# Sequential Sampling Equilibrium 

## Duarte Gonçalves*


#### Abstract

This paper introduces an equilibrium framework based on sequential sampling in which players face strategic uncertainty over their opponents' behavior and acquire informative signals to resolve it. Sequential sampling equilibrium delivers a disciplined model featuring an endogenous distribution of choices, beliefs, and decision times, that not only rationalizes well-known deviations from Nash equilibrium, but also makes novel predictions supported by existing data. It grounds a relationship between empirical learning and strategic sophistication, and generates stochastic choice through randomness inherent to sampling, without relying on indifference or choice mistakes. Further, it provides a rationale for Nash equilibrium when sampling costs vanish.


Keywords: Game Theory; Sequential Sampling; Information Acquisition; Response Time; Bayesian Learning; Strategic Uncertainty; Statistical Decision Theory.
JEL Classifications: C70, D83, D84, C41.

[^0]
## 1. Introduction

When faced with a choice, such as which TV show to watch, snack to buy, or savings account to open, decision-makers spend time and cognitive effort to resolve uncertainty about which alternative is best. This gives rise to a fundamental trade-off between speed, how long it takes to choose, and accuracy, the likelihood of choosing the best alternative.

Rooted in Wald's seminal work (1947), sequential sampling emphasizes this trade-off between speed and accuracy. It models the decision-maker's reasoning as sampling informative signals, reflecting the notion that individuals draw upon memory to guide their choices, even in novel settings-a premise supported by neurological evidence. ${ }^{1}$ Initially adopted by cognitive sciences to explain reaction times in perception problems (Ratcliff, 1978; Luce, 1986), sequential sampling stands as a cornerstone for understanding the relationship between choice and decision time in various domains, from consumer behavior to risky or intertemporal choice. ${ }^{2}$ Its success stems from its ability to rationalize established empirical regularities, particularly a time-revealed indifference ${ }^{3}$ whereby slower choices are associated with weaker preference intensity and greater choice randomness (Fudenberg et al., 2018; Alós-Ferrer et al., 2021).

When making decisions in strategic settings like online auction bidding, protest participation decisions, or stock market transactions, individuals often grapple with uncertainty about others' behavior, and spend time and cognitive effort to address this strategic uncertainty. And since the cost and benefit of reasoning may vary across situations, the time and effort committed to resolving this uncertainty is itself endogenous. These observations resonate with experimental evidence from strategic environments: Existing work has shown that in dominance-solvable games, faster decisions are associated with less strategically sophisticated actions (Agranov et al., 2015; Rubinstein, 2016; Recalde et al., 2018), and that stronger incentives entail longer decision times and greater strategic sophistication (Alaoui and Penta, 2016; Alaoui et al., 2020; Alós-Ferrer and Bruckenmaier, 2021; Esteban-Casanelles and Gonçalves, 2022). Moreover, in binary-action games longer decision times tend to be associated with indifference (Schotter and Trevino, 2021; Frydman and Nunnari, 2023). While these findings agree with our understanding of individual decision-making, they are difficult to reconcile with existing equilibrium concepts.

[^1]This paper introduces an equilibrium framework based on sequential sampling in which players face strategic uncertainty and acquire informative signals to resolve it. Players have a prior belief about others' distribution of actions and, before taking an action, they acquire costly signals about others' behavior. These signals are assumed to be informative of others' behavior and therefore to have informational value. Then, strategic uncertainty drives information acquisition and each player optimally trades-off the cost and benefit to sample. As optimal sequential sampling renders players' action distributions dependent on their opponents', a sequential sampling equilibrium corresponds to a fixed-point, a consistent distribution of actions of all players. This delivers a disciplined model featuring an endogenous distribution of choices, beliefs, and decision times, that not only rationalizes empirical patterns relating choices and decision times and well-known deviations from Nash equilibrium, but also makes novel predictions supported by existing data. Moreover, it provides a rationale for Nash equilibrium as a limit case when sampling costs vanish.

The solution concept builds on an individual decision-making model of sequential sampling in a rich environment of choice under uncertainty. Players effectively act as decision-makers: each takes as given others' uncertain behavior, characterized by an unknown distribution. Sequential sampling serves as a stylized model of stepwise reasoning about others' behavior, occurring prior to making a decision. ${ }^{4}$ Each player faces an optimal stopping problem, trading-off informational gains and costs. Upon stopping, players choose an action to maximize their expected payoff, given their posterior beliefs. Optimal sequential sampling yields stochastic choice through the randomness inherent to sampling, without relying on indifference or choice mistakes. The actions chosen upon stopping depend on posterior beliefs, informed by the sampled observations, whose distribution depends on others' behavior. For simplicity, we focus on the case in which players can sample at a cost from their opponents' choice distribution and defer the discussion of more general information structures.

