Beeps: Online Appendix

Jeffrey C. Ely *

July 25, 2016

1 Calculations From The Introduction

Here are the calculations for the expected waiting time under the random beep with probability z. First, consider the probability that there has been an arrival conditional on no beep prior to date t. This is the probability that of at least one *silent* arrival conditional on the event that there was no *noisy* arrival. Notice that by the memory-less property of the Poisson process, conditional on any number of noisy arrivals, the distribution of silent arrivals is Poisson with parameter $\lambda(1 - z)$.[1] That is, the probability of a silent arrival is independent of the number of noisy arrivals. In particular the probability of at least one silent arrival conditional on no noisy arrival is just the unconditional probability of a silent arrival, i.e.

$$1 - e^{-\lambda(1-z)t}$$

*Charles E. and Emma H. Morrison Professor of Economics, Northwestern University. jeff@jeffely.com

[1]For example, suppose that a hidden Poisson process operates over a time period of length t and an observer is told that *at least k* arrivals occurred. Then the conditional probability distribution is that the total number of arrivals is $k + x$ where x is Poisson with the same parameter.

Therefore the time it takes for the agent's beliefs to reach p^* in the absence of a beep is

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

and the probability that this time is reached without a beep is

$$e^{-\lambda z \omega}.$$

Finally, the expected time before the first beep conditional on that beep occurring prior to the agent's beliefs reaching the threshold is

$$\frac{1}{\lambda z} - \frac{\omega}{e^{\lambda z \omega} - 1}$$

Since the agent stops working either because his beliefs reach the threshold or because a beep is heard earlier than ω, the total expected waiting time is the weighted average

$$\omega e^{-\lambda z \omega} + \left(\frac{1}{\lambda z} - \frac{\omega}{e^{\lambda z \omega} - 1}\right)\left(1 - e^{-\lambda z \omega}\right)$$

and after some algebra this simplifies to

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

2 Formal Derivation Of The HJB Equation

Let us approximate the optimized continuous time discounted payoff for the principal by discretizing the time dimension into intervals of length h and the summation

$$J(t, \mu_t) = \mathbf{E} \sum_{s=0}^{\infty} e^{-r(t+sh)} u(\mu_{t+sh}) \cdot h + O(h^2).$$

Here $J(t, \mu_t)$ gives the principal's maximal expected total discounted continuation pay-off beginning at time t when the agent's beliefs at instants $\{t + sh\}_{s \geq 0}$ are given by μ_{t+sh}. By the principal of optimality

$$J(t, \mu_t) = \max_{\substack{q \in \Delta(\Delta S) \\ \mathbf{E}q = \mu_t}} \left[e^{-rt} u(\nu_t) h + e^{-rt} J(t + h, f(\nu_t)) \right] + O(h^2).$$

where q denotes a lottery whose realization is ν_t. The optimal policy is stationary so $J(t, \mu) = J(t', \mu)$ and we can write

$$J(t, \mu) = e^{-rt} V(\mu)$$

and

$$e^{-rt} V(\mu) = \max_{\substack{q \in \Delta(\Delta S) \\ \mathbf{E}q = \mu}} \mathbf{E}_q \left[e^{-rt} u(\nu) h + e^{-r(t+h)} V(f(\nu)) \right] + O(h^2).$$

Dividing through by e^{-rt},

$$V(\mu) = \max_{\substack{q \in \Delta(\Delta S) \\ \mathbf{E}q = \mu}} \mathbf{E}_q \left[u(\nu) h + e^{-rh} V(f(\nu)) \right] + O(h^2).$$

Now a first-order approximation

$$e^{-rh} V(f(\nu)) = e^{-rh} \left[V(\nu) + V'(\nu) \frac{d\nu}{dt} h \right] + O(h^2),$$

so

$$V(\mu) = \max_{\substack{q \in \Delta(\Delta S) \\ \mathbf{E}q = \mu}} \mathbf{E}_q \left\{ u(\nu) h + e^{-rh} \left[V(\nu) + V'(\nu) \frac{d\nu}{dt} h \right] \right\} + O(h^2).$$

I claim that at an optimum q^* of the maximization above, $\mathbf{E}_{q^*} V(\nu) = V(\mu)$. To see why, for each $\nu \in \Delta(S)$, let $q^*(\nu)$ be a maximizer for the optimization that defines $V(\nu)$. Then

we have

$$u(v)h + e^{-rh}V(f(v)) \leq \mathbf{E}_{q^*(v)}\left[u(v')h + e^{-rh}V(f(v'))\right] = \max_{\substack{q \in \Delta(\Delta S) \\ \mathbf{E}q = v}} \mathbf{E}_q\left[u(v')h + e^{-rh}V(f(v'))\right]$$

since the left-hand side is a feasible value for the right-hand side optimization, taking q to be the degenerate lottery. Therefore

$$V(\mu) = \mathbf{E}_{q^*}\left[u(v)h + e^{-rh}V(f(v))\right] \leq \mathbf{E}_{q^*}\mathbf{E}_{q^*(v)}\left[u(v')h + e^{-rh}V(f(v'))\right] = \mathbf{E}_{q^*}V(v)$$

but also

$$V(\mu) \geq \mathbf{E}_{q^*}V(v)$$

since the compound lottery in the middle expression above is feasible for the optimization that defines $V(\mu)$. We can thus re-arrange as follows

$$\left(1 - e^{-rh}\right)V(\mu) = \max_{\substack{q \in \Delta(\Delta S) \\ \mathbf{E}q = \mu}} \mathbf{E}_q\left[u(v)h + e^{-rh}V'(v)\frac{dv}{dt}h\right] + O(h^2).$$

Dividing through by h and then taking $h \to 0$ we obtain by l'Hopital's rule ,

$$rV(\mu_t) = \max_{\substack{q \in \Delta(\Delta S) \\ \mathbf{E}q = \mu}} \mathbf{E}_q\left[u(v) + V'(v)\frac{dv}{dt}\right]$$

or

$$rV = \operatorname{cav}\left[u + V' \cdot \frac{dv}{dt}\right].$$

3 Calculations For The Basic Model

3.1 Value Function For Delayed Beep

First, lets calculate the continuation value when the agent's belief is at the threshold p^*. The principal collects a flow payoff of 1 until a beep arrives. He discounts payoffs at rate r and beeps arrive at rate λ. This yields the following expected discounted value $V(p^*)$.

