Deception Under the Veil of Noise*

Jawwad Noor and Fernando Payró

April 17, 2026

Abstract

We study a dynamic predator–prey game in which a predator can conceal its movement under naturally occurring environmental noise. In the safe state, forest noise is i.i.d., whereas in the dangerous state the predator contributes additional noise as it approaches the prey. The prey updates her beliefs about danger from the realized noise sequence and chooses whether to remain vigilant. We characterize equilibrium patterns of noise generated in the forest and show that a marker for deception is a hot-hand effect, whereby streaks persist with increasing probability.

JEL Classification: D01, D9.

Keywords: deception, optimal stopping, endogenous information, belief biases.

1 Introduction

Self-interest motivates individuals to exploit and manipulate others to their advantage. Such exploitation is typically more effective if the target is unaware that it is happening, or at the very least uncertain of it. In this paper we study how rational deceivers might take advantage of a natural smoke screen to hide their intentions.

We analyze a dynamic predator-prey game where a prey decides whether or not to be vigilant, based on the noise in the forest.[1] If the state is safe then the noise from the forest is i.i.d., but if the state is dangerous, then some of that noise is optimally made by a predator as it steps toward the prey in order to attack it. In each period, nature randomly chooses whether or not to make noise, the predator decides whether to take a step conditional on what nature does. Nature's noise

***Noor is at the Department of Economics, Boston University, 270 Bay State Road, Boston MA 02215. Email: jnoor@bu.edu. Payró is at the Department of Economics and Economic History, Universitat Autonoma de Barcelona, Barcelona School of Economics and Center for the Study of Organizations and Decisions in Economics, Campus de la UAB, Edifici B, Bellaterra, 08193. Email: fernando.payro@uab.cat.** **The authors are grateful to Marc Rysman, to seminar audiences at National University of Singapore, Management University of Singapore, Hitotsubashi University and Johns Hopkins University, and to the audience at Barcelona Summer Forum 2025 and Econometric Society World Congress 2025 for valuable feedback. The usual disclaimer applies. Payr\'o gratefully acknowledges financial support from the Ministerio de Economía y Competitividad and Feder (PID2020-116771GB-I00 and PID2023-147183NB-I00), the Generalitat de Catalunya through the research grant 2021 SGR 00194 and the Spanish Agencia Estatal de Investigación (AEI), from the Severo Ochoa Programme for Centres of Excellence in R\&D (Barcelona School of Economics CEX2024-001476-S)", funded by MCIN/AEI/10.13039/501100011033.**

[1] We use the term "noise" to mean "sound", rather than some sort of random error term.

acts as a smoke screen: if nature has made noise, the predator can take a step under the veil of that noise, and if nature is silent, the predator's step will make noise. The predator's stepping strategy impacts the noise coming out of the environment, which the prey uses to update her beliefs about whether the forest is safe or dangerous, that is, whether there is a predator in the forest or not.

The setup contains features that fit several economic scenarios, and also speaks to psychology. For instance:

1. Selling: A seller (of a product, a political narrative, a religious ideology, etc.) may seek to sway a potential buyer by reporting the quality of the product along different dimensions (forest noise). While the seller may know the true quality (natural noise), it may manipulate the information (turn silence into noise) by highlighting only a subset of facts, using engaging or misleading language to describe its product and competing products, etc. The buyer (the prey) may be unsure if the seller (the predator) is honest (safe state) or willing to misrepresent the product (dangerous state). When called on to buy the product (an attack) the seller's success depends on the extent of positively framed information delivered (steps) but also on how vigilant the buyer is at that time.

Interpreting the seller and the buyer as being the same person leads to a description of self-deception. A person's reasoning may either be objective or biased by a latent desire or emotion – these are the states and the hidden desire is the predator. Reasoning generates arguments for or against an action (forest noise, with noise representing a favorable argument). The hidden desire for doing the action may lead arguments against the action to be reframed or misinterpreted so as not to disfavor the action (turning silence into noise). At some point the desire urges a call to action (the attack). This models the psychological process of motivated reasoning and gives rise to psychological phenomena such as confirmation bias (Rabin and Schrag 1999).

2. Hostile takeovers: In order to gain control of a target company (prey) without the approval or consent of the target's board of directors, and in order to prevent triggering a defensive response (vigilance), an acquiring firm (predator) has to buy the target company's stock in a concealed fashion until it makes a bid to take over the company (attack). For instance, in an effort to take over Hermes, LVMH bought 4.9% of Hermes at a time (e.g. directly on the open market and indirectly through financial entities), since French corporation law required holdings above 5% to be disclosed.[2] The target company observes market activity involving its stock (forest noise), but the acquiring company knows which trades are normal market activity (natural noise) and which ones are part of the take-over (steps).

4. Hiding from High Frequency Traders (HFTs): HFTs use sophisticated strategies, including algorithms that execute trades and cancel orders in fractions of a second and pattern recognition software, in order to take advantage of small arbitrage opportunities but also to sniff out and profit from larger orders placed by institutional investors (like Fidelity). In order to avoid being detected and exploited by HFT's, institutional investors often employ sophisticated execution algorithms to break down large orders into smaller slices that match the market. In this setup, both parties seek to deceive each other. This requires an extension of our model (which we leave to future research).

Our main interest in this paper is in determining what traces a patient deceiver might leave of its presence. Specifically, we are interested in how the pattern of noise coming out of the forest differs between the safe and dangerous states of the world. A natural conjecture is that perfectly patient deceivers may perfectly mimic nature, taking steps only under the cover of nature's noise, and attack when in optimal striking distance. This in fact corresponds with a so-called "creeping" hostile take-over, where the acquiring company gradually buys shares on the open market until it accumulates a controlling interest. We refer to equilibria with such a strategy as a *Moriarty*

[2]LVMH failed because Hermes became aware that very little of its stock was being traded in the market.

equilibrium. [3]

Our results show that, in accordance with intuition, when the prey is unsuspecting (that is, prior belief on danger is low), then a Moriarty equilibrium can exist. But, surprisingly, when the prey's prior beliefs are high then equilibrium *necessitates manipulation.* Intuitively, by manipulating with very high probability, the prey expects a lot of noise from a dangerous forest, and therefore when a silence occurs (which it inevitably does because the probability of manipulation turns out to be strictly less than 1 in equilibrium), the prey's beliefs about danger plunge. With the prey thereby lulled into a false sense of security, the predator perpetuates the prey's beliefs by adopting a Moriarty strategy from this point on.

In terms of the observable implications of deception, we show that a deceiver will either perfectly mimic the safe environment (in which case it leaves no trace of its presence), or it will create noise patterns that are *necessarily* non-i.i.d. We show that in an intuitive class of equilibria that we study, there must be instances of a *hot hand effect*: a streak persists with increasing probability.

This paper overlaps in spirit with standard cheap talk (Crawford and Sobel 1982) and Bayesian persuasion (Kamenica and Gentzkow 2011). However, in these models, signals are never a source of utility, whereas in our model signals are payoff relevant for the predator: each step increases the predator's payoff, but each step creates noise, which is a signal for the prey. Unlike Bayesian persuasion, our paper does not assume commitment. Moreover, these models are not optimal stopping problems. Our paper also overlaps in spirit with models of reputation (see Mailath and Samuelson 2006 for a review), which are repeated games where a player is able to increase its payoff by mimicking a "commitment type" that always plays a deterministic action. In our model, in order to manage the beliefs of the prey, the predator benefits from mimicking nature's stochastic behavior (albeit only up to the point that the prey does not become vigilant). Unlike the reputation literature, our model is an optimal stopping problem and not a repeated game.[4] Moreover, while the focus of the reputation literature is to characterize the payoffs that are achievable through reputation, our focus is on characterizing the stochastic behavior of a predator in equilibrium.

Finally, there are some related papers on strategic information manipulation in dynamic environments. Anderson and Smith (2013) study dynamic deception under noisy monitoring, where an informed player's actions are only imperfectly observed by uninformed rivals. Thus the rivals know that a "deceiver" took an action, but they do not know what the action is. In our model the prey can never be sure that there exists a deceiver in the forest. Frankel and Kartik (2019) study how strategic manipulation ("gaming the system") can reduce the informativeness of a test. The deceiver here is a "gaming" type with low ability that can perfectly mimic a non-gaming type with high ability. The test design is given, and in particular there gamer is not playing against a designer of the test. In our model, the predator plays against a prey that takes actions – she chooses how attentive to be – which affect the payoff of the predator. Our game generically has equilibria where the deceiver does not perfectly mimic nature. More peripherally, Hörner and Skrzypacz (2016) survey a literature on dynamic learning, experimentation, and information design in strategic settings. In our paper, the signal process is not an exogenously noisy observation of action, but is itself endogenously shaped by concealed behavior under environmental noise, which distinguishes our setting from the experimentation and optimal stopping literatures.

More generally, the main departure from these literatures is that our main object of interest is

[3] This is named after Professor James Moriarty, a fictional character and criminal mastermind in Sir Arthur Conan Doyle's "Sherlock Holmes" series.

[4] For this reason we are able to directly assume that the predator is perfectly patient rather than having to study a limit as impatience vanishes.

the equilibrium time-series structure of observables, in particular when deception generates serial dependence in the signal process. For this reason, the paper's contribution is best understood as identifying how strategic concealment under background noise can leave observable traces.

2 Model:

2.1 The Game

We describe a predator-prey game where the prey has to decide, based on the noise in the forest, whether it is safe enough to relax, and a predator has to decide how to (noisily) approach the prey. Formally:

- States: The state space $\Omega = \{s, d\}$ consists of a "safe" state and a "dangerous" state.
- Signals: In any period $t = 1, 2, ...,$ there is either silence (denoted $x_t = 0$) or sound/noise (denoted $x_t = 1$) that comes out of the forest. The signal space in period $1 \leq t \leq \infty$ is $X^t = \{0, 1\}^t$, where $X^0 = \emptyset$. For any $t > 0$ and $x^t \in X^t$, denote the total amount of noise in x^t by

$$\sum x^t := \sum_{i=1}^{t} x_i^t.$$

Denote by $x^t y^s$ the concatenation of the sequences x^t, y^s.