A sequential sampling equilibrium determines an endogenous joint distribution of actions, beliefs, and stopping time. Equilibrium emerges as a consistency condition on the distribution over players' actions arising from the fact that signals sampled are informative. We show a sequential sampling equilibrium always exists. The proof follows from the novel observation that each player's optimal stopping time is uniformly bounded with respect to opponents' distribution of actions-which renders this into a computationally tractable finite-horizon problem. While our equilibrium definition is in terms of action distributions, a player's optimal sequential sampling

[^2]together with the equilibrium distribution of opponents' actions pins down an endogenous distribution of actions, beliefs, and stopping time. Furthermore, as a player is uncertain about others' (equilibrium) distribution of actions, and the latter corresponds to but one of the possible action distributions the player entertains, equilibrium stopping time does not correspond to the players' expected costs of information.

Sequential sampling equilibrium provides a rationale for the relationship between higher incentives, longer decision times, and more sophisticated play. While sequential sampling equilibrium does allow for players to choose non-rationalizable with positive probability, we prove that when scaling up players' payoffs, only $k$-rationalizable actions are chosen with positive probability at any equilibrium, where the order of rationalizability depends monotonically on the scaling factor. This connection between empirical learning and strategic sophistication arises directly from the fact that higher payoffs induce longer decision times, leading players to sample sufficiently to learn and choose only k-rationalizable actions.

We then turn to binary action games and examine how relative payoffs influence the joint distribution of choices and decision time. We establish comparative statics results for sequential sampling. First, increasing the payoffs to a given action leads to that action being chosen more often and faster and the other less often and slower, a finding that generalizes beyond two-action settings. Second, an increase in the underlying probability that an action is optimal leads to an analogous result. These results allow us to prove two behavioral implications of sequential sampling equilibrium: that the frequency with which an action is chosen increases in its payoffs, and that the opponent chooses the best response to that action more often and faster. If the former corresponds to a well-documented deviation from Nash equilibrium in experimental literature (e.g. Goeree and Holt, 2001), the latter provides a novel prediction on how time relates to choice, which we find to be borne out by existing experimental evidence.

Sequential sampling equilibrium also has implications for players' equilibrium beliefs. Experimental evidence has suggested that beliefs about others' behavior are often biased (Costa-Gomes and Weizsäcker, 2008), appear stochastic (Friedman and Ward, 2022), and depend on own incentives even when others' behavior is held fixed (Esteban-Casanelles and Gonçalves, 2022). All these patterns are implied by sequential sampling, where beliefs upon stopping will typically be biased toward the prior, stochastic, as they depend on the realized observations players sample, and payoff-dependent, given these affect when players stop sampling. To go beyond these general properties, we consider the case of Beta-distributed priors and prove that sequential sampling
equilibrium provides a foundation for time-revealed indifference observed in binary action games. Specifically, we show that the longer the decision time, the closer is the player to being indifferent between taking either action-a game-theoretic counterpart to Fudenberg et al. (2018). Furthermore, we uncover monotone comparative statics on how beliefs respond to payoff changes. In a nutshell-and recalling that stopping beliefs are stochastic-when payoffs to a player's action increase, the opponent's equilibrium beliefs shift (in a first-order stochastic dominance) toward assigning a higher probability to that action being chosen.

Sequential sampling also provides a new rationale for Nash equilibrium, based on costly information acquisition. While optimal stopping implies that conditional on stopping observations are neither independent nor identically distributed, it is possible to use martingale theory to show that sequential sampling equilibria nevertheless converge to Nash equilibria. However, not all Nash equilibria can be reached through this approach: those involving weakly dominated actions cannot, while pure strategy Nash equilibria that don't involve such actions can.

We also explain how this solution concept can be seen as a steady state of a dynamic process. Specifically, we show how sequential sampling equilibria coincide with the steady states of the distribution of play of short-lived players who sequentially sample from data on past play, and obtain global asymptotic convergence results for a generic class of $2 \times 2$ games. This parallels the role of Nash equilibria in scenarios where these short-lived players have frictionless access to the entirety of past data (Fudenberg and Kreps, 1993; Fudenberg and Levine, 1998).

Finally, we conclude with a discussion of variations to the model, including extensions to incomplete information games and more general information structures. It is straightforward to adjust the solution concept to Bayesian games by having samples include information on the realized actions as well as the state. This can be interpreted as making inferences about a specific context by also relying on information about behavior in similar settings. An analogous result to that of convergence to Nash equilibrium ensues: limit points of Bayesian sequential sampling equilibria as sampling costs vanish are Bayesian Nash equilibria. Furthermore, we provide an extension to more general information structures, accommodating situations in which recollections are noisy, some players' signals are silent about a subset of their opponents, or where in general players' ability to distinguish between opponents' actions or types is limited. In the latter case, indistinguishable action profiles constitute an analogy class and it is shown that limit points of a sequence of equilibria with vanishing costs are analogy-based expectations equilibria (Jehiel, 2005; Jehiel and Koessler, 2008).

To summarize, sequential sampling equilibrium constitutes a flexible equilibrium framework for analyzing strategic interaction. It provides a rationale for standard solution concepts, accounts for several behavioral patterns that have been documented in experiments, and makes novel predictions not just regarding choices that individuals make in strategic settings, but for timedstochastic choice data, the joint distribution of choices, beliefs, and decision times.

### 1.1. Related Literature

This paper is related to two broad literatures: sequential sampling and information acquