Opportunity Hunters: A Model of Competitive Sequential Inspections*

Ran Eilat† Zvika Neeman‡ Eilon Solan§

December, 2023

Abstract

We introduce a new type of games, called "opportunity-hunting games," in which two players compete to discover an uncertain event ("opportunity") that occurs at an unobserved and random point in time. Players can inspect whether the event has already occurred again and again, but each inspection is costly. Varying the parameters of the model spans the range from games where competition between the players to be the first to identify the opportunity is the dominant force, to games in which free-riding on the other player's effort is the dominant force. We characterize the game's unique symmetric Markov Perfect Equilibrium.

Keywords: opportunities, sequential inspection, opportunity-hunting games.

JEL Classification Codes: C72, C73, D83.

*We thank Daniel Bird, Yeon-Koo Che, Laura Doval, Kfir Eliaz, Johannes Hörner, John Levy, Elliot Lipnowski, Qingmin Liu, Ady Pauzner, Alessandro Pavan, Jacopo Perego, Andrea Prat, Steven Schochet and seminar participants in Columbia University, Tel-Aviv University, Bar-Ilan University, the University of Glasgow, the 10th Israeli IO Workshop and the 24th ACM Conference on Economics and Computation (EC'23) for helpful comments. Eilat and Neeman acknowledge the support of the Israel Science Foundation, Grant #1792/23. Solan acknowledges the support of the Israel Science Foundation, Grant #211/22.

†Department of Economics, Ben-Gurion University of the Negev, eilatr@bgu.ac.il

‡School of Economics, Tel Aviv University, zvika@tauex.tau.ac.il

§The School of Mathematical Sciences, Tel Aviv University, eilons@tauex.tau.ac.il

1 Introduction

According to the Oxford Dictionary, an *opportunity* is a set of circumstances that makes it possible to do something. Identifying an opportunity, however, is oftentimes difficult. Indeed, opportunities do not make appointments – the exact moment in time in which they appear is uncertain, identifying them is costly, and the window of time to act upon them is narrow. When multiple players are interested in the same opportunity, a competition to find it emerges.

There are many economic situations in which players compete to identify an opportunity and seize it quickly. For example, consider two firms vying to capture a certain market. The firms can capture the market only if it is ripe, and checking for this condition incurs an unrecoverable cost (say, advertising costs), which is paid by the firm who moves first. If the inspection reveals that the market is not ripe, then the firms may inspect again, and again, at later dates.

The dynamics of such competition are influenced by the characteristics of the market in question. In some markets, being the first mover enhances the chances of successfully capturing the ripe market. This advantage is evident in industries like technology and fashion, where the ability to introduce innovative products or set new trends can lead to significant market share. On the other hand, there are markets where the second mover has the upper hand. In these cases, the second firm can learn from the experience of the first mover, enabling it to improve upon the initial attempt. A notable example is the retail sector, where Walmart, as a follower, managed to offer comparable products at lower prices, ultimately becoming a dominant force in the industry.

The advantage of the first- or second-mover, as well as the cost associated with inspecting the market's ripeness, play crucial roles in shaping the nature of the competition between the two firms. Understanding these dynamics is essential for devising optimal strategies in scenarios where firms compete to seize opportunities swiftly while considering the market conditions and associated costs.

Other natural contexts in which a competition for opportunities emerges include innovation, patenting, launching of new products, asset sales, headhunting and dating. What these examples all have in common is that the optimal timing of action depends on identifying an unobservable event while considering the strategy of the other player. Additionally, players may act more than once in their attempt to identify this event. These two features distinguish the hunting for opportunities and add an important new dimension to what has been learned from the literature on strategic timing games such as "preemption games" and "war of attrition games." What characterizes the competition for opportunities? what do its equilibria look like? What are the lessons for the design of the search for opportunities?

MODEL AND OUTLINE OF RESULTS. To address these questions we introduce the following *opportunity-hunting* game. An event ("an opportunity") occurs at some unobserved and random point in time. There are two players who wish to discover the event ("seize the opportunity") once it occurs. At each point in time each player can inspect whether the event has already occurred at a cost of $c > 0$ per inspection. The inspections are publicly observable. If a player discovers that the event has already occurred, she obtains the payoff v_1, the other player obtains the payoff v_2, and the game ends; otherwise, the players can inspect again, and again, later in time. The player who inspects at t is referred to as the "leader" at t; the other player is referred to as the "follower" at t. If the prize is never found, then the payoff to both players is zero (minus inspection costs, if any). If the two players inspect simultaneously, then they are randomly assigned to be leader and follower. By varying the parameters of the model, it is possible to span the range from games in which the competition between the players to identify the opportunity is the dominant force, to games in which free-riding on the other player's effort is the dominant force.

We characterize the symmetric Markov Perfect Equilibrium (MPE) of the game and show that it is unique. We show that depending on one parameter – the follower prize v_2 – equilibrium takes one of two forms. When the follower prize v_2 is smaller than some threshold value, the equilibrium involves frequent synchronized inspections by the players. Intuitively, when the benefit to the follower is small, both players compete over being the first to discover the prize. This competition pushes them to inspect as early as possible, but not so early that the probability that the prize has already appeared is too small. Consequently, the two players synchronize their inspections at the earliest mutually beneficial time. We show that the players' discount rate does not affect the rate of inspection in this type of equilibrium.

By contrast, when the benefit to the follower is large, the incentive of players to "free-ride" on each other's inspections is more intense. In this case, the players cannot synchronize their inspections in equilibrium, because each player always prefers to slightly delay her inspection so she benefits from the other player's inspection without incurring the associated cost. It follows that in this case the unique symmetric MPE exhibits diffusion in which the players inspect randomly at different times. Notably, in this equilibrium, changing the value of v_2 has no effect on the players' equilibrium payoffs. In particular, the rent that is generated by a higher value of v_2 is completely dissipated. This implies that from a designer's perspective, it is better to concentrate efforts on decreasing the cost of inspection c rather than increasing the reward v_2.

CONTRIBUTION. This paper makes three main contributions. Firstly, it introduces opportunity hunting games that capture the strategic forces inherent in many real world situations, and offers a characterization of their symmetric MPEs. These MPEs are unique and take on a simple form.

Secondly, the model offers a unified framework that integrates two qualitatively different and well-known types of competition: one where preemption is the primary motivation and another where free-riding dominates. One parameter, the follower's prize v_2, determines which type of competition prevails. These two modes of competition lead to the two previously mentioned equilibria. Thus, from an applied perspective, our model provides a simple unified framework that allows economists to study markets or industries in which the type of competition is unknown in advance or changes dynamically over time.

Lastly, from a methodological perspective, our work also offers three more specialized contributions. First, we provide useful definitions of histories and strategies in continuous time games. These definitions, which are natural in our context, can also be applied in other continuous time games. Second, the characterization of equilibrium hinges on a condition that requires players to be indifferent among all the inspection times in the support of the equilibrium distribution, which is expressed as a functional equation. While this in itself is not new, in our model the functional equation that characterizes the equilibrium also depends on the equilibrium payoff itself.[1] Our solution method demonstrates how to find the equilibrium despite this endogeneity. Third, we show how a result from the Theory of Distributions (Hörmann and Steinbauer (2009)) can be used to prove that the equilibrium is unique. Specifically, for the case where v_2 is large, this result implies that if our functional equation has a continuous solution, then this solution is unique.

Related Literature

The game we consider falls into the general category of stochastic games with a partially observable state. See, e.g., Davis and Varaiya (1973) and Hansen, Bernstein and Zilberstein (2004). Such games are usually intractable, and have been studied only in the case of two-players, often with zero-sum payoffs. Papers in this literature are mostly concerned with proving the existence of a value or equilibrium for the game and not in the qualitative properties of equilibrium or in comparative statics, which is our focus here.

Numerous studies have examined war of attrition games and preemption games as sim-

[1] In contrast, in war of attrition games the equation that characterizes the equilibrium strategy describes an indifference condition that does not directly involve the players' equilibrium payoffs. See, e.g., Hendricks, Weiss and Wilson (1988).

ple timing (or stopping) games, where players must choose the optimal moment to act.[2] The "classical" war of attrition game studied by Hendricks, Weiss and Wilson (1988) involves two players engaged in a cost-intensive struggle, with the last player standing declared the winner. The symmetric Nash equilibrium in this game entails a mixture of quitting times. Fudenberg and Tirole (1985) analyze a preemption game where firms compete to adopt a new technology. The first adopter benefits, but early adoption comes at a higher cost. Fudenberg and Tirole (1985) characterize the equilibria of the game, demonstrating that firms either adopt at different deterministic times or simultaneously at one of many possible times in some interval, depending on the parameters. More recent papers in this literature include Argenziano and Schmidt-Dengler (2014), Anderson, Smith and Park (2017), and Smirnov and Wait (2022).

Subsequent works in this literature have focused on the incorporation of private information into one-shot timing games, in diverse contexts. Notable examples include Hopenhayn and Squintani (2011), Murto and Valimaki (2011), Awaya and Krishna (2021), Bobtcheff, Levy and Mariotti (2021), Shahanaghi (2022), and Cetemen and Margaria (2023). In contrast to this line of research, our extension of the basic model introduces the possibility of repeated actions.

Our work is also related to the substantial body of literature on experimentation. Papers in this literature study games in which players who face uncertainty about the value of different choices engage in strategic learning over time to reduce this uncertainty. In some of these papers the focus is on the rate of experimentation. See, e.g., Bonatti and Hörner (2011) and Décamps and Mariotti (2004). In other papers, experimentation is modeled as a multiple-armed bandit problem. See, e.g., Keller, Rady and Cripps (2005), Rosenberg, Solan and Vieille (2007), Rosenberg, Salomon and Vieille (2013), and Hörner, Klein and Rady (2022). Unlike in our paper, where the focus is on *when* to next check whether the event has already occurred, in this literature, the focus is either on how much to experiment or on which experiment to conduct next (i.e., which arm to play). Accordingly, papers in this literature have studied the interplay between private and public learning, and the implications of the possibility of free-riding on learning by others. In a recent paper, Frick and Ishii (2023) study a model in which a continuum of agents receive stochastic opportunities, and have to decide whether to adopt an innovation or wait and learn from others' experience. They show how equilibrium behavior depends on the features of the learning environment.

Our model is also related to the literature on contests and the design of prize schemes.

[2]When the number of times players can act is bounded, a timing game can be viewed as a sequence of standard timing games, where players can act only once. See, e.g., Laraki and Solan (2005). In our case, the number of possible actions for each player is potentially infinite.

See, Glazer and Hassin (1988), Moldovanu and Sela (2001), Che and Gale (2003), and Moldovanu, Sela and Shi (2007). However, this literature has a different focus and the games studied in this literature typically have a one-dimensional strategy space.

Finally, Ball and Knopfle (2023) study optimal inspection policies in a dynamic Principal-Agent model with partial observability, and derive conditions for the optimality of deterministic vs. random inspections.

ORGANIZATION. The rest of the paper is organized as follows. In Section 2 we present the model. In Section 3 we analyze the one-player problem, which serves as a benchmark for the rest of the analysis. The two-player game is presented and analyzed in Section 4, and its symmetric Markov Perfect Equilibria are presented in Section 5. In Section 6 we discuss alternative modelling assumptions and equilibrium concepts. Section 7 concludes. All proofs are relegated to the Appendix.

2 Model

We consider a continuous-time two-player game, which we call an *opportunity-hunting* game. A prize appears at a random time according to an exponential distribution with parameter $\lambda > 0$. That is, for every time $t \geq 0$, the probability that the prize has appeared by time t is $1 - e^{-\lambda t}$. After the prize appears, it remains hidden from the players until they actively search for it. The first player to discover the prize obtains a payoff $v_1 > 0$; the other player obtains a payoff $v_2 \geq 0$.[3,4] Both players discount future payoffs at the rate r (i.e., a unit prize at t generates a current payoff e^{-rt}). If the prize is never found, then the payoff to both players is zero (minus inspection costs, if any).

At each point in time each player can check whether the prize has already appeared. A player who inspects at time t is referred to as the *leader* at t. The other player is referred to as the *follower* at t. The players' inspection times are commonly observed.[5] The cost of checking is $c > 0$ per inspection. We assume that the cost of inspection is smaller than the

[3]The assumption that v_2 is non-negative ensures that equilibrium payoffs are non-negative, which is important for some of the results presented below. There are, of course, interesting applications where $v_2 < 0$. The analysis of such cases is different and falls outside the scope of this paper.

[4]We thus implicitly assume that the player who discovers the prize immediately "claims" it. As explained in Section 6.3 below, this assumption is not crucial for our results.

[5]On the one hand, this assumption simplifies the analysis because it implies that the players are symmetrically informed. On the other hand, it complicates the analysis because it expands the strategy space. This assumption implies that we consider "closed-loop" equilibria (see, e.g., Fudenberg and Tirole (1991), p. 130). In such equilibria, each player can observe, and therefore also respond to, deviations from equilibrium of other players. By contrast, an "open-loop" equilibrium is an equilibrium where deviations by other players are unobserved. In Section 6.3 we discuss the implications of assuming that inspection times are unobserved.

leader's prize, i.e., $c < v_1$.[6] If, upon inspection, a player discovers the prize, then the game ends, and the two players' payoffs are as described above. Otherwise, the game continues and players can inspect again and again. If the two players inspect at the exact same time, then they are randomly assigned to be leader (who incurs the cost of inspection and obtains v_1 if she discovers the prize) or follower (who does not incur the cost of inspection, and obtains the payoff v_2 if the prize is discovered).[7]

In the subsequent analysis, our primary focus will be on symmetric Markov perfect equilibria of the game, which we formally define in Section 2.2. In these equilibria, both players employ Markov strategies, which possess a simple structure, as explained below. Notice however that Markov perfect equilibria have to be robust against deviations to *any* strategy, not just Markov strategies. Thus, to characterize the equilibrium we have to present a comprehensive definition of a general strategy in the game.

Because time is continuous, a strategy cannot simply be a function from time and history into actions. The problem lies in the difficulty of mapping such strategies into distributions of game plays.[8] Therefore, we adopt the following approach: A strategy instructs a player when to inspect next, as a function of the times in which past inspections took place and the identity of the player who inspected at each time. In principle, the decision of when to inspect next also depends on the outcomes of past inspections. However, since the game ends once the prize is discovered, we can assume that all prior inspections have failed to find the prize. Additionally, a player may choose not to inspect at all, which we denote as an inspection at time $t = \infty$.

In order to correctly define a strategy in this way, we need to address a few technical subtleties, which we take care of in the next section (where we also discuss the different approaches taken by the literature). Readers who wish to skip the technical details can proceed directly to Section 2.2 without any disruption to overall coherence.

[6] The case in which $c \geq v_1$ is trivial because the players would never inspect.

[7] In Section 6.3 we discuss the implications of the alternative assumption that when both players inspect simultaneously, they both incur the cost of inspection.

