The appendix includes proofs that support the results in the paper and descriptions of numerical optimization routines, when applicable. It proceeds in four sections. Section 1 develops the model. Section 2 introduces properties of search over rugged landscapes characterized by a path of Brownian motion. I then consider optimal implemented strategies for the incumbent and replacement in Section 3. Finally, Section 4 derives optimal experimental strategies for the incumbent in each of the five cases described in the paper.

1 Setup

An organization consists of a principal (the organization), p, and two agents (managers), a_1 and a_2, facing an uncertain task. The outcome of the task y is a stochastic function of a strategy x. The function that maps strategies to outcomes is defined by the sum of a deterministic function $-(1-x)^2$ and an unknown error term unique to each agent $\epsilon_i(x)$ for $i \in \{1, 2\}$. The error term $\epsilon_i(x)$ is modeled as a path of scaled standard Brownian motion $\sigma B_i(x)$ where $Cov(B_j(x), B_k(x)) = 0$ for $j \neq k$.

The strategy x can take any real value, i.e., $x \in \mathbb{R}$. However, the Brownian path is "tied down" at a known value of 0 for $x = 0$, i.e., $B_i(0) = 0$, $\forall i$.

Taken together,

$$y = -(1-x)^2 + \sigma B_i(x) \qquad \text{where } B_i(0) = 0$$

The principal assigns a_1 to the task. a_1 has the opportunity to try an experimental strategy $\tilde{x}$ and observe its outcome $\tilde{y}$. The principal observes $\tilde{y}$, makes inference about $B_1(x)$, and decides whether to keep a_1 or replace a_1 with a_2. Whichever agent is chosen then selects an implemented strategy x and observes the implemented outcome y.

The principal is risk-averse with a utility function u_p defined by mean-variance preferences over implemented outcomes. Specifically, $E[u_p] = E(y) - Var(y)$. Agents are tenure-motivated. They prefer to be kept rather than be replaced. However, all else equal, they prefer a strategy that produces a desirable outcome for the principal to one that produces an undesirable outcome.

The crux of the model is that these career concerns create an agency problem for the principal. a_1's experimental strategy $\tilde{x}$ will be biased in order to minimize the probability that a_1 is replaced with a_2, where that probability is represented by π, instead of maximizing $E[u_p]$.

I define tie-breaking rules, for ease of formal exposition. When choosing between two strategies x_i and x_j where $x_i < x_j$, that produce the same expected utility, p prefers the strategy that produces a higher expected outcome, i.e., higher $E(y)$. If x_i and x_j produce outcomes with the same expected outcome, then p prefers

x_i. I also assume that when selecting between the incumbent a_1 and a replacement a_2 that produce the same expected utility p prefers a_2 (although all results hold straightforwardly if p replaces a_1 with probability greater than or equal to $\frac{1}{2}$).

1.1 Information and Commitment

Prior to a_1's experiment, all actors have the same information about the task. What is unknown are the realizations of $\epsilon_i(x)$ for $i \in \{1, 2\}$.

I analyze cases in which p can commit to a *replacement rule* and in which p cannot commit to one. If p cannot commit, then it replaces a_1 whenever a_2 is at least as desirable as a_1. If p can commit, then it defines a set of experimental strategy-outcome pairs $(\tilde{x}, \tilde{y})$ which will lead p to replace a_1 with a_2. One type of commitment involves never replacing a_1, i.e., providing a tenure guarantee.

I also analyze cases in which p can observe $\tilde{x}$ and in which p cannot observe $\tilde{x}$. After a_1's experiment, a_1 learns $(\tilde{x}, \tilde{y})$. If the experiment is *observable*, then p observes $(\tilde{x}, \tilde{y})$ as well. If it is *tacit*, then p only observes $\tilde{y}$. If $\tilde{x}$ is tacit, then p bases its replacement decision solely on $\tilde{y}$ and prior knowledge about the task.

There are five cases that I analyze in the model, which are summarized in Table 1 in the paper:

Case 1: Observable Strategies with Commitment

Case 2: Observable Strategies with No Commitment

Case 3: Tacit Strategies with No Commitment

Case 4: Tacit Strategies with Outcome-based Commitment

Case 5: Guaranteed Tenure

1.2 Equilibrium Concept

I find subgame perfect Bayesian Nash equilibria for four of the conditions: $\tilde{x}$ observable and p can commit (Case 1), $\tilde{x}$ observable and p cannot commit (Case 2), and $\tilde{x}$ tacit and p can commit (Cases 4 and 5). In each of these conditions, information is perfect. That is, all decision-makers have full knowledge of all prior decisions in the game.

When p cannot commit and $\tilde{x}$ is tacit (Case 3), p is unaware of $\tilde{x}$, so I use the perfect Bayesian equilibrium concept, which requires that all actors be sequentially rational conditional on information and beliefs, even for decisions never reached in equilibrium.

I use backwards induction to derive the equilibria in the model. First, I find the optimal implemented strategies for a_1 and a_2, where the optimal implemented strategy for a_1 is conditional on $(\tilde{x}, \tilde{y})$. Then, I define the experimental strategies and outcomes for which p replaces a_1. Then, I find the equilibrium $\tilde{x}$, i.e., $\tilde{x}$ that minimizes the probability of replacement, π.

2 Properties of Search

2.1 Search on Known Landscapes

Before analyzing the principal-agent dynamics, it is valuable to build some intuition about the organization's search problem. First, note that if there were no uncertainty, i.e., $\epsilon_i(x) = 0$, $\forall i, x$, utility is maximized at $x^* = 1$.

$$
\begin{aligned}
E[u_p(y)] &= -(1-x)^2 \\
\frac{d}{dx}E[u_p(y)] &= 2(1-x) \\
\frac{d^2}{d^2x}E[u_p(y)] &= -2 && \text{Therefore, function is concave} \\
2(1-x^*) &= 0 && \text{Function maximized at } \frac{d}{dx}E[u_p(y)] = 0 \\
x^* &= 1
\end{aligned}
$$

2.2 Adding Rugged Uncertainty

Next, I define some important characteristics of standard Brownian motion. Standard Brownian motion, generally denoted $B(t)$, is a continuous Gaussian stochastic process with stationary and independent increments such that $t \geq 0$, $B(0) = 0$, $E[B(t)] = 0$, and $Cov(B(t_1), B(t_2)) = E[B(t_1)B(t_2)] = min(t_1, t_2)$. Thus, $Var(B(t)) = E[B(t)^2] = t$. The probability distribution of $\sigma B(t)$ is Normal with mean 0 and variance $t\sigma^2$.