$$
\begin{aligned}
\int_0^\infty \int_0^t e^{-rs} \lambda e^{-\lambda t} ds dt &= \int_0^\infty \left[r^{-1} e^{-rs} |_{s=0}^{s=t} \right] \lambda e^{-\lambda t} dt \\
&= \int_0^\infty \frac{1 - e^{-rt}}{r} \lambda e^{-\lambda t} dt \\
&= r^{-1} \left[1 - \int_0^\infty e^{-rt} \lambda e^{-\lambda t} dt \right] \\
&= r^{-1} \left[1 - \lambda \int_0^\infty e^{-(r+\lambda)t} dt \right] \\
&= r^{-1} \left[1 + \left(\frac{\lambda}{r+\lambda} e^{-(r+\lambda)t} |_{t=0}^{t=\infty} \right) \right] \\
&= r^{-1} \left[1 - \frac{\lambda}{r+\lambda} \right] \\
&= \frac{1}{r+\lambda}
\end{aligned}
$$

Next consider any $\mu \leq p^*$. Let $\tau(\mu)$ be the time remaining before the belief reaches p^*. The principal collects a flow payoff of 1 until the belief for a duration $\tau(\mu)$ then earns the continuation value $V(p^*)$. The overall payoff is thus

$$
V(\mu) = \int_0^{\tau(\mu)} e^{-rt} dt + e^{-r\tau(\mu)} \frac{1}{r+\lambda} = \frac{1}{r} \left[1 - e^{-r\tau(\mu)} + e^{-r\tau(\mu)} \frac{r}{r+\lambda} \right] \tag{1}
$$

3.2 Solving the HJB Equation

We begin with the function in brackets, i.e.,

$$u + V' \cdot \frac{dv}{dt} \tag{2}$$

For $\mu \leq p^*$, the flow payoff is constant $u \equiv 1$, the value function is given by Equation 1 and we have

$$1 + \frac{1}{r}\left[re^{-r\tau(\mu)} \cdot \frac{d\tau}{d\mu} - re^{-r\tau(\mu)}\frac{r}{r+\lambda} \cdot \frac{d\tau}{d\mu}\right]\frac{dv}{dt} = 1 + \left[e^{-r\tau(\mu)} - e^{-r\tau(\mu)}\frac{r}{r+\lambda}\right]\frac{d\tau}{d\mu} \cdot \frac{dv}{dt}$$

and noting that $\frac{d\tau}{d\mu} = -(\frac{dv}{dt})^{-1}$, this reduces to

$$1 - e^{-r\tau(\mu)} + e^{-r\tau(\mu)}\frac{r}{r+\lambda}$$

which is simply rV over the range $\mu \leq p^*$. For $\mu > p^*$, the flow payoff is constant $u \equiv 0$, and the value function is linear with negative slope

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

Thus

$$u + V' \cdot \frac{dv}{dt} = \frac{-V(p^*)}{1 - p^*}\lambda\,(1 - \mu)$$

which is negative for $\mu < 1$ and zero at $\mu = 1$.

Finally we consider the concavification. We can analyze the two intervals $[0, p^*]$ and $[p^*, 1]$ separately. Since the bracketed function equals rV which is already concave on the interval $[0, p^*]$, it is its own concavification on that interval. Turning to the interval $[p^*, 1]$, the linear function which connects the points $(p^*, V(p^*))$ and $(1, 0)$ is concave and no smaller than the bracketed function over this interval. Therefore the concavifi-

6

cation (which is the pointwise smallest such function) must have these two points in its graph. And any function with smaller values elsewhere would not be concave. Thus the linear function is the concavification on this interval.

Finally, rV is just the union of these two functions, and it is concave over the full domain $\mu \in [0, 1]$. It is therefore the concavification of the function in brackets.

4 Proofs For The General Single Agent Model

Proof of Lemma 1. For any current state s and subsequent state s' let $\chi(s, s')$ equal the probability conditional on the process being in state s at time l that the process is in state s' at time $l + h$ where h is the discrete period length. Let χ be the matrix of these transition probabilities. Then if the agent begins period t with belief μ and obtains no information, he will begin period $t + 1$ with belief

$$f(\mu) = \mu \cdot \chi.$$

□

Proof of Theorem 1. It remains only to show that the operator

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

is a contraction mapping on the complete metric space of bounded real-valued functions with the sup metric. The Blackwell sufficient conditions for T to be contraction mapping (with modulus δ) are as follows.

1. If $V \geq V'$ then $TV \geq TV'$,

2. If c is a constant then $T(V + c) \leq TV + \delta c$.

The first follows from the observation that if $g \geq g'$, then any concave function which is weakly larger than g is also weakly larger than g' and therefore $\operatorname{cav} g \geq \operatorname{cav} g'$. The second follows from the observation that for any function h and any constant k, a concave function y is weakly larger than h if and only if $y + k$ is weakly larger than $h + k$. Therefore $\operatorname{cav}(h + k) = (\operatorname{cav} h) + k$. □

5 Calculations For the Monopoly Example

It was asserted that a convex u implies that full-disclosure is optimal and that a concave u implies that no disclosure is optimal. These claims are immediate consequences of Theorem 1. In particular consider any convex u. If V is convex then so is $(1 - \delta)u + \delta \circ f$. The operator T applied to any V is hence the concavification of convex function, i.e. linear (and hence convex.) Thus the set of all convex functions V is closed under the operator T when u is convex. Thus its fixed point is in the (sup-norm) closure of of the set of all convex functions. The set of all convex functions is closed in the sup-norm topology. The optimal value function is thus the concavification of some convex function and therefore it is linear. A linear value function implies that full-disclosure is optimal. An analogous argument shows that the optimal value function is (strictly) concave when u is (resp. strictly) concave. A (strictly) concave value function implies that no disclosure is (resp. strictly) optimal.

Turning to the mixed demand case, let $W(\nu)$ represent the value of the no-disclosure policy beginning with belief ν. It satisfies the recursive equation (recall that we have set λ and r both to 1.)

$$W(\mu) = u(\mu) - W'(\mu) \cdot \frac{d\mu}{dt} \tag{3}$$

and it inherits the shape of u: it is first strictly concave and then strictly convex as μ

8

ranges from 0 to 1.[2] Refer to Figure 1 throughout.

Figure 1: The mixed-demand monopoly example.

We will show that the optimal policy is partially informative: at belief $\mu = 1$ the monopoly continuously sends a partially informative message which sends the belief either to an interior target v^* or keeps it at 1. Once the belief is sent to v^*, no further information is disclosed and the principal earns continuation value $W(v^*)$.