[8] A standard example that illustrates this difficulty is a strategy in which in each point in time, a player inspects with probability one-half, independently of past play. It is impossible to embed such a sequence in a well defined probability space (see also Judd (1985)). Another example is the following strategy (see, Kamada and Rao (2023)): Inspect at time t if no player inspected until t *and* $t = 1/n$ for some natural number n; otherwise, do not inspect. It can be shown that there is no play of the game that is consistent with both players employing this strategy.

2.1 On the Definition of Strategies

In this section, we formally define the notions of history, strategy, and play. The challenge is to define these notions in a way that is, on the one hand, both general and natural, and on the other hand, ensures that a pair of strategies induces a unique distribution over the set of plays. This ensures that the expected payoffs associated with any pair of strategies are well defined.

In principle, a player may inspect as often as she likes. For example, a player may inspect at times $\frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \frac{4}{5}, \ldots, 1\frac{1}{2}, 1\frac{2}{3}, 1\frac{3}{4}, 1\frac{4}{5}, \ldots$. Such an inspection policy is likely to be suboptimal, but a definition of a strategy must accommodate it.

We begin by defining the notion of history.

Definition (History). *A* history *is a 3-tuple $\langle \alpha^*, \{t_\alpha\}_{\alpha\leq\alpha^*}, \iota\rangle$ where: (i) α^* is a countable ordinal;*[9] *(ii) $\{t_\alpha\}_{\alpha\leq\alpha^*}$ is an increasing sequence of non-negative real numbers, such that $t_\alpha = \lim_{\beta<\alpha} t_\beta$ for every limit ordinal $\alpha \leq \alpha^*$; and (iii) ι is a function that assigns two nonempty sets of players, ι_a and a singleton set ι_i, to every successor ordinal $\alpha \leq \alpha^*$.*

The interpretation of a history is as follows. The players' inspection times before t_{α^*} are given by

$$\{t_\alpha : \alpha \leq \alpha^*, \alpha \text{ is a successor ordinal}\}.$$

For each time t_α in this set, the players in $\iota_a(\alpha)$ are the players who attempted an inspection at that time, and the player in the singleton set $\iota_i(\alpha)$ is the player who actually inspected at that time.[10,11]

Example. To illustrate the definition above, consider a history in which a player inspects at times $\frac{1}{2}, \frac{2}{3}, \frac{3}{4}, \ldots$, and then at times $1\frac{1}{2}, 1\frac{2}{3}, 1\frac{3}{4}, \ldots$. To describe this history, we set $\alpha^* = 2\omega$. Then, $t_0 = 0, t_1 = \frac{1}{2}$, $t_2 = \frac{2}{3}$, $t_3 = \frac{3}{4}, \ldots$, $t_\omega = 1$, $t_{\omega+1} = 1\frac{1}{2}$, $t_{\omega+2} = 1\frac{2}{3}$, $t_{\omega+3} = 1\frac{3}{4}, \ldots$, and $t_{2\omega} = 2$. In this description, $0, \omega$ and 2ω are limit ordinals. All other ordinals, namely $1, 2, 3, \ldots, \omega+1, \omega+2, \omega+3, \ldots$, are successor ordinals, each associated with a specific inspection time. In this example, since only one player inspects during this history, $\iota_a = \iota_i$ is the singleton set containing this player for every successor ordinal. ■

[9] Ordinal numbers are linearly ordered labels that include the natural numbers and have the property that every set of ordinals has a least element. This facilitates the definition of an ordinal number ω that is greater than every natural number, along with ordinal numbers $\omega+1$, $\omega+2, \ldots$, which are greater than ω. Similarly, the ordinal number 2ω is greater than ω plus every natural number, and so forth.

[10] By definition, no player inspects at times t_α for *limit* ordinals α. It is possible to also consider a more general set of strategies where players may also inspect at times that are associated with limit ordinals. Doing so would not affect our results.

[11] The requirement that α^* is countable is without loss of generality, because between each t_α and $t_{\alpha+1}$ there is at least one rational number, and the set of rational numbers is countable.

Next, we define the notion of strategy. A strategy is a measurable function from histories to distributions over future inspection times. For measurability to be well defined, we endow both of these spaces with appropriate sigma-algebras. The sigma-algebra over the space of distributions on $[0,\infty]$ is the weak-$*$ topology. The sigma-algebra over the space of histories is the one inherited from $\cup_\alpha \mathbb{R}^\alpha$, where the union is over all countable ordinals. Formally,

Definition (Strategy). *A* strategy σ_i *of player* i *is a measurable function that assigns to every history* $h = \langle \alpha^*, \{t_\alpha\}_{\alpha \le \alpha^*}, \iota \rangle$ *a probability distribution over* $(t_{\alpha^*}, \infty]$*. An atom of* $\sigma_i(h)$ *at* ∞ *is interpreted as if the player assigns positive probability to never inspecting.*

A *play* is an object that indicates how the game was played from beginning to end. In our model, the definition of play coincides with that of history, with one difference: when $\langle \alpha^*, \{t_\alpha\}_{\alpha \le \alpha^*}, \iota \rangle$ is a history, $\{t_\alpha\}_{\alpha \le \alpha^*}$ is a collection of non-negative real numbers. In a play, we allow t_{α^*} to be equal to ∞. This is because along the play the players may continue inspecting ad infinitum and never discover the prize. We endow the space of plays with the sigma-algebra inherited from $\cup_\alpha (\mathbb{R} \cup \{\infty\})^\alpha$, where the union is over all countable ordinals.

The next result ensures that every pair of strategies induces a well defined distribution over plays. This implies that the players' expected payoffs from any pair of strategies are well defined.

Lemma 1. *Every pair of strategies* (σ_1, σ_2) *induces a unique probability distribution over plays.*

NOTE ON THE LITERATURE. The literature on continuous-time games has proposed various approaches for modeling histories and strategies. In the context of our model, our approach offers several advantages compared to the existing body of work.

For example, in the literature on optimal control (e.g., Cardaliaguet (2007)) a player's strategy specifies a grid of time points. Starting at time zero, and for any time point on the grid, the player "commits to" a plan of action, which may depend on the history of play, to be played until the next time point on the grid. However, applying this method in our model would only permit the characterization of ε-equilibria.

Other papers (e.g., Bergin and MacLeod (1993)) have employed the concept of "inertia" in a strategy. Specifically, once an action is played, the player cannot switch to another action within a small interval of time.[12] Papers in this literature study the limit of equilibria as the length of the inertia interval decreases to zero. In contrast, our approach does not rely on the implicit assumption that a player can act only finitely many times in any finite time

[12] Alternatively, a player cannot repeat the action within the mentioned time interval.

interval, and does not need to address the issue of whether the equilibrium correspondence is continuous as the length of the inertia interval vanishes.[13]

Stinchcombe (1992) proposes a definition for strategy that, similar to ours, employs the concept of ordinals. A key distinction between our definition and Stinchcombe's (1992) is that, in our framework, actions are instantaneous and occur at times associated with each ordinal. In contrast, in Stinchcombe (1992), a player maintains the same action between any two times associated with consecutive ordinals. This distinction makes our definition more useful for games in which the payoff from actions is "discrete" (e.g., a cost of inspection).

Finally, Fudenberg and Tirole (1985) analyze the limit of a sequence of games that are played on a discrete time grid. However, in those cases where solving the game in continuous time is simpler, as in our case, our approach is more straightforward.

2.2 Markov Perfect Equilibrium

Suppose that the last inspection has occurred at time t. Since the prize appears according to an exponential distribution,[14] which is memoryless, and since payoffs are exponentially discounted, if an inspection at time t fails to discover the prize, then the subgame that starts at time $t+$ (i.e., immediately after t) is equivalent to the whole game. Because inspection times are observable, the players' beliefs coincide after every possible history of the game.

In this type of environments, it is natural for players to employ *Markov strategies*. Generally, Markov strategies are strategies that are measurable with respect to payoff-relevant histories (see, e.g., Fudenberg and Tirole (1991)). In our model, this implies that Markov strategies may only depend on the time of the last (failed) inspection.

Definition (Markov Strategy). *A Markov strategy is a cumulative probability distribution function (CDF) F over $(0,\infty]$, with $F(0)=0$.*

The interpretation of a Markov strategy is that a player selects the first inspection time according to a distribution F; and for each time $t>0$ at which an inspection fails to discover the prize, the process of inspection "restarts" at t, and the player selects again a new inspection time according to the distribution F shifted to $(t,\infty]$. The requirement that $F(0)=0$ reflects the fact that a player cannot inspect at the exact same time as the previous inspection. The difference $1-\lim_{t\to\infty}F(t)$ is interpreted as the probability of never inspecting.

[13] Kamada and Rao (2023) study an environment in which focusing on standard "inertial" strategies may be inadequate. In their framework, the action of each player can depend on the realization of a stochastic process and on the history of play at each point in time. They identify restrictions on strategies and histories that guarantee that each strategy profile induces a unique path of play.

[14] Specifically, for every time $t'>t$, the probability that the players assign to the event that the prize has already appeared by time t' is determined by Bayes' rule and is given by $1-e^{-\lambda(t'-t)}$.

In terms of the general definition of strategy presented in Section 2.1, a Markov strategy is a strategy σ that assigns for every history of past inspections h the same cumulative probability distribution over the time of the next inspection by the player. Formally,

$$\sigma(h)(t_{\alpha^*}+t)=\begin{cases} F(t), & \forall t \geq 0, \\ 0, & \forall t < 0, \end{cases}$$

where t_{α^*} is the last inspection time in the history h.[15]

We say that a pair of Markov strategies constitutes a Markov Perfect Equilibrium (MPE) if they are best responses to each other at the beginning of the game, and after every history. Namely, there is no other strategy, *Markov or not*, that yields a higher payoff to the player when played against the opponent's strategy.

However, the next result shows that to verify that a given pair of Markov strategies is an MPE, it is sufficient to verify that a player cannot benefit by deviating to another *Markov* strategy. While this property is well known in discrete time models, in our model the argument is more subtle because time is continuous.

Proposition 1. *A pair of Markov strategies* (σ_1, σ_2) *is an MPE if and only if no player has another Markov strategy which is a profitable deviation.*

Remark 1. *To establish Proposition 1, a comprehensive definition of a general strategy is required. Such a definition is developed in Section 2.1.*

In the analysis below we restrict attention to symmetric Markov Perfect Equilibria of the game (in Section 6.1 we discuss MPEs that are not symmetric and equilibria that are not MPE).

Finally, note that if a pair of Markov strategies is an MPE, then they are best responses to each other *at any time* t (and not only immediately after the last failed inspection). This is because if a player has a profitable deviation at some time t, then she also has a profitable deviation immediately after the last failed inspection before time t.

3 Benchmark: The One-Player Problem

We begin the analysis by considering a simpler problem where a single player searches for a prize. This serves as a benchmark for the analysis of the two-player game. We maintain all

[15] More precisely, as explained in Section 2.1, the time t_{α^*} may be an accumulation point of previous inspection times in the history $h = \langle \alpha^*, \{t_\alpha\}_{\alpha \leq \alpha^*}, \iota \rangle$.

the model's assumptions, except for the presence of only one player, who receives a payoff of v_1 upon discovering the prize.

For every $t > 0$, let $w(t)$ be the (single-) player's expected payoff under a pure Markov strategy in which the player inspects at time t (and because the strategy is Markov, inspects again t units of time after every failed inspection):

$$w(t) = e^{-rt}\left(-c + (1 - e^{-\lambda t})v_1 + e^{-\lambda t}w(t)\right), \quad \forall t > 0. \tag{1}$$

The three terms in parentheses on the right-hand side of Eq. (1) capture the expected payoff from inspection at time t: c is the cost of inspection; $(1 - e^{-\lambda t})v_1$ is the expected payoff upon discovering the prize; and $e^{-\lambda t}w(t)$ is the expected payoff upon failing to discover the prize and "restarting" the process of inspection.

Simplifying Eq. (1) yields:

$$w(t) = \frac{(1 - e^{-\lambda t})v_1 - c}{e^{rt} - e^{-\lambda t}}. \tag{2}$$

The next lemma establishes two basic properties of the player's payoff function $w(\cdot)$.

Lemma 2. *The function $w(\cdot)$ is strictly quasi-concave and has a unique maximizer.*

We denote the unique maximizer of $w(\cdot)$ by $\tau^* \in (0, \infty)$. Intuitively, when deciding when to inspect, the player balances two opposite effects of slightly postponing (or advancing) the time of inspection. On the one hand, postponement increases the probability that the prize has already appeared. On the other hand, discounting lowers the benefit from finding the prize. Overall, the sum of the two effects is the total gain from postponement. This gain from postponement is given by the derivative of the function $w(t)$ with respect to t. At the optimal inspection time τ^* this derivative is equal to zero.

Inspection every τ^* units of time is optimal among all Markov strategies. Our next result records this observation and shows that this Markov strategy is in fact optimal among *all* strategies.

Proposition 2. *The player has an optimal inspection strategy, which is Markov. This strategy instructs the player to inspect at time τ^*, and τ^* units of time after every failed inspection, until the prize is discovered. The player's expected payoff from this strategy is given by $w(\tau^*)$.*

Proposition 2 asserts the existence of an optimal strategy, the fact that it is Markov, and relates it to the maximizer of the function $w(\cdot)$. Intuitively, existence can be deduced from the proof of Proposition 1 (invoked for the case where the other player's Markov strategy is to never inspect), which implies that if an optimal strategy exists, its payoff can be approximated by a Markov strategy. Because the strategy outlined in the proposition is optimal among all Markov strategies, it follows that it is optimal among all strategies as well. However, in the proof, we adopt a more direct approach, leveraging recent findings in the literature by Jasso-Fuentes, Menaldi and Prieto-Rumeau (2020) and Stachurski and Zhang (2021).

Intuitively, when the value of the prize v_1 increases, the optimal inspection rate increases as well. This is because the effective cost of postponement is greater. Similarly, when the inspection cost c increases, the optimal inspection rate decreases. We record these two observations in the next lemma.

Lemma 3. *The optimal inspection rate τ^* is decreasing in v_1 and increasing in c.*

4 The Two-Player Game

We now shift the focus back to the two-player game. The addition of another player introduces a strategic element into the players' considerations. In this section we derive results that facilitate the characterization of the symmetric MPE of the game.

In Section 4.1 we introduce two notions, the follower gain and the delay gain, which play a crucial role in the analysis of the game. In Section 4.2, we use these notions to derive qualitative properties of the equilibria of the game. In Section 4.3 we describe an indifference condition that is necessary for symmetric MPE, and finally in Section 4.4 we discuss the properties of a symmetric MPE *if* it is continuous.

4.1 The Follower Gain and the Delay Gain

Fix a symmetric MPE (not necessarily continuous), and denote the players' symmetric payoff in this MPE by v_0. What does Player 1 gain by slightly postponing her inspection? The answer depends on the time t and whether or not Player 2 inspects in the meantime.