The process by which these signals are generated depends on the state:

In the safe state, nature determines in each period t whether there will be silence (denoted $\overline{x}_t = 0$) or noise (denoted $\overline{x}_t = 1$). The space of natural noise sequences up to period t is X^t. We will denote a natural noise sequence by $\overline{x}^t = (\overline{x}_1, ..., \overline{x}_t)$ to distinguish it from a forest noise sequence $x^t = (x_1, ..., x_t)$. The distribution over X^∞ generated by nature is the i.i.d. distribution with bias $\theta \in (0, 1)$:

$$\pi(\overline{x}^t) = \theta^{\sum \overline{x}^t}(1-\theta)^{t-\sum \overline{x}^t} \text{ for any } \overline{x}^t \in X^t.$$

In the safe state, nature's noise sequence $\overline{x}^t$ is the noise sequence that comes out of the forest:

$$x^t = \overline{x}^t.$$

In the dangerous state, in addition to the noise being generated by nature, there is a predator that chooses in each period t whether to take a step (denoted $\hat{x}_t = 1$) or not (denoted $\hat{x}_t = 0$) towards the prey. Each step creates noise that is indistinguishable from nature's noise. The space of stepping sequences up to period t is X^t, but we will use the notation $\hat{x}^t = (\hat{x}_1, ..., \hat{x}_t)$ to distinguish a stepping sequence from a natural noise $\overline{x}^t = (\overline{x}_1, ..., \overline{x}_t)$ and forest noise sequence $x^t = (x_1, ..., x_t)$. If nature produces sequence $\overline{x}^t \in X^t$ and the predator takes steps $\hat{x}^t \in X^t$, then the noise from the forest is given by

$$x^t = \max\{\overline{x}^t, \hat{x}^t\},$$

where the max is taken component-wise, that is, $x_i^t = \max\{\overline{x}_i^t, \hat{x}_i^t\}$ for each $i \leq t$. Note that noise is not additive – when both nature and the predator make noise, then the noise from the forest is $x_t = 1$. The upshot is that when $x_t = 1$ it is not possible to tell if this is made by nature or the predator or both. Observe also that the presence of the predator can only make the forest more noisy, not less. This rules out possibilities such as when the presence of a predator scares away other animals in the forest, leading to a conspicuous silence.

The distribution ν over forest noise-sequences X^∞ in the dangerous state is determined endogeneously by the predator's equilibrium stepping strategy, as we describe below.

• Strategies: In each period the prey decides whether to be inattentive or vigilant (denoted $e_t = 1$ and $e_t = 0$ respectively), the predator decides whether or not to step towards the prey, $\hat{x}_t \in \{0, 1\}$, and it also decides whether to attack the prey or wait (denoted $a_t = 1$ and $a_t = 0$, respectively). The order in which these decisions are made in a given period t is given in the time-line. We assume that the prey's choice of inattention is not observed by the predator when it decides whether to attack, but nature's realization is fully observed when it decides whether to take a step. We also assume that the game ends after the attack occurs and the payoffs of the prey and predator are realized.

t_a • prey: inattentive or vigilant — — — t_b • predator: attack or wait

— — — t_c • nature: noise/silence — — — t_d • predator: step or stay — — —

More generally, the prey is assumed to choose a mixed strategy

$$x^t \mapsto \xi(x^t),$$

which describes at each history of forest noises x^t the probability $\xi(x^t) \in [0, 1]$ of being inattentive. The predator is assumed to choose a mixed strategy

$$x^t \mapsto \alpha(x^t),$$

which describes the probability $\alpha \in [0, 1]$ of attacking the prey at each history of forest noises. Further, we assume the predator and prey can interact at infinity conditionally on the predator not attacking on finite time and thus, ξ and α also describe probabilities for histories x^∞.

The stepping decision is conditioned on the history of forest noises x^t but also on the realization of nature's $\overline{x}_t \in \{0, 1\}$, leading to a probability of taking a step $\rho(x^t, \overline{x}_t) \in [0, 1]$. Thus the predator chooses a strategy

$$(x^t, \overline{x}_t) \mapsto \rho(x^t, \overline{x}_t).$$

Unlike the prey, the predator gets to observe both x^t and $\overline{x}^t$ but the predator strategies condition actions only on x^t and $\overline{x}_t$. The following elucidates what this means. Note that at $t = 1$ the forest noise history is null, and therefore the first stepping decision depends only on $\overline{x}_1$. The outcome of this decision deterministically generates forest noise x_1. That is, x_1 is a function of $\overline{x}_1$ and the realization of $\rho(\emptyset, \overline{x}_1)$. In period 2, the stepping decision is a function that depends on x_1 (and $\overline{x}_2$ of course) but x_1 is determined by $\overline{x}_1$ and the realization of the stepping strategy adopted in period 1. This shows that the dependence of $\rho(x^t, \overline{x}_t)$ on x^t is to be understood as dependence on $\overline{x}^t$ and the realization of the stepping strategy up to period t. The dependence of the attack strategy $\alpha(x^t)$ on x^t is similarly interpreted.

Our key insight about deception is that *rational deceivers will always exploit any opportunity provided by circumstances to enhance their agenda.* This intuitive observation provides traction in our analysis. More formally, in our model, it will be a *dominant strategy* for the predator to take a step whenever nature makes noise, that is, $\rho(x^t, 1) = 1$. This is because whatever belief the prey may have about this, the predator only gains by taking the step deterministically since the cover provided by noise is perfect. In order to save substantially on notation, we restrict attention to the set of strategies that satisfy $\rho(x^t, 1) = 1$, and define $\rho(x^t) := \rho(x^t, 0)$, the probability of taking a step *if nature is silent*. We adopt the terminology that the predator engages in *manipulation* if it steps when nature is silent, and it exhibits *restraint* otherwise.

Finally, we make a useful (and notation-saving) observation about the dominant strategy feature $\rho(x^t, 1) = 1$. Since $\overline{x}_t = 1 \implies \hat{x}_t = 1$ we see that $\hat{x}^t \geq \overline{x}^t$, and it follows that

$$x^t = \max\{\overline{x}^t, \hat{x}^t\} = \hat{x}^t.$$

That is, in the dangerous state, all the noises made by the forest entail steps made by the predator, whether under the cover of nature's noise or via manipulation. In particular,

$$\sum x^t = \sum \hat{x}^t, \tag{1}$$

so that the total noise in the forest also reveals the total number of steps taken by the predator as it advances towards the prey. We use this observation below to simplify exposition.

• Choice: The stepping strategy ρ generates a distribution ν_ρ over forest noise sequences X^∞ by the following formula:

$$\nu_\rho(x^t) = \prod_{i=1}^{t} (\theta + (1-\theta)\rho(x^i))^{x_i} ((1-\theta)(1-\rho(x^i)))^{1-x_i}.$$

Thus, if a prey conjectures ρ then it has a conjecture about the distribution ν_ρ over X^∞ in the dangerous state. Given the knowledge that in the safe state the distribution is π over X^∞, it can use Bayes's rule to form posterior beliefs about danger. Letting $\mu = (\mu(s), \mu(d)) \in \Delta(\{s, d\})$ denote the prior over the state space $\{s, d\}$, the posterior at x^t conditional on not having been attacked is

$$\mu(d|x^{t-1}) = \frac{\mu(d)\nu_\rho(x^{t-1})}{\mu(d)\nu_\rho(x^{t-1}) + \mu(s)\pi(x^{t-1})}.$$

Since there is no possibility of an attack in the safe state, it is clear that the posterior on danger is 1 if the prey observes an attack.

The prey gets $k \in (0, 1)$ units of utility if it is vigilant. When it is vigilant, it can never get caught if attacked. If it is inattentive, then it gets utility 0 if it is captured by the predator, and utility 1 if it is not captured or not attacked. The probability $\varphi(\sum \hat{x}^t)$ of capture at x^t depends on the total number of steps $\sum \hat{x}^t$ taken by the predator, which we saw in (1) is conveniently given by $\sum x^t$, is strictly increasing in the number of steps. Thus the expected utility at x^t for the prey that plays strategy ξ is

$$U(\xi(x^t); x^t) = \xi(x^t)\left[\mu(s|x^t) + \mu(d|x^t)[\alpha(x^t)(1 - \varphi(\sum x^t)) + (1 - \alpha(x^t)\right] + (1 - \xi(x^t))k$$

$$= \xi(x^t)\left[1 - \mu(d|x^t)\alpha(x^t)\varphi(\sum x^t)\right] + (1 - \xi(x^t))k.$$

The prey is assumed to be myopic: it makes a choice of $\xi(x^t)$ at each history to maximize expected utility received at that history alone.

The predator gets utility 1 if it captures the prey and 0 otherwise. There is no utility from stepping per se, other than for the fact that steps impact the probability of capture. The predator is also completely patient. Given a conjecture about ξ, the predator is assumed to choose (ρ, α) that solves the following dynamic programming problem:

$$V(x^t) = \max \left\{ \begin{array}{c} \varphi(\sum x^t)\xi(x^t), \\ V(x^t 1), \\ \theta V(x^t 1) + (1-\theta) V(x^t 0) \end{array} \right\}.$$

The first term on the right-hand side is the expected utility of an attack. Here we use the assumption that at the time of attack the predator does not know if the prey is vigilant or inattentive, and therefore the probability of a successful capture is $\varphi(\sum x^t)\xi(x^t)$. The second term corresponds to manipulation (taking a step if nature is silent): the predator steps whenever nature makes noise since it is a dominant strategy to do so, and manipulation implies that the predator steps also when nature is silent. It follows that there is noise in the forest in that period with probability 1, deterministically leading to the continuation sequence $x^t 1$. Given that the predator is completely patient, the (discounted) utility is then just $V(x^t 1)$. The third term corresponds to exercising restraint: with probability θ nature makes noise and the predator steps (leading to the continuation sequence $x^t 1$) and with probability $1-\theta$ nature is silent and the predator does not step (leading to continuation $x^t 0$).

For perspective, we note that the literature on reputation is also interested in the case of completely patient agents: while it is common for papers to assume impatience (discount factor $\delta < 1$), the papers study equilibria that arise as $\delta \to 1$. The reason they do not assume $\delta = 1$ directly is because utility from infinite-horizon sequences of action is not well-defined in that case. Since we have an optimal stopping problem, we can assume $\delta = 1$ directly.

The only assumption we maintain on primitives throughout (and henceforth without explicit reference) is that:

Assumption 1 *The capture probability $N \mapsto \varphi(N)$ is weakly increasing but not eventually constant,*[5] *and satisfies $\varphi(0) = 0$ and $\varphi(\infty) = \lim_{N\to\infty} \varphi(N) > 1-k$.*

The assumption implies that it is impossible to capture the prey if the predator does not take a step, and the more steps it takes, the more likely it is to capture the prey, with a maximum probability $\varphi(\infty)$ that is “sufficiently” high This maximum probability cannot be achieved with a finite number of steps however, that is, φ is not eventually flat.