Define $X(v^*)$ by

$$X(v^*) = u(1) - \frac{X(v^*) - W(v^*)}{(1 - v^*)}. \tag{4}$$

Note that this formula, rearranged, gives the initial value from a policy which, at rate $1/(1 - v^*)$ sends the belief to v^* (and otherwise keeps it fixed at 1.)

$$X(v^*) = \frac{u(1)}{1 + \frac{1}{(1-v^*)}} + \frac{\frac{1}{(1-v^*)}W(v^*)}{1 + \frac{1}{(1-v^*)}} = \frac{(1-v^*)u(1) + W(v^*)}{2 - v^*}$$

---

[2]The function W can be computed in closed form:

$$W(\mu) = \frac{6\mu^{1/2} + 8\mu^{5/4} + 3\mu^2}{72}.$$

I claim that there is a unique $v^* \in (0, 1)$ satisfying

$$W'(v^*) = \frac{X(v^*) - W(v^*)}{1 - v^*}.$$

To prove this, first note that for all v^* in the convex range of W, $X(v^*) \geq W(1) > W(v^*) + W'(v^*)(1 - v^*)$. Let us prove these inequalities in turn.

Consider a policy which discloses nothing until the belief reaches $\mu \in (v^*, 1)$ whereupon it sends a binary message which sends the belief to 1 or v^* with probabilities p and $(1 - p)$ respectively. After that it reverts to non-disclosure forever. The value of this policy is

$$W(1) + e^{-\tau(\mu)} \left[pW(1) + (1 - p)W(v^*) - W(\mu) \right]$$

where $\tau(\mu)$ is the time it takes to reach belief μ. If W is convex to the right of v^* and by the law of total probability, this is greater than $W(1)$. By induction therefore the value of a policy which sends the same binary signal every time the belief reaches v^* is greater than $W(1)$. Since this is true for all $\mu \in (v^*, 1)$, it is true by continuity for the policy that generates $X(v^*)$ (i.e. taking limits as μ approaches 1), hence $X(v^*) \geq W(1)$.

The second inequality follows immediately from strict convexity of W on $[v^*, 1]$. Thus, if there exists a solution to the displayed equation it must occur at a value of v^* at which W is (strictly) concave. Moreover by continuity there must exist at least one solution because $W'(0)$ is infinite (because $u'(0)$ is infinite) and $X(0)$ is finite. Rearrange the equation to obtain

$$(2 - v^*)W'(v^*) + W(v^*) - u(1) = 0.$$

Note that the left-hand side is proportional to the first derivative of X evaluated at v^*.

$$X'(v^*) = \frac{(2 - v^*)W'(v^*) + W(v^*) - u(1)}{(2 - v^*)^2}$$

and thus the derivative vanishes at a value v^* that satisfies the displayed equation. When W is strictly concave, the function $W'(\mu)(1 - \mu)$ is strictly decreasing and therefore can intersect at most once where $X'(\mu)$ is zero.

Let $L(\mu)$ be the linear function connecting the points $(v^*, W(v^*))$ and $(1, X(v^*))$ and consider the candidate value function

$$V(\mu) = \begin{cases} W(\mu) & \text{if } \mu \leq v^* \\ \\ L(\mu) & \text{if } \mu > v^*. \end{cases}$$

The function V is concave because we have shown that W is concave on $[0, v^*]$.

We will show that it solves the HJB equation. Define the function U as the bracketed expression in the HJB equation:

$$U(\mu) = u(\mu) - V'(\mu) \cdot \frac{d\mu}{dt}.$$

For $\mu < v^*$, since $V(\mu) = W(\mu)$, i.e. the no information value, it satisfies the differential equation in Equation 3. Therefore $U(\mu) = V(\mu)$. Also, for $\mu = 1$, the definition of $X(v^*)$ (Equation 4) ensures that $U(1) = V(1)$ because $\frac{X(v^*) - W(v^*)}{(1 - v^*)}$ is the slope of L and hence equals $V'(1)$ and $\frac{d\mu}{dt}(1) = -1$. We will now show that for $\mu \in (v^*, 1)$ we have $U(\mu) \leq V(\mu)$. It will then follow that the concave function V is the smallest concave function no smaller than U, i.e. the concavification.

Since

$$V(v^*) = u(v^*) + V'(v^*)\frac{d\mu}{dt}(v^*)$$

11

and

$$V(1) = u(1) + V'(1)\frac{d\mu}{dt}(1)$$

and V is linear on $[v^*, 1]$, as is $V'\frac{d\mu}{dt}$, it follows that over this range

$$V = l + V' \cdot \frac{d\mu}{dt}$$

where l is the linear function whose graph is the chord connecting the points $(v^*, u(v^*))$ and $(1, u(1))$ on the graph of u. Moreover, by definition $U(\mu) = u(\mu) + V'(\mu)\frac{d\mu}{dt}(\mu)$. So if we can show that $l \geq u$ over the range $[v^*, 1]$ then it follows that $U \leq V$.

Consider the graph of u. Since u is concave then convex, there exists a point $\bar{v} \in (0, 1)$ such that for all $\mu \geq \bar{v}$, the chord connecting points $(\mu, u(\mu))$ and $(1, u(1))$ is above the graph of u, and for all $\mu < \bar{v}$, the chord intersects the graph in three points. In fact, the point $\bar{v}$ is the unique point satisfying[3]

$$u'(\bar{v}) = \frac{u(1) - u(\bar{v})}{1 - \bar{v}}.$$

Direct calculation[4] reveals that $v^* > \bar{v}$. Thus, $l \geq u$ over $[v^*, 1]$.

We have shown $U \leq V$. And since the two functions coincide over $[0, v^*]$ and at 1, there can be no smaller concave function that is pointwise larger than U. Thus V is the concavification of U and thus solves the HJB equation.

---

[3]In fact if we were to consider a static version of this problem, we would apply the methods of Kamenica and Gentzkow (2011) and consider the concavification of u. The concavification would coincide with u to the left of $\bar{v}$, and would be linear to the right. Thus, $\bar{v}$ would be the target belief of a partially informative policy in the static problem.

[4]Specifically, $\bar{v} \approx 0.031$ while $v^* \approx 0.058$.

6 Additional Examples

6.1 Beeps In Discrete Time

In discrete time the email beeps example can be described as follows. The set of states is $S = \{0, 1\}$ indicating whether or not an email has arrived to the inbox. The discrete time transition probability from state 0 to 1 is the probability within a period of length h that at least one email arrives and is given by $M = 1 - e^{-\lambda h}$ yielding the following law of motion for the agent's beliefs when uninformed:

$$f(v_t) = v_t + (1 - v_t)M$$

The principal's indirect utility function is

$$u(v) = \begin{cases} 1 & \text{if } v \leq p^* \\ 0 & \text{otherwise} \end{cases}$$

and he maximizes his expected discounted utility where the discount factor is $e^{-r\Delta}$ given a continuous time discount rate r.