FOLLOWER GAIN. Suppose that Player 1 slightly postpones her inspection at some time t (while keeping the times of future inspections unchanged), and Player 2 inspects in the meantime. We call Player 1's (undiscounted) gain from this postponement the *follower gain* at t, and define it formally as follows.

Definition (Follower Gain). *The follower gain at t is the quantity*

$$(1-e^{-\lambda t})(v_2 - v_1) + c. \tag{3}$$

Intuitively, the follower gain at t is a player's gain from a postponement of inspection at time t, which turns her into the follower rather than the leader. Accordingly, the first term in the sum in Eq. (3) is the difference between the second and first prizes, multiplied by the probability that the prize has appeared by time t. The second term is the cost, which is saved for sure.

The following lemma collects some useful observations about the follower gain that immediately arise from the definition. Its proof is omitted.

Lemma 4. *The follower gain satisfies the following properties:*

1. *If v_2 is smaller (greater) than v_1, then the follower gain is decreasing (increasing) in t.*

2. *If $v_2 \geq v_1 - c$, then the follower gain is positive for all $t > 0$.*

3. *If $v_2 < v_1 - c$ then the follower gain is equal to zero at the time*

$$t^\circ \equiv -\frac{1}{\lambda}\ln\left(\frac{v_1 - v_2 - c}{v_1 - v_2}\right). \tag{4}$$

The follower gain is positive for all times $t < t^\circ$, and is negative for all times $t > t^\circ$.

Notably, t° is the time in which a player is indifferent between being the leader and follower. Such a time exists only if $v_2 < v_1 - c$.

DELAY GAIN. Suppose instead that Player 1 slightly postpones her inspection at time t (while keeping the times of future inspections unchanged), but Player 2 *does not* inspect in the meantime. Player 1's payoff from being the leader at t is given by

$$L(t) \equiv e^{-rt}\left(-c + (1 - e^{-\lambda t})v_1 + e^{-\lambda t} v_0\right). \tag{5}$$

The expression on the right-hand side of Eq. (5) is similar to the right-hand side of Eq. (1), except that $w(t)$ is replaced by the equilibrium payoff v_0. It is noteworthy that v_0, and so also $L(t)$, are determined endogenously in equilibrium. This key feature distinguishes our model of repeated inspections from previous stopping games that were analyzed in the literature, where the payoff from taking an action is exogenously fixed.

Thus, the discounted gain from postponing the inspection from t to $t+\varepsilon$, conditional on the other player not inspecting during $[t, t+\varepsilon]$ and keeping the times of future inspections unchanged, is given by the difference $L(t+\varepsilon) - L(t) \approx \varepsilon L'(t)$. We refer to the derivative of $L(t)$ with respect to t as the *discounted delay gain*. We define the *delay gain* at t to be the discounted delay gain divided by e^{-rt}. Consequently, we have that,

Definition (Delay Gain). *The delay gain at t is the quantity*

$$L'(t)/e^{-rt} = rc - rv_1 + (\lambda + r)e^{-\lambda t}v_1 - (\lambda + r)e^{-\lambda t}v_0. \tag{6}$$

The following lemma describes two properties of the delay gain. The first property follows directly from the definition. The second property relates the delay gain and the optimal inspection time in the one-player problem. This result is not immediate and its proof appears in the appendix.

Lemma 5. *The delay gain satisfies the following properties:*

1. *If $v_0 < v_1$, then the delay gain is decreasing in t, and is negative for sufficiently large t.*

2. *If $v_0 = w(\tau^*)$, then τ^* is the unique time at which the delay gain is equal to zero.*

4.2 Qualitative Properties of Equilibrium

In this section, we develop qualitative properties of the equilibrium that facilitate the characterization of equilibrium in Section 5.

The first property follows immediately from our definition of the follower gain and the delay gain.

Observation 1 ("Opposite Signs"). *Let F be a symmetric MPE strategy. If t belongs to the support of F, then the delay gain and the follower gain cannot both be positive at t and cannot both be negative at t.*

Intuitively, if both the delay gain and follower gain are positive at a particular time t in the support of a symmetric MPE strategy, then the continuity of these quantities suggests the existence of a small interval around t where both gains remain positive. Consequently, it is beneficial for a player to postpone all inspections conducted within this interval to its upper end. This deviation benefits the player regardless of whether the other player conducts inspections during the same interval.

Likewise, when both the delay gain and follower gain are negative at a specific time t, similar reasoning suggests that it is more favorable for a player to slightly advance all inspections around t. This contradicts the optimality of the strategies employed.

PROBABILITY OF INSPECTION. In principle, the players may employ strategies that refrain from inspection with a positive probability. However, as shown by the next result, such strategies cannot be part of a symmetric MPE.

Proposition 3 (Inspection with Probability 1). *Let F be a symmetric MPE strategy. Then,* $\lim_{t\to\infty} F(t) = 1$.

The intuition behind this result is as follows. Suppose there exists a symmetric MPE strategy F such that $\lim_{t\to\infty} F(t) < 1$, so that the event that the players never inspect has a positive probability. If both players employ the strategy F, then there exists a sufficiently large time s that is such that: (i) The probability that neither player inspects before time s is positive, (ii) Given that no inspection occurred by time s, the conditional probability of an inspection at or after time s is extremely small, and (iii) The probability that the prize appeared before time s is close to 1. Thus, conditional on reaching time s, the payoff to Player 1 (evaluated at time s) from following the equilibrium strategy is negligibly small. On the other hand, the expected payoff from inspecting at time s (evaluated at s) is close to $v_1 - c > 0$. This is a contradiction.

ATOMS. Whether or not a symmetric MPE is a continuous probability distribution is a-priori unknown. However, *if* a symmetric MPE has an atom, then the next result characterizes its point in time.

Proposition 4 (Atoms). *Let F be a symmetric MPE strategy. If F has an atom at t, then $t = t^\circ$, where t° is defined in Eq. (4).*

To see the intuition for this result, suppose F is a symmetric MPE with an atom of size 1 at t.[16] The probability that Player 1 is the leader at t is one-half. By slightly postponing or advancing the inspection, Player 1 becomes the follower or leader, respectively, with probability 1. Because such deviations cannot be profitable in equilibrium, the follower gain at t must be equal to zero. Thus, by Lemma 4 it must be the case that $t = t^\circ$.

Lemma 4 also implies that if $v_2 > v_1 - c$ then the follower gain is always positive. Thus, the argument above implies the following corollary.

Corollary 1. *If $v_2 > v_1 - c$, then a symmetric MPE strategy has no atoms.*

BOUNDS OF THE SUPPORT. Fix a symmetric MPE strategy. Denote the equilibrium payoff under this MPE by v_0, and the lower and upper bounds of its support by $\underline{t}$ and $\overline{t}$, respectively. Our next results describe properties of the bounds of the support.

Proposition 5 (Lower Bound). *Let F be a symmetric MPE strategy whose support has a minimum of $\underline{t}$. Then, the players' equilibrium payoff is given by $v_0 = w(\underline{t})$, where $w(\cdot)$ is the single player payoff function given by Eq. (2).*

[16]In the proof, we extend this argument to an atom of any positive mass

The proof of this result is the following. Suppose first that there is no atom at $\underline{t}$. In equilibrium, a player has to be indifferent between all the inspection times in the support of her strategy. In particular, a player has to be willing to inspect at $\underline{t}$, knowing that the other player has not inspected until that time. In fact, inspecting at $\underline{t}, 2\underline{t}, \ldots$ ensures that the other player will never get to inspect. This makes the player's payoff identical to her payoff from inspecting every $\underline{t}$ units of time in the one-player game, which is given by $w(\underline{t})$. If there is an atom at $\underline{t}$, then Proposition 4 implies that $\underline{t} = t^{\circ}$. Hence, the player is indifferent between being the leader or the follower at $\underline{t}$, and the conclusion of Proposition 5 follows.

The next corollary follows immediately.

Corollary 2. *Let F be a symmetric MPE strategy with an equilibrium payoff v_0. Then, v_0 cannot exceed the leader's prize minus the cost, i.e., $v_0 < v_1 - c$.*

Indeed, by Proposition 5, the equilibrium payoff is equal to the payoff of a single player who inspects at $\underline{t}$. This payoff is (weakly) smaller than the optimal payoff of a single player, which is smaller than $v_1 - c$ because of the potential costs associated with wasted inspections and the time required for the prize to appear.

Finally, the supremum of the support satisfies the following property.

Proposition 6 (Upper Bound). *Let F be a symmetric MPE strategy. Then, the supremum of the support of F is either $\bar{t} = \infty$ or $\bar{t} = t^{\circ} > 0$ (defined in Eq. (4)).*

To see the intuition for this result, suppose that $\bar{t}$ is finite and smaller than t°, and that the game has reached time $\bar{t} - \varepsilon$ for some small $\varepsilon > 0$. If a player adheres to the original strategy, symmetry of the MPE implies that she would end up becoming the leader and follower with probability one-half each. However, by postponing all inspections until slightly after $\bar{t}$, the player would become the follower with probability one. Because $\bar{t} < t^{\circ}$, by Lemma 4 the follower gain is positive, and therefore this deviation is profitable. An analogous argument applies for the case in which $\bar{t}$ is finite and larger than t°.

4.3 The Indifference Condition

Let F be a symmetric MPE strategy (not necessarily continuous). In equilibrium a player must be indifferent between all the inspection times in the support of the strategy F. Thus, the following indifference condition must hold for every t in the support of F:

$$v_0 = (1 - F(t))e^{-rt}\left(-c + (1 - e^{-\lambda t})v_1 + e^{-\lambda t}v_0\right) + \int_{s=0}^{t} e^{-rs}\left((1 - e^{-\lambda s})v_2 + e^{-\lambda s}v_0\right) \mathrm{d}F(s). \quad \text{(IND)}$$

To understand the indifference condition (IND), note that the right-hand side of (IND) is a player's expected payoff from inspecting at t, provided that (i) the other player employs F, and (ii) the players would keep employing F in the future. Indeed, if Player 2 does not inspect before or at t, which occurs with probability $1-F(t)$, then Player 1's expected payoff is given by $-c+(1-e^{-\lambda t})v_1+e^{-\lambda t}v_0$ (see the explanation that follows Eq. (1)). And, if Player 2 inspects at time $s<t$, then Player 1's payoff is given by $(1-e^{-\lambda s})v_2+e^{-\lambda s}v_0$. The first term describes Player 1's expected payoff if Player 2 discovers the prize at s, and the second term describes her expected payoff if Player 2 does not discover the prize, and the game "restarts" at s.

The indifference condition (IND) is a necessary condition that has to be satisfied for all the times t in the support of F.[17] Although the structure of this support is endogenously determined in equilibrium, we show that it nevertheless has a simple form.

Proposition 7. *The support of a symmetric MPE strategy F is convex.*

The proof of this "no holes in the support" property hinges on Corollary 2, which ensures that $v_1>v_0$. By Lemma 5, this implies that the delay gain, as given by Eq. (6), is decreasing in t *in equilibrium*, which is necessary in order to prove the proposition.[18]

4.4 Properties of Continuous Equilibria

If the MPE strategy is continuous, then it is possible to derive from Eq. (IND) tighter conditions that facilitate the characterization of the MPE.

Suppose that F is a continuous symmetric MPE strategy that is supported on an interval $[\underline{t},\bar{t}]$,[19] and denote its density by f. Differentiate both sides of the indifference condition (IND) according to t and rearrange to obtain the following equation, which holds for all $t\in(\underline{t},\bar{t})$:

$$f(t)\left[(1-e^{-\lambda t})(v_2-v_1)+c\right]=-(1-F(t))\left[rc-rv_1+(\lambda+r)e^{-\lambda t}v_1-(\lambda+r)e^{-\lambda t}v_0\right]. \quad (7)$$

Eq. (7) is a first-order linear differential equation. By Proposition 5, in an MPE we have that $v_0=w(\underline{t})$. Hence, Eq. (7) is parametrized by the lower bound of the support $\underline{t}$ and it has to hold on the interval $(\underline{t},\bar{t})$.

[17] Notice that satisfaction of the indifference condition is not sufficient for equilibrium, and that the indifference condition says nothing about times that lie outside this support.

[18] Thus, if t_1 and t_2 are in the support of F, but the open interval (t_1,t_2) is not in the support of F, then it is beneficial for a player to either postpone inspection from t_1 to $t_1+\varepsilon$ or to advance inspection from t_2 to $t_2-\varepsilon$ for some small $\varepsilon>0$, which is a contradiction.

[19] If $\bar{t}=\infty$, then F is supported on $[\underline{t},\bar{t})$.

It is noteworthy that the term in the square brackets on the left-hand side of Eq. (7) is the follower gain at t, given by Eq. (3). And, the term in the square brackets on the right-hand side of Eq. (7) is the delay gain at t, given by Eq. (6). Indeed, in line with Observation 1, Eq. (7) implies that these two terms cannot be both positive or negative at any time t in the support of F.

The next lemma characterizes the solution F of Eq. (7) given the support $[\underline{t}, \bar{t}]$ and under the initial condition $F(\underline{t}) = 0$.

Lemma 6. *Given values $\underline{t} < \bar{t}$, the solution of Eq. (7) under the initial condition $F(\underline{t}) = 0$ is*

$$F(t) = 1 - e^{G(t) - G(\underline{t})}, \quad \forall t \in [\underline{t}, \bar{t}]. \tag{8}$$

The function $G(t)$ is defined as follows:[20]

1. *If $v_2 = v_1 - c$, then $G(t) = \frac{D}{B\lambda} e^{\lambda t} + \frac{E(w(\underline{t}))}{B} t$.*
2. *If $v_2 \neq v_1 - c$ and $A + Be^{-\lambda t} > 0$ for all $t \in [\underline{t}, \bar{t})$, then $G(t) = \frac{D}{A} t + \frac{1}{\lambda} \left(\frac{D}{A} - \frac{E(w(\underline{t}))}{B} \right) \ln \left(A + Be^{-\lambda t} \right)$.*

where $A \equiv c - v_1 + v_2$, $B \equiv v_1 - v_2$, $D \equiv r(c - v_1)$, and $E(v_0) \equiv (\lambda + r)(v_1 - v_0)$.

Given a support $[\underline{t}, \bar{t}]$, Eq. (8) describes a *candidate* for a continuous symmetric MPE. To show that this candidate is indeed a continuous symmetric MPE, we need to show that it is a CDF, and that there are no profitable deviations outside the support $[\underline{t}, \bar{t}]$.