Our equilibrium notion is that of a perfect Bayesian equilibrium in mixed strategies:

Definition 1 *(Equilibrium) An equilibrium is defined by strategies (ρ, α, ξ) and a system of beliefs $\mu = \{x^t \mapsto \mu(d|x^t)\}$ such that:*

1. *μ is derived from Bayes' Rule using (ρ, α).*
2. *Given ρ, α, the prey's strategy ξ maximizes $U(\cdot; \mu(d|x^t))$ for all x^t,*
3. *Given e, the predator's strategy (ρ, α) solves its dynamic programming problem.*

In the sequel we restrict attention to a subclass of equilibria where ξ satisfies natural properties and characterize the set of manipulation strategies ρ that can be supported in this subclass. We further restrict attention to patient predators by assuming $\delta = 1$.

[5]That is, there does not exist M such that $\varphi(\cdot)$ is constant on $\{M, M+1, ...\infty\}$.

A strong intuition one might have about patient predators is that they may advance only under the cover of nature's noise, and attack only at perfect striking distance. We refer to this as a Moriarty strategy:

Definition 2 *The predator strategies* (ρ, α) *constitute a* Moriarty strategy *if*

$$\rho(x^t) = 0 \text{ for all } x^t,$$

$$\alpha(x^t) = 1 \iff t = \infty.$$

An equilibrium (ρ, α, ξ) *is a* Moriarty equilibrium *if* (ρ, α) *are a Moriarty strategy.*

A Moriarty strategy is defined by a strategy$\rho = 0$ of never manipulating, and of attacking only at $t = \infty$. Along any sequence x^∞ the probability of capture is maximized by attacking at $t = \infty$.

We conclude by connecting our model to our motivating applications in the Introduction. In the hostile take-over example, forest noise corresponds to the indicator function for whether or not the target company's stock was bought in the market. Nature corresponds to regular buying activity ($\overline{x}_i = 1$) that does not involve an intent of take-over. The acquiring firm's step ($\hat{x}_i = 1$) corresponds to buying activity with the intent of take-over. If either nature or the acquiring firm buys ($\max\{\overline{x}_i, \hat{x}_i\} = 1$), the target company sees market activity in its stock (forest noise $x_i = 1$). An attack is a bid to take over the target company. Vigilance on the part of the target company corresponds to defensive maneuvers, making the acquisition harder or less desirable.[6] The probability of a successful or desirable takeover may be less than 1 even once a majority of shares are acquired due to defensive maneuvers following a takeover bid. A Moriarty strategy corresponds to what is known as a "creeping take-over", where the acquiring company quietly acquires a stake in the company until it makes a take-over bid.

Next, consider the selling example. Imagine each t to correspond to a dimension of a product. Nature decides whether the dimension is "good" (corresponding to $\overline{x}_i = 1$) or not, and the seller decides whether to "present it as good" or not. If the dimension is good then it will be observed as such regardless of how the seller presents it. If the dimension is bad then it will be perceived as good if and only if the seller presents it as good. An attack corresponds to asking the buyer if they would like to buy the product at the current time. Vigilance on the part of the buyer is a commitment not to buy from this seller. The assumption that φ is increasing in the number of steps implies that the probability of buying is increasing only in the number of good dimensions, without regard to the number of bad dimensions. This could describe a buyer who is only looking for sufficient reasons for buying. A buyer who is evaluating the product based on both reasons for and against it is better captured by an extension of our model where, for instance, the probability of success is a function of the average noise $\frac{\sum x^t}{t}$.

Finally, the HFT example corresponds to an extension of the model where the predator is not sure of the prey's presence, and the prey seeks to deceive the predator in this regard.

[6]There are several examples of defensive responses. Elections of board members can be staggered so that it takes several years for a new owner to replace the board of directors. Its most valuable assets can be transferred or sold in order to make it less attractive to the acquiring company. The company can buy back its own shares to reduce the number available on the market and consolidate ownership among trusted shareholders. Key executives can be promised large financial rewards if they lose their jobs due to a take-over, increasing costs for the acquiring company.

3 Alert Equilibria

3.1 Definition and Result

In this section we study the following class of equilibria.

Definition 3 *(Alert Equilibrium) For any $N < \infty$, an N-*alert equilibrium *is an equilibrium (ρ, α, ξ) where ξ satisfies:*

$$\sum x^t \leq N \implies \xi(x^t) = 1$$

$$\sum x^t > N \implies 0 < \xi(x^t 1) < \xi(x^t) < \xi(x^t 0) < 1.$$

In an alert equilibrium, the prey is alert beyond some total noise $N < \infty$. Since stepping-when-noise is a dominant strategy for the predator, the probability of capture at the end of any history x^t is $\varphi(\sum x^t)$. Thus, for low values of total noise $\sum x^t$, the probability of capture is low, and the prey may well choose to be completely inattentive, $\xi(x^t) = 1$. The definition above requires that there is some noise level N up until which the prey is completely inattentive. The substantive assumption here is that the threshold depends on total noise $\sum x^t$ and is otherwise independent of any other details of the history x^t. For instance it rules out that the prey may tolerate more noise if the sequence x^t has many 0's in it.

Beyond the threshold N, the definition requires that the prey is never willing to ever be completely inattentive, and neither does it ever strictly prefer to be completely vigilant: $\sum x^t > N \implies \xi(x^t) \in (0, 1)$. While there may exist equilibria with complete vigilance, it is very natural to imagine scenarios where a prey never finds itself requiring complete vigilance, since it is never in the predator's benefit to let the prey get to that point (since the probability of capture is then 0).

Finally, the definition comes with the very natural restriction that silence makes the prey more likely to be inattentive, while noise makes it less likely. However another intuition one may have is that too much silence could make the prey nervous, leading it to become more vigilant. But this intuition makes more sense in an environment where the predator's presence can create silence – for instance its presence may scare away other animals that would otherwise make noise in the forest. In our setting the predator can only add to the noise in the environment.

The main result of this section characterizes the set of manipulation strategies in this class of equilibria. Say that a manipulation strategy ρ is *supported* by an alert equilibrium if there exists an alert equilibrium (ρ, α, ξ) with ρ as its manipulation strategy. Our first result establishes existence, and characterizes the set of manipulation strategies supported. Proofs of all results can be found in the appendix.

Theorem 1 *An N-alert equilibrium exists if and only if $\varphi^{-1}(1-k) - 1 < N < \infty$. Moreover, for any such N, a manipulation strategy ρ is supported by an N-alert equilibrium if and only if it satisfies*

(a) for all x^t, $\rho(x^t) < 1$,
(b) for all $\sum x^t > N$, $\rho(x^t) = 0$,
(c) for all $\sum x^t \leq N$,

$$\inf_{x^t : \sum x^t = N+1} \prod_{i=1}^{t} (1 + \frac{(1-\theta)}{\theta} \rho(x^i))^{x_i} (1 - \rho(x^i))^{1-x_i} > \frac{1-k}{k - (1 - \varphi(N+1))} \frac{\mu(s)}{\mu(d)}.$$

The first claim in the theorem is that N-alert equilibria does not always exist. To interpret the lower bound, observe that it can be written as $k > 1 - \varphi(N+1)$, which states that the utility k of being vigilant is greater than the probability of not being captured if attacked at noise $N+1$.[7] When this fails, then for any x^t with $\sum x^t = N+1$ the prey has a preference for being inattentive, and we get $\xi(x^t) = 1$. This is impossible for an N-alert equilibrium, which requires $\xi(x^t) < 1$ at such x^t. When the condition holds, then for any x^t with $\sum x^t = N+1$ the prey prefers to be vigilant $\xi(x^t) = 0$. But the prey can be made indifferent if the probability of an attack is small enough, and consequently it is possible to obtain $0 < \xi(x^t) < 1$. In fact, the result shows that an N-alert equilibrium can be constructed.

The theorem goes on to characterize equilibrium manipulation behavior through three conditions. The first condition tells us that manipulation is never deterministic. This feature is not limited to alert equilibria: we show in Lemma 6 that the game with a patient predator ($\delta = 1$) has no equilibrium with deterministic manipulation at any history ($\rho(x^t) = 1$ for some x^t). The reason is as follows. If manipulation is deterministic at x^t, then regardless of whether nature makes noise or not in that period, the predator will take a step (it takes a step when nature makes noise since that is a dominant strategy, and it takes a step when noise is silent since we are hypothesizing deterministic manipulation). In equilibrium, the prey expects that with probability 1 there is noise in the current period. Then the patient predator has a profitable deviation: if it chooses restraint, then with probability $1-\theta$ there will be silence, and the prey will erroneously conclude with certainty that the state must be safe. The predator can then manipulate arbitrarily many times and attack with increasingly better odds of success.

The second condition in the theorem tells us that "eventually" there can never be any manipulation. The third condition tells us what values of ρ are permissible "early" in the game.[8]

3.2 Properties

Since Theorem 1 does not make transparent what kinds of ρ are possible, we highlight some key implications in the following propositions.

We first note that Moriarty equilibria are incompatible with N-alert equilibria, although full restraint $\rho = 0$ is possible under some conditions.

Proposition 1 *The following hold for any $\varphi^{-1}(1-k) - 1 < N < \infty$.*
(a) There exists an N-alert equilibrium with $\rho = 0$ iff $\frac{\mu(d)}{1-\mu(d)} > \frac{1-k}{\varphi(N+1)-(1-k)}$.
(b) There does not exist an N-alert Moriarty equilibrium.

Part (a) says that total restraint $\rho = 0$ is possible only as long as the prey's prior on danger is high enough. The high prior along with the predator's restraint is consistent with the prey being alert in equilibrium. To see that restraint is not possible when the prey's prior on danger is low, observe the following. Since the prey becomes alert beyond N, its posterior on danger must be higher than the prior. This is possible only if it anticipated some manipulation. The predator on the other hand is indifferent between all manipulation strategies for $\sum x^t \leq N$, because N will come about ν_ρ-almost surely at any ρ, consequently yielding (what we show in the proof must be)

[7] By assumption 1 , $\lim_{n\to\infty} \varphi(n) > 1-k$, and so there always exist N satisfying $k > 1 - \varphi(N+1)$.

[8] Note that the right hand side expression $\frac{1-k}{k-(1-\varphi(N+1))} \frac{\mu(s)}{\mu(d)}$ is strictly positive given that N satisfies $N > \varphi^{-1}(1-k) - 1$ by hypothesis.

its payoff $\varphi(N)$ regardless of ρ. It follows that in equilibrium it is willing to adopt the manipulation strategy that the prey anticipates.

While restraint is possible, it is not possible that the equilibrium is Moriarty. In a Moriarty equilibrium, the predator never attacks in finite time, and consequently the prey always remains inattentive at all sequences ($\xi(x^t) = 1$ for all x^t), which is ruled out by N-alert equilibria, which require that the prey becomes alert at $\sum x^t = N$. Another incompatibility is that (as we see in the proof of Theorem 1) in any N-alert equilibria, with ν-probability one, the prey is always alert at $T = \infty$ and so an attack at $T = \infty$ yields zero utility, and the predator is strictly better off attacking at some $t < \infty$. Indeed, in any N-alert equilibrium there is a strictly positive probability of attack at each finite time such that $\sum x^t > N$.

It is intuitive that patient predators will play a Moriarty strategy. It is also intuitive that prey who are aware of the possibility of deception might be alert. The take-away is that one of these two intuitions must be relinquished. For instance, in the context of hostile takeovers, if we take a creeping takeover as evidence of a Moriarty equilibrium, then we must believe that target firms enjoy a false sense of security. If instead we believe that target firms are alert, then we must expect predators to be noisy (whether because a little noise does not affect their outcome, or because they were never patient to begin with).

Next, we explore whether noise may be i.i.d. in the dangerous state.

Proposition 2 *For any $\varphi^{-1}(1-k) - 1 < N < \infty$ and any N-alert equilibrium, if ρ is constant for all $\sum x^t \leq N$ then $\rho = 0$.*

If ρ is constant for $\sum x^t \leq N$, then for such sequences, ν_ρ must be an i.i.d. distribution with bias $\theta + (1-\theta)\rho$. The content of Proposition 2 is that we can never have an i.i.d. environment for $\sum x^t \leq N$ unless $\rho = 0$. While it is easy to imagine that the environment can be non-i.i.d. when a predator is present, it is not at all obvious that it *must* be non-i.i.d. in an equilibrium where the predator manipulates. We take this as one of the main empirical take-aways: *in a deceptive environment, the predator either leaves no trace of its presence ($\rho = 0$) or the occurrence of noise must be non-i.i.d. in early periods.*

The next result gives us some sense of the non-i.i.d. nature of the forest noise.

Proposition 3 *For any $\varphi^{-1}(1-k) - 1 < N < \infty$ and in any N-alert equilibrium, it must be that for any $\sum x^t \leq N$,*

$$\inf_n \nu(1|x^t 0^n) = \theta.$$

Early in the game, after a sequence of 0's the ν-probability of a further 0 will typically be less than $1-\theta$ (strictly less so when the probability of manipulation is strictly positive, $\rho(x^t 0^n) > 0$). Proposition 3 states that the ν-probability of a further 0 must not be bounded away from $1-\theta$, and in fact must eventually increase to $1-\theta$. This is a hot hand effect of sorts for a sequence of silences.

4 Inattentive Equilibria

4.1 Definition and Results

In this section we consider a counterpart of an alert equilibrium for $N = \infty$. This is a class of equilibria where, except for a set of sequences that have ν_ρ-measure 0, the prey is completely inattentive along x^∞.

Definition 4 *(Inattentive Equilibrium) An* inattentive equilibrium *is an equilibrium* (ρ, α, ξ) *where* ξ, ν_ρ *satisfy:*

$$\nu_\rho(\{x^\infty : \xi(x^t) = 1 \textit{ for all } t \leq \infty\}) = 1$$

Say that a manipulation strategy ρ is *supported* by an inattentive equilibrium if there exists an inattentive equilibrium (ρ, α, ξ) with ρ as its manipulation strategy.

Theorem 2 *An inattentive equilibrium exists. A strategy* ρ *can be supported by an inattentive equilibrium if and only if* ρ *is such that* $\rho < 1$ *and*

$$\lim_{t\to\infty} \mu(d|x^t) \leq 1 - k,\ \ \nu_\rho\text{-}a.s.$$

The result establishes existence and characterizes the set of all manipulation strategies that can be supported by an inattentive equilibrium. In words, a strategy ρ arises in an inattentive equilibrium if and only if never it manipulates deterministically, and if it almost surely induces a limiting posterior on danger that is sufficiently low. When the posterior satisfies $\mu(d|x^\infty) \leq 1 - k$ then the prey finds it optimal to be inattentive even if, conditional on the state being dangerous, an attack is sure to happen and sure to lead to capture with probability $\lim_N \varphi(N)$. Since the Theorem does not say more transparently what kind of ρ are supported by an inattentive equilibrium, we proceed to get a sense of the kinds of ρ that are permitted and ruled out.

We first observe that a Moriarty equilibrium is not always guaranteed to exist.

Proposition 4 *An inattentive Moriarty equilibrium exists iff* $\mu(d) \leq 1 - k$.

If the prey's prior belief about danger is sufficiently low, then $\rho = 0$ will keep the beliefs there and the prey will always be inattentive, including at $T = \infty$. The predator will reach its highest possible payoff in this case by attacking at $T = \infty$. While it is intuitive that a Moriarty strategy can arise in equilibrium when the prey's beliefs about danger are low, what is less obvious is that when the prey's beliefs are too high the predator *must* engage in manipulation. In fact it must manipulate with high probability (though not probability 1), and in fact must manipulate with even higher probability after noise. If it does so, then almost surely there will eventually be some silence. The reason is that this is the *only* way to bring the prey's posterior down: since a high probability manipulation strategy implies that the odds of silence are very low, the occurrence of silence leads the prey's posterior beliefs about danger to plunge. Once the posterior is sufficiently low the predator can then set $\rho = 0$ to keep beliefs frozen and to have the prey available and fully inattentive at $T = \infty$ (although this is not the unique strategy since the predator is completely patient).

Next we explore whether noise may be i.i.d. in the dangerous state.

Proposition 5 *In any inattentive equilibrium, if* ρ *is constant for all* x^t *then* $\rho = 0$.

In the case of alert equilibria, we saw (in Proposition 2) that it is impossible for the predator to engage in i.i.d. manipulation. The proposition states that this continues to be the case here. The reason is different however. In an inattentive equilibrium an attack only takes place at $T = \infty$, which means that the prey has infinitely many periods of signals in which to learn the state. If ρ is constant and strictly positive then the environment in the dangerous state is i.i.d. but a bias that is different from that in the safe state, implying that the prey will eventually learn that the true state. But this is not optimal for the predator.

It is not possible to provide a counterpart to Proposition 3 because ρ is practically unrestricted up to any finite time period. Indeed, ρ must be such that the prey's beliefs are less than $1-k$ only *eventually*. In the next subsection we consider a refinement that selects a unique ρ, and proceed to analyze it.

4.2 Swift Manipulation

Consider manipulation strategies that push beliefs below $1-k$ swiftly:

Definition 5 *A manipulation strategy ρ is swift if*
(a) ρ is more swift than any ρ': for any x^t,

$$\mu_{\rho'}(d|x^t) \leq 1-k \implies \mu_{\rho}(d|x^t) \leq 1-k.$$

(b) $\rho \leq \overline{\rho}$ whenever $\overline{\rho}$ is also more swift than ρ'.

Take any manipulation strategy ρ' and suppose it is the case that for any history x^t, when ρ' gets beliefs below $1-k$ at x^t then so does ρ. Thus, ρ is more swift than ρ'. Part (a) of the definition requires ρ to be more swift than all manipulation strategies ρ'. If there is more than one such swiftest strategy, part (b) requires ρ to be the swiftest strategy that also uses the least possible amount of manipulation. We establish that there exists a strategy satisfying the definition, and moreover it is unique. We also show that it is supported by a unique inattentive equilibrium.

Proposition 6 *There exists a unique swift manipulation strategy and it is given by*

$$\rho(x^t) = \begin{cases} 1 - \frac{(1-k)\mu(s|x^t)}{k\mu(d|x^t)} & \text{if } \mu(d|x^t) > 1-k \\ 0 & \text{otherwise} \end{cases}.$$

Furthermore, ρ is supported by a unique inattentive equilibrium.

Since there is a unique swift inattentive equilibrium, we can now explore the nature of noise produced by ν in this equilibrium.

Proposition 7 *Suppose $\mu(d) > 1-k$. In an inattentive equilibrium with a swift ρ, the conditional probability of noise after a sequence of noise, $\nu(1|1^n)$, is strictly increasing in n and satisfies $\lim_{n\to\infty} \nu(1|1^n) = 1$.*

Thus, a swift manipulation strategy generates increasing probabilities of noise conditional on observing noise. The reason is that the longer the streak of noise, the higher the likelihood of continued noise must be in order for a subsequent silence to sharply lower the posterior below $1-k$. The pattern of increasing likelihood of continuation of a streak is an instance of the hot-hand effect. [9]

[9] In the evidence, a hot hand effect can occur after a streak at any point in a sequence, and not just a streak starting at time 1, as in the proposition. This is readily rationalized: in our model the prey knows that if a predator is at all present then it is present at time 1, but a natural generalization of our model would be where the entry period of the predator is random.

5 Conclusion

This paper develops a dynamic model of deception in which a patient predator strategically manages the prey's beliefs as it makes a noisy advance towards it. We characterize the full set of equilibria defined by intuitive restrictions on the prey's attention, and identify conditions under which nontrivial deception is sustainable. When the prey begins with low concern about danger, a Moriarty equilibrium—where the predator indefinitely postpones attack and perfectly mimics a safe environment—can arise. In contrast, when the prey enters the interaction with sufficiently high prior beliefs of danger, manipulation becomes unavoidable: any equilibrium must involve the predator generating noise so as to "reset" the prey's fears. When prey are alert then a Moriarty equilibrium is not possible.

A central result of the paper is that the equilibrium distribution of observable signals is *necessarily* non-i.i.d. when the deceiver chooses not to perfectly mimic nature. Depending on the prey's initial vigilance, different forms of serial dependence arise. When the prey is alert, equilibrium behavior implies a hot hand effect of sorts: following long sequences of silence, the probability that silence continues cannot remain uniformly low, and must eventually increase to its natural benchmark. When the prey is inattentive, the equilibrium exhibits a sharper hot-hand effect: conditional on observing a long streak of noise, the probability of continued noise increases and converges to one.

While there is an extensive literature in psychology and economics documenting belief biases (see Benjamin 2019 for a review), there is little research exploring why they might exist in the first place.[10] The current paper suggests a possible evolutionary reason for why people misperceive randomness, specifically by believing in the gambler's fallacy. One of the main results of this paper is that a predator can be detected by non-i.i.d. noise in the environment that takes the form of a hot hand effect. A misspecified belief about randomness in the *opposite* direction (that is, in the direction of a gambler's fallacy) would help the prey recognize danger faster. Intuitively, if the prey believes that randomness requires streaks to break with high probability, then observing a long streak (which is a likely consequence of a hot hand effect) would lead the prey to sharply update beliefs in the direction of danger, thereby improving the odds of survival. We leave it to future research to explore this message.