Unlike in most contexts where Brownian drift is a function of time, in this model the Brownian motion $B(x)$ is also permitted to take negative values, as well, i.e., $x \in \mathbb{R}$. Allow $B(x)$ where $x < 0$ to be the reflection of a path of standard Brownian motion over the y-axis, i.e., $B(x) = B'(-x)$ for $x < 0$, and $Cov(B(x), B'(x)) = 0$.

Once $(\tilde{x}, \tilde{y})$ is known, the distribution of $B(x)$ depends on the realization of $\epsilon(\tilde{x})$. For example, assume that two strategy-outcome pairs (x_l, y_l) and (x_r, y_r) are known, where $x_l < x_r$. For all $x < x_l$, $B(x)$ is distributed Normal with mean $\epsilon(x_l)$ and variance $|x_l - x|\sigma^2$. For all $x > x_r$, $B(x)$ is distributed Normal with mean $\epsilon(x_r)$ and variance $|x_r - x|\sigma^2$. For all $x_l < x < x_r$, $B(x)$ is distributed Normal with mean $\epsilon(x_l) + (x - x_l)\frac{\epsilon(x_r)-\epsilon(x_l)}{x_r - x_l}$ and variance $\frac{(x_r-x)(x-x_l)}{x_r-x_l}\sigma^2$. The region of the Brownian path bounded by two known points is called a *Brownian bridge*.

3 Optimal Implemented Strategies (x^*)

3.1 Optimal Implemented Strategy for Replacement (a_2)

If a_1 is replaced, then a_2 seeks to maximize $E[u_p]$. Let x^* represent the optimal implemented strategy for a_2. As $Cov(B_1(x), B_2(x)) = 0$, knowledge of $B_1(\tilde{x})$ provides a_2 has no additional information about $B_2(x)$.

x^* is greater than or equal to 0, because for $x < 0$, $\frac{d}{dx}E[u_p] > 0$, as $E(y)$ is increasing in x and $Var(y)$ is decreasing x.

For $x \geq 0$:

$$\begin{aligned} E[u_p] &= -(1-x)^2 - x\sigma^2 \\ \frac{d}{dx}E[u_p] &= 2(1-x) - \sigma^2 \\ \frac{d^2}{d^2x}E[u_p] &= -2 && \text{Function is concave} \end{aligned}$$

For $\sigma^2 < 2$, the right-hand side $\frac{d}{dx}E[u_p]$ is positive at $x = 0$. Therefore, $x^* > 0$.

$$\begin{aligned} 2(1-x^*) - \sigma^2 &= 0 && \text{Function maximized at } \frac{d}{dx}E[u_p] = 0 \\ x^* &= 1 - \frac{\sigma^2}{2} \\ E[u_p] &= -\sigma^2(1 - \frac{\sigma^2}{4}) && (1) \end{aligned}$$

For $\sigma^2 \geq 2$, the right-hand side $\frac{d}{dx}E[u_p]$ is negative at $x = 0$, so $x^* = 0$, and $E[u_p] = -1$.

3.2 Optimal Implemented Strategy for Incumbent (a_1), Conditional on Experimental Strategy and Outcome, $(\tilde{x}, \tilde{y})$

If a_1 is retained, I group a_1's implemented strategy, x, into one of four categories:

Retain $\tilde{x}$: If $x = \tilde{x}$, then the experimental strategy is *retained*.

Backslide: If $x = 0$, then the implemented strategy *backslides* to the known strategy-outcome pair of $(0, -1)$.

Explore: If $x > \tilde{x}$ for $\tilde{x} > 0$ or $x < \tilde{x}$ for $\tilde{x} < 0$, a_1 *explores* outside of the support of known strategy-outcome pairs.

Triangulate: If $0 < x < \tilde{x}$ or $\tilde{x} < x < 0$, then a_1 *triangulates* between known strategy-outcome pairs.

I begin by characterizing x^* conditional on $(\tilde{x}, \tilde{y})$ for $\tilde{x} \geq 0$ for various ranges of σ^2.

First, I establish some useful results for the four categories of implemented strategies for $\tilde{x} > 0$.

3.2.1 Backsliding, $x = 0$

If a_1 backslides, then $E[u_p] = -1$. Note that $\frac{dE[u_p]}{d\epsilon_1(\tilde{x})} = 0$, conditional on $\tilde{x}$, because the expected utility associated with backsliding is constant. Define the second-period strategy if a_1 backslides as $x^B = 0$.

3.2.2 Retaining $x = \tilde{x}$

If a_1 retains $\tilde{x}$, then $E[u_p] = \tilde{y} = -(1-\tilde{x})^2 + \epsilon_1(\tilde{x})$. Further, $\frac{dE[u_p]}{d\epsilon_1(\tilde{x})} = 1$, conditional on $\tilde{x}$. Define the second-period strategy if a_1 retains $\tilde{x}$ as $x^R = \tilde{x}$.

3.2.3 Exploring, $x > \tilde{x}$

If a_1 explores, then x maximizes $E[u_p]$ conditional on $(\tilde{x}, \tilde{y})$ and $x > \tilde{x}$. Let μ equal the slope of the Brownian bridge spanning the two known strategy-outcome pairs, i.e., $\frac{\epsilon_1(\tilde{x})}{\tilde{x}}$.

$$
\begin{aligned}
E[u_p] &= -(1-x)^2 + \mu\tilde{x} - (x-\tilde{x})\sigma^2 \\
\frac{d}{dx}E[u_p] &= 2(1-x) - \sigma^2 \\
\frac{d^2}{d^2x}E[u_p] &= -2 && \text{Function is concave}
\end{aligned}
$$

For $\sigma^2 < 2$ and $\tilde{x} \geq 1 - \frac{\sigma^2}{2}$ or for $\sigma^2 \geq 2$, $\frac{d}{dx}E[u_p]$ is less than or equal to zero, so exploration is strictly dominated by retaining $x = \tilde{x}$.

Let x^E represent the optimal x, conditional on exploring:

$$
\begin{aligned}
2(1-x^E) - \sigma^2 &= 0 && \text{Function maximized at } \frac{d}{dx}E[u_p] = 0 \\
x^E &= 1 - \frac{\sigma^2}{2} \\
E[u_p(y|x^E)] &= -\sigma^2(1 - \frac{\sigma^2}{4}) + \tilde{x}(\mu + \sigma^2) && (2) \\
\frac{dE[u_p]}{d\epsilon_1(\tilde{x})} &= 1 && \text{As } \mu\tilde{x} = \epsilon_1(\tilde{x})
\end{aligned}
$$

3.2.4 Triangulating, $x \in (0, \tilde{x})$

If a_1 triangulates between strategies with known outcomes, i.e., $x \in (0, \tilde{x})$:

$$
\begin{aligned}
E[u_p] &= -(1-x)^2 + \mu x - \frac{x(\tilde{x}-x)}{\tilde{x}}\sigma^2 \\
\frac{d}{dx}E[u_p] &= 2(1-x) + \mu - \left(1 - \frac{2x}{\tilde{x}}\right)\sigma^2 \\
\frac{d^2}{d^2x}E[u_p] &= -2 + \frac{2\sigma^2}{\tilde{x}} && \text{Function is concave for } \tilde{x} > \sigma^2
\end{aligned}
$$

For $\tilde{x} > \sigma^2$, the utility function is concave. Further, the right-hand side $\frac{d}{dx}E[u_p]$ is negative at $x = 0$ for $\mu \leq \sigma^2 - 2$. Therefore, for $\mu \leq \sigma^2 - 2$, backsliding dominates triangulating.