In Section 3.1 I derive the value function and optimal policy analytically. Here we will follow a sequence of diagrams to visualize the derivation and gain intuition. Recall the Bellman equation:

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

Consider as an initial guess, a linear V. We can trace through the operations on the right-hand side. The first step is to compose V with the transition mapping f. Since f is linear and $f(v) > v$, this composition has the effect of flattening V by rotating its graph to the left as illustrated in the following figure.
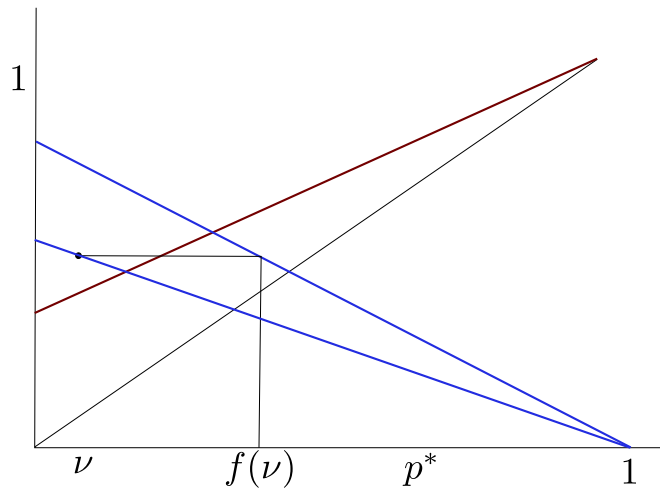
Next, we take the convex combination with the step function u yielding the discontinuous decreasing function below.
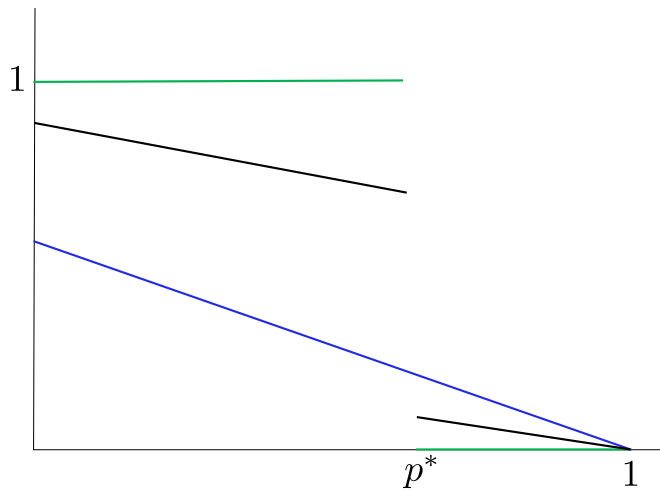
Lastly, we concavify the resulting function as illustrated in the next figure,
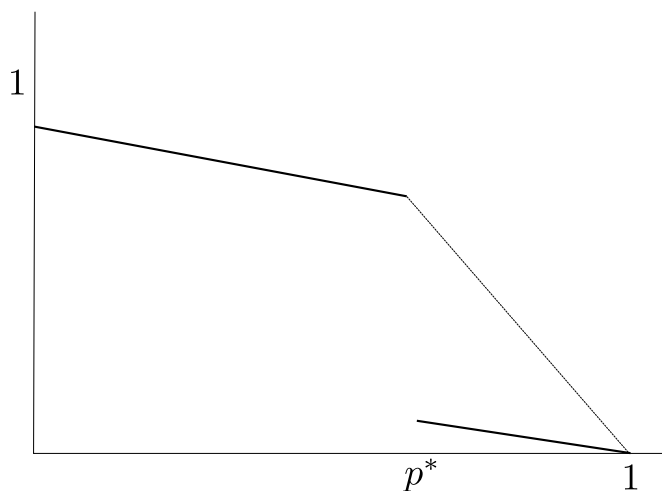
and we have the first iteration of the right-hand side operator. Notice that the function obtained differs from the initial candidate value function which is therefore not a fixed point and not the optimal value. In fact, since the beep-on mechanism discussed in the introduction yields a linear value function, we have shown that beep-on is not an optimal mechanism.[5]

Let us take stock of this first iteration and its implications for the optimal policy. Recall that the concavification represents the optimal lottery over interim beliefs, i.e. the optimal message distribution. At beliefs along a segment where the concavification differs from the underlying function, the optimal policy is to randomize between the beliefs at the endpoints of the segment. Thus in the interval $(p^*, 1]$, the principal wants to send the agent to either $\mu = p^*$ or $\mu = 1$ with the appropriate probabilities. At beliefs along a segment where the concavifcation and the underlying function coincide it is optimal to send no message, as is here between $\mu = 0$ and $\mu = p^*$.

The kink at p^* is a remnant of the discontinuity in the flow payoff u. It is easy to see that this kink will re-appear at every step of the iteration, as well as the linear segment

---

[5]We did not discuss how to interpret beep-on when the agent begins with a prior greater than zero. A fitting story is the following. The agent arrives to his office in the morning with a belief μ_0 that there is an email already there waiting for him. If indeed there is an email it will beep when his computer boots up, i.e. with probability μ_0. If it does not, then his belief jumps to 0 and beep-on continues from there. Thus, the value at μ_0 is just $(1 - \mu_0)V(0)$.

15

from p^* to 1. What subsequent iterations add are additional kinks, first at the point $f^{-1}(p^*)$ in the second iteration, then at $f^{-2}(p^*)$, etc. This occurs when we compose the drift mapping f with the previous iteration, shifting the kinks successively leftward. As we continue to iterate these are the qualitative features of the fixed point to which we converge.[6] The optimal value function is represented in Figure 2.
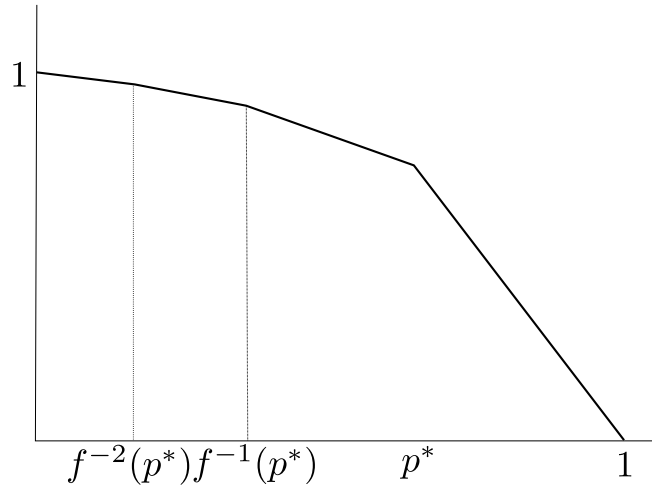

Figure 2: The optimal value function for the Beeps problem.