A Preliminary Lemmas

Lemma 1 *(a) Beliefs are given by*

$$\mu(d|x^{t-1}) = \frac{\mu(d)\prod_{i=1}^{t}(1+\frac{(1-\theta)}{\theta}\rho(x^i))^{x_i}((1-\rho(x^i))^{1-x_i}}{\mu(d)\prod_{i=1}^{t}(1+\frac{(1-\theta)}{\theta}\rho(x^i))^{x_i}(1-\rho(x^i))^{1-x_i}+\mu(s)},$$

or equivalently

$$\mu(d|x^{t-1}) = \frac{\mu(d|x^{t-1})[(1+\frac{(1-\theta)}{\theta}\rho(x^t))^{x_t}(1-\rho(x^t))^{1-x_t}]}{\mu(d|x^{t-1})[(1+\frac{(1-\theta)}{\theta}\rho(x^t))^{x_t}(1-\rho(x^t))^{1-x_t}]+\mu(s|x^{t-1})}.$$

(b) For any x^t, the prey is inattentive $\xi(x^t)=1$ iff

$$\mu(d|x^t) \leq \frac{1-k}{\alpha(x^t)\varphi(\sum x^t)}.$$

[10] Noor and Payro (2025) model the gambler's fallacy and analyze an evolutionary horserace between a Bayesian population and a population of agents who exhibit the gambler's fallacy. They show that the latter always survive.

Proof. (a) Since $\nu_\rho(x^t) = \prod_{i=1}^t (\theta + (1-\theta)\rho(x^i))^{x_i}((1-\theta)(1-\rho(x^i)))^{1-x_i}$,we see that

$$\mu(d|x^t) = \frac{\mu(d)\nu_\rho(x^t)}{\mu(d)\nu_\rho(x^t) + \mu(s)\pi(x^t)}$$

$$= \frac{\mu(d)\prod_{i=1}^t (\theta + (1-\theta)\rho(x^i))^{x_i}((1-\theta)(1-\rho(x^i)))^{1-x_i}}{\mu(d)\prod_{i=1}^t (\theta + (1-\theta)\rho(x^i))^{x_i}((1-\theta)(1-\rho(x^i)))^{1-x_i} + \mu(s)\prod_{i=1}^t \theta^{x_i}(1-\theta)^{1-x_i}}$$

$$= \frac{\mu(d)\prod_{i=1}^t (\theta + (1-\theta)\rho(x^i))^{x_i}((1-\rho(x^i))^{1-x_i}}{\mu(d)\prod_{i=1}^t (\theta + (1-\theta)\rho(x^i))^{x_i}(1-\rho(x^i))^{1-x_i} + \mu(s)\prod_{i=1}^t \theta^{x_i}}$$

$$= \frac{\mu(d)\prod_{i=1}^t (1 + \frac{(1-\theta)}{\theta}\rho(x^i))^{x_i}((1-\rho(x^i))^{1-x_i}}{\mu(d)\prod_{i=1}^t (1 + \frac{(1-\theta)}{\theta}\rho(x^i))^{x_i}(1-\rho(x^i))^{1-x_i} + \mu(s)}.$$

Observe also that

$$\mu(d|x^t) = \frac{\mu(d|x^{t-1})\nu_\rho(x_t|x^{t-1})}{\mu(d|x^{t-1})\nu_\rho(x_t|x^{t-1}) + \mu(s|x^{t-1})\pi(x_t|x^{t-1})}$$

$$= \frac{\mu(d|x^{t-1})[(\theta + (1-\theta)\rho(x^t))^{x_t}((1-\theta)(1-\rho(x^t)))^{1-x_t}]}{\mu(d|x^{t-1})[(\theta + (1-\theta)\rho(x^t))^{x_t}((1-\theta)(1-\rho(x^t)))^{1-x_t}] + \mu(s|x^{t-1})\theta^{x_t}(1-\theta)^{1-x_t}}$$

$$= \frac{\mu(d|x^{t-1})[(1 + \frac{(1-\theta)}{\theta}\rho(x^t))^{x_t}(1-\rho(x^t))^{1-x_t}]}{\mu(d|x^{t-1})[(1 + \frac{(1-\theta)}{\theta}\rho(x^t))^{x_t}(1-\rho(x^t))^{1-x_t}] + \mu(s|x^{t-1})}.$$

(b) Using the entry condition we see that

$$U(1;x^t) \geq U(0;x^t)$$

$$\iff \left[1 - \mu(d|x^t)\alpha(x^t)\varphi(\sum x^t)\right] \geq k$$

$$\iff \mu(d|x^t) \leq \frac{1-k}{\alpha(x^t)\varphi(\sum x^t)},$$

as desired. ∎

Lemma 2 *There is no equilibrium such that $\alpha(\overline{x}^t) = 1$ for some $\overline{x}^t$ with $t < \infty$.*

Proof. If $\alpha(\overline{x}^t) = 1$ for some $\overline{x}^t$ with $t < \infty$, then any continuation sequence $\overline{x}^t z^r$ is impossible in the dangerous state, that is, $\nu(\overline{x}^t z^r) = 0$ for all z^r. It follows that $\mu(d|x^t z^r) = 0$ and it is optimal for prey to set $\xi(x^t z^r) = 1$. But then $\alpha(\overline{x}^t) = 1$ cannot be optimal: while $\alpha(\overline{x}^t) = 1$ yields utility $\xi(\overline{x}^t)\varphi(\sum \overline{x}^t)$, deviating to $\alpha(\overline{x}^t) = 0$ followed by a manipulation and subsequent attack in the next period leads to utility $\varphi(\sum \overline{x}^t 1) > \xi(\overline{x}^t)\varphi(\sum \overline{x}^t)$. ∎

Next we show that it is not possible to support an equilibrium where the predator wishes to attack only after some N is reached, and beyond N deterministically manipulates at at least one sequence.

Lemma 3 *Take any N s.t. $\xi(x^t) = 1$ if $\sum x^t \leq N$. Then there is no equilibrium such that for all x^t,*

$$\sum x^t \leq N \implies \alpha(x^t) = 0,$$

$$\sum x^t > N \implies \alpha(x^t) > 0,$$

$$\exists \overline{x}^t \text{ s.t. } \sum \overline{x}^t > N \text{ and } \rho(\overline{x}^t) = 1.$$

Proof. Take any N s.t. $\xi(x^t) = 1$ of $\sum x^t \leq N$. We show that there exists no ξ that induces the desired predator strategies (α, ρ). It is necessary that these desired strategies satisfy the following conditions: for any x^t s.t. $\sum x^t \leq N$, since waiting is optimal,

$$\varphi(N) \leq \max\{V(x^t 1), \theta V(x^t 1) + (1-\theta) V(x^t 0)\}$$

and for any x^t s.t. $\sum x^t > N$, since attack happens with positive probability,

$$V(x^t) = \xi(x^t)\varphi(\sum x^t)$$

and for some $\overline{x}^t$with $\sum \overline{x}^t > N$, since the predator deterministically manipulates,

$$V(\overline{x}^t 1) = \xi(\overline{x}^t)\varphi(\sum \overline{x}^t) \geq \theta V(\overline{x}^t 1) + (1-\theta) V(\overline{x}^t 0).$$

Observe that since $\rho(\overline{x}^t) = 1$ by hypothesis, any sequence $\overline{x}^t 0 z^r$ that has silence following $\overline{x}^t$ is impossible in the dangerous state, that is, $\nu(\overline{x}^t 0 z^r) = 0$ for all z^r. It follows that $\mu(d|x^t 0 z^r) = 0$ and it is optimal for prey to set $\xi(x^t 0 z^r) = 1$. Indeed, $V(x^t 0 z^r) = \varphi(\sum x^t 0 z^r)$. Observe in particular that

$$\begin{aligned} &\max\{V(\overline{x}^t 0 z^r 1), \theta V(\overline{x}^t 0 z^r 1) + (1-\theta) V(\overline{x}^t 0 z^r 0)\} \\ &= \max\{\varphi(\sum \overline{x}^t 0 z^r 1), \theta \varphi(\sum \overline{x}^t 0 z^r 1) + (1-\theta)\varphi(\sum \overline{x}^t 0 z^r)\} \\ &= \varphi(\sum \overline{x}^t 0 z^r) \end{aligned}$$

and so the hypothesis requires that for all z^r,

$$\varphi(\sum \overline{x}^t 0 z^r) = V(\sum \overline{x}^t 0 z^r)$$

$$\geq \max\{V(\overline{x}^t 0 z^r 1), \theta V(\overline{x}^t 0 z^r 1) + (1-\theta) V(\overline{x}^t 0 z^r 0)\} = \varphi(\sum \overline{x}^t 0 z^r 1).$$

It follows that $\varphi(\sum \overline{x}^t 01) \geq \varphi(\sum \overline{x}^t 01^2) \geq \varphi(\sum \overline{x}^t 01^3) = \ldots$ and thus for all m,

$$\varphi(\sum \overline{x}^t 01) \geq \varphi(\sum \overline{x}^t 01^{m+1}).$$

Denoting $\theta := \lim_{n\to\infty} \varphi(n)$ and given assumption 1, this leads to the contradiction that $\theta > \varphi(\sum \overline{x}^t 01) \geq \varphi(\sum \overline{x}^t 01^{m+1}) \to \theta$. ∎

B Proof of Theorem 1

For any N-alert prey strategy ξ, define

$$X^- := \{x^t : \sum x^t \leq N\}$$

$$\partial X^- := \{x^t : \sum x^t = N\}$$

$$X^+ := \{x^t : \sum x^t > N\}.$$

Lemma 4 *The following hold in any N-alert equilibrium such that $x^t \in \partial X^- \implies \alpha(x^t) < 1$:*

(i) $x^t \in \partial X^- \implies V(x^t) = \varphi(\sum x^t) = \theta V(x^t 1) + (1-\theta)V(x^t 0) = V(x^t 1)$, $\xi(x^t 1) = \frac{\varphi(\sum x^t)}{\varphi(\sum x^t 1)}$, *and* $\xi(x^t 0) = 1$.

(ii) $x^t \in \partial \underline{X} \implies x^t 0 \in \partial \underline{X}$.

(iii) $x^t \in X^+ \implies \mu(d|x^t) > \frac{1-k}{\varphi(\sum x^t)}$.