Remark 2. *To better understand the economic interpretation of Eq. (7), suppose that Player 2 inspects according to the continuous distribution F. What does Player 1 gain from delaying her inspection at time t by Δt? Conditional on the event that Player 2 did not inspect until t, Player 1 switches from being the leader to being the follower with probability $\Delta t \cdot f(t)/(1 - F(t))$, and she gains the follower gain. With the complementary probability, Player 1 inspects at time $t + \Delta t$, and gains $\Delta t \cdot$(delay gain). Taken together, this implies that the expected gain from a small delay at time t is given by*

$$H(t, \Delta t) = \frac{\Delta t \cdot f(t)}{1 - F(t)} \cdot (\textit{follower gain}) + \left(1 - \frac{\Delta t \cdot f(t)}{1 - F(t)} \right) \cdot \Delta t \cdot (\textit{delay gain}).$$

A necessary condition for the continuous strategy F to be a symmetric MPE strategy is that $\lim_{\Delta t \to 0} (H(t, \Delta t)/\Delta t) = 0$. This yields Eq. (7).

[20] Notice that the case in which $v_2 \neq v_1 - c$ and $A + Be^{-\lambda t}$ is not positive for all $t \in [\underline{t}, \bar{t})$ is not covered. However, this case turns out to not be relevant for our analysis.

5 Markov Perfect Equilibria in the Two-Player Game

In this section, we characterize the symmetric MPE of the two-player game. As previously explained, the game encompasses two strategic considerations. On the one hand, players compete to discover the prize, while on the other hand, the presence of another player introduces the possibility of "free-riding" on the other player's inspection.

We demonstrate that the MPE takes one of two forms, depending on which consideration dominates, and this, in turn, depends on the parameters of the game. The first type of MPE involves frequent synchronized inspections, while the second exhibits slow diffusion, in which players inspect randomly at different times.

Recall that τ^* is the optimal inspection time in the one-player game (see Proposition 2). The following threshold value plays a crucial role in the characterization of equilibrium:

$$\widetilde{v}_2 \equiv v_1 - \frac{c}{1 - e^{-\lambda \tau^*}}. \tag{9}$$

Simple algebraic manipulation shows that $v_2 \geq \widetilde{v}_2$ if and only if $t^\circ \geq \tau^*$, with equality when $v_2 = \widetilde{v}_2$. When $v_2 \leq \widetilde{v}_2$ we say that the follower prize is *small*, and when $v_2 > \widetilde{v}_2$ we say that the follower prize is *large*.

5.1 The Case of a Small Follower Prize: Synchronized Inspection

Suppose that the follower prize is small, namely, $v_2 \leq \widetilde{v}_2$. Intuitively, in this scenario, each player prefers to discover the prize and collect v_1, rather than waiting to collect the small amount v_2. We show that in this case the player synchronize their inspections. Formally, the unique symmetric MPE strategy has an atom of probability one at t°.

Theorem 1. *Suppose that $v_2 \leq \widetilde{v}_2$. The unique symmetric MPE strategy of the two-player game is given by*

$$F^S(t) = \begin{cases} 0, & t < t^\circ, \\ 1, & t \geq t^\circ, \end{cases} \tag{10}$$

where t° is defined in Eq. (4).

In the equilibrium described in Theorem 1, both players synchronize their inspections and inspect at times $t^\circ, 2t^\circ, \ldots$. Namely, the players inspect simultaneously at the exact times in which they are indifferent between being leader and follower.

To see why this is an equilibrium, recall that by Proposition 5, the players' payoff under this strategy is $w(t^\circ)$. A player cannot benefit from inspecting earlier because $v_2 \leq \widetilde{v}_2$ implies that $t^\circ \leq \tau^*$. This means that by advancing the inspection to $t < t^\circ$, the player's payoff would be equal to the payoff of the player in the one-player game who inspects at the same times. Lemma 2 implies that this payoff is quasiconcave in the inspection time, and so such a deviation is not profitable.[21] And, deviating and inspecting after time t° would turn the player into a follower at t°, which would not affect her payoff.

The intuition for why F^S is the *unique* symmetric MPE strategy hinges on the observation that when $v_2 \leq \widetilde{v}_2$, in equilibrium there is no time t in which the players inspect and the follower and delay gains have opposite signs. Therefore, by Observation 1, the support of a symmetric MPE strategy must be concentrated at t° where the follower gain is zero, implying that it coincides with F^S. Despite the positive delay gain at t°, a player cannot profit by delaying her inspection as this would make her a follower with probability one.[22]

The following corollary is an immediate implication of Theorem 1.

Corollary 3. *If $v_2 \leq \widetilde{v}_2$, then the players inspect more frequently compared to the one-player benchmark.*

The next observation summarizes the effects of the model's parameters on the equilibrium frequency of inspection and payoff.

Observation 2. *Suppose that $v_2 \leq \widetilde{v}_2$. Then,*[23]

1. *A larger value of the leader's prize v_1 or the arrival rate λ increase both the frequency of inspection and the MPE payoff. The inspection cost c has the opposite effect.*

2. *A larger value of the follower's prize v_2 decreases the frequency of inspection, but increases the MPE payoff.*

3. *The players' discount rate r does not affect the frequency of inspection.*

As intuitively expected, a larger value of v_1 induces a lower value of t°, leading to a higher frequency of inspection. The effect on the equilibrium payoff is more subtle: on the one hand, a higher value of v_1 increases the players' payoff, but on the other hand, faster inspection decreases it. It can be shown that the first effect dominates, and that the players' equilibrium payoff increases in the value of v_1. Analogous arguments demonstrate that the

[21] This argument makes use of the observation that if a player prefers to deviate to an earlier inspection time once, then the player would also want to deviate every time she finds herself in the same situation.

[22] Recall that by definition, a player enjoys the delay gain only if the other player does not inspect in the meantime.

[23] The statements below hold true so long as the condition that $v_2 \leq \widetilde{v}_2$ is preserved, which implies that the MPE is described by Theorem 1.

prize's rate of arrival λ has a similar effect, and the cost of inspection c has the opposite effect.

To see the effect of v_2, note that a larger value of v_2 makes being the leader relatively less attractive, and so decreases the frequency of inspection. By Proposition 5, this also increases the players' equilibrium payoff because it delays inspections.[24] While this second result holds true for small values of v_2, this is not the case for large values of v_2, as shown in the next section.

Finally, perhaps surprisingly, the rate of inspection is unaffected by the players' discount rate r. This is because in equilibrium, the players inspect when they are indifferent between being the leader and the follower, and the point in time in which this happens is independent of the discount rate.

5.2 The Case of a Large Follower Prize: Diffused Inspection

When the follower's prize v_2 is large, both players benefit from free-riding on each other. This has a significant effect on the symmetric MPE of the game.

Suppose that $v_2 > \widetilde{v}_2$. The quantity τ° defined below will be useful in the characterization of the MPE:

$$\tau^\circ = \begin{cases} t^\circ, & \text{if } v_2 < v_1 - c, \\ \infty, & \text{if } v_2 \geq v_1 - c. \end{cases} \tag{11}$$

The unique symmetric MPE of the two-player game is characterized in the next theorem. Unlike the MPE described in Theorem 1 this MPE exhibits diffused inspections.

Theorem 2. *Suppose that $v_2 > \widetilde{v}_2$. Then, there is a unique symmetric MPE strategy. This MPE strategy is characterized by the continuous CDF given by Eq. (8) that is supported on the interval $[\underline{t}, \overline{t}] = [\tau^*, \tau^\circ]$.*

To prove Theorem 2 we show that the conditions required by Lemma 6 are satisfied, and that the solution F described by Eq. (8) is indeed a CDF. We then show that inspections outside the support of F are not profitable deviations.

To prove that F is the *unique* symmetric MPE, we first show that the lower bound of the support of any symmetric MPE must be equal to τ^*, and the upper bound is at most t°. This pins down the equilibrium payoff $v_0 = w(\tau^*)$. We then employ a general result in the Theory of Distributions. The right-hand side of (IND) is of the form $(1 - F(t))g(t) + \int_{r=0}^{t} h(r)\mathrm{d}F(r)$, where g and h are two smooth functions, and the left-hand side is a constant.

[24] It is noteworthy that unlike v_1 and c, the value of v_2 does not affect the function $w(\cdot)$. Thus, it affects the players' equilibrium payoff only through its effect on the timing of inspection.

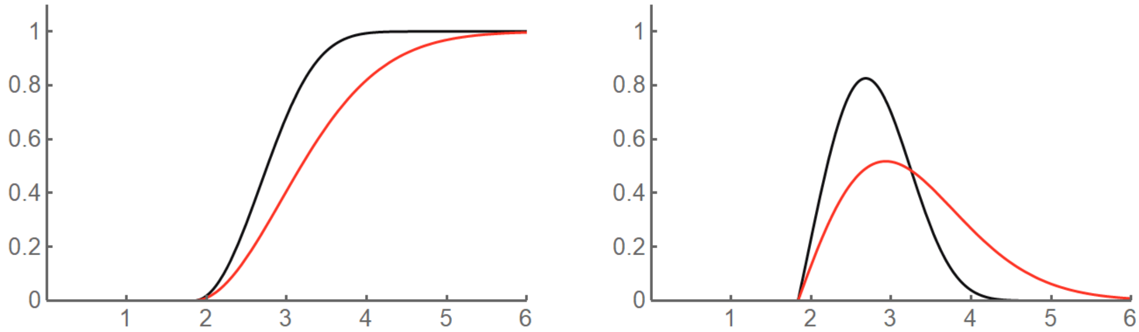

(a) CDFs of symmetric MPEs

(b) Densities of symmetric MPEs

Figure 1: CDFs and densities of symmetric MPE strategies for different parameter values. The black curves correspond to the parameters $v_1 = 10$, $v_2 = 7.5$, $c = 2$, $r = 0.75$, $\lambda = 0.25$, and $\widetilde{v}_2 = 4.566$. The red curves correspond to the same parameters, with the only difference that $v_2 = 9.5$.

Simple calculations show that the functions g and h can be equal only a finite number of times on the interval $[\underline{t}, \bar{t}]$. According to Theorem 2.24 in Hörmann and Steinbauer (2009), on any interval in which g and h are not equal, if a continuous solution to (IND) exists, then this solution is unique among all possible solutions, continuous or not. Because the distribution F in Eq. (8) is differentiable, it follows that (IND) does not admit any non-differentiable solutions.

Figure 1 depicts a symmetric MPE strategy and its density for different parameter values. It is noteworthy that the density of the distribution of inspection times in an MPE is first increasing and then decreasing in t. As can be seen in the figure, a higher value of the follower prize v_2 "stretches" the density of the MPE strategy to the right (towards larger values). By contrast, the lower bound of the support of the MPE CDF is always equal to τ^*, for any value of $v_2 > \widetilde{v}_2$. This last fact, together with Proposition 5, implies the following result.

Corollary 4. *Suppose that $v_2 > \widetilde{v}_2$. Then, the players' symmetric MPE payoff v_0 is equal to the payoff in the one-player problem.*

It is also noteworthy that by Corollary 4, as long as $v_2 > \widetilde{v}_2$, the players' equilibrium payoff is *independent of the value of* v_2. Intuitively, this independence is due to the fact that as v_2 increases, the players intensify free-riding on each other by delaying their inspections. This is illustrated in Figure 1 that shows that the symmetric MPE strategy with a larger v_2 first-order stochastically dominates the symmetric equilibrium strategy with a smaller v_2.

To see the effect of the other model parameters on the MPE payoff recall that $v_0 = w(\tau^*) \equiv \max_{t \geq 0} w(t)$ where the function $w(\cdot)$ is given by Eq. (2). Computing the derivative of $w(\cdot)$ with respect to each one of the model's parameters gives rise to the following observation.

Observation 3. *Suppose that $v_2 > \tilde{v}_2$. Then,*[25]

1. *A larger value of the leader's prize v_1 or the arrival rate λ increase the MPE payoff.*

2. *The follower's prize v_2 does not affect the MPE payoff.*

3. *A larger value of the inspection cost c decreases the MPE payoff.*

5.3 Social Welfare

Suppose that the players' utilities are quasilinear, so that utility is transferable. Thus, the socially optimal inspection policy is given by the solution to the one-player problem, where the prize is given by $v_1 + v_2$.

Throughout the analysis presented above, the value of the follower's prize v_2 played a crucial role in the characterization of equilibrium. The next proposition demonstrates the effect of the value of v_2 on social welfare.

Proposition 8. *There exists a unique threshold $\hat{v}_2 = \hat{v}_2(v_1, c, \lambda, r) \leq \tilde{v}_2$ such that:*
(i) when $v_2 < \hat{v}_2$, the players inspect "too frequently" relative to the socially optimal policy;
(ii) when $v_2 > \hat{v}_2$, the players inspect "too rarely" relative to the socially optimal policy;
(iii) when $v_2 = \hat{v}_2$, the players inspect according to the socially optimal policy.

The key idea is that when v_2 is small, the competition between the two players is fierce. The players inspect too frequently in an attempt to be the first player to find the prize, which erodes their joint welfare. By contrast, when v_2 is large the players try to free-ride on each other, and therefore they inspect too rarely. There is a unique value of v_2 for which the two effects exactly offset each other.

An implication of the fact that $\hat{v}_2 \leq \tilde{v}_2$ is that an economist who lacks direct access to the model's parameters but can still observe the equilibrium dynamics in play can still infer from observation of diffused inspections that the players inspect too rarely compared to the social optimum.

Our analysis also offers insights for the design of regulatory policy:

Observation 4. *Consider a social planner who needs to allocate a resource amount $v > c$ between a leader's prize (v_1) and a follower's prize (v_2). Then, there exists an allocation v_1 and v_2 that satisfies $v_1 + v_2 = v$, and induces the socially optimal rate of inspection in equilibrium.*

[25] The statements below hold true so long as the condition that $v_2 > \tilde{v}_2$ is preserved, which implies that the MPE is described by Theorem 2.

To see this note that with the prize scheme ($v_1 = v, v_2 = 0$), in equilibrium the players inspect too frequently compared to the socially optimal inspection level. Conversely, with the prize scheme ($v_1 = c + \varepsilon, v_2 = v - c - \varepsilon$) for some small $\varepsilon > 0$, the temptation to free-ride on the other player's action governs the players' equilibrium behavior, causing them to inspect too rarely. Continuity of the quantities τ^*, t°, and $\widetilde{v}_2$ in v_1 and in v_2 implies the existence of a value of the follower prize for which the equilibrium inspection frequency coincides with the socially optimal level.[26]

6 Discussion

6.1 Other Equilibria

Thus far, our analysis has focused on the symmetric Markov perfect equilibria of the game. Symmetric MPEs possess two appealing characteristics. Firstly, they consist of Markov strategies, which, intuitively, require less sophistication from the players as they allow the players to only condition their behavior on payoff-relevant variables. Secondly, since the game and the players are ex-ante symmetric, studying symmetric equilibria is a natural choice. Consequently, symmetric MPEs serve as a natural starting point for analyzing the game.