Also, the left hand side $\frac{d}{dx}E[u_p]$ is positive at $x = \tilde{x}$ for $\mu \geq 2(\tilde{x}-1)-\sigma^2$. Therefore, for $\mu \geq 2(\tilde{x}-1)-\sigma^2$, retaining $\tilde{x}$ dominates triangulating.

For $\tilde{x} > \sigma^2$ and $\mu \in (\sigma^2 - 2, 2(\tilde{x}-1)-\sigma^2)$, a_1 prefers to triangulate. Let x^T indicate the optimal x if a_1 triangulates.

$$
\begin{aligned}
2(1-x^T) + \mu - \left(1 - \frac{2x^T}{\tilde{x}}\right)\sigma^2 = 0 && \text{Function maximized at } \frac{d}{dx}E[u_p] = 0 \\
x^T = \frac{2+\mu-\sigma^2}{2(1-\frac{\sigma^2}{\tilde{x}})}
\end{aligned}
$$

Next, I derive properties of $E[u_p]$ conditional on $\tilde{x}$. First, I place bounds on $\frac{dE[u_p]}{d\epsilon_1(\tilde{x})}$. Then, I derive the expected utility associated with $\tilde{x}$ using a Taylor series expansion.

I show that $\frac{dE[u_p]}{d\epsilon_1(\tilde{x})} \in [0, 1]$:

$$
\begin{aligned}
x^T &= \frac{2 + \frac{\epsilon_1(\tilde{x})}{\tilde{x}} - \sigma^2}{2(1-\frac{\sigma^2}{\tilde{x}})} \\
E[u_p] &= -(1-x^T)^2 + \frac{\epsilon_1(\tilde{x})x^T}{\tilde{x}} - \frac{x^T(\tilde{x}-x^T)}{\tilde{x}}\sigma^2 \\
\frac{dx^T}{d\epsilon_1(\tilde{x})} &= x' = \frac{1}{2(\tilde{x}-\sigma^2)} > 0 && \text{As } \tilde{x} > \sigma^2 \\
\frac{dE[u_p]}{d\epsilon_1(\tilde{x})} &= 2(1-x^T)x' + \frac{x^T}{\tilde{x}} + \frac{\epsilon_1(\tilde{x})x'}{\tilde{x}} - (1 - \frac{2x^T}{\tilde{x}})x'\sigma^2 \\
\frac{d^2E[u_p]}{d^2\epsilon_1(\tilde{x})} &= \frac{x'}{\tilde{x}} > 0 && \text{As } \tilde{x} > 0 \text{ and } x' > 0
\end{aligned}
$$

As:

- $E[u_p]$ is convex with respect to $\epsilon_1(\tilde{x})$;
- $\frac{dE[u_p]}{d\epsilon_1(\tilde{x})} = 0$ for $\epsilon_1(\tilde{x}) = \tilde{x}(\sigma^2 - 2)$; and

- $\frac{dE[u_p]}{d\epsilon_1(\tilde{x})} = 1$ for $\epsilon_1(\tilde{x}) = \tilde{x}(2(\tilde{x}-1)-\sigma^2)$

It follows that:

$$\frac{dE[u_p]}{d\epsilon_1(\tilde{x})} \in [0,1] \qquad\qquad \forall \epsilon_1(\tilde{x}) \in [\tilde{x}(\sigma^2-2), \tilde{x}(2(\tilde{x}-1)-\sigma^2)]$$

Next, I derive a simplified version of $E[u_p]$ conditional on triangulating using a Taylor series expansion. At $\epsilon_1(\tilde{x}) = \tilde{x}(\sigma^2-2)$, $E[u_p] = -1$ and $\frac{dE[u_p]}{d\epsilon_1(\tilde{x})} = 0$. I can therefore re-write $E[u_p]$ as:

$$E[u_p] = \frac{(\epsilon_1(\tilde{x}) + \tilde{x}(2-\sigma^2))^2}{4\tilde{x}(\tilde{x}-\sigma^2)} - 1 \tag{3}$$

3.3 Characterizing the Optimal Implemented Strategy, x^*

I next derive the optimal x, defined x^*, conditional on $(\tilde{x}, \tilde{y})$ and σ. I first examine x^* conditional on $\tilde{x} \geq 0$ and three regions of σ^2: low uncertainty, or $\sigma^2 < \frac{2}{3}$; medium uncertainty, or $\sigma^2 \in [\frac{2}{3}, 2)$, and high uncertainty, or $\sigma^2 \geq 2$. Then, I examine x^* conditional on $\tilde{x} < 0$.

3.3.1 Low uncertainty: $\sigma^2 < \frac{2}{3}$ and $\tilde{x} \geq 0$

For $\tilde{x} < 1 - \frac{\sigma^2}{2}$, retaining $\tilde{x}$ is strictly dominated by exploring [see section 3.2.3]. Therefore, a_1 chooses between exploring, triangulating, and backsliding.

Because the derivative of $E[u_p]$ conditional on x^T is increasing in $\epsilon_1(\tilde{x})$ and bounded by 0 and 1, there is at most one point of indifference between exploring and triangulating.

I define the regions in which a_1 prefers exploration to backsliding and vice versa by setting equation 2 equal to -1 (the expected utility from backsliding).

Let $x^{E,B}$ indicate the optimal x, conditional on exploring or backsliding.

$$x^{E,B} = \begin{cases} x^B = 0 & \text{if } \mu < \frac{\sigma^2(1-\frac{\sigma^2}{4})-1}{\tilde{x}} - \sigma^2 \\ x^E = 1 - \frac{\sigma^2}{2} & \text{if } \mu \geq \frac{\sigma^2(1-\frac{\sigma^2}{4})-1}{\tilde{x}} - \sigma^2 \end{cases} \tag{4}$$

I find the minimum positive value for which $\tilde{x}$ triangulates for some $\epsilon_1(\tilde{x})$ by setting the bound at which a_1 is indifferent between backsliding and triangulating [$\mu = \sigma^2 - 2$, as derived in section 3.2.4] equal to the bound at which a_1 is indifferent between backsliding and exploring (see equation 4) and solving for $\tilde{x}$. This value is $\frac{(\sigma^2-2)^2}{8(1-\sigma^2)}$.