Proof. (i) The restriction $\alpha(x^t) < 1$ for $x^t \in \partial X^-$ implies that the predator weakly prefers to wait rather than attack at x^t. Since $\rho < 1$, waiting is always optimal with restraint, and so the value of waiting equals the utility of waiting with restraint. Thus, at any $x^t \in \partial X^-$, it must be that

$$\xi(x^t)\varphi(\sum x^t) \leq \theta \xi(x^t 1)\varphi(\sum x^t 1) + (1-\theta)\xi(x^t 0)\varphi(\sum x^t)$$

$$\text{and } \theta \xi(x^t 1)\varphi(\sum x^t 1) + (1-\theta)\xi(x^t 0)\varphi(\sum x^t) \geq \xi(x^t 1)\varphi(\sum x^t 1)$$

$$\iff \varphi(\sum x^t) \leq \theta \xi(x^t 1)\varphi(\sum x^t 1) + (1-\theta)\varphi(\sum x^t) \text{ and } \xi(x^t 0)\varphi(\sum x^t) \geq \xi(x^t 1)\varphi(\sum x^t 1)$$

$$\iff \xi(x^t 1) \geq \frac{\varphi(\sum x^t)}{\varphi(\sum x^t 1)} \text{ and } \xi(x^t 0)\frac{\varphi(\sum x^t)}{\varphi(\sum x^t 1)} \geq \xi(x^t 1)$$

$$\iff \xi(x^t 1) \geq \frac{\varphi(\sum x^t)}{\varphi(\sum x^t 1)} \geq \xi(x^t 0)\frac{\varphi(\sum x^t)}{\varphi(\sum x^t 1)} \geq \xi(x^t 1) \text{ (since } \xi(x^t 0) \leq 1)$$

$$\iff \xi(x^t 1) = \frac{\varphi(\sum x^t)}{\varphi(\sum x^t 1)} \text{ and } \xi(x^t 0) = 1.$$

The desired assertions follow.

(ii) By the preceding,

$$x^t \in \partial X^- \implies \xi(x^t 0) = 1 \implies x^t 0 \in \partial X^-.$$

(iii) Prove the contrapositive. For any x^t, if $\mu(d|x^t) \leq \frac{1-k}{\varphi(\sum x^t)}$ then for any $\alpha(x^t) \in [0,1]$,

$$\mu(d|x^t) \leq \frac{1-k}{\varphi(\sum x^t)} \leq \frac{1-k}{\alpha(x^t)\varphi(\sum x^t)},$$

and thus the prey strictly prefers to be inattentive, $\xi(x^t) = 1$ (by lemma 1(b)). Since the ξ is N-alert. it follows that $x^t \in X^-$.

■

Lemma 5 *Suppose there is N such that $V(x^t) = \xi(x^t)\varphi(\sum x^t)$ for all $\sum x^t > N$. Then, for any $\sum x^t > N$,*

$$V(x^t) = \theta V(x^t1) + (1-\theta)V(x^t0) \iff \xi(x^t1) = [\xi(x^t) - (1-\theta)\xi(x^t0)]\frac{\varphi(\sum x^t)}{\theta\varphi(\sum x^t1)}$$

and

$$\theta V(x^t1) + (1-\theta)V(x^t0) \geq V(x^t1) \iff \xi(x^t0)\varphi(\sum x^t) \geq \xi(x^t1)\varphi(\sum x^t1).$$

Proof. Using definitions we see that under the hypothesis,

$$\xi(x^t)\varphi(\sum x^t) = \theta V(x^t1) + (1-\theta)V(x^t0)$$

$$\iff \xi(x^t)\varphi(\sum x^t) = \theta\xi(x^t1)\varphi(\sum x^t1) + (1-\theta)\xi(x^t0)\varphi(\sum x^t))$$

$$\iff \xi(x^t1) = [\xi(x^t) - (1-\theta)\xi(x^t0)]\frac{\varphi(\sum x^t)}{\theta\varphi(\sum x^t1)},$$

and

$$\theta V(x^t1) + (1-\theta)V(x^t0) \geq V(x^t1)$$

$$\iff \theta\xi(x^t1)\varphi(\sum x^t1) + (1-\theta)\xi(x^t0)\varphi(\sum x^t) \geq \xi(x^t1)\varphi(\sum x^t1)$$

$$\iff \xi(x^t0)\varphi(\sum x^t) \geq \xi(x^t1)\varphi(\sum x^t1).$$

■

Lemma 6 *The equilibrium (α, ρ, ξ) is N-alert with $N < \infty$ iff*

$$N > \varphi^{-1}(1-k) - 1$$

$$\alpha(x^t) = \begin{cases} 0 & \sum x^t < N \\ \alpha \in [0,1) & \sum x^t = N \\ \frac{1-k}{\mu(d|x^t)\varphi(\sum x^t)} & N < \sum x^t < \infty \\ 1 & t = \infty \end{cases},$$

$$\sum x^t = N+1 \implies \xi(x^t) = \frac{\varphi(N)}{\varphi(N+1)},$$

$$\sum x^t \geq N+1 \implies \xi^*(x^t) < \xi^*(x^t0) < \min\{1, \frac{1}{(1-\theta)}\xi(x^t)\}$$

$$\sum x^t \geq N+1 \implies \xi(x^t1) = [\xi(x^t) - (1-\theta)\xi(x^t0)]\frac{\varphi(\sum x^t)}{\theta\varphi(\sum x^t1)} < \xi(x^t),$$

$$\xi(x^\infty) = \begin{cases} 0 & \text{if } \sum x^\infty > N \\ 1 & \text{otherwise} \end{cases},$$

$$\rho < 1,$$

$$\sum x^t > N \implies \rho(x^t) = 0,$$

$$\inf_{x^t:\sum x^t = N+1} \prod_{i=1}^{t}(1 + \frac{(1-\theta)}{\theta}\rho(x^i))^{x_i}(1-\rho(x^i))^{1-x_i} > \frac{1-k}{\varphi(N+1)-(1-k)}\frac{\mu(s)}{\mu(d)}.$$

Proof. We omit the proof of sufficiency. To prove necessity, suppose (α, ρ, ξ) is an alert equilibrium with $N < \infty$.

Step 1: Restrictions on α and N.

It is optimal for the predator to set $\alpha(x^t) = 0$ when $\sum x^t < N$, since $\delta = 1$ and the prey is fully available at $\sum x^t = N$, which is reached almost surely (this is reached almost surely by π and therefore also by the ν induced by ρ). In equilibrium we cannot have $\alpha(x^t) = 1$ at any x^t (Lemma 2), and so $\alpha(x^t) < 1$ at $\sum x^t = N$. But then in order to support $0 < \xi(x^t) < 1$ for all $\sum x^t > N$, given lemma 1(b), the prey must be made indifferent by setting $\alpha(x^t)$ at the unique value for which

$$\mu(d|x^t) = \frac{1-k}{\alpha(x^t)\varphi(\sum x^t)}.$$

Such $\alpha(x^t)$ exists and lies in $(0, 1)$, since $\sum x^t > N \implies \mu(d|x^t) > \frac{1-k}{\varphi(\sum x^t)}$ by Lemma 4(iii). Clearly, $\alpha(x^\infty) = 1$ is a weakly dominant strategy for the predator.

Finally, observe that

$$\mu(d|x^t) > \frac{1-k}{\varphi(N+1)} \implies 1 > \frac{1-k}{\varphi(N+1)} \implies N > \varphi^{-1}(1-k) - 1.$$

Step 2: Restrictions on ξ and ρ for $\sum x^t > N$.

As we saw in Step 1, $\sum x^t > N$ implies $\alpha(x^t) \in (0, 1)$. This is possible only if the predator is indifferent between attacking or waiting. This implies that

$$\sum x^t > N \implies V(x^t) = \xi(x^t)\varphi(\sum x^t).$$

By Lemma 3, it must be that $\rho < 1$ in equilibrium, implying also that

$$\sum x^t > N \implies V(x^t) = \theta V(x^t 1) + (1-\theta)V(x^t 0).$$

By Lemma 5:

$$\xi(x^t 1) = [\xi(x^t) - (1-\theta)\xi(x^t 0)]\frac{\varphi(\sum x^t)}{\theta\varphi(\sum x^t 1)}.$$

The first expression is one of the desired conditions on ξ in this step. Since $\xi(x^t 1) > 0$, this expression implies that it must be that $\xi(x^t) - (1-\theta)\xi(x^t 0) > 0$, which is equivalent to $\frac{1}{(1-\theta)}\xi(x^t) > \xi(x^t 0)$. Since ξ is N-alert, $\xi < 1$ and $\xi(x^t 0) > \xi(x^t)$. Thus, we obtain another desired condition on ξ:

$$\xi^*(x^t) < \xi^*(x^t 0) < \min\{1, \frac{1}{(1-\theta)}\xi(x^t)\}.$$

Note that under these restrictions (and assumption 1),

$$\xi(x^t 1) = [\xi(x^t) - (1-\theta)\xi(x^t 0)]\frac{\varphi(\sum x^t)}{\theta\varphi(\sum x^t 1)} < \theta\xi(x^t)\frac{\varphi(\sum x^t)}{\theta\varphi(\sum x^t 1)} \leq \xi(x^t),$$

verifying that the property $\xi(x^t 1) < \xi(x^t)$ of alert equilibria is satisfied.

The final desired restriction, namely, $\sum x^t = N + 1 \implies \xi(x^t) = \frac{\varphi(N)}{\varphi(N+1)}$, comes from Lemma 4(i) and the fact that $\sum x^t = N \implies \alpha(x^t) < 1$ (Lemma 2).

Step 3: Restriction ρ for $\sum x^t > N$.

To establish that $\sum x^t > N \implies \rho(x^t) = 0$, we show $\xi(x^t0)\varphi(\sum x^t) > \xi(x^t1)\varphi(\sum x^t1)$ for all such x^t. Lemma 5 then implies $\rho(x^t) = 0$. Since ξ is N-alert, we have $\xi(x^t0) > \xi(x^t)$, and consequently the expression $\xi(x^t1) = [\xi(x^t) - (1-\theta)\xi(x^t0)]\frac{\varphi(\sum x^t)}{\theta\varphi(\sum x^t1)}$ from the previous step implies $\xi(x^t1) = [\xi(x^t) - (1-\theta)\xi(x^t0)]\frac{\varphi(\sum x^t)}{\theta\varphi(\sum x^t1)} < [\xi(x^t) - (1-\theta)\xi(x^t)]\frac{\varphi(\sum x^t)}{\theta\varphi(\sum x^t1)} = \xi(x^t)\frac{\varphi(\sum x^t)}{\varphi(\sum x^t1)}$ and thus

$$\xi(x^t1) < \xi(x^t)\frac{\varphi(\sum x^t)}{\varphi(\sum x^t1)}.$$

But then

$$\xi(x^t1) < \xi(x^t)\frac{\varphi(\sum x^t)}{\varphi(\sum x^t1)} \text{ and } \xi(x^t0) > \xi(x^t)$$

$$\implies \xi(x^t0)\varphi(\sum x^t) > \xi(x^t)\varphi(\sum x^t) > \xi(x^t1)\varphi(\sum x^t1).$$

That is, $\xi(x^t0)\varphi(\sum x^t) > \xi(x^t1)\varphi(\sum x^t1)$, as was to be shown.