However, the game also has asymmetric equilibria. For example, consider the case in which $v_2 > v_1$. Then, the following pair of strategies is an equilibrium. Player 1 inspects at the optimal frequency in the single-player problem, i.e., every τ^* units of time; Player 2 never inspects. To see that this is an equilibrium of the game, notice that given that Player 2 never inspects, Player 1 essentially faces a single-player problem and so responds optimally. And, because $v_2 > v_1$, Player 2 cannot do better than free-riding on Player 1's inspection.

The game also has symmetric equilibria that are not Markov strategies. An example of such an equilibrium can be illustrated as follows: Assume that $v_1 = 1$, $v_2 = 0$, $c = 0.05$, $r = \lambda = 0.5$. Computation shows that in this case, $\tau^* \approx 0.646$. Consider the following symmetric strategies: Both players inspect every τ^* units of time until the prize is discovered or until one of the players deviates. If a player deviates, then the players revert to playing the MPE strategies characterized in Theorem 1, where they inspect every $t^\circ \approx 0.1$ units of time. Notice that because $v_2 = 0$, the players' continuation payoff after reverting to the strategies characterized by Theorem 1 is zero.

[26] Specifically, denote by $\overline{v}_2$ the value of the follower's prize for which the equilibrium inspection frequency coincides with the socially optimal level. Then, $\overline{v}_2$ has to satisfy the equation $t^\circ(v_1 = v - \overline{v}_2, v_2 = \overline{v}_2) = \tau^*(v)$, and thus can be computed numerically.

To understand why this is an equilibrium, observe that the most profitable deviation for a player is to inspect at $\tau^* - \varepsilon$ for some small $\varepsilon > 0$. However, calculations reveal that the gain from such a deviation is not significant enough to warrant the deviation, taking into account the fact that in the continuation game the players will inspect "too frequently," every $t^\circ \approx 0.1$ units of time. These strategies are not Markov strategies because the players' actions depend on components of history (specifically, whether a player has deviated before or not), which are not payoff relevant.[27]

6.2 Equilibrium Payoffs and MaxMin Values

The maxmin value of a player is the highest payoff that the player can secure in the game, regardless of the other players' strategies. In opportunity hunting games, the symmetric MPE of the game generates an equilibrium payoff that coincides with the maxmin value.

Proposition 9. *The players' payoffs in the symmetric MPEs characterized in Theorems 1 and 2 are equal to the players' maxmin values.*

The proof of Proposition 9 is constructive. It shows that by inspecting every τ^* units of time (in the case where $v_2 > \widetilde{v}_2$) or every t° units of time (in the case where $v_2 \leq \widetilde{v}_2$) a player can secure at least the payoffs $w(\tau^*)$ or $w(t^\circ)$, respectively. These payoffs precisely match the players' symmetric MPE payoffs. In fact, they are equal to the payoff that Player 1 would receive if Player 2 never inspects. In the proof, we establish that if Player 2 does inspect sometimes, then Player 1's payoff from this strategy can only be higher.

Proposition 9 implies that the players' payoffs in the symmetric MPEs characterized in Theorems 1 and 2 are weakly lower than the payoffs in any other equilibrium of the game, if such an equilibrium exists.

6.3 Modelling Assumptions

We now turn to discuss three of our modelling assumptions and examine some natural alternatives.

Observable vs. Unobservable Actions

Throughout our analysis, we have assumed that the players' inspection times are observable. This implies that both players are symmetrically informed at all times. While this assump-

[27] There also exist non-Markovian equilibria in which the players' strategies are history-dependent on the path of play, and not only after deviations as in the example above. When $v_2 < \widetilde{v}_2$, a simple example for such an equilibrium is where Player 1 inspects at $t^\circ, 3t^\circ, 5t^\circ, \ldots$ and Player 2 inspects at $2t^\circ, 4t^\circ, 6t^\circ, \ldots$

tion is applicable in certain environments, there are scenarios where it is less suitable.

In the parameter range in which the follower prize is small (i.e., $v_2 < \widetilde{v}_2$) our analysis and the characterization of the symmetric MPE are unchanged under the alternative assumption that inspection times are unobservable. The reason is that, as stated in Theorem 1, when the follower prize is small, the players inspect simultaneously at t°, which is the point in time in which a player is indifferent between being the leader and the follower. This also implies that a player would still inspect at the same time as long as the strategy of the other player is to inspect at t°, even if the inspection time itself is not observable.

This reasoning no longer holds in the parameter range in which the follower prize is large, because in this case the fact that the symmetric MPE involves mixed strategies implies that the players would be asymmetrically informed about past inspection times, which affects their strategic considerations. The characterization of MPE in this environment remains an open question.

Payoffs upon Simultaneous Inspection

In the model, we assumed that if the players inspect simultaneously, they are randomly assigned the roles of leader (incurring the inspection cost c and obtaining the payoff v_1 if the prize has already appeared) and follower (not incurring the inspection cost and obtaining the payoff v_2 if the prize has already appeared). Assuming instead that if the two players inspect at the exact same time, then each of them pays half the cost of inspection, i.e., $c/2$, and they are randomly assigned to receive the prizes v_1 and v_2, yields identical results.

An alternative and plausible assumption is that if both players perform their inspections simultaneously, then *both* players incur the *entire* inspection cost c, while the allocation of the leader and follower prizes remains random. In the parameter range where the follower prize is large ($v_2 > \widetilde{v}_2$), the MPE characterized in Theorem 2 remains valid. This is because this MPE involves continuous strategies, and therefore the probability that both players inspect simultaneously is zero.

However, in the parameter range where the follower prize is small ($v_2 < \widetilde{v}_2$), the MPE characterized in Theorem 1 is no longer an MPE under the alternative assumption that both players bear the entire inspection cost upon simultaneous inspection. This is because a player can gain a payoff that is bounded away from zero by slightly advancing their inspection time. This reasoning suggests that a symmetric MPE strategy cannot include an atom. Therefore, in the absence of any other special adjustments to the model (such as the use of a correlation device), no symmetric MPE exists in this case.[28]

[28]The relationship between the type of tie-breaking assumption employed and the existence or non-existence

Decoupling the Discovery and the Claiming of the Prize

In the model, we have assumed that the first player to discover the prize claims it immediately and receives a payoff of v_1, while the other player receives a payoff of v_2. While this assumption holds true in certain scenarios, there are instances where a player may be able to defer the claiming of the prize despite knowing of its presence. This becomes particularly appealing when $v_2 > v_1$. In such cases, the allure of being the follower rather than the leader prompts players to vie for the follower position even after learning of the prize's existence.

The model can be adapted to incorporate this consideration in a straightforward manner. First, players engage in a competition to discover the prize, bearing a cost of c per inspection as in the original model. Once the prize is discovered, a "competition" to claim the prize ensues. In this competition, the act of claiming the prize carries no cost; the first player to act receives the payoff v_1, whereas the other player receives the payoff v_2. If both players act simultaneously, the player who discovered the prize is granted priority.

An equilibrium in this adapted game can be characterized using the results presented above. In the case where $v_1 \geq v_2$, it is optimal for the player who discovers the prize to immediately claim it. This implies that the analysis of this variation of the game coincides with our previous analysis.

However, when $v_2 > v_1$, the player who discovers the prize does not want to be the first to claim it. In this case, because it is commonly known that the prize has already appeared, the two players initiate a war of attrition over who will be the leader and who will be the follower. In this war of attrition, the cost of not acting is due to discounting. Let $\overline{v}$ denote the players' expected payoff in a symmetric MPE of this war of attrition subgame.[29] Thus, the analysis of the first stage competition between the players in the adapted game coincides with the analysis presented in Section 5.2 where the leader and follower prizes are both equal to $\overline{v}$.

---

of equilibrium is well-established. For further discussion on this matter and the potential adaptations introduced by the literature, see Fudenberg and Tirole (1991).

[29] The existence of a symmetric MPE in this subgame can be established, for example, by adapting the analysis presented in the previous sections to the case where the prize always exists. In this symmetric MPE, both players inspect according to a continuous distribution with the CDF $F(t) = 1 - e^{-\frac{r(v_1-c)}{c-v_1+v_2}(t-t^{SG})}$, which is supported over the interval $[t^{SG}, \infty)$, where t^{SG} is the time in which the subgame begins. The players' equilibrium payoff, discounted to the time in which the subgame begins, is given by $v_1 - c$.

7 Conclusions

In this paper, we analyze a game of hunting for "opportunities" – events that occur at uncertain points in time and identifying them is costly. We characterize the unique symmetric Markov Perfect Equilibrium in this game. The qualitative form of this equilibrium depends on the model's parameters, and attains one of two forms: one featuring simultaneous and synchronous inspections, and another involving inspections dispersed randomly over time.

In addition to its theoretical contribution as the first paper (to the best of our knowledge) to solve opportunity hunting games, our model also offers applied economists a practical framework for analyzing competitive scenarios centered around opportunities. Importantly, our framework does not require the economist to commit beforehand to the type of competition (e.g. first- vs second- mover advantage) or to the dominant strategic force guiding players' behavior within the model. Instead, these can be estimated directly from the available data.

Finally, our analysis makes methodological contributions to the literature on continuous-time timing games in which agents can act multiple times. We hope that these contributions will prove valuable for the analysis of other models sharing a similar structure.

Bibliography

Anderson, A., L. Smith, and A. Park. 2017. "Rushes in large timing games." *Econometrica*, 85(3): 871–913.

Argenziano, R., and P. Schmidt-Dengler. 2014. "Clustering in n-player preemption games." *Journal of European Economic Association*, 12(2): 368–396.

Awaya, Y., and V. Krishna. 2021. "Startups and Upstarts: Disadvantageous Information in R&D." *Journal of Political Economy*, 129(2): 534–569.

Ball, I., and J. Knopfle. 2023. "Should the Timing of Inspections be Predictable?" Working Paper.

Bergin, James, and W Bentley MacLeod. 1993. "Continuous time repeated games." *International Economic Review*, 21–37.

Bobtcheff, C., R. Levy, and T. Mariotti. 2021. "Negative results in science: Blessing or (winner's) curse?" Working Paper.

Bonatti, A., and J. Hörner. 2011. "Collaborating." *American Economic Review*, 101(2): 632–663.

Cardaliaguet, Pierre. 2007. "Differential games with asymmetric information." *SIAM journal on Control and Optimization*, 46(3): 816–838.

Cetemen, D., and C. Margaria. 2023. "Exit Dilemma: The Role of Private Learning on Firm Survival." Working Paper.

Che, Y-K., and I. Gale. 2003. "Optimal design of research contests." *American Economic Review*, 93: 646–671.

Davis, M.H.A., and P. Varaiya. 1973. "Dynamic programming conditions for partially observable stochastic systems." *SIAM Journal on Control*, 11(2): 226–261.

Décamps, J.P., and T. Mariotti. 2004. "Investment timing and learning externalities." *Journal of Economic Theory*, 118(1): 80–102.

Frick, Mira, and Yuhta Ishii. 2023. "Innovation Adoption by Forward-Looking Social Learners." *Working Paper*.

Fudenberg, D., and J. Tirole. 1985. "Preemption and rent-equalization in the adoption of new technology." *The Review of Economic Studies*, 52(3): 383–401.

Fudenberg, D., and J. Tirole. 1991. *Game Theory.* MIT Press.

Glazer, Amihai, and Refael Hassin. 1988. "Optimal contests." *Economic Inquiry*, 26(1): 133–143.

Hansen, E.A., D.S. Bernstein, and S. Zilberstein. 2004. "Dynamic programming for partially observable stochastic games." *American Association for Artificial Intelligence (AAAI-04)*, 4: 709–715.

Hendricks, K., A. Weiss, and C. Wilson. 1988. "The War of Attrition in Continuous Time with Complete Information." *International Economic Review*, 29(4): 663–680.

Hopenhayn, H.A., and F. Squintani. 2011. "Preemption games with private information." *The Review of Economic Studies*, 78(2): 667–692.

Hörmann, G., and R. Steinbauer. 2009. *Lecture Notes on the Theory of Distributions.* Universitaet Wien.

Hörner, J., N. Klein, and S. Rady. 2022. "Over-coming free riding in bandit games." *The Review of Economic Studies*, 89(4): 1948–1992.

Jasso-Fuentes, H., J.-L. Menaldi, and T. Prieto-Rumeau. 2020. "Discrete-time control with non-constant discount factor." *Mathematical Methods in Operation Research*, 92: 377–399.

Judd, Kenneth L. 1985. "The Law of Large Numbers with a Continuum of IID Random Variables." *Journal of Economic Theory*, 35(1): 19–25.

Kamada, Y., and N. Rao. 2023. "Strategies in Stochastic Continuous-Time Games." UTMD

Working Paper.

Keller, R., S. Rady, and M. Cripps. 2005. "Strategic experimentation with exponential bandits." *Econometrica*, 73(1): 39–68.

Laraki, R., and E. Solan. 2005. "The value of zero-sum stopping games in continuous time." *SIAM Journal on Control and Optimization*, 43(5): 1913–1922.

Moldovanu, B., and A. Sela. 2001. "The Optimal Allocation of Prizes in Contests." *American Economic Review*, 91: 542–558.

Moldovanu, B., A. Sela, and X. Shi. 2007. "Contests for Status." *Journal of Political Economy*, 115(2): 338–363.

Murto, P., and J. Valimaki. 2011. "Learning and information aggregation in exit games." *The Review of Economic Studies*, 78: 1426–1461.

Rosenberg, D., A. Salomon, and N. Vieille. 2013. "On games of strategic experimentation." *Games and Economic Behavior*, 82: 31–51.

Rosenberg, D., E. Solan, and N. Vieille. 2007. "Social learning in one arm bandit problems." *Econometrica*, 75(6): 1591–1611.

Shahanaghi, S. 2022. "Competition and Errors in Breaking News." Working Paper.

Smirnov, V., and A. Wait. 2022. "General timing games with multiple players." SSRN Discussion Paper.

Stachurski, J., and J. Zhang. 2021. "Dynamic programming with state-dependent discounting." *Journal of Economic Theory*, 192: 105190.

Stinchcombe, Maxwell B. 1992. "Maximal strategy sets for continuous-time game theory." *Journal of Economic Theory*, 56(2): 235–265.

Appendix: Proofs

Proof of Lemma 1

We will construct the unique probability distribution over plays induced by (σ_1, σ_2) by transfinite induction. The distribution of t_0 is determined by $\sigma_1(\emptyset)$ and $\sigma_2(\emptyset)$, where $\emptyset$ denotes the empty history.

Let α_* be a limit ordinal, and suppose that the distribution over histories with $\alpha < \alpha_*$ is uniquely defined. Then $t_{\alpha_*} = \sup_{\alpha<\alpha_*} t_\alpha$.

Let $\alpha_* + 1$ be a successor ordinal, and suppose that the distribution over histories with $\alpha \le \alpha_*$ is uniquely defined. Denote by h the random variable of the history up to t_{α_*}. The conditional distribution of t_{α_*+1} given h is determined by $\sigma_1(h)$ and $\sigma_2(h)$.