For $\tilde{x} \geq 1 - \frac{\sigma^2}{2}$, retaining $\tilde{x}$ dominates exploring. Therefore, a_1 chooses between retaining $\tilde{x}$, triangulating, and backsliding. As $\tilde{x} > \sigma^2$ if $\tilde{x} \geq 1 - \frac{\sigma^2}{2}$ and $\sigma^2 < \frac{2}{3}$, there are values of $\epsilon_1(\tilde{x})$ for which triangulating dominates retaining $\tilde{x}$ and backsliding, as defined in section 3.2.4.

Using these results, I consider three cases: In Case 1, $\tilde{x} \leq \frac{(\sigma^2-2)^2}{8(1-\sigma^2)}$. In Case 2, $\tilde{x} \in (\frac{(\sigma^2-2)^2}{8(1-\sigma^2)}, 1-\frac{\sigma^2}{2})$. In Case 3, $\tilde{x} \geq 1-\frac{\sigma^2}{2}$.

Case 1 $\tilde{x} \leq \frac{(\sigma^2-2)^2}{8(1-\sigma^2)}$:

For these $\tilde{x}$, a_1 chooses between backsliding and exploring. Therefore,

$$x^* = x^{E,B}$$

Case 2 $\tilde{x} \in (\frac{(\sigma^2-2)^2}{8(1-\sigma^2)}, 1-\frac{\sigma^2}{2})$:

For these $\tilde{x}$, a_1 chooses between backsliding, triangulating, and exploring.

I leave the value of μ conditional on $\tilde{x}$ that defines indifference between triangulating and exploring, which I denote $\mu^{T=E}$, to be defined implicitly. However, note that it is strictly less than the boundary representing indifference between retaining $\tilde{x}$ and triangulating [$\mu = 2(\tilde{x}-1)-\sigma^2$, as defined in section 3.2.4], as exploring strictly dominates retaining $\tilde{x}$ for $\tilde{x} < 1-\frac{\sigma^2}{2}$. The μ that defines indifference between triangulating and backsliding is $\sigma^2 - 2$, as defined in section 3.2.4.

In summary,

$$x^* = \begin{cases} x^B = 0 & \text{if } \mu < \sigma^2 - 2 \\ x^T & \text{if } \mu \in [\sigma^2 - 2, \mu^{T=E}] \\ x^E = 1 - \frac{\sigma^2}{2} & \text{if } \mu > \mu^{T=E} \end{cases}$$

Case 3 $\tilde{x} \geq 1-\frac{\sigma^2}{2}$:

For these $\tilde{x}$, a_1 chooses between backsliding, triangulating, and retaining $\tilde{x}$. Therefore,

$$x^* = \begin{cases} x^B = 0 & \text{if } \mu < \sigma^2 - 2 \\ x^T & \text{if } \mu \in [\sigma^2 - 2, 2(\tilde{x}-1)-\sigma^2] \\ x^R = \tilde{x} & \text{if } \mu > 2(\tilde{x}-1)-\sigma^2 \end{cases} \quad (5)$$

3.3.2 Medium uncertainty: $\sigma^2 \in [\frac{2}{3}, 2)$ and $\tilde{x} \geq 0$

What differentiates the medium uncertainty case from the low uncertainty case is that in the medium uncertainty case a_1 never triangulates for $\tilde{x} < 1 - \frac{\sigma^2}{2}$.

As before, I consider 3 cases. In Case 1, $\tilde{x} < 1 - \frac{\sigma^2}{2}$. In Case 2, $\tilde{x} \in [1 - \frac{\sigma^2}{2}, \sigma^2]$. In Case 3, $\tilde{x} > \sigma^2$.

Case 1 $\tilde{x} < 1 - \frac{\sigma^2}{2}$:

$x^{E,B}$ strictly dominates triangulating for $\tilde{x} < 1 - \frac{\sigma^2}{2}$. Therefore,

$$x^* = x^{E,B}$$

Case 2 $\tilde{x} \in [1 - \frac{\sigma^2}{2}, \sigma^2]$:

For $\tilde{x} \in [1 - \frac{\sigma^2}{2}, \sigma^2]$, retaining $\tilde{x}$ dominates exploring. For $\tilde{y} > -1$ [or, equivalently, $\mu > \frac{(1-\tilde{x})^2-1}{\tilde{x}}$], retaining $\tilde{x}$ also dominates backsliding. Otherwise, backsliding dominates retaining $\tilde{x}$.

Let $x^{R,B}$ indicate the the optimal x conditional on retaining $\tilde{x}$ or backsliding.

$$x^{R,B} = \begin{cases} x^B = 0 & \text{if } \tilde{y} \leq -1 \\ x^R = \tilde{x} & \text{if } \tilde{y} > -1 \end{cases}$$

As the minimum value for which triangulating dominates retaining $\tilde{x}$ and backsliding is σ^2 [see section 3.2.4], for $\tilde{x} \in [1 - \frac{\sigma^2}{2}, \sigma^2]$

$$x^* = x^{R,B}$$

Case 3 $\tilde{x} > \sigma^2$:

For $\tilde{x} > \sigma^2$, a_1 chooses between backsliding, triangulating, and retaining $\tilde{x}$. Thus, x^* is defined as in equation 5.

3.3.3 High uncertainty: $\sigma^2 \geq 2$ and $\tilde{x} \geq 0$

For $\sigma^2 \geq 2$, retaining $\tilde{x}$ strictly dominates exploring for all x. Therefore, I consider just two cases. In Case 1, $\tilde{x} \leq \sigma^2$. In Case 2, $\tilde{x} > \sigma^2$.

Case 1 $\tilde{x} \leq \sigma^2$:

a_1 selects between retaining $\tilde{x}$ and backsliding. Therefore,

$$x^* = x^{R,B}$$

Case 2 $\tilde{x} > \sigma^2$:

a_1 chooses between backsliding, triangulating, and retaining $\tilde{x}$. Thus, x^* is defined as in equation 5.

3.3.4 High uncertainty: $\sigma^2 \geq 2$ and $\tilde{x} < 0$

Next, I briefly characterize x^* for $\tilde{x} < 0$. Note that for $\tilde{x} < 0$, a_1 gains no additional information about $B_1(x)$ for $x > 0$. Thus, none of the equilibria reported in the paper support negative $\tilde{x}$. As such, I present the results with less explication than I do for $\tilde{x} > 0$.