Now consider what happens to the optimal value function when we shorten the period length. In the shorter time interval the belief moves less between periods and f approaches the identity mapping. This has two implications. First, the number of kinks multiplies and in the limit the value function is smooth to the left of p^*. Second, the slope of the linear segment just to the left of p^* approaches the slope of the linear segment from p^* to 1. In the limit therefore, the value function is differentiable at p^* and indeed at every belief, see Figure 3.

It follows that to the left of p^*, it is uniquely optimal to send no message. In the discrete-time approximation, the linear segments between kinks allowed for a multiplicity of optimal policies ranging from randomization across the whole segment to no

---

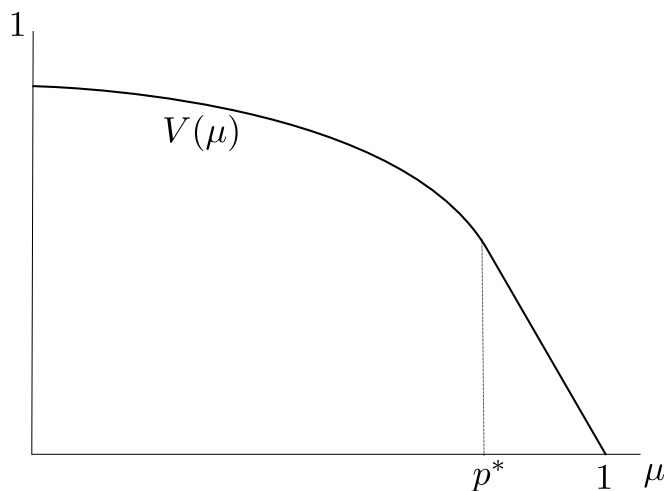[6]The number of kinks will be the maximum index k such that $f^k(0) \le p^*$.

Figure 3: The continuous time value function.

message at all. In continuous time, the strict concavity implies that any non-degenerate lottery is suboptimal. To the right of p^*, it remains optimal to randomize between p^* and 1.

6.2 Ergodic Process

In the beeps example the state $s = 1$ is absorbing. When on the other hand the process is ergodic a new issue must be addressed. If the agent's belief reaches $\mu = 1$, it will begin to drift back to the interior, enabling further information revelation by the principal. How does the principal optimally incorporate this possibility?

To address this, consider now full-support transition probabilities so that the process admits a unique invariant distribution μ^* and let's assume $\mu^* > p^*$.[7]

With μ^* as the invariant distribution, the law of motion for beliefs is no longer monotonic. Beliefs greater than μ^* move downward and beliefs below μ^* move upward. The

---

[7]If $\mu^* \leq p^*$ the problem is simpler and less interesting. Eventually the beliefs will reach p^*. Once there the principal can cease sending messages and the agent will remain at p^* and work forever. The only problem to solve is how to get the agent to start working as quickly as possible when he begins away from p^*. It can be shown that this is accomplished by following exactly the strategy from the original beeps example.

mapping f crosses the 45-degree line at μ^*:
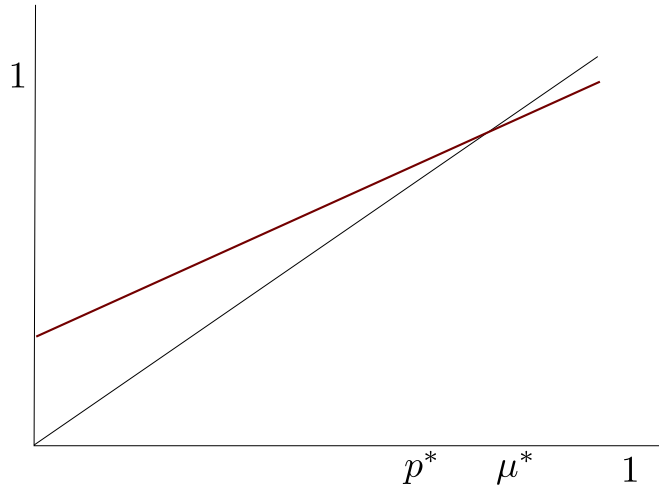
To aid the analysis, it helps to make a general observation about *absorbing sets* of beliefs. Say that an interval $\mathcal{I} \subset [0,1]$ is *absorbing under f* if $f(\mu) \in \mathcal{I}$ for all $\mu \in \mathcal{I}$. According to the following lemma, if u is linear over an interval that is absorbing under f, then the value function must also be linear over that interval.

Lemma 1. *Suppose that $\mathcal{I}$ is absorbing under f, and that u is linear over $\mathcal{I}$. Then V is also linear on $\mathcal{I}$.*

Proof. Consider any candidate value function W which is linear over $\mathcal{I}$. Since f is linear and u is linear over $\mathcal{I}$, the formula

$$(1 - \delta)u + \delta\left(W \circ f\right) \tag{5}$$

must be linear over $\mathcal{I}$ because it is the convex combination of two functions which are linear over that interval. (That the composition is linear follows from the assumption that $\mathcal{I}$ is absorbing and W is linear on $\mathcal{I}$.)

The concavification of Equation 5 must be linear over $\mathcal{I}$. Thus, iteration starting with W must always stay within the set of functions which are linear over $\mathcal{I}$, i.e. the set

18

of such functions is invariant under the value. Since the value mapping is a contraction, iteration converges globally to a fixed point, the fixed point must belong to any invariant set of functions. □

When $\mu^* > p^*$ the interval $(p^*, 1)$ is absorbing under f. Therefore the value function is linear there. It follows that the value function has the same shape as in the absorbing beeps problem. The optimal policy is therefore identical. In terms of implementation there is only one novelty: at $\mu = 1$ beliefs are now trending downward, i.e. the agent knows that eventually there will be a transition from state 1 to state 0. According to the optimal policy, the instant beliefs move into the interior the principal is randomizing between the two endpoints p^* and 1. This is achieved by a random message that reveals state changes but with false positives. As soon as a transition occurs the message is sent, but also each instant a transition does not occur the message is still sent but with a probability less than 1 calibrated so that the message induces interim belief p^*. Since the message has false positives but not false negatives, the agent remains at $\mu = 1$ as long as no message is heard.

In particular, the value $V(1)$ is positive now because the agent will periodically switch back to working when his beliefs jump down to p^*.

6.3 Three States

When there are three states, $S = \{0, 1, 2\}$ and two actions, the threshold is no longer a point but a line segment through the simplex of beliefs ΔS. On one side of the line the agent takes action 0 and on the other side he takes action 1. The belief dynamics operate in a 2-dimensional simplex and can therefore be significantly more complicated. In this section I analyze a simple extension of the beeps problem to 3 states to illustrate. This example is special because the belief dynamics have a monotonicity property. Once the

beliefs leave the region where the agent takes a given action they never return (absent information from the principal). With two states the belief dynamics always have this property. With three states the beliefs may cycle and move up and down the steps of the principal's payoff function u on their path toward the invariant distribution. Such problems are considerably more complex to solve. In independent work Renault, Solan and Vieille (2014) analyze the many-state problem in more detail.