Proof of Proposition 1

We show that for every $\varepsilon > 0$, every Markov strategy has an ε-best response, which is itself a Markov strategy.

We introduce the following notation. For every pair of strategies (σ_1, σ_2) and every history $h = \langle \alpha^*, \{t_\alpha\}_{\alpha \le \alpha^*}, \iota \rangle$, denote by $U(\sigma_1, \sigma_2; h)$ the expected payoff of Player 1 under (σ_1, σ_2) in the subgame that starts at h. The quantities $(U(\sigma_1, \sigma_2; h))_h$ are related through the following recursive equation:

$$\begin{aligned} U(\sigma_1,\sigma_2;h) &= \int_{s_1,s_2\in(t_{\alpha^*},\infty]} \Bigg[\left(1_{s_1<s_2} + \frac{1}{2}\cdot 1_{s_1=s_2}\right) e^{-r(s_1-t_{\alpha^*})} \\ &\quad \cdot \left(-c + (1-e^{-\lambda(s_1-t_{\alpha^*})})v_1 + e^{-\lambda(s_1-t_{\alpha^*})}U(\sigma_1,\sigma_2;h')\right) \\ &\quad + \left(1_{s_1>s_2} + \frac{1}{2}\cdot 1_{s_1=s_2}\right) e^{-r(s_1-t_{\alpha^*})} \\ &\quad \cdot \left((1-e^{-\lambda(s_2-t_{\alpha^*})})v_2 + e^{-\lambda(s_2-t_{\alpha^*})}U(\sigma_1,\sigma_2;h')\right)\Bigg] \sigma_1(h)(\mathrm{d}s_1)\sigma_2(h)(\mathrm{d}s_2), \end{aligned} \tag{12}$$

where h' is the history after the current one, so it is created by adding to h either $(t_{\alpha^*+1} = s_1, \iota_a(\alpha^*+1) = \iota_i(\alpha^*+1) = \{1\})$ (if $s_1 < s_2$), or $(t_{\alpha^*+1} = s_2, \iota_a(\alpha^*+1) = \iota_i(\alpha^*+1) = \{2\})$ (if $s_1 > s_2$), or $(t_{\alpha^*+1} = s_1, \iota_a(\alpha^*+1) = \{1,2\}), \iota_i(\alpha^*+1) = \{1\})$ (if $s_1 = s_2 < \infty$ and Player 1 actually inspected), and $(t_{\alpha^*+1} = s_2, \iota_a(\alpha^*+1) = \{1,2\}), \iota_i(\alpha^*+1) = \{2\})$ (if $s_2 = s_1 < \infty$ and Player 2 actually inspected).

Denote

$$\begin{aligned} \widetilde{U}(\sigma_1,\sigma_2,h) &:= \int_{s_1,s_2\in(t_{\alpha^*},\infty]} \Bigg(\left(1_{s_1<s_2} + \frac{1}{2}\cdot 1_{s_1=s_2}\right) e^{-r(s_1-t_{\alpha^*})}\left(-c + (1-e^{-\lambda(s_1-t_{\alpha^*})})v_1\right) \\ &\quad + \left(1_{s_1>s_2} + \frac{1}{2}\cdot 1_{s_1=s_2}\right) e^{-r(s_1-t_{\alpha^*})}\left((1-e^{-\lambda(s_2-t_{\alpha^*})})v_2\right)\Bigg) \sigma_1(h)(\mathrm{d}s_1)\sigma_2(h)(\mathrm{d}s_2), \end{aligned}$$

the unconditional expected payoff of Player 1 due to the first inspection done after t_{α^*} (normalized to time t_{α^*}), and by

$$\widetilde{P}(\sigma_1,\sigma_2,h) := \int_{s_1,s_2\in(t_{\alpha^*},\infty]} e^{-r(\min\{s_1,s_2\}-t_{\alpha^*})}\sigma_1(h)(\mathrm{d}s_1)\sigma_2(h)(\mathrm{d}s_2),$$

the expected discounted time (normalized to time t_{α^*}) until the next inspection. Finally, define

$$\Lambda(\sigma_1,\sigma_2,h) := \widetilde{U}(\sigma_1,\sigma_2,h)/\widetilde{P}(\sigma_1,\sigma_2,h),$$

to be the normalized expected payoff to Player 1 until the next inspection. When σ_1 and σ_2 are Markov, $\Lambda(\sigma_1,\sigma_2,h)$ is independent of h, and coincides with Player 1's expected payoff under (σ_1, σ_2).

Player 1's expected payoff under (σ_1, σ_2) is a convex combination of $(\Lambda(\sigma_1, \sigma_2, h))_h$, where the weight of $\Lambda(\sigma_1, \sigma_2, h)$ is given by the unconditional probability that the first inspection done after history h is successful.

Fix now an $\varepsilon > 0$ and a Markov strategy σ_2 of Player 2, and let σ_1 be an $(\varepsilon/2)$-best response of Player 1 to σ_2. Denote $u := \sup_h \Lambda(\sigma_1, \sigma_2, h)$, where the supremum is over all histories h. Since the expected payoff under (σ_1, σ_2) is a convex combination of $(\Lambda(\sigma_1, \sigma_2, h))_h$, it follows that this expected payoff is at most u.

Let h_0 be a history such that $\Lambda(\sigma_1, \sigma_2, h_0) \geq u - \frac{\varepsilon}{2}$. Finally, let σ_1' be the Markov strategy that is defined by $\sigma_1(h_0)$: after each inspection, the distribution of the next inspection time of Player 1 is according to $\sigma_1(h_0)$, shifted to the time of the last inspection.

Since σ_1' and σ_2 are Markov strategies, the expected payoff under (σ_1', σ_2) is $\Lambda(\sigma_1', \sigma_2, h)$ (and this quantity is independent of h). By the choice of h_0, this expected payoff is at least $u - \varepsilon/2$, and hence at least the expected payoff under (σ_1, σ_2) plus ε.

Finally, now let (σ_1, σ_2) be a pair of Markov strategies. We have showed that if Player 1 has a profitable deviation to some strategy, then she has a profitable deviation to a Markov strategy. Thus, to verify that a strategy pair is an MPE, it suffices to verify that no player has a profitable deviation to a Markov strategy. The proposition follows.

Proof of Lemma 2

Substituting $x = e^{-\lambda t}$ in $w(\cdot)$, and differentiating $w(\cdot)$ with respect to x, yields:

$$\frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{-c + (1-x)v_1}{x^{-(r/\lambda)} - x}\right) = \frac{(v_1 - c)x^{(r/\lambda)+1} + \left(-v_1 \cdot \frac{\lambda + r}{\lambda}\right)x + (v_1 - c)(r/\lambda)}{x^{(r/\lambda)+1}\left(x^{-(r/\lambda)} - x\right)^2}. \tag{13}$$

The denominator on the right-hand side of Eq. (13) is positive. Hence, the sign of the derivative is determined solely by the sign of the numerator on the right-hand side. It is easy to verify that this sign is positive at $x = 0$ and negative at $x = 1$. And, because the derivative of the numerator with respect to x is given by $-\frac{1}{\lambda}(\lambda + r)\left(\left(1 - x^{\frac{r}{\lambda}}\right)v_1 + cx^{\frac{1}{\lambda}r}\right)$, which is negative for all $x \in [0,1]$, the numerator is decreasing in x. Let x^* denote the (unique) value for which the numerator is equal to zero. Thus, $w(x)$ is increasing for all $x < x^*$ and decreasing for all $x > x^*$. The result follows from the fact that $x = e^{-\lambda t}$ is decreasing in t.

Proof of Proposition 2

We first prove that the player has an optimal strategy which is Markov. We then show that it is given by inspecting at times $k\tau^*$, $k \in \mathbb{N}$, until the prize is found.

To prove that the player has an optimal strategy which is Markov, we use a result due to Jasso-Fuentes, Menaldi and Prieto-Rumeau (2020) or Stachurski and Zhang (2021), which provide conditions under which a Markov decision problem with general state space and state/action-dependent discount factor admits the dynamic programming principle.

Our first goal is, then, to provide an alternative representation of the decision problem, where time is discrete. Suppose the set of states is $\mathbb{N} \times [0,\infty]$: the first coordinate counts the number of inspections that were already made by the player, and the second coordinate corresponds to time. The initial state is $(0,0)$. The set of actions of the player at state (k,t) for $t < \infty$ is $[t,\infty]$; the interpretation of action $a \in (t,\infty)$ is that the player inspects at time a, the action $a = \infty$ corresponds to no future inspection, and the action $a = t$ is interpreted as doing another inspection at time t. The payoff that corresponds to action a is $-c + (1 - e^{-\lambda(a-t)})v_1$, with probability $1 - e^{-\lambda(a-t)}$ the prize is found and the game terminates, and with probability $e^{-\lambda(a-t)}$ the prize is not found, and the game continues to state $(k+1,a)$, with a state/action-dependent discount factor $e^{-r(a-t)}$. The states (k,∞) for $k \in \mathbb{N}$ are absorbing; the player has no available actions, and the payoff is 0.

The payoff function is continuous, the action sets are compact, and the transitions are continuous, hence by Jasso-Fuentes, Menaldi and Prieto-Rumeau (2020) or Stachurski and Zhang (2021) the dynamic-programming principle applies, and the player has an optimal strategy σ, which is pure and Markov. The reader can verify that σ is also optimal in the original problem.

The definition of τ^* implies that the optimal pure Markov strategy is to inspect at times $k\tau^*$ for $k \in \mathbb{N}$, which completes the proof.

Proof of Lemma 3

As in the proof of Lemma 2, substitute $x = e^{-\lambda t}$ in $w(\cdot)$ and recall that the optimal sampling time τ^* is determined by the value x^* that solves $\frac{\mathrm{d}w(\cdot)}{\mathrm{d}x}(x) = 0$. By Eq. (13) we have that $\frac{\lambda \cdot (x^*)^{\frac{r+\lambda}{\lambda}} + r}{x^*} = \frac{v_1}{v_1 - c}(\lambda + r)$. The right-hand side is decreasing in v_1 and increasing in c, and since $x^* \in (0,1]$, the left-hand side is decreasing in x^*. It follows that x^* is increasing in v_1, and decreasing in c, as claimed.

Proof of Lemma 5

Part (1): Follows immediately from the definition of the delay gain.

Part (2): By Proposition 2, in the single player problem, the optimal strategy is to inspect every τ^* units of time. By the Dynamic Programming Principle of Optimality (which holds because of the argument presented in the proof of Proposition 2), it follows that in the one-player problem, if the continuation payoff following a failed inspection is fixed at $w(\tau^*)$, then the optimal inspection time is τ^* (possibly not unique).

Therefore, the derivative according to t of the expression $e^{-rt}\left(-c+(1-e^{\lambda t})v_1+e^{\lambda t}w(\tau^*)\right)$, which is similar to the expression on the right-hand side of Eq. (1) where $w(t)$ is replaced by $w(\tau^*)$, is equal to zero at $t=\tau^*$. It follows that if $w(\tau^*)=v_0$, then $L'(\tau^*)=0$ and so also the delay gain $L'(\tau^*)/e^{-r\tau^*}=0$.

Uniqueness follows from the fact that $v_0=w(\tau^*)<v_1$, and therefore the expression in Eq. (6) is monotonic in t and can be equal to zero for at most one value of t.

Proof of Proposition 3

Assume to the contrary that $F(\infty):=\lim_{t\to\infty}F(t)<1$. Let $\varepsilon>0$ be sufficiently small such that $\varepsilon<\min\{v_1-c,\frac{v_1}{v_1+v_2}\}$. Let $s>0$ be sufficiently large such that $\frac{F(\infty)-F(s)}{1-F(s)}<\varepsilon$. In words, the conditional probability that a player will inspect after time s, provided she did not inspect up to time s, is smaller than ε. Such s exists since $F(\infty)<1$. Assume, without loss of generality, that s is sufficiently large so that $(1-e^{-\lambda s})v_1>\varepsilon(v_1+v_2)$.

The probability that no player inspects before time s is $(1-F(s))^2>0$. All subsequent calculations assume that no player inspected before time s and are discounted to time s.

Since the probability that a player inspects after time s is $\frac{F(\infty)-F(s)}{1-F(s)}<\varepsilon$, the expected payoff of a player who follows F is smaller than $\varepsilon\cdot(v_1+v_2)$. Since $v_2\geq 0$, the payoff of a player who inspects at s is at least $(1-e^{-\lambda s})v_1$.[30] Yet, by our choice of s we have $(1-e^{-\lambda s})v_1>\varepsilon(v_1+v_2)$, which contradicts the assumption that F is an equilibrium.

Proof of Proposition 4

The idea of the proof is as follows: when both players inspect at t, the identity of the leader and the follower is determined randomly. If there is an atom of size $\widehat{q}$ at t, then the probability that at least one of the players inspects at t is $1-(1-\widehat{q})^2$, and hence the probability that

[30] With probability $1-e^{-\lambda s}$, the player discovers the prize at s, and obtains the payoff v_1; with the complimentary probability, she obtains the equilibrium payoff v_0, which is positive because $v_2\geq 0$.

a specific player will be the leader is half this quantity, which is $\widehat{q}-\frac{\widehat{q}^2}{2}=\widehat{q}(1-\frac{\widehat{q}}{2})$, which is smaller than $\widehat{q}$ and larger than $\widehat{q}(1-\widehat{q})$. By slightly delaying the inspection to $t+\varepsilon$ (and ignoring the probability of an inspection between t and $t+\varepsilon$, which goes to 0 with ε), the player lowers the probability to be the leader because of this atom to $\widehat{q}(1-\widehat{q})$ while by slightly advancing the inspection to $t-\varepsilon$ (again, ignoring the probability of an inspection between $t-\varepsilon$ and t), the player increases the probability to be the leader because of this atom to $\widehat{q}$. In an equilibrium, the player cannot profit by delaying the inspection at t or by advancing it in time, which implies that both eventualities yield her the same payoff. It turns out that this condition pins down t.

We turn this idea into a formal proof. All calculation we do next are discounted to t and conditional that no player inspected before t. The payoff of a player who inspects alone at t is $\gamma_1 := -c+(1-e^{-\lambda t})v_1+e^{-\lambda t}v_0$, and the payoff of the other player is $\gamma_2 := (1-e^{-\lambda t})v_2+e^{-\lambda t}v_0$, where v_0 in both equations denotes the equilibrium payoff. The payoff of the players if they both inspect at t is $\gamma_{12} := \frac{1}{2}\gamma_1+\frac{1}{2}\gamma_2$. Denote by γ_0 the payoff of each player if no-one inspects at t.

Denote the size of the atom at t by $\widehat{q} := F(t)-F(t_-)$, and the conditional probability of inspecting at t if no inspection was done by that player before t by $q := \frac{\widehat{q}}{1-F(t_-)}$. Discounted to t, a player's payoff under F, conditional that no player inspected before t, is $g_1 = q^2\gamma_{12}+q(1-q)\gamma_1+q(1-q)\gamma_2+(1-q)^2\gamma_0$.