I begin by considering $\sigma^2 \geq 2$. Then, I consider $\sigma^2 < 2$.

For $\sigma^2 \geq 2$, a_2 chooses between retaining $\tilde{x}$ if $\tilde{y}$ is sufficiently high, triangulating between $\tilde{x}$ and 0, and backsliding to $x = 0$.

$$E[u_p] = -(1-x)^2 + \mu x + \frac{x(\tilde{x}-x)}{\tilde{x}}\sigma^2$$

$$\frac{d}{dx}E[u_p] = 2(1-x) + \mu + \left(1 - \frac{2x}{\tilde{x}}\right)\sigma^2$$

$$\frac{d^2}{d^2x}E[u_p] = -2 - \frac{2\sigma^2}{\tilde{x}} \qquad \text{Function is concave for } \tilde{x} < -\sigma^2$$

I consider four cases. In Case 1, $\tilde{x} \leq -\sigma^2$. In Case 2, $\tilde{x} > -\sigma^2$ and $\mu \geq -2 - \sigma^2$. In Case 3, $\tilde{x} > -\sigma^2$ and $\mu \in (\sigma^2 - 2(1-\tilde{x}), -2 - \sigma^2)$. In Case 4, $\tilde{x} > -\sigma^2$ and $\mu \leq \sigma^2 - 2(1-\tilde{x})$.

Case 1 $\tilde{x} \geq -\sigma^2$:

a_1 backslides to $x = 0$ if $\tilde{y} < -1$ and retains $\tilde{x}$ otherwise.

Case 2: $\tilde{x} < -\sigma^2$ and $\mu \geq -2 - \sigma^2$:

The left-hand side $\frac{d}{dx}E[u_p] > 0$ for $x = 0$, so a_1 backslides to $x = 0$.

Case 3: $\tilde{x} < -\sigma^2$ and $\mu \in (\sigma^2 - 2(1-\tilde{x}), -2 - \sigma^2)$

The left-hand side $\frac{d}{dx}E[u_p] < 0$ for $x = 0$ and $\frac{d}{dx}E[u_p] > 0$ for $x = \tilde{x}$, so a_1 triangulates between $\tilde{x}$ and 0. Specifically,

$$2(1-x) + \mu + \left(1 - \frac{2x}{\tilde{x}}\right)\sigma^2 = 0 \qquad \text{Function maximized at } \frac{d}{dx}E[u_p] = 0$$

$$x^* = \frac{2 + \mu + \sigma^2}{2(1 + \frac{\sigma^2}{\tilde{x}})}$$

Case 4: $\tilde{x} < -\sigma^2$ and $\mu \leq -2 - \sigma^2$

$\frac{d}{dx}E[u_p] < 0$ for $x = \tilde{x}$, so a_2 retains $x = \tilde{x}$.

3.3.5 Low/Medium uncertainty: $\sigma^2 < 2$ and $\tilde{x} < 0$

For $\sigma^2 < 2$, the same applies, except that a_1 will explore rather than backslide to $x = 0$. The optimal x conditional on exploring is $x^E = 1 - \frac{\sigma^2}{2}$ [see section 3.2.3] which generates $E[u_p] = -\sigma^2(1 - \frac{\sigma^2}{4})$.

I first find the maximum value for which there exists $\epsilon_1(\tilde{x})$ such that triangulating dominates either exploring or retaining $\tilde{x}$ by setting the bound of $\tilde{y}$ at which a_1 retains $\tilde{x}$ equal to $-\sigma^2(1 - \frac{\sigma^2}{4})$ and solving for $\tilde{x}$. This bound is $-\frac{\sigma^2}{2} - \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}}$.

I analyze five cases, depending on whether $\tilde{x}$ exceeds this bound and the level of $\tilde{y}$.

Case 1 $\tilde{x} \geq -\frac{\sigma^2}{2} - \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}}$ and $\tilde{y} \geq -\sigma^2(1 - \frac{\sigma^2}{4})$:

$x^* = x^R = \tilde{x}$.

Case 2 $\tilde{x} \geq -\frac{\sigma^2}{2} - \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}}$ and $\tilde{y} < -\sigma^2(1 - \frac{\sigma^2}{4})$

$x^* = x^E = 1 - \frac{\sigma^2}{2}$.

Case 3 $\tilde{x} < -\frac{\sigma^2}{2} - \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}}$ and $\tilde{y} > \tilde{x}(\sigma^2 - 2(1-\tilde{x})) - (1-\tilde{x})^2$

$x^* = x^R = \tilde{x}$.

Case 4 $\tilde{x} < -\frac{\sigma^2}{2} - \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}}$ and $\tilde{x}(\sigma^2 - 2(1-\tilde{x})) - (1-x)^2 \geq \tilde{y} \geq \tilde{y}^{E=T}$, where $\tilde{y}^{E=T}$ implicitly defines the $\tilde{y}$ such that $E[u_p]$ if a_1 triangulates equals $E[u_p]$ if a_1 explores:

$x^* = x^T$.

Note that at $\tilde{y}^{E=T}$, $\epsilon_1\tilde{x} \geq \tilde{x}(-2 - \sigma^2) > 0$ as exploring dominates backsliding.

Case 5 $\tilde{x} < -\frac{\sigma^2}{2} - \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}}$ and $\tilde{y} < \tilde{y}^{E=T}$:

a_1 explores and $x^* = x^E = 1 - \frac{\sigma^2}{2}$.

4 Optimal Experimental Strategies ($\tilde{x}^*$)

4.1 Case 1: Experimental Strategy with Observable Strategy and Commitment. Or, Finding the Optimal $\tilde{x}$ for p, $\tilde{x}^{**}$

I begin by demonstrating the optimal $\tilde{x}$ for p, which I denote $\tilde{x}^{**}$. This serves as a baseline against which I compare deviations when career concerns are introduced.

Note that this outcome can be achieved if $\tilde{x}$ is observable and p can commit to a replacement rule in the second period using the following rubric: p tells a_1 that a_1 will be replaced if $\tilde{x} \neq \tilde{x}^{**}$. p then retains a_1

when the expected utility associated with retaining a_1 exceeds the expected utility associated with replacing a_1 with a_2.

I find $\tilde{x}^{**}$ using backwards induction. That is, conditional on x^*, I optimize $\tilde{x}$ in order to maximize $E[u_p]$.

