$$(1-\delta)u + \delta\,(u \circ f)$$

Concave u Iteration

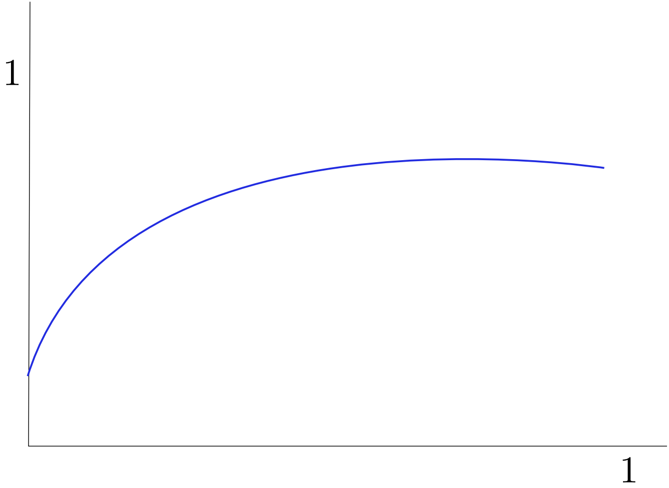

$(1-\delta)u + \delta\,(u \circ f)$ still concave.

Concave u Iteration

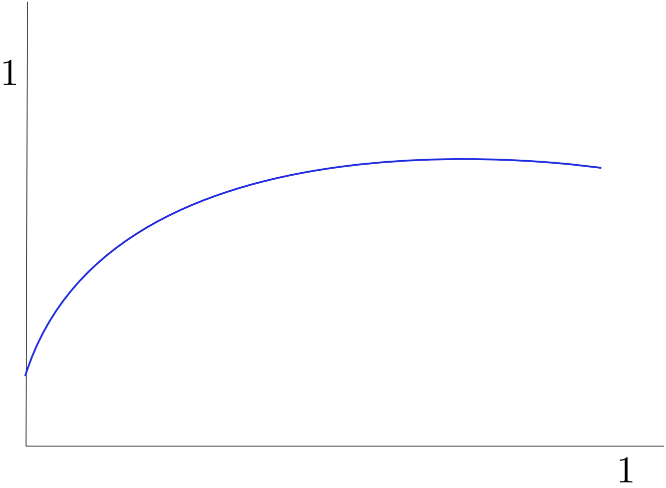

Convex hull of a concave function is the function itself.

Concave u Iteration

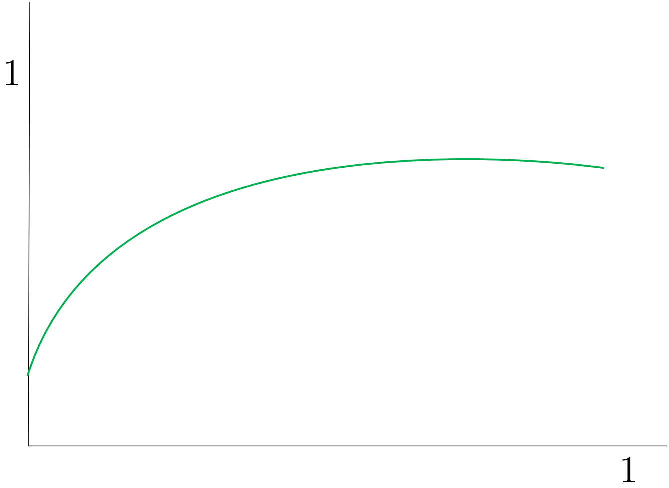

etc.

Convex u Iteration

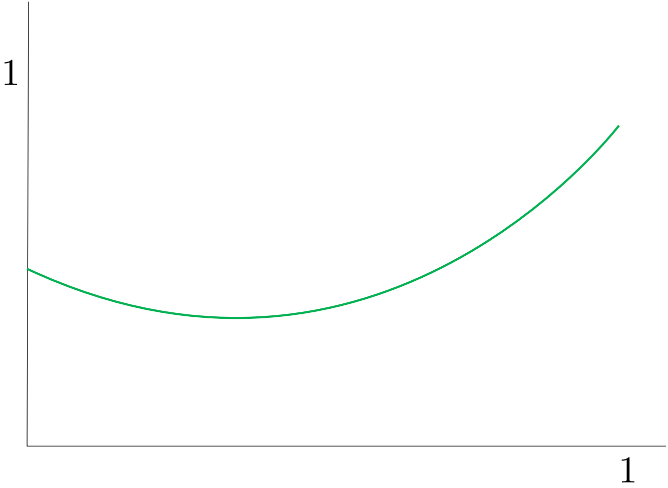

Convex u.