For every $\varepsilon>0$ denote by F_-^ε the CDF that moves the atom at t to $t-\varepsilon$, and by F_+^ε the CDF that moves the atom at t to $t+\varepsilon$. Denote by g_-^ε the payoff of a player who adopts the CDF F_-^ε until the first inspection and afterwards follows F, while her opponent follows F (discounted to $t-\varepsilon$, conditional that the prize was not inspected before $t-\varepsilon$). Then

$$g_- := \lim_{\varepsilon\to 0} g_-^\varepsilon = q\gamma_1+q(1-q)\gamma_2+(1-q)^2\gamma_0. \tag{14}$$

For every $\varepsilon>0$ denote by g_+^ε the payoff of a player who adopts the CDF F_+^ε until the first inspection and afterwards follows F, while her opponent follows F (discounted to $t+\varepsilon$, conditional that the prize was not inspected before t). Then

$$g_+ := \lim_{\varepsilon\to 0} g_+^\varepsilon = q\gamma_2+q(1-q)\gamma_1+(1-q)^2\gamma_0. \tag{15}$$

In particular, $g_1=\frac{1}{2}g_-+\frac{1}{2}g_+$. Since F is an equilibrium, $g_-, g_+ \le g_1$. Hence, $g_1=\frac{1}{2}g_-+\frac{1}{2}g_+ \le \frac{1}{2}g_1+\frac{1}{2}g_1=g_1$. It follows that $g_+=g_-=g_1$, and hence $\gamma_1=\gamma_2$, or, equivalently, $e^{-\lambda t}=\frac{v_1-v_2-c}{v_1-v_2}$. The result follows. The equality $\gamma_1=\gamma_2$ further implies that the player is indifferent between being the leader or the follower at t.

Proof of Proposition 5

The proof appears in the text.

Proof of Proposition 6

By Proposition 3, $F(\bar{t}) = 1$. Suppose by way of contradiction that $\bar{t}$ is finite and different from t°. Thus, by Proposition 4 the follower gain at $\bar{t}$ is different from zero. We will derive a contradiction, whether the follower gain is positive or negative.

If the follower gain at $\bar{t}$ is positive, then fixing the strategy of Player 2, it is beneficial for Player 1 to postpone all inspections in the interval $[\bar{t}-\varepsilon, \bar{t}]$ to $\bar{t}+\delta$, for any small $\delta > 0$ (i.e., immediately after $\bar{t}$). By doing so, Player 1 ensures that she becomes the follower conditional on having reached $\bar{t}-\varepsilon$. This improves her payoff because the follower gain is bounded away from zero, while the cost of discounting that is induced by the additional delay is negligible.[31]

Analogously, if the follower gain at $\bar{t}$ is negative, then it is beneficial to advance all inspections in the interval $[\bar{t}-\varepsilon, \bar{t}]$ to $\bar{t}-\varepsilon$. This way, conditional on reaching $\bar{t}-\varepsilon$, Player 1 increases the probability of being the leader from 1/2 to 1 while the probability to find the prize is only negligibly affected.

Proof of Proposition 7

Recall that the support of F is a closed set, so its complement is an open set. Suppose by contradiction that the support of F is not convex. Then there is an interval $(\hat{t}_0, \hat{t}_1)$ that is disjoint from the support of F, yet $\hat{t}_0$ and $\hat{t}_1$ lie in the support of F. It follows that the integral on the right-hand side of (IND) is the same for every $t \in [\hat{t}_0, \hat{t}_1)$, because $\mathrm{d}F(r)$ vanishes inside this interval.

We distinguish between two cases, where $\hat{t}_1 \neq t^\circ$ and where $\hat{t}_0 \neq t^\circ$. Since $\hat{t}_1 \neq \hat{t}_0$, at least one of these cases holds.

Case 1: Suppose that $\hat{t}_1 \neq t^\circ$. Since there is no atom at $\hat{t}_1$, both $1-F(t)$ and the integral on the right-hand side of (IND) are constant on $[\hat{t}_0, \hat{t}_1]$ (and not only on $[\hat{t}_0, \hat{t}_1)$). Denote the coefficient of $(1-F(t))$ in (IND) by

$$h(t) := e^{-rt}\left(-c + (1-e^{-\lambda t})v_1 + e^{-\lambda t}v_0\right), \quad \forall s \in [\hat{t}_0, \hat{t}_1]. \tag{16}$$

The function h is differentiable. Since F is a symmetric MPE strategy, for every $t \in (\hat{t}_0, \hat{t}_1)$

[31] This cost arises due to the delay in future inspections.

the left-hand side of (IND) is larger than or equal to the right-hand side of (IND). It follows that $h(t) \le h(\hat{t}_0) = h(\hat{t}_1), \quad \forall t \in [\hat{t}_0, \hat{t}_1]$. Hence, $h'(\hat{t}_0) \le 0$ and $h'(\hat{t}_1) \ge 0$.

The derivative of h is

$$h'(t) = e^{-rt}\left((r+\lambda)e^{-\lambda t}(v_1 - v_0) - r(v_1 - c)\right). \tag{17}$$

The term e^{-rt} does not change the sign of the product in Eq. (17), and the term in the parenthesis is decreasing with t (since $v_1 > v_0$ by Corollary 2). This contradicts the fact that $h'(\hat{t}_0) \le 0$ and $h'(\hat{t}_1) \ge 0$.

Case 2: Suppose that $\hat{t}_0 \ne t^\circ$. For every $t \le \hat{t}_1$ denote

$$f(t) := e^{-\lambda t} v_1 + (1 - e^{-\lambda t}) v_0, \qquad g(t) := e^{-r(\hat{t}_1 - t)}\left(e^{-\lambda \hat{t}_1} v_1 + (1 - e^{-\lambda \hat{t}_1}) v_0\right). \tag{18}$$

Consider the situation faced by a player at $\hat{t}_0$, conditional on the other player not inspecting before, and discount payoffs to $\hat{t}_0$. Since there is no atom at $\hat{t}_0$, if the player inspects at $\hat{t}_0$, her payoff is $f(\hat{t}_0)$. Since $(\hat{t}_0, \hat{t}_1)$ is disjoint of the support of F, the probability that the other player inspects in the interval $[\hat{t}_0, \hat{t}_1)$ is zero. Hence, if the player inspects at $\hat{t}_1$, her payoff is $g(\hat{t}_0)$.[32]

In equilibrium the player must be indifferent between inspection at times $\hat{t}_0$ and $\hat{t}_1$, and therefore $f(\hat{t}_0) = g(\hat{t}_0)$. But, this is impossible when $\hat{t}_0 \ne \hat{t}_1$. To see this, notice that by Eq. (18) we have that $f(\hat{t}_1) = g(\hat{t}_1)$. However, $f'(t) = -\lambda e^{-\lambda t}(v_1 - v_0) < 0, \quad \forall t \le \hat{t}_1$, and $g'(t) = r e^{-r(\hat{t}_1 - t)}\left(e^{-\lambda \hat{t}_1} v_1 + (1 - e^{-\lambda \hat{t}_1}) v_0\right) > 0, \quad \forall t \le \hat{t}_1$. Therefore, there does not exist a time $t < \hat{t}_1$ such that $f(t) = g(t)$.

Proof of Lemma 6

When $A = 0$, we necessarily have $B = c$, and Eq. (7) reduces to $f(t) \cdot c e^{-\lambda t} = (1 - F(t)) \cdot \left(-D - E(v_0) e^{-\lambda t}\right)$. This equation can be written as $-\frac{f(t)}{1-F(t)} = \frac{-D - E(v_0) e^{\lambda t}}{c e^{\lambda t}}$, whose solution is $F(t) = 1 - e^{G(t) - G(\underline{t})}$, where $G(t) = \frac{D}{c\lambda} e^{\lambda t} + \frac{E(w(\underline{t}))}{c} t$.

When $A \ne 0$, Eq. (7) can be equivalently presented as

$$-\frac{f(t)}{1 - F(t)} = \frac{D + E(w(\underline{t})) \cdot e^{-\lambda t}}{A + B e^{-\lambda t}} = \frac{E(w(\underline{t}))}{B} + \left(D - \frac{A E(w(\underline{t}))}{B}\right) \frac{1}{A + B e^{-\lambda t}}, \quad \forall t \in (\underline{t}, \bar{t}),$$

[32] If $\hat{t}_1 = t^\circ$, the player is indifferent between being the leader and follower, and this equation still holds.

and therefore in the interval $[\underline{t},\bar{t}]$ we have

$$\int\left(-\frac{f(t)}{1-F(t)}\right)\mathrm{d}t=\int\left(\frac{E(w(\underline{t}))}{B}+\left(D-\frac{AE(w(\underline{t}))}{B}\right)\frac{1}{A+Be^{-\lambda t}}\right)\mathrm{d}t.$$

Using the facts that $\int\frac{1}{A+Be^{-\lambda t}}\mathrm{d}t=\left(\frac{t}{A}+\frac{1}{A\lambda}\ln\left(A+Be^{-\lambda t}\right)\right)+K$ (when $A,\lambda\neq 0$ and $A+Be^{-\lambda t}>0$ in the interval of integration) and $\int-\frac{f(t)}{1-F(t)}\mathrm{d}t=\ln(1-F(t))+K$, we obtain:

$$\begin{aligned}\ln(1-F(t))&=\frac{E(w(\underline{t}))}{B}t+\left(D-\frac{AE(w(\underline{t}))}{B}\right)\left(\frac{t}{A}+\frac{1}{A\lambda}\ln\left(A+Be^{-\lambda t}\right)\right)+K\\&=\frac{D}{A}t+\frac{1}{\lambda}\left(\frac{D}{A}-\frac{E(w(\underline{t}))}{B}\right)\ln\left(A+Be^{-\lambda t}\right)+K,\quad\forall t\in[\underline{t},\bar{t}].\end{aligned}$$

Denoting $G(t)=\frac{D}{A}t+\frac{1}{\lambda}\left(\frac{D}{A}-\frac{E(w(\underline{t}))}{B}\right)\ln\left(A+Be^{-\lambda t}\right)$, we obtain Eq. (8). Thus, when $A\neq 0$, the only continuous solution of Eq. (7) is given by Eq. (8).

Proof of Theorem 1

Denote $x(t)=e^{-\lambda t}$. The proof consists of four steps. We first argue that F^S is a symmetric MPE.

Step 1: F^S defined in Eq. (10) is a symmetric MPE strategy.

By Proposition 5, the payoff under F^S is $w(t^\circ)$. As mentioned before, the condition $v_2\leq\tilde{v}_2$ implies that $t^\circ\leq\tau^*$, and because by Lemma 2 $w(\cdot)$ is quasiconcave and attains its maximum at τ^*, it follows that no player can profit by deviating and inspecting at some time $t<t^\circ$ (a deviation that yields $w(t)$). And, deviating and inspecting after time t° only increases the probability that the player is the follower at t°. This is not profitable because, by Proposition 4, the follower gain at t° is zero.

To prove that F^S is the unique symmetric MPE, we fix a symmetric MPE strategy $F\neq F^S$, and show that (i) the support of F is included in $\left[\min\{t^\circ,t^D\},\max\{t^\circ,t^D\}\right]$, where t^D is the value that sets the delay gain equal to zero (that is, such that $D+E(v_0)\cdot x\left(t^D\right)=0$), (ii) $t^D\leq t^\circ$, and (iii) $t^D\geq t^\circ$. It follows that the probability mass in F must be concentrated at t°, which implies that it coincides with F^S. A contradiction.

Let $\underline{t}$ and $\bar{t}$ denote the lower and upper bounds, respectively, of the support of an MPE strategy F, and let $v_0=w(\underline{t})$ denote the equilibrium payoff (see Proposition 5). Recall that $A+Bx(t^\circ)=0$, that is at t° the follower gain vanishes.

Step 2: The support of F is included in the interval $\left[\min\{t^\circ,t^D\},\max\{t^\circ,t^D\}\right]$.

By Corollary 2, and by Lemma 5, the delay gain is decreasing in t. Since $v_2 \le \widetilde{v}_2 < v_1$, by Lemma 4, the follower gain is decreasing in t as well. Therefore, for every $t < \min\{t^\circ, t^D\}$ both $A + Bx(t)$ and $D + E(v_0) \cdot x(t)$ are positive, and for every $t > \max\{t^\circ, t^D\}$ both $A + Bx(t)$ and $D + E(v_0)) \cdot x(t)$ are negative. By Observation 1, these two quantities must have different signs at any t in the support of F, and the claim follows.

Step 3: $t^D \le t^\circ$.

Suppose by way of contradiction that $t^D > t^\circ$. Since $F \ne F^S$, and since F assigns probability 0 to the interval $[0, t^\circ)$, we have $\bar{t} > t^\circ$. By Lemma 4, there is no atom at $\bar{t}$. Moreover, $\bar{t}$ is finite.[33] It follows that there is some time $\bar{t} - \varepsilon$ that slightly precedes $\bar{t}$, such that the players inspect with a positive probability between $\bar{t} - \varepsilon$ and $\bar{t}$ (formally, $F(\bar{t}) > F(\bar{t} - \varepsilon)$).

Fix $\varepsilon > 0$ sufficiently small, such that $\bar{t} - \varepsilon > t^\circ$. Since by Lemma 4 the follower gain is decreasing in t, it is negative on $(\bar{t} - \varepsilon, \bar{t}]$, and by Observation 1, the delay gain is positive on this interval. Suppose a player considers advancing all inspections from the interval $(\bar{t} - \varepsilon, \bar{t}]$ to $\bar{t} - \varepsilon$, *conditional on no player inspecting before* $\bar{t} - \varepsilon$. This has two effects on the player's payoff. The first effect is the delay gain. Because the length of the interval $(\bar{t} - \varepsilon, \bar{t}]$ is ε, this gain is in the order of magnitude of ε. The second effect is the follower gain. By advancing the inspection to $\bar{t} - \varepsilon$, the player ensures that she becomes the leader rather than the follower. Because the follower gain is negative on $(\bar{t} - \varepsilon, \bar{t}]$, switching from being the follower on $(\bar{t} - \varepsilon, \bar{t}]$ to being the leader at $\bar{t} - \varepsilon$, yields a positive payoff that does not vanish as ε becomes small. Consequently, for ε sufficiently small, the overall effect of advancing the inspection is dominated by the follower gain, which is positive. Thus, the deviation considered above is profitable, contradicting the assumption that F is an MPE.

Step 4: $t^D \ge t^\circ$.