Consider a variation of the email problem in which the agent wishes to check as soon as the probability is sufficiently large, at least p^*, that there are at least *two* emails to read. Let $s \in S = \{0,1,2\}$ denote the number of unread emails currently in the inbox, where $s = 0$ and $s = 1$ indicate the exact number and $s = 2$ indicates 2 or more. Maintain the assumption that email arrives at rate λ. The following diagram illustrates the situation. The simplex is the set of possible beliefs for the agent and the line depicts the threshold above which the agent checks. The arrows show the flow of the agent's beliefs when the principal withholds information. The constant arrival of email implies that the beliefs move up and to the right.
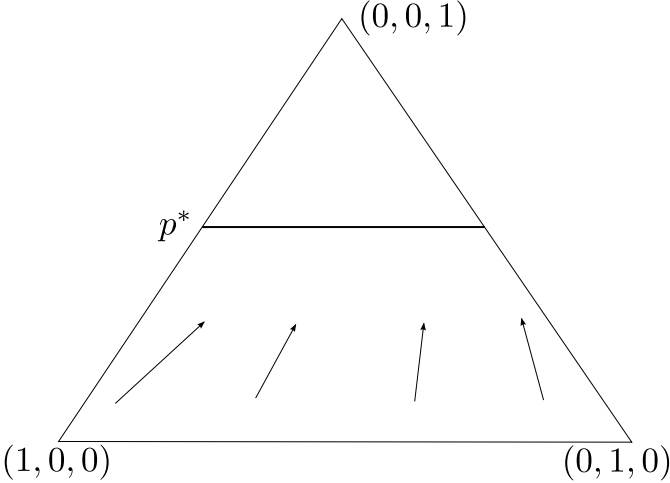
This analysis of this problem follows similar lines as the two-state email beep problem. In the diagram below the lines represent points which lead by iterations of f to the threshold line. If we begin with a candidate value function which is linear and equal

20

to zero at the $s = 2$ vertex, iterations lead to a piecewise linear value function with kinks along these segments.[8] The continuous time limit value function will therefore be linear along these line segments but strictly concave along rays toward the $s = 2$ vertex in the region below the threshold. It will be linear above the threshold and equal to zero at the $s = 2$ vertex. Thus, the optimal mechanism is a delayed beep signaling a past arrival of the second email.
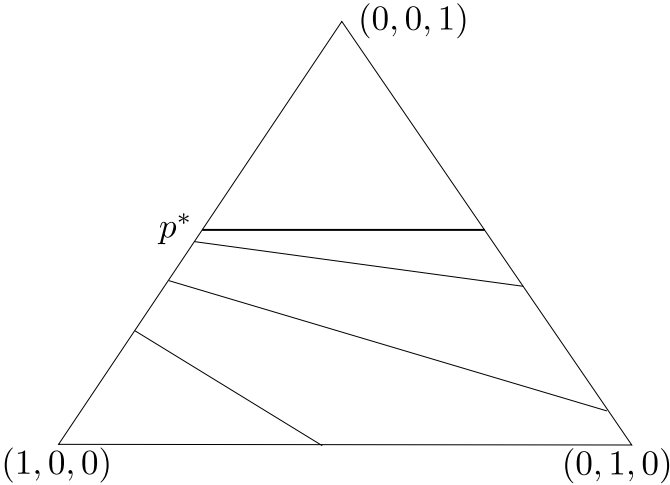
The novelty that arises with three states concerns the evolution of beliefs and continuation values along the threshold. At the threshold the principal's policy is designed to maximize the probability that the agent continues working. As usual this is accomplished by sending the agent either to the threshold or to the most distant point in the shirk region (with the appropriate probability), in this case the vertex $(0, 0, 1)$ where the agent is certain that two emails are waiting. This signal allows the agent to increase his belief that a single email has arrived and thus as long as the agent remains at the threshold, this belief will continue to trend upward, converging toward the right face of the simplex. Note that the right face, where the agent is certain that at least 1 email is waiting, is isomorphic to the original 1-dimensional beeps problem because the agent is simply waiting to find out if one more email arrives.

---

[8]These are not the level sets. As we will show below, the value declines as we move to the right along the p^* line. This pattern will thus be preserved at all of its inverse images under f.

The following diagram shows the path of the agent's beliefs. The beliefs will follow this path until they reach the threshold, then remain on the path until a beep sounds. As long as there is no beep the beliefs will converge asymptotically to the right face.
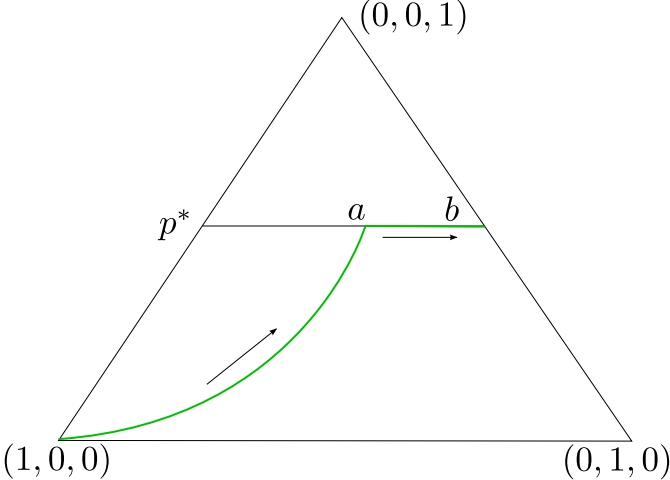
It follows that the length of the optimal delay must change as time passes. To see this, first consider the delay length at the point a where the beliefs first touch the threshold. Let's determine the delay length that keeps the agent on the threshold. Let t_a denote the length of time it takes for beliefs to reach a from the vertex $(1,0,0)$. If the beep has delay t_a then when the agent is at a and hears no beep his updated beliefs continue to attach probability p^* to the presence of 2 emails. The absence of a beep tells the agent that a second email did not arrive t_a moments ago. The key question is what does this information tell the agent about the conditional distribution over the remaining states $s = 0$, $s = 1$. At that point in the past he was at the point $(1,0,0)$. In particular he is certain that not even the first email had arrived. Learning that a second email did not arrive at that moment gives him no information about the other states since the simultaneous arrival of two emails has probability zero.