Step 4: Restriction on ρ for $\sum x^t \leq N$.

Since $V(x^t) = \theta V(x^t1) + (1-\theta)V(x^t0)$ for all x^t (implied by $\rho < 1$ which is required by equilibrium), and since N is reached with probability 1, we have

$$\sum x^t \leq N \implies V(x^t) = \sum_{x^t z^s : \sum x^t z^s = N} \pi(z^s)\xi(x^t z^s)\varphi(N)$$

$$= \varphi(N) \sum_{x^t z^s : \sum x^t z^s = N} \pi(z^s) = \varphi(N).$$

Thus

$$\sum x^t \leq N \implies V(x^t) = \varphi(N).$$

It follows that

$$\sum x^t < N \implies \theta V(x^t1) + (1-\theta)V(x^t0) = V(x^t1) > \varphi(\sum x^t)$$

which implies that the predator strictly prefers to wait and is otherwise indifferent between manipulation and restraint whenever $\sum x^t < N$. Consequently, any ρ is a best response to ξ. In particular, ρ can arise in equilibrium iff it is consistent with ξ.

We need ρ to be such that $\sum x^t > N \implies \xi(x^t) < 1$. By Lemma 4(iii),

$$\sum x^t > N \implies \frac{1-k}{\varphi(\sum x^t)} < \mu(d|x^t)$$

$$\implies \frac{1-k}{\varphi(\sum x^t)} < \frac{\nu(x^t)\mu(d)}{\nu(x^t)\mu(d) + \pi(x^t)\mu(s)}$$

$$\implies \frac{\nu(x^t)}{\pi(x^t)} > \frac{1-k}{\varphi(\sum x^t) - (1-k)} \frac{\mu(s)}{\mu(d)}$$

$$\implies \frac{\prod_{i=1}^t (\theta + (1-\theta)\rho(x^i))^{x_i}((1-\theta)(1-\rho(x^i)))^{1-x_i}}{\pi(x^t)} > \frac{1-k}{\varphi(\sum x^t) - (1-k)} \frac{\mu(s)}{\mu(d)}$$

$$\implies \prod_{i=1}^{t}(1+\frac{(1-\theta)}{\theta}\rho(x^i))^{x_i}(1-\rho(x^i))^{1-x_i} > \frac{1-k}{\varphi(\sum x^t)-(1-k)}\frac{\mu(s)}{\mu(d)}.$$

Since $\sum x^t \geq N+1 \implies \rho(x^t)=0$ (as in step 2) it is sufficient for ρ to satisfy this only for $\sum x^t = N+1$, since then for all continuations z^s,

$$\frac{1-k}{\varphi(\sum x^t z^s)} \leq \frac{1-k}{\varphi(\sum x^t)} < \mu(d|x^t) = \mu(d|x^t z^s).$$

Therefore, ρ must satisfy:

$$\prod_{i=1}^{t}(1+\frac{(1-\theta)}{\theta}\rho(x^i))^{x_i}(1-\rho(x^i))^{1-x_i} > \frac{1-k}{\varphi(N+1)-(1-k)}\frac{\mu(s)}{\mu(d)} \text{ for all } \sum x^t = N+1$$

$$\iff \inf_{x^t:\sum x^t=N+1}\prod_{i=1}^{t}(1+\frac{(1-\theta)}{\theta}\rho(x^i))^{x_i}(1-\rho(x^i))^{1-x_i} > \frac{1-k}{\varphi(N+1)-(1-k)}\frac{\mu(s)}{\mu(d)}.$$

Note that the condition "$N > \varphi^{-1}(1-k)-1$" implies that $\varphi(N+1)-(1-k) > 0$ and thus $\frac{1-k}{\varphi(N+1)-(1-k)}\frac{\mu(s)}{\mu(d)} > 0$.

Step 5: Restriction on $\xi(x^\infty)$.

Consider the set of sequences that consist of x^∞ such that $\sum x^\infty > N$ (this set has ν_ρ-measure 1). For any such x^∞ there is a finite t such that $\sum x^t = N+1$, in which case (by step 2), we must have $\xi(x^t) \in (0,1)$. This is possible only if $\mu(d|x^t) > \frac{1-k}{\varphi(\sum x^t)} > 1-k$. Since $\rho(x^t z^m)=0$ for every continuation z^m of x^t (by step 3), it follows that forest noise is uninformative and so $\mu(d|x^\infty) > 1-k$ and thus $\xi(x^\infty)=0$ for all x^∞ is a best response.

The (ν_ρ-measure 0) set of sequences consists of x^∞ such that $\sum x^\infty \leq N$. For these sequences, $\mu(d|x^\infty) \leq \frac{1-k}{\varphi(N+1)}$ and thus $\xi(x^\infty)=1$. ■

C Proof of Propositions 1, 2 and 3

The proof for Proposition 1(b) is a corollary of lemma 6, where we see that the attack strategy is N-alert equilibrium is not a Moriarty attack strategy. The proof of part (a) is that if $\rho=0$ then for all x^t,

$$\prod_{i=1}^{t}(1+\frac{(1-\theta)}{\theta}\rho(x^i))^{x_i}(1-\rho(x^i))^{1-x_i} = 1.$$

Thus, condition (iii) in the theorem can be satisfied with $\rho=0$ iff $\frac{1-k}{\varphi(N+1)-(1-k)}\frac{\mu(s)}{\mu(d)} < 1$, that is, if the probability of danger is high enough.

The proof of Proposition 2 is as follows: If $\rho(\cdot)$ is constant for all $\sum x^t \leq N$, then

$$\min_{x^t:\sum x^t=N+1}\prod_{i=1}^{t}(1+\frac{(1-\theta)}{\theta}\rho(x^i))^{x_i}(1-\rho(x^i))^{1-x_i}$$

$$= \min_{x^t:\sum x^t=N+1}(1+\frac{(1-\theta)}{\theta}\rho)^{N+1}(1-\rho)^{t-N-1}$$

$$= \frac{(1+\frac{(1-\theta)}{\theta}\rho)^{N+1}}{(1-\rho)^{N+1}} \min_{x^t:\sum x^t = N+1} (1-\rho)^t.$$

This is decreasing in t. But then $\min_{x^t:\sum x^t=N+1}(1-\rho)^t$ can be made arbitrarily small. Since $\frac{1-k}{\varphi(N+1)-(1-k)}\frac{\mu(s)}{\mu(d)} > 0$, condition (iii) in Theorem 1 cannot be satisfied with a constant nonzero ρ.

To prove Proposition 3, assume by way of contradiction there exists an N-alert equilibrium such that for some x^t with $\sum x^t \leq N$ we have $\inf_n \rho(x^t 0^n) > 0$. We show that this implies there exists y^t with $\sum y^t \leq N$ such that $\prod_{i=1}^t (1+\frac{(1-\theta)}{\theta}\rho(y^i))^{y_i}(1-\rho(y^i))^{1-y_i} < \frac{1-k}{\varphi(N+1)-(1-k)}\frac{\mu(s)}{\mu(d)}$, an impossibility according to part (c) of Theorem 1.

To prove this, let $y^{t+n} = x^t 0^n$ for any n and observe that

$$\prod_{i=1}^{t+n}(1+\frac{(1-\theta)}{\theta}\rho(y^i))^{y_i}(1-\rho(y^i))^{1-y_i} = \prod_{i=1}^{t}(1+\frac{(1-\theta)}{\theta}\rho(x^i))^{x_i}(1-\rho(x^i))^{1-x_i}\prod_{j=1}^{n}(1-\rho(x^t 0^j)).$$

By part (a) of Theorem 1 $\rho(x^t 0^j) \in [0,1)$ for all j. Moreover, by hypothesis, $\rho(x^t 0^j) \in (0,1)$ for infinitely many j. Hence, $\prod_{j=1}^n (1-\rho(x^t 0^j)) \leq (1-\inf_j \rho(x^t 0^j))^n \to 0$ since $\inf_n \rho(x^t 0^n) > 0$. Therefore there exists some n for which $y^{t+n} = x^t 0^n$ is such that $\prod_{i=1}^t(1+\frac{(1-\theta)}{\theta}\rho(y^i))^{y_i}(1-\rho(y^i))^{1-y_i} < \frac{1-k}{\varphi(N+1)-(1-k)}\frac{\mu(s)}{\mu(d)}$, a contradiction. Thus

$$\inf_n \rho(x^t 0^n) = 0.$$

Given that $\nu_\rho(x^t) = \prod_{i=1}^t (\theta + (1-\theta)\rho(x^i))^{x_i}((1-\theta)(1-\rho(x^i)))^{1-x_i}$ for any x^t, we see that

$$\nu(1|x^t 0^n) = \frac{\nu_\rho(x^t 0^n 1)}{\nu_\rho(x^t 0^n)} = \theta + (1-\theta)\rho(x^t 0^n 1).$$

Therefore $\inf_n \nu(1|x^t 0^n) = \theta + (1-\theta)\inf_n \rho(x^i) = \theta$, completing the proof of the proposition.

D Proof of Propositions 6 and 7

Proposition 6 is established in the following three lemmas and Proposition 7 is established in the final lemma.

Lemma 7 *ρ is more swift than any ρ' iff*

$$\mu(d|x^t) \leq 1-k \implies \rho(x^t) = 0$$

$$\mu(d|x^t) > 1-k \implies \rho(x^t) \geq 1 - \frac{(1-k)\mu(s|x^t)}{k\mu(d|x^t)}.$$

Proof. Consider ρ as in the statement of the lemma.

Step 1. $\mu(d) \leq 1-k$ implies $\mu(d|x^t) \leq 1-k$ for all x^t.

This follows from the fact that $\rho = 0$ whenever beliefs are below $1-k$.

Step 2. $x^t \neq 1^t \implies \mu(d|x^t 0) \leq 1-k$.

If the prior satisfies $\mu(d) \leq 1-k$ then this follows from Step 1. If $\mu(d) > 1-k$ then we show that a single instance of silence pushes below $1-k$. In particular we show that $\mu(d|x^t 0) \leq 1-k$ for any x^t. Observe that for any x^t,

$$\mu(d|x^t 0) \leq 1-k$$
$$\iff \frac{(1-\theta)(1-\rho(x^t))\mu(d|x^t)}{(1-\theta)(1-\rho(x^t))\mu(d|x^t)+(1-\theta)\mu(s|x^t)} \leq 1-k$$
$$\iff \rho(x^t) \geq 1-\frac{(1-k)\mu(s|x^t)}{k\mu(d|x^t)},$$

as desired.