4.1.1 Equilibrium Replacement Rule for the Principal

The principal p retains a_1 when the the expected utility from x^* is higher if a_1 is in charge of implementation than if a_2 is in charge. Otherwise, p replaces a_1. Next, I define the regions of $(\tilde{x}, \tilde{y})$ for which this is the case. In the first case, $\tilde{x} < 0$. In the second case, $\tilde{x} \geq 0$ and $\sigma^2 \geq 2$. In the third case, $\tilde{x} \geq 0$ and $\sigma^2 < 2$.

Negative $\tilde{x}$ For $\tilde{x} < 0$, the values of the $\epsilon_1(\tilde{x})$ for which a_1 explores [backslides] for $\sigma^2 < 2$ [$\sigma^2 \geq 2$] represent strategy-outcome pairs for which p is indifferent between a_1 and a_2 and, as such, for which p replaces a_1.

High Uncertainty For $\tilde{x} \geq 0$ and $\sigma^2 \geq 2$, the values of $\epsilon_1(\tilde{x})$ for which a_1 backslides also make p indifferent between a_1 and a_2. Therefore, p replaces a_1 for these strategy-outcome pairs, as well. For other strategy-outcome pairs, a_1 is not replaced.

Low or Medium Uncertainty Next, I examine $\tilde{x} \geq 0$ and $\sigma^2 < 2$.

For $\tilde{x} \leq 1 - \frac{\sigma^2}{2}$: p replaces a_1 if $\mu \leq -\sigma^2$. This boundary is derived by setting the expected utility from a_2 [equation 1] to the expected utility from a_1 exploring [equation 2] and noting that $x^* = x^E$ for $\mu = -\sigma^2$.

For $\tilde{x} \in (1 - \frac{\sigma^2}{2}, \frac{\sigma^2}{2} + \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}})$: p replaces a_1 if $\tilde{y} \leq E[u_p(a_2)] = -\sigma^2(1 - \frac{\sigma^2}{4})$ (or, equivalently, $\mu \leq \frac{(1-\tilde{x})^2 - \sigma^2(1-\frac{\sigma^2}{4})}{\tilde{x}}$). This boundary reflects the fact that $x^* = x^R$ for $\tilde{y} = -\sigma^2(1 - \frac{\sigma^2}{4})$.

The right-most limit of this region (i.e., $\frac{\sigma^2}{2} + \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}}$), is defined by setting the expected utility from a_2 to the boundary between $x^* = x^R$ and $x^* = x^T$ (see equation 5) and solving for $\tilde{x}$. Note that this limit is greater than or equal to 1 for all σ^2.

Finally, for $\tilde{x} \geq \frac{\sigma^2}{2} + \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}}$, p replaces a_1 if:

$$\epsilon_1(\tilde{x}) \leq (2 - \sigma^2)(\sqrt{\tilde{x}(\tilde{x} - \sigma^2)} - \tilde{x})$$

Which is derived by setting the expected utility from a_2 to the expected utility from triangulating (see equation 3).

4.1.2 Finding the Optimal Experimental Strategy

I find the optimal experimental strategy (from the perspective of p), denoted $\tilde{x}^{**}$, using genetic optimization, implemented using the `rgenoud` package in `R`. The function optimized is a weighted average of the probabilities of $\tilde{y}$ conditional on $\tilde{x}$ multiplied by the respective expected utilities associated with $\tilde{y}$.

Inspection of the graphs that relate $\tilde{y}$ to $E[u_p]$ indicate that the problem is very likely to be single-peaked for $\tilde{x} > 0$ and single-peaked for $\tilde{x} < 0$, which indicates that genetic optimization is very likely to find a global optimum.

4.2 Case 3: Experimental Strategy with Tacit Strategy and No Commitment

Next, I examine how $\tilde{x}$ changes in equilibrium if p cannot commit to a replacement rule. In order to find the equilibrium, I identify an experimental strategy and replacement rule for which (1) the experimental strategy minimizes the probability of replacement conditional on the replacement region and (2) the replacement region the optimal response of p to the experimental strategy.

Let $\underline{\tilde{y}}$ indicate a level of $\tilde{y}$ for which p replaces a_1 for all $\tilde{y}$ less than or equal to it. Let $\underline{\tilde{y}}^*$ indicate the equilibrium level of $\underline{\tilde{y}}$. Let $\tilde{x}^*$ indicate the $\tilde{x}$ that minimizes π conditional on $\underline{\tilde{y}}$.

Note that p never replaces a_1 for:

- $\sigma^2 < 2$ and $y + (1 - \tilde{x})^2 = \epsilon_1(\tilde{x}) \geq 0$
- $\sigma^2 \geq 2$ and $\tilde{x} \in (0, 2)$

In these regions, expected utility for x^* always exceeds the expected utility from a replacement. In equilibrium, therefore:

- $\underline{\tilde{y}}^* + (1 - \tilde{x}^*)^2 < 0$ for $\sigma^2 < 2$
- $\underline{\tilde{y}}^* + (1 - \tilde{x}^*)^2 < 0$ for $\sigma^2 \geq 2$ if $\tilde{x}^* \in (0, 2)$

First, I show that $\underline{\tilde{y}}^* \in [-1, 0]$. If $\underline{\tilde{y}}^*$ were less than -1, then the $\tilde{x}$ that minimizes π would be 0. This is not supported in equilibrium, because p replaces a_1 if $\tilde{x} = 0$. If $\underline{\tilde{y}}^*$ were greater than 0, then p would replace a_1 when a_1 strictly dominates a_2, which is also not supported in equilibrium.

Next, I show that the optimal experimental strategy $\tilde{x}^*$ lies between 0 and 1 if $\underline{\tilde{y}}^* \in [-1, 0]$

Note first that $\tilde{x}^* > 0$ for $\underline{\tilde{y}}^* \geq -1$. For $\tilde{x} \leq 0$, $P(\tilde{y} \leq -1) \geq \frac{1}{2}$. Because $P(\tilde{y} \leq -1)$ approaches $\frac{1}{2}$ from below as $\tilde{x}$ approaches 0 from above, $\tilde{x} \leq 0$ is strictly dominated.

Next, I find the $\tilde{x}$ that minimizes the probability of replacement, π, conditional on $\underline{\tilde{y}}$ for $\tilde{x} > 0$.

$$Z = \frac{\underline{\tilde{y}} + (1-\tilde{x})^2}{\sigma\sqrt{\tilde{x}}} \qquad Z < 0 \text{ for } \sigma^2 < 2, \text{ or } \sigma^2 \geq 2 \text{ and } \tilde{x} \in (0,2)$$

$$\pi = \Phi(Z)$$

$$\frac{dZ}{d\tilde{x}} = Z' = \frac{-\underline{\tilde{y}} + 3\tilde{x}^2 - 2\tilde{x} - 1}{2\sigma\tilde{x}^{\frac{3}{2}}}$$

$$\frac{d\pi}{d\tilde{x}} = \phi(Z)Z' \tag{6}$$

$$\frac{d^2\pi}{d^2\tilde{x}} = -Z\phi(Z)(Z')^2 \qquad \pi \text{ convex for } Z \leq 0 \tag{7}$$

For $\tilde{x} > 2$ and $\sigma^2 \geq 2$, $Z' > 0$, so π is strictly increasing in this region. π is therefore quasi-convex with respect to $\tilde{x}$ for all σ^2.