Thus, the absence of a beep tells him that his beliefs at the time t_a ago were correct, and thus that his current beliefs should be updated from those prior beliefs based only on the information that a time period of length t_a has passed during which he learns

nothing about arrivals. By construction that updated belief assigns exactly p^* to $s = 2$.

By contrast, consider a point like b, further to the right. At this point he attaches higher probability to the presence of a single email. Suppose the principal continued to use a beep with delay t_a. What does the agent believe conditional on hearing no beep? He learns that as of t_a ago, the second email had not arrived. At that point in the past he assigned positive probability to the arrival of the first email. Of course the absence of the second email is information: it makes it less likely that the first email arrived. But nevertheless he will continue to assign positive probability to the event that a first email had arrived t_a ago. To obtain his current beliefs he will update that posterior based on knowing that t_a time has passed. His updated belief that a second has arrived during that time will be larger when starting from a positive probability of a first email than when starting with probability zero. The latter starting point would lead him to p^*, so using the delay length t_a would put him above p^*.

Therefore, in order to keep the agent on the threshold, the delay length must be shorter the more time has passed. In particular if the second email arrives at time t' then the delay before beeping must be shorter than if the second email arrived at time $t < t'$. What happens to this delay length asymptotically as time increases? Since the beliefs are approaching probability 1 that exactly 1 email is waiting, the length of time it takes for the probability of a second email to equal p^* converges to the length of time it would take if the agent were certain at the outset that the first email had already arrived. This is just the length of time for the arrival of a single email and that is the length t^* from the 1-dimensional problem.

We can understand in these terms why the continuation value must decline as we move toward the right face. Because the delay is shortening but the arrival rate of email is constant, it follows that beeps are arriving more quickly and thus the agent is jumping sooner on average to the upper vertex.

6.4 Three Actions

The solutions for each of the examples considered thus far are special in at least two senses. First, as was shown by Renault, Solan and Vieille (2014) the optimal mechanism is *greedy*; in particular it is nearly identical to the optimal mechanism in a static version of the problem. Second, the optimal policy is monotonic in that there is an initial phase of silence leading eventually to a phase of random messages. These features are typical of problems with two states and two actions. In this subsection I present the simplest 3 action problem in which the optimal mechanism is very different than the static problem and in particular consists of an early phase of messages followed by a duration of silence and then ultimately a final message phase.[9]

Consider the following indirect utility function with 3 steps:

$$u(v) = \begin{cases} 5/4 & \text{if } v = 0 \\ 1 & \text{if } v \in (0, 1/2] \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

Here is a story that goes with it. When the agent is certain there is no email waiting he can work with full concentration. When there is a small chance of an email he is tempted to check but he resists the temptation. The effort spent on willpower makes him somewhat less productive. Finally when his beliefs cross the threshold (here $p^* = 1/2$) he succumbs to temptation and stops working.

The interval $(1/2, 1]$ is absorbing and u is linear there and so by Lemma 1 the value function has the familiar linear segment and the optimal mechanism randomizes between 1/2 and 1 when the beliefs are anywhere in between. To analyze the interval

---

[9] Renault, Solan and Vieille (2014) also present an example in which the optimal mechanism is non-greedy involving two actions but three states. Two-state problems are generally easier to analyze so the example below is simpler.

$[0, 1/2)$ note that one feasible mechanism for the principal is to use the optimal mechanism from the basic beep example. That would yield a value function which is differentiable at the threshold $1/2$. Even though this is not optimal for the present problem, it places a lower bound on the optimal value function. In particular, the optimal value function cannot have a kink at $1/2$. The value function is in fact smooth and strictly concave over some interval $(p^{**}, 1/2]$ so that no message is optimal there. However, unlike the original beeps problem, $p^{**} \neq 0$ and on the interval $[0, p^{**}]$ the value function is again linear. Intuitively, randomizing between 0 and p^{**} allows the principal to stay at the point 0, earning the high flow payoff of $5/4$ for some duration, whereas following the original beeps solution the beliefs would spend only an instant there.

With these observations in hand we can now show formally how the optimal mechanism changes in the three action version. Indeed, if we take the original beeps value function V as a candidate continuous-time value function for the three-action problem, the HJB equation tells us to compute its time-derivative and add it to the three-step u in Equation 6. We obtain exactly the same graph except that the height at $\mu = 0$ jumps up by $1/4$. When we concavify:
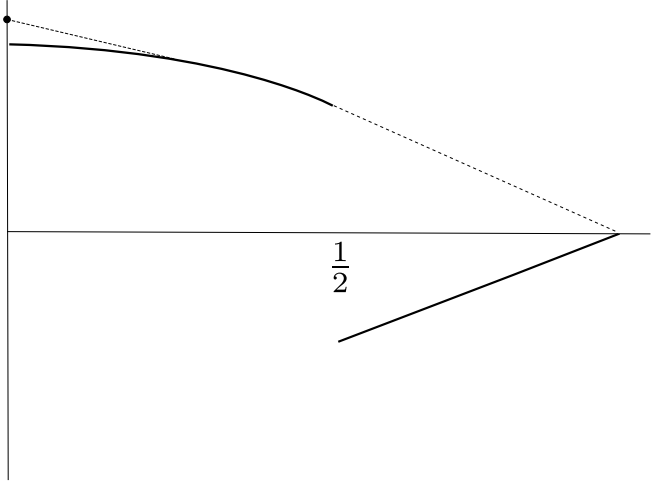
We see that we do not recover V but instead there is now a linear segment over an initial interval. This gives us a hint that the optimal value function will have a similar shape.

25

And indeed we can reduce the functional fixed-point problem in the HJB equation to a parametric equation with a single unknown p^{**}. For suppose we pick a threshold p^{**} where the mechanism switches from randomizing to silence. Knowing that the value function will be smooth at that point tells us what the slope of the initial linear segment must be. It must be the slope of the email value function at p^{**}, namely $V'(p^{**})$. This therefore tells us a candidate value for $V(0)$:

$$V(p^{**}) + p^{**}V'(p^{**})$$

Now when we take the resulting candidate value function and feed it into the right-hand side of the HJB equation we obtain a new value for $V(0)$:

$$u(0) + V'(p^{**})\lambda$$

We can solve the model by picking the p^{**} that equates these and thus produces a fixed point.

$$V(p^{**}) + p^{**}V'(p^{**}) = 5/4 + V'(p^{**})\lambda.$$

A distinguishing feature of this example is that the optimal policy uses both false negatives and false positives. In the initial phase, the signals that move from 0 to p^{**} are false positives: conditional on the agent receiving this signal he changes his behavior but there is a probability $1 - p^{**}$ that this signal was received even if no email arrived. On the other hand, once the agent reaches $\mu = 1/2$ the optimal policy reverts to the false negatives from the original email problem.