Convex u Iteration

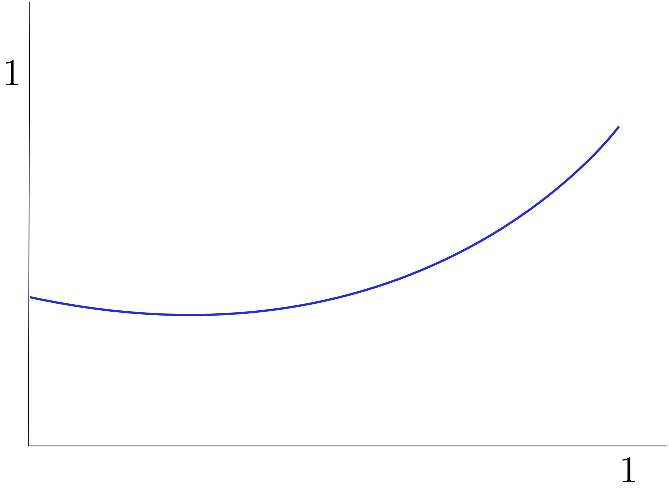

$(1-\delta)u + \delta\,(u \circ f)$ still convex.

Convex u Iteration

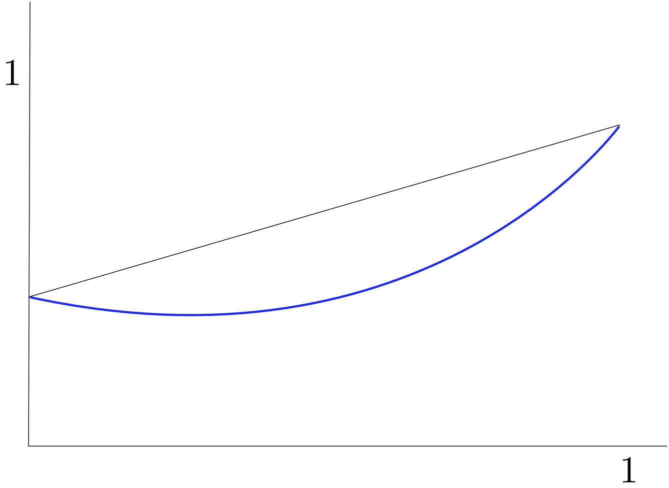

The convex hull of a convex function is linear.

Convex u Iteration

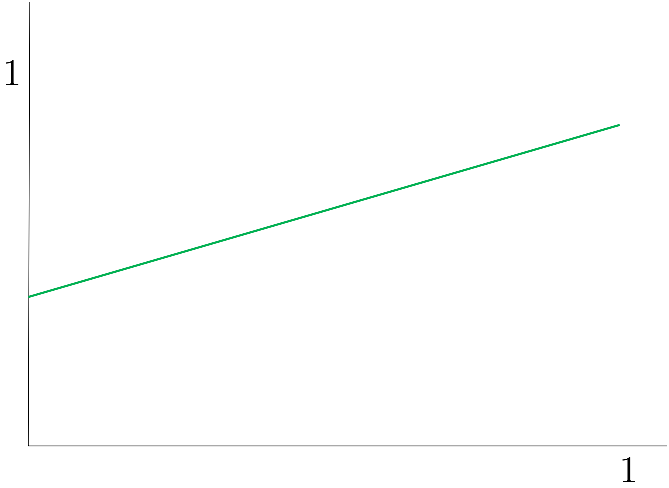

etc.

Intuition?

- The expected value of the agent's posterior at time t is independent of the policy.
- Indeed it equals the posterior in the null policy (beep off.)
- So the choice of policy is just a gamble over paths whose "expected value" is the no-beep path.
- Wherever the function u is convex the principal is "risk-loving" over paths.
- Wherever it is concave he is "risk-averse"

Suggests that the continuous time optimum always coincides with the static optimum.

Counterexample

$$u(v) = \begin{cases} 1 & \text{if } v = 0 \\ 3/4 & \text{if } v \in (0, 1/2] \\ 0 & \text{otherwise} \end{cases}$$

Ongoing Work

- Solving the model directly in continuous time, HJB equation:

$$rV = \text{cav}\left[u + V' \cdot \frac{df}{dt}\right]$$

- Extensions
  - Long run agent (dynamic incentives)
  - Multiple agents (higher-order beliefs)
  - Action affects the state
  - Action generates information
  - Agent has private information