As the right-hand side $\frac{d\pi}{d\tilde{x}} < 0$ at $\tilde{x} = 0$ and $\frac{d\pi}{d\tilde{x}} > 0$ at $\tilde{x} = 1$ for $\underline{\tilde{y}} \in [-1, 0]$, $\tilde{x}^* \in (0,1]$, $\forall \sigma^2$.

As π is convex with respect to $\tilde{x}$ for $\tilde{x} \in (0,1]$, I solve for $\tilde{x}^*$ by finding $\tilde{x}$ that satisfies $\frac{d\pi}{d\tilde{x}} = 0$:

$$Z(\tilde{x}^*) = \frac{\underline{\tilde{y}} + (1-\tilde{x}^*)^2}{\sigma\sqrt{\tilde{x}^*}}$$

$$\phi\left(Z(\tilde{x}^*)\right)\frac{dZ(\tilde{x}^*)}{d\tilde{x}} = 0 \qquad \text{Function maximized at } \frac{d\Phi(Z(\tilde{x}^*))}{d\tilde{x}} = 0$$

$$\frac{dZ(\tilde{x}^*)}{d\tilde{x}} = 0 \qquad \text{As } \phi(Z) > 0\ \forall Z$$

$$\frac{-\underline{\tilde{y}} + 3(\tilde{x}^*)^2 - 2\tilde{x}^* - 1}{2\sigma(\tilde{x}^*)^{\frac{3}{2}}} = 0$$

$$\tilde{x}^* = \frac{1}{3}\left(\sqrt{4 + 3\underline{\tilde{y}}} + 1\right) \tag{8}$$

As $\underline{\tilde{y}}^* \in [-1, 0]$, $\tilde{x}^* \in [\frac{2}{3}, 1]$ in equilibrium.

Next, I derive $\underline{\tilde{y}}^*$ conditional on $\tilde{x}^*$.

In equilibrium, if $\sigma^2 < 2$, p replaces a_1 if $\tilde{y} \leq -(1-\tilde{x}^*)^2 - \tilde{x}^*\sigma^2$ for $\tilde{x}^* \in (0, 1-\frac{\sigma^2}{2})$ and if $\tilde{y} \leq -\sigma^2(1-\frac{\sigma^2}{4})$ for $\tilde{x}^* \in [1-\frac{\sigma^2}{2}, 1]$. If $\sigma^2 \geq 2$, p replaces a_1 if $\tilde{y} \leq -1$.

First, I show that an equilibrium is supported for $\underline{\tilde{y}}^* = -\sigma^2(1-\frac{\sigma^2}{4})$ for $\sigma^2 < 2$ and $\underline{\tilde{y}}^* = -1$ for $\sigma^2 \geq 2$.

For $\sigma^2 < 2$ and $\underline{\tilde{y}}^* = -\sigma^2(1-\frac{\sigma^2}{4})$, $\tilde{x}^* = \frac{1}{3}\left(\sqrt{4 - 3\sigma^2(1-\frac{\sigma^2}{4})} + 1\right)$, which lies within $[1-\frac{\sigma^2}{2}, 1]$ and therefore is an equilibrium.

Next, I show that the above equilibrium is unique for $\sigma^2 < 2$. For $\tilde{x} \in (0, 1-\frac{\sigma^2}{2})$ and $\sigma^2 < 2$, $\underline{\tilde{y}}^*$ in equilibrium is $-(1-\tilde{x}^*)^2 - \tilde{x}^*\sigma^2$. Further, after rearranging Equation 8, $\underline{\tilde{y}}^*$ also must equal $3\tilde{x} - 2\tilde{x} - 1$. Setting these equations to each other, I find that $\tilde{x} = 1-\frac{\sigma^2}{4}$, which exceeds $1-\frac{\sigma^2}{2}$, generating a contradiction.

For $\sigma^2 \geq 2$ and $\underline{\tilde{y}} = -1$, $\tilde{x}^* = \frac{2}{3}$, which lies within $(0,1]$ and is also supported in equilibrium.

As a result, the unique equilibrium replacement region is:

$$\underset{\sim}{y}^* = \begin{cases} -\sigma^2(1-\frac{\sigma^2}{4}) & \text{For } \sigma^2 < 2 \\ -1 & \text{For } \sigma^2 \geq 2 \end{cases} \tag{9}$$

The unique best response of a_1, defined $\tilde{x}^*$, is:

$$\tilde{x}^* = \begin{cases} \frac{1}{3}\left(\sqrt{4-3\sigma^2(1-\frac{\sigma^2}{4})}+1\right) & \text{For } \sigma^2 < 2 \\ \frac{2}{3} & \text{For } \sigma^2 \geq 2 \end{cases} \tag{10}$$

4.3 Case 2: Experimental Strategy with Observable Strategy and No Commitment

I next demonstrate that the equilibrium replacement region and experimental strategy remains the same if $\tilde{x}$ is observable.

The major difference between the two cases is that in Case 2 p can vary $\underset{\sim}{y}$ in response to $\tilde{x}$. Therefore, I identify the experimental strategy that minimizes the probability of replacement, conditional on a replacement region that is a rational response to $\tilde{x}$.

Let $\underset{\sim}{y}^*(\tilde{x})$ indicate the level of $\tilde{y}$ conditional on $\tilde{x}$ for which p replaces a_1 for all $\tilde{y}$ less than or equal to it in equilibrium. Let $\tilde{x}^*$ indicate the $\tilde{x}$ that minimizes π conditional on $\underset{\sim}{y}^*(\tilde{x})$. Let $\underset{\sim}{y}^*$ represent the equilibrium $\underset{\sim}{y}^*(\tilde{x}^*)$.

First, I focus on the regions of $\tilde{x}$ for which $\underset{\sim}{y}^*(\tilde{x})$ does not vary with $\tilde{x}$, in which cases many of the results from Section 4.2 hold. I show that $\tilde{x}^*$ defined in Equation 10 minimizes the probability of replacement.