Suppose by way of contradiction that $t^D < t^\circ$. Recall that by the definition of t^D we have $x\left(t^D\right) = \frac{-D}{E(v_0)} = \frac{r}{\lambda + r} \frac{v_1 - c}{v_1 - v_0}$. By Lemma 5, $x(\tau^*) = \frac{r}{\lambda + r} \frac{v_1 - c}{v_1 - w(\tau^*)}$. As before, the fact that $v_2 \le \widetilde{v}_2$, implies that $\tau^* \ge t^\circ > t^D$. We therefore have $x(\tau^*) = e^{-\lambda \tau^*} < e^{-\lambda t^D} = x\left(t^D\right)$, or, equivalently,

$$\frac{r}{\lambda + r} \frac{v_1 - c}{v_1 - w(\tau^*)} < \frac{r}{\lambda + r} \frac{v_1 - c}{v_1 - v_0}.$$

It follows that $w(\tau^*) < v_0 = w(\underline{t})$. This is a contradiction because by definition τ^* is the maximizer of $w(\cdot)$.

[33]To see this, suppose that $\bar{t} = \infty$. Fix some $\delta > 0$ and consider a time interval $(x, x + \delta)$, parametrized by $x > t^\circ$, in which inspection occurs with a positive probability under F. Because the follower gain is decreasing in t, it is negative for all $t \in (x, x + \delta)$. By Eq. (6), the delay gain is negative on the interval $(x, x + \delta)$, for every large enough x. However, this is a contradiction to Observation 1.

Proof of Theorem 2

The proof consists of six steps. In the first three steps we show that $\tau^* < t^\circ$ and that the strategy F defined in Eq. (8) with $[\underline{t}, \bar{t}] = [\tau^*, t^\circ]$ is a symmetric MPE. In the last three steps, we show that this is the unique symmetric MPE.

Step 1: If $t^\circ < \infty$ then $A < 0$.

Since t° is finite, and since by Proposition 4 we have $e^{-\lambda t^\circ} = \frac{v_1-(v_2+c)}{v_1-v_2}$, it follows that the ratio $\frac{v_1-(v_2+c)}{v_1-v_2}$ is in $(0,1)$. If $v_1 < v_2$ this ratio is larger than 1, if $v_1 = v_2$ this ratio is undefined, and if $v_1 \in (v_2, v_2+c]$ this ratio is nonpositive. Therefore $v_1 > v_2 + c$, or, equivalently, $A < 0$.

Step 2: $\tau^* < t^\circ$.

If $t^\circ = \infty$ the claim holds trivially. If $t^\circ < \infty$, then by Step 1 we have $A < 0$. Because $v_2 > \widetilde{v}_2 = v_1 - \frac{c}{1-e^{-\lambda\tau^*}}$, it follows that $e^{-\lambda\tau^*} > 1 - \frac{c}{v_1-v_2} = -\frac{A}{B}$. Thus, $\tau^* = -\frac{1}{\lambda}\ln(e^{-\lambda\tau^*}) < -\frac{1}{\lambda}\ln(-B/A) = t^\circ$.

Step 3: The strategy F defined by Eq. (8) with $[\underline{t}, \bar{t}] = [\tau^*, t^\circ]$ is a symmetric MPE.

If $v_2 > v_1 - c$, then $A + Be^{-\lambda t} = (v_2 - v_1 + c)(1-e^{-\lambda t}) + ce^{-\lambda t}$ is positive for all t. If $v_2 < v_1 - c$, then, since $v_2 > \widetilde{v}_2$, the quantity $A + Be^{-\lambda t}$ is positive for all t in the support $[\tau^*, t^\circ)$. This is because by Lemma 4 the follower gain is decreasing and is equal to zero at $t = t^\circ$. Therefore, the conditions of Lemma 6 are satisfied, and the function F given by Eq. (8) is a solution of (IND).

To show that this solution is a CDF, we need to show that f is non-negative on the support $[\underline{t}, \bar{t}]$ and that $F(\bar{t}) = 1$. By Lemma 5 and Corollary 2, the delay gain, $D + E(v_0)e^{-\lambda t}$, is negative for all t in the support. Therefore, Eq. (7) implies that the derivative f is non-negative on the support. Next, inspection of Eq. (8) reveals that when $\bar{t} = t^\circ$, we have that $F(\bar{t}) = F(t^\circ) = 1$.[34]

It remains to show that no player wants to deviate and inspect at a time outside the support of F. To see this, recall that by Proposition 5, the payoff under F is equal to $w(\underline{t})$. Since $w(\cdot)$ is maximized at τ^*, no player can profit from inspecting before $\underline{t} = \tau^*$. And, because (i) the equilibrium payoff is positive (see footnote 3), and (ii) $F(t^\circ) = 1$, no player can profit by inspecting after time t°.

We now turn to show that the equilibrium described above is unique.

Step 4: When $v_2 > \widetilde{v}_2$, the lower bound $\underline{t}$ of the support of any symmetric MPE is equal to τ^*.

[34] If $A = 0$ then $t^\circ = \infty$, and so $F(t^\circ) = 1$; if $A > 0$ then $t^\circ = \infty$, and so $F(t^\circ) = 1$; and, if $A < 0$ then t° is determined such that $A + Be^{\lambda t^\circ} = 0$, and so again $F(t^\circ) = 1$.

Suppose that the lower bound $\underline{t}$ of the support of a symmetric MPE is smaller than τ^*. By Lemma 4 the follower gain at $\underline{t}$ is positive. We show that when $\underline{t} < \tau^* < t^\circ$, the delay gain is also positive at $\underline{t}$, which implies a contradiction to Observation 1. To see why the delay gain is positive, recall that by Lemma 5, if the continuation payoff upon a failed inspection was $w(\tau^*)$, then the delay gain would have been zero at τ^* and positive at $\underline{t}$. However, the continuation payoff following a failed inspection is $v_0 = w(\underline{t}) < w(\tau^*)$, which implies that a small delay in inspection (conditional on the other player not inspecting in the meantime) is even *more* profitable than in the case in which the continuation payoff is $w(\tau^*)$. Thus, the delay gain is also positive at $\underline{t}$ when the continuation payoff is $v_0 = w(\underline{t})$.

We have therefore established that $\underline{t} \geq \tau^*$. Suppose by way of contradiction that $\underline{t} > \tau^*$. Proposition 5 implies that the payoff under the symmetric MPE is equal to $w(\underline{t})$. However, by deviating to a Markov strategy that inspects at τ^*, a player can guarantee $w(\tau^*) > w(\underline{t})$. A contradiction.

Step 5: When $v_2 > \widetilde{v}_2$, the upper bound $\underline{t}$ of the support of any symmetric MPE is no greater than t°.

By Lemma 4, at $t = t^\circ$ the follower gain is zero. By Step 4 and Lemma 5, at $t = \tau^*$ the delay gain is zero. By Step 2 we have $\tau^* < t^\circ$. Since both follower gain and delay gain are decreasing, the upper bound of the support of the MPE strategy cannot be greater than t°.

Step 6: The strategy F defined in Eq. (8) with $[\underline{t}, \bar{t}] = [\tau^*, t^\circ]$ is the *unique* symmetric MPE.

Eq. (IND) holds for every t in the support of F. By Proposition 7 the support of F is convex, hence Eq. (IND) holds for every $t \in [\underline{t}, \bar{t}] = [\tau^*, \bar{t}]$.

By Proposition 5, the equilibrium payoff under F is $v_0 = w(\underline{t}) = w(\tau^*)$. The function F defined in Eq. (8) is a continuous solution of Eq. (IND). The Theory of Distributions states that if the integral equation (IND) has a solution that is a continuous distribution, then it is the unique solution (see, e.g., Theorem 2.24 in Hörmann and Steinbauer (2009)). Therefore, F is the only solution of (IND) for $\bar{t}$. Inspecting the formula of F, we see that if $\bar{t} < t^\circ$, then $F(\bar{t}) < 1$, which is a contradiction. This implies that $\bar{t}$ must coincide with t°, and the proof is complete.

Proof of Proposition 8

By Theorem 1, when $v_2 \leq \widetilde{v}_2$, the symmetric MPE involves synchronous inspection.

When $v_2 = 0$, the socially optimal inspection time is τ^* because the social gain from discovering the prize is v_1, while under the equilibrium inspection time is $t^\circ < \tau^*$. Hence, for $v_2 = 0$, in equilibrium, the players inspect too frequently relative to the social optimum.

When $v_2 = \widetilde{v}_2$ we have $t^\circ = \tau^*$, and hence the equilibrium inspection time is τ^*. Lemma 3 implies that the optimal inspection time for one player with a prize that is equal to $v_1 + v_2$ is more frequent than τ^*.

The MPE inspection time t° increases continuously in v_2 (see Eq. (4)), and the optimal inspection time for one player with a prize that is equal to v_1+v_2 decreases continuously in v_2 (see Lemma 3). Therefore, by the discussion above, these two curves (of the MPE inspection time t°, and of the optimal inspection time for one player) intersect for some value $\widehat{v}_2 < \widetilde{v}_2$. It follows that when $v_2 = \widehat{v}_2$ the two players inspect synchronously at the socially optimal frequency.

By Theorem 2, for values $v_2 \geq \widetilde{v}_2$ the players start inspecting at τ^*, where τ^* is the optimal sampling time for a single player when the prize is v_1. As for the case where $v_2 = \widetilde{v}_2$, this implies that the players inspect too rarely relative to the social optimum.

Proof of Proposition 9

The proof of the proposition relies on the following lemma.

Lemma 7. *Let $g(t) = \frac{1-e^{-\lambda t}}{e^{rt}-e^{-\lambda t}}$. The function $g(t)$ is decreasing in t, for every $t > 0$, $r > 0$, and $\lambda > 0$.*

Proof. We will show that the derivative of $g(t)$ is negative for all $t > 0$, $r > 0$, and $\lambda > 0$. That is:

$$g'(t) = \frac{\lambda e^{-\lambda t}(e^{rt} - e^{-\lambda t}) - (re^{rt} + \lambda e^{-\lambda t})(1-e^{-\lambda t}))}{(e^{rt} - e^{-\lambda t})^2} < 0, \quad \forall r > 0, t > 0, \lambda > 0.$$

Because the denominator is positive, it suffices to demonstrate that the numerator is negative. This will hold as soon as

$$re^{rt}(1-e^{-\lambda t}) > \lambda e^{-\lambda t}(e^{rt} - 1), \quad \forall t > 0, r > 0, \lambda > 0. \tag{19}$$

To see why this inequality holds, note that when $r = 0$, both sides of Eq. (19) are equal to 0. Hence, it is sufficient to show that the derivative of the left-hand side of Eq. (19) with respect to r is greater than the derivative of the right-hand side. This is equivalent to proving that:

$$tr + 1 > \frac{\lambda t e^{-\lambda t}}{1 - e^{-\lambda t}}, \quad \forall r > 0,\ t > 0,\ \lambda > 0.$$

The left-hand side is minimized at $r = 0$. Thus, it suffices to show that $1 - e^{-\lambda t} \geq \lambda t e^{-\lambda t}$, or equivalently $e^{\lambda t} \geq 1 + \lambda t$, for all $t > 0, \lambda > 0$. This holds true because $e^x \geq 1 + x$ for all $x \in \mathbb{R}$. $\square$

We turn to the proof of Proposition 9 and distinguish between two cases.

Case 1: $v_2 > \widetilde{v}_2$. In this case, the equilibrium payoff in the symmetric MPE, as characterized by Theorem 2, is given by $w(\tau^*)$. We will show that Player 1 can secure this payoff regardless of the strategy of Player 2.

Suppose that Player 1 adopts the Markovian strategy that inspects at time τ^* with probability 1. If Player 2 never inspects, Player 1's payoff is $w(\tau^*)$. By the one-shot deviation principle, we only need to show that a single inspection by Player 2 at time $t < \tau^*$ weakly increases Player 1's payoff.[35] Therefore, it suffices to show that $e^{-rt}\big((1-e^{-\lambda t})v_2 + e^{-\lambda t}w(\tau^*)\big) \geq w(\tau^*), \quad \forall t < \tau^*$, or, equivalently,

$$\frac{1-e^{-\lambda t}}{e^{rt}-e^{-\lambda t}}v_2 \geq w(\tau^*), \quad \forall t < \tau^*.$$

Because $v_2 > \widetilde{v}_2$, and because $\frac{1-e^{-\lambda t}}{e^{rt}-e^{-\lambda t}}$ is decreasing in t (as per Lemma 7), it suffices to establish that

$$\frac{1-e^{-\lambda\tau^*}}{e^{r\tau^*}-e^{-\lambda\tau^*}}\widetilde{v}_2 \geq w(\tau^*). \tag{20}$$

Substituting $\widetilde{v}_2 = v_1 - \frac{c}{1-e^{-\lambda\tau^*}}$ and $w(\tau^*) = \frac{(1-e^{-\lambda\tau^*})v_1 - c}{e^{r\tau^*}-e^{-\lambda\tau^*}}$ into Eq. (20) and simplifying reveals that the two sides of Eq. (20) are equal. This concludes the proof for the case when $v_2 > \widetilde{v}_2$.

Case 2: $v_2 \leq \widetilde{v}_2$. In this case, the equilibrium payoff in the symmetric MPE, as characterized by Theorem 1, is given by $w(t^\circ)$. As before, we will demonstrate that Player 1 can secure this payoff regardless of the strategy of Player 2.

Suppose that Player 1 adopts the Markovian strategy that inspects at time t° with probability 1. As in Case 1, by the one-shot deviation principle, we need to show that a single inspection of Player 2 at $t < t^\circ$ weakly increases Player 1's payoff. That is, it suffices to show that $e^{-rt}\big((1-e^{-\lambda t})v_2 + e^{-\lambda t}w(t^\circ)\big) \geq w(t^\circ) \quad \forall t < t^\circ$, or equivalently,

$$w(t^\circ) \leq \frac{(1-e^{-\lambda t})}{e^{rt}-e^{-\lambda t}}v_2 \quad \forall t < t^\circ. \tag{21}$$

To see why Eq. (21) holds, recall that at time t° the follower gain is 0. Therefore, $w(t^\circ)$ is equal to Player 1's payoff if Player 2 is the one who inspects at t° (rather than Player 1).

[35] To see why the one-shot deviation principle applies, suppose that the game was zero-sum: Player 2 gains c for every inspection of Player 1, she loses v_1 if Player 1 finds the prize, and she loses v_2 if she (Player 2) finds the prize. In this case, the one-shot deviation principle implies that if Player 2 has no profitable one-shot deviation, then she does not have any profitable deviation. Since the game is zero-sum, this implies that if any one-shot deviation of Player 2 does not decrease Player 1's payoff, then no deviation of Player 2 decreases Player 1's payoff.

Hence, $w(t^\circ) = e^{-rt^\circ}\big((1-e^{-\lambda t^\circ})v_2 + e^{-\lambda t^\circ} w(t^\circ)\big)$, or equivalently,

$$w(t^\circ) = \frac{(1-e^{-\lambda t^\circ})}{e^{rt^\circ} - e^{-\lambda t^\circ}} v_2. \tag{22}$$

Taken together, Eq. (22) and Lemma 7 imply that the inequality (21) is satisfied for all $t < t^\circ$.