Step 3 Prove the result

Take any ρ'. Take any x^t and suppose $\mu'(d|x^t) \leq 1-k$. Consider two cases involving the prior belief:

case i: $\mu(d) \leq 1-k$.

Then $\mu(d|x^t) \leq 1-k$ by Step 1.

case ii: $\mu(d) > 1-k$.

Since noise increases the posterior on danger, we note that if $x^t = 1^t$ then $\mu'(d|x^t) \geq \mu'(d) > 1-k$, contradicting $\mu'(d|x^t) \leq 1-k$. Therefore it must be that $x^t_i = 0$ for some $i \leq t$. But then by Step 2 we have $\mu(d|x^t) \leq 1-k$, as desired. ■

Lemma 8 *ρ is swift iff*

$$\rho(x^t) = \begin{cases} 1-\frac{(1-k)\mu(s|x^t)}{k\mu(d|x^t)} & \text{if } \mu(d|x^t) > 1-k \\ 0 & \text{otherwise} \end{cases}.$$

Proof. Given the characterization of the swiftest equilibria in the previous lemma, we see that the the swiftest equilibria with the smallest ρ must be as desired. ■

Lemma 9 *There exist a unique inattentive equilibria where ρ is swift.*

Proof. Lemma 12 characterizes the unique equilibrium involving a swift ρ. Since there is only one swift strategy, it follows that there is only inattentive equilibrium with a swift strategy. ■

Lemma 10 *Suppose $\mu(d) > 1-k$ and ρ is swift. Then $\nu(1|1^n)$ is strictly increasing in n and $\lim_{n\to\infty}\nu(1|1^n) = 1$.*

Proof. Since π has full support, we have $\mu(s|1^n) \in (0,1)$ for all n. By the definition of ρ and since $\mu(d) > 1-k$, we also have $\rho(1^n) > 0$ for all n. But $\rho(1^n) > 0$ for all n then implies that for any history 1^n, further noise must strictly decrease the belief in safety: $\mu(s|1^n) > \mu(s|1^n 1)$. Thus

$$\rho(1^n 1) = 1-\frac{(1-k)\mu(s|1^n 1)}{k\mu(d|1^n 1)} > 1-\frac{(1-k)\mu(s|1^n)}{k\mu(d|1^n)} = \rho(1^n).$$

Thus $\rho(1^n 1) > \rho(1^n)$ for all n. Since $\lim_{n\to\infty}\mu(s|1^n) = 0$ we also have that $\lim_{n\to\infty}\rho(1^n) = 1$. We obtain the desired conclusions once we recall that $\nu(1|1^n) = \theta + (1-\theta)\rho(1^n 1)$ (as in the proof of Proposition 3). ■

E Proof of Theorem 2

Theorem 2 is established by the following two lemmas.

Lemma 11 *A strategy ρ can be supported by an inattentive equilibrium iff ρ is such that $\rho < 1$ and for all x^∞*

$$\lim_{t\to\infty} \mu(d|x^t) \leq 1-k,\ \nu\text{-a.s.}$$

where ν is the measure on X^∞ induced by ρ.

Proof. Sufficiency: Given the ν, by the Martingale Convergence Theorem, $\mu(d|X^t)$ converges to some random variable with support $[0,1]$. For the predator it is then optimal to attack only at ∞ if $\xi(x^\infty) = 1$, ν-a.s.Therefore, the predator is indifferent between all ρ that induces $\xi(x^\infty) = 1$, ν-a.s. Since $\alpha = 0$ for all finite sequences, $\xi(x^t) = 1$ is a best response for the prey for all $t < \infty$. But $\xi(x^\infty) = 1$ ν-a.s. requires

$$\lim_{t\to\infty} \mu(d|x^t) \leq 1-k,\ \nu\text{-a.s.}$$

Necessity: if that condition holds, then $\xi(x^\infty) = 1$ ν-a.s., and then optimal to attack only at ∞. Hence, it is optimal for the predator to set $\alpha(x^t) = 0$ for all finite $t < \infty$ and for prey to set $\xi(x^t) = 1$ for all $t \leq \infty$. ∎

The next lemma uses swift manipulation strategies (see Definition 5 and Proposition 6).

Lemma 12 *There exists an inattentive equilibrium (ρ, α, ξ) where ρ is a swift manipulation strategy. In this equilibrium,*

$$\alpha(x^t) = 1 - \xi(x^t) = 0 \textit{ if } t < \infty$$
$$\alpha(x^\infty) = 1$$
$$\xi(x^\infty) = 1 \textit{ if } x^\infty \neq 1^\infty$$
$$\xi(x^\infty) = 0 \textit{ if } x^\infty = 1^\infty.$$

Proof. The predator wishes to induce the prey to be inattentive at any x^∞ with $\sum x^\infty = \infty$. This can be ensured by the following procedure. Define $\rho(x^t)$ inductively for all $x^t \in X \cup \{\emptyset\}$ by

$$\rho(x^t) = \begin{cases} 1 - \frac{(1-k)\mu(s|x^t)}{k\mu(d|x^t)} & \text{if } \mu(d|x^t) > 1-k \\ 0 & \text{otherwise} \end{cases}.$$

If the prior is low, $\mu(d) \leq 1-k$, then the strategy involves no manipulation, leading to limiting beliefs that equal the prior. If the prior is high, $\mu(d) > 1-k$, then the predator manipulates with "high" probability. Intuitively, at any x^t, if $\mu(d|x^t) > 1-k$ then $\rho(x^t)$ is such that the posterior after a silence satisfies $\mu(d|x^t0) = 1-k$.[11] That is, ρ is such that silence in the next period drives beliefs down to $1-k$. Once this happens, then the predator never manipulates.

[11] To verify this observe that:

$$\mu(d|x^t 0) = 1-k$$
$$\iff \frac{(1-\theta)(1-\rho(x^t))\mu(d|x^t)}{(1-\theta)(1-\rho(x^t))\mu(d|x^t) + (1-\theta)\mu(s|x^t)} = 1-k$$
$$\iff \rho(x^t) = 1 - \frac{(1-k)\mu(s|x^t)}{k\mu(d|x^t)}.$$

By the martingale convergence theorem, given ρ, it must be that posteriors converge ν_ρ-almost surely, that is, $\lim_{t\to\infty}\mu(d|x^t)$ exists ν_ρ-a.s. Since $\nu_\rho(\{x^\infty : x^\infty = 1^\infty \text{ or } \sum x^\infty < \infty\}) = 0$, it follows that ν_ρ-almost surely, a silence appears and infinitely many noises are made along x^∞. Since a silence must appear ν_ρ-almost surely, with such a manipulation procedure, the predator is assured that $\mu(d|x^t) \leq 1-k$ hold eventually for large t almost surely. Indeed, $\mu(d|x^\infty) \leq 1-k$ ν_ρ-a.s. Since $\sum x^\infty = \infty$ ν_ρ-a.s., the predator is thus able to capture the prey with probability $\varphi(\infty) = 1$.

Given the preceding we therefore see that ρ as defined above and the following strategies constitute an equilibrium:

$$\alpha(x^t) = 1 - \xi(x^t) = 0 \text{ if } t < \infty$$

$$\alpha(x^\infty) = 1$$

$$\xi(x^\infty) = 1 \text{ if } x^\infty \neq 1^\infty$$

$$\xi(x^\infty) = 0 \text{ if } x^\infty = 1^\infty.$$

An implication of the result is that the predator exhibits a hot hand effect: manipulation is positive as long as there is a lot of noise, but stops if there are a lot of zeros. ■

F Proof of Propositions 4 and 5

The proof of Proposition 4 is as follows. If $\mu(d) \leq 1-k$ then a swift manipulation strategy requires $\rho = 0$. But then Lemma 12 establishes the existence of an inattentive Moriarty equilibrium in this case. If on the other hand we have $\mu(d) > 1-k$ then suppose by way of contradiction that an inattentive Moriarty equilibrium exists. Since $\rho = 0$, then posterior beliefs equal the prior μ, that is, $\mu(d|x^t) = \mu(d)$ for all x^t. But then $\lim_{t\to\infty}\mu(d|x^t) = \mu(d) > 1-k$, ν-a.s. and by Lemma 11 we are not in an inattentive equilibrium, a contradiction.

To prove Proposition 5, suppose $\rho(\cdot)$ is constant at some $\overline{\rho} \in [0,1)$. If $\overline{\rho} = 0$ then Proposition 4 confirms the possibility of an inattentive equilibrium, as desired. To see that an equilibrium does not exist otherwise, suppose that $\overline{\rho} \in (0,1)$. Then ν is i.i.d. with bias $\theta + (1-\theta)\overline{\rho}$. By standard Bayesian consistency results, $\lim_{t\to\infty}\mu(d|x^t) = 1$, ν-a.s. But since $k \in (0,1)$, we see that $\lim_{t\to\infty}\mu(d|x^t) > 1-k$, ν-a.s. and so by Lemma 11 there does not exist an inattentive equilibrium that supports $\overline{\rho}$, establishing the contradiction.

References

[1] Anderson, A. and L. Smith (2013): "Dynamic Deception," *American Economic Review*, 103(7), 2811–2847.

[2] Benjamin, D. (2019). Errors in Probabilistic Reasoning and Judgment Biases. *Handbook of Behavioral Economics* (edited by Doug Bernheim, Stefano DellaVigna, and David Laibson). Elsevier Press.

[3] Crawford, V. P., & Sobel, J. (1982). Strategic Information Transmission. *Econometrica,* 50(6), 1431–1451. https://doi.org/10.2307/1913390

[4] Frankel, A. and N. Kartik (2019): "Muddled Information," Journal of Political Economy, 127(4), 1739–1796.

[5] Hörner, J. and A. Skrzypacz (2016): "Learning, Experimentation, and Information Design," in Advances in Economics and Econometrics: Theory and Applications, 63–98.

[6] Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6), 2590–2615.

[7] Mailath, G. & Samuelson, L. (2015). Reputations in Repeated Games. *Handbook of Game Theory with Economic Applications* (edited by H. Peyton Young and Shmuel Zamir), Elsevier, Volume 4 , Pages 165-238, https://doi.org/10.1016/B978-0-444-53766-9.00004-5.

[8] Noor, J. & Payro, F. (2025). An Axiomatic Model of the Law of Small Numbers. mimeo.

[9] Rabin, M. and J Schrag (1999). A Model of Confirmatory Bias. *The Quarterly Journal of Economics*, 114(1), 37-82. https://www.jstor.org/stable/2586947