Case 1 $\tilde{x} \in [1-\frac{\sigma^2}{2}, 1]$ and $\sigma^2 < 2$; or $\tilde{x} \in [0,1]$ and $\sigma^2 \geq 2$:

Equations 9 and 10 characterize the unique equilibrium

Case 2 $\tilde{x} \in (1, \frac{\sigma^2}{2} + \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}}]$ and $\sigma^2 < 2$; or $\tilde{x} \in (1,2)$ and $\sigma^2 \geq 2$:

$\underset{\sim}{y}^*$ is the same as for $\tilde{x} \in [1-\frac{\sigma^2}{2}, 1]$. Therefore, it follows from the convexity of π in equation 7, that $\tilde{x} \in [1-\frac{\sigma^2}{2}, 1]$ strictly dominates $\tilde{x} \in (1, \frac{\sigma^2}{2} + \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}}]$.

Case 3 $\tilde{x} \in [2, \sigma^2]$ and $\sigma^2 \geq 2$:

$\pi > \frac{1}{2}$, as $\underset{\sim}{y}^* = -1 > E[\tilde{y}|\tilde{x}] = -(1-x)^2$. Therefore, $\tilde{x}$ in this region is dominated by $\tilde{x} = \frac{2}{3}$.

As such, for the regions of $\tilde{x}$ for which $\underset{\sim}{y}^*(\tilde{x})$ does not vary with $\tilde{x}$, the equilibrium defined in equations 9 and 10 still holds.

Next, I show that equilibrium $\tilde{x}$ in equation 10 leads to a lower π than any other $\tilde{x}$. As such, this equilibrium is unique.

Case 4 $\tilde{x} \leq 0$:

The probability of replacement for $\tilde{x} \leq 0$ is strictly greater than $\frac{1}{2}$, as any $\epsilon_1(\tilde{x}) \leq 0$ leads p to replace a_1. As $\tilde{x} \to 0$ from the right, the probability of replacement approaches $\frac{1}{2}$ from below [see section 4.1.1]. As such, $\tilde{x} \leq 0$ is strictly dominated.

Case 5 $\tilde{x} \in (0, 1 - \frac{\sigma^2}{2})$ and $\sigma^2 < 2$:

a_1 is replaced if $\mu \leq -\sigma^2$, which is derived in section 4.1.1. I next show that the probability of replacement is strictly decreasing in $\tilde{x}$ in this region, indicating that the $\tilde{x}$ that minimizes π is greater than $1 - \frac{\sigma^2}{2}$.

$$\pi = \Phi\left(\frac{-\sigma^2 \tilde{x}}{\sigma\sqrt{\tilde{x}}}\right) = \Phi\left(-\sigma\sqrt{\tilde{x}}\right)$$
$$\frac{d\pi}{d\tilde{x}} = \phi(-\sigma\sqrt{\tilde{x}})\left(-\frac{\sigma}{2\sqrt{\tilde{x}}}\right) < 0$$

Case 6 $\tilde{x} > \frac{\sigma^2}{2} + \frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}}$ and $\sigma^2 < 2$:

I demonstrate that π is increasing in this region, which means that $\tilde{x}^*$ cannot fall in it, either.

$$Z = \frac{(2-\sigma^2)(\sqrt{\tilde{x}(\tilde{x}-\sigma^2)} - \tilde{x}) + (1-\tilde{x})^2}{\sigma\sqrt{\tilde{x}}} \qquad \text{From equation 3}$$
$$\pi = \Phi(Z)$$
$$\frac{dZ}{d\tilde{x}} > 0 \qquad \text{For } \tilde{x} > 1$$
$$\frac{d\pi}{d\tilde{x}} = \phi(Z)\frac{dZ}{d\tilde{x}} > 0 \qquad \text{For } \tilde{x} > 1$$

As π is strictly increasing for $\tilde{x} > 1$, and $\frac{\sqrt{(\sigma^2-1)^2+1}}{\sqrt{2}} > 1$, $\tilde{x}$ in this region is not supported in equilibrium, either.

Case 7 $\tilde{x} > \sigma^2$ and $\sigma^2 \geq 2$:

p replaces a_1 if $\mu \leq \sigma^2 - 2$, as derived in section 4.1.1 and 4.1.2. This means that $\pi > \frac{1}{2}$, so $\tilde{x}$ in this region is strictly dominated.

As π is strictly decreasing for $\tilde{x} < 1 - \frac{\sigma^2}{2}$ and is either increasing or greater than $\frac{1}{2}$ for $\tilde{x} > 1$, the equilibrium defined in section 4.2 is also the unique equilibrium if $\tilde{x}$ is observable.

4.4 Cases 4 and 5: Experimental Strategy with Tacit Strategy and Commitment

If p can commit to a replacement rule, however, it can induce a more desirable experimental strategy. I define two classes of replacement rules to which p can commit.

The first is guaranteed tenure (Case 5). Under guaranteed tenure, $\pi = 0$ for all $\tilde{x}$. As such, a_1 selects $\tilde{x}$ to maximize $E[u_p]$. I find $\tilde{x}^*$ under guaranteed tenure the same optimization routine described in section 4.1.2. Note, however, that $\tilde{x}^*$ under guaranteed tenure deviates from $\tilde{x}^{**}$ because under guaranteed tenure p foregoes the option of replacing an undesirable a_1.

The second rule is outcome-based replacement (Case 4). Under outcome-based replacement, p commits to a value of $\tilde{y}$, which I denote as $\underline{\tilde{y}}^c$, such that p replaces a_1 if $\tilde{y} \leq \underline{\tilde{y}}^c$. By calibrating $\underline{\tilde{y}}^c$, p can induce a desirable $\tilde{x}$. However, for $\underline{\tilde{y}}^c > \underline{\tilde{y}}^*$, p is committing itself to replace a_1 when a_1 is more desirable than a_2 and for $\underline{\tilde{y}}^c < \underline{\tilde{y}}^*$ p is committing itself to retain a_1 when a_1 is less desirable than a_2. I find $\underline{\tilde{y}}^c$ using numerical optimization within a box constraint that $\underline{\tilde{y}}^c \in [-1, 6\sigma]$. $\underline{\tilde{y}}^c < -1$ is strictly dominated by $\underline{\tilde{y}}^c = -1$. The probability that $\tilde{y} \geq 6\sigma$ is minuscule for all $\tilde{x}$ and, as such, $E[u_p] \approx -\sigma^2(1 - \frac{\sigma^2}{2})$ for $\sigma^2 < 2$ and -1 for $\sigma^2 \geq 2$ at $\underline{\tilde{y}}^c = 6\sigma$.

For $\sigma^2 \geq 2$, guaranteed tenure is always the best commitment, because $E[u_p]$ conditional on $\tilde{x}$ is always weakly greater than -1, which is p's expected utility from a_2. For $\sigma^2 < 2$, I compare the two replacement rules numerically to find the dominant replacement regime.