7 Appendix For The Multi-Agent Section

7.1 Construction Of The Mechanism

The two-agent mechanism is constructed by the following logic. Let's consider mechanisms in which agent 1 has a beep with a delay t^*, and agent 2 withdraws immediately if the state change happens after t^*. Within mechanisms with these properties, if the state changes after date t^*, i.e. $\phi > t^*$ then agent 1 will be the last to withdraw and the date of the bank run is $\phi + t^*$. It remains only to specify the date at which 2 withdraws if the state change happens before t^*. Let us define the function $\psi(\phi)$ to give the date at which agent 2 withdraws as a function of the state change date ϕ. The only requirement of the random variable ψ is that its distribution be exponential with parameter λ, to satisfy incentive compatibility for agent 2.

Then, conditional on $\phi \leq t^*$ agent 2 switches to run at $\tilde{r}_2 = \psi(\phi)$, agent 1 switches to run at $\tilde{r}_1 = \phi + t^*$ and a bank run occurs at the following time

$$\max\{\tilde{r}_1, \tilde{r}_2\} = \max\{\phi + t^*, \psi(\phi)\}$$

with conditional expected value (conditional on the state change occurring prior to t^*)

$$\int_{s=0}^{\infty} 1 - G(s)ds \tag{7}$$

where

$$G(s) = \text{Prob}\left(\{\tilde{r}_1 \leq s \text{ and } \tilde{r}_2 \leq s\} \mid \phi < t^*\right)$$

is the conditional CDF of $\max\{\tilde{r}_1, \tilde{r}_2\}$.

From the theory of copulae in probability theory, and in particular Sklar's Theorem and the Fréchet-Hoeffding Theorem (see Nelsen (1999)) we can minimize G pointwise,

and therefore maximize the expected delay length in Equation 7 by adopting as the the joint distribution for ϕ and ψ (whose marginal distributions we are given) the *Fréchet-Hoeffding lower bound.* In this case the lower bound is achieved by setting

$$\psi(\phi) = -\frac{1}{\lambda} \log \left(\frac{F(\phi)}{p^*} \right) + t^* \tag{8}$$

for all $\phi \leq t^*$ where F is the (marginal) CDF of ϕ, (i.e. the CDF of the exponential distribution with parameter λ.) To verify, conditional on $\phi \leq t^*$ the distribution of $\psi - t^*$ is

$$
\begin{aligned}
\text{Prob}\,(\psi - t^* \leq t \mid \phi \leq t^*) &= \text{Prob}\left(-\frac{1}{\lambda} \log \left(\frac{F(\phi)}{p^*} \right) \leq t \mid \phi \leq t^* \right) \\
&= \text{Prob}\left(F(\phi) \geq p^* e^{-\lambda t} \mid \phi \leq t^* \right) \\
&= \frac{\text{Prob}\left(\{F(\phi) \geq p^* e^{-\lambda t}\} \cap \{\phi \leq t^*\} \right)}{\text{Prob}\,(\phi \leq t^*)} \\
&= \frac{\text{Prob}\left(\{F(\phi) \geq p^* e^{-\lambda t}\} \cap \{F(\phi) \leq p^*\} \right)}{p^*} \\
&= \frac{\text{Prob}\left(F^{-1}(p^* e^{-\lambda t}) \leq \phi \leq F^{-1}(p^*) \right)}{p^*} \\
&= \frac{F\left(F^{-1}(p^*) \right) - F\left(F^{-1}(p^* e^{-\lambda t}) \right)}{p^*} \\
&= \frac{p^* \left(1 - e^{-\lambda t} \right)}{p^*} \\
&= 1 - e^{-\lambda t},
\end{aligned}
$$

i.e. exponential with parameter λ. Moreover, because ψ is a strictly decreasing function of ϕ, the corresponding random variables are perfectly negatively correlated. And since ψ gives the run time of agent 2 and $\phi + t^*$ gives the run time of agent 1, this makes the two run times maximally negatively correlated.

Under the proposed mechanism, the run times for the two agents are deterministic

functions of the state change date, ϕ. In particular,

$$\tilde{r}_1(\phi) = \phi + t^*$$

$$\tilde{r}_2(\phi) = \begin{cases} \phi & \text{if } \phi \geq t^* \\ -\frac{1}{\lambda} \log \left(\frac{F(\phi)}{p^*} \right) + t^* & \text{if } \phi < t^* \end{cases}$$

The graphs of these functions are depicted in Figure 4. Highlighted in blue is the graph of $\max \tilde{r}_i$, the date at which a bank run occurs.
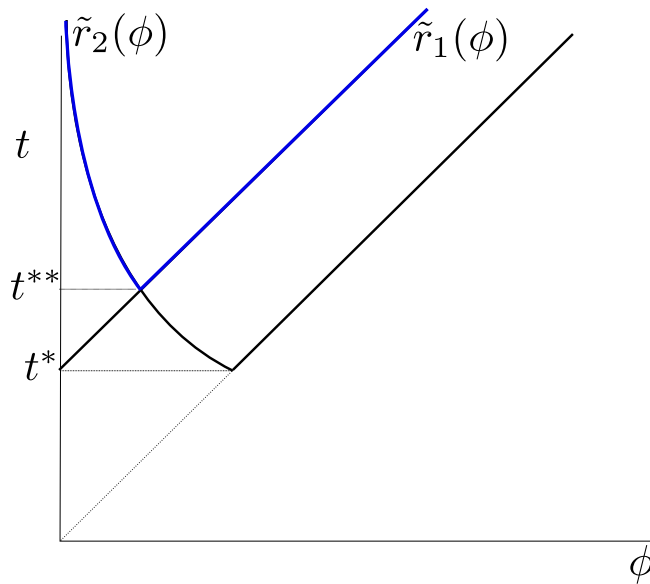

Figure 4: Beep times $\tilde{r}_i$ as a function of the state change time ϕ for the optimal multi-agent mechanism.

First, notice that a bank run cannot occur prior to the date t^{**}. It is the unique intersection of the two curves above, $t^{**} = \tilde{r}_1(\phi^{**}) = \tilde{r}_2(\phi^{**})$. Since this occurs prior to t^* it

is on the auxiliary clock segment of agent 2's mechanism and so

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

$$F(t^{**} - t^*) = p^* e^{-\lambda t^{**}}$$

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

$$t^{**} = t^* + \frac{1}{\lambda} \log(1 + p^*) \tag{9}$$

References

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian persuasion." *American Economic Review*, 101(6).

Nelsen, Roger B. 1999. *An introduction to copulas.* Springer.

Renault, Jérôme, Eilon Solan, and Nicolas Vieille. 2014. "Optimal Dynamic Information Provision." *arXiv preprint arXiv:1407.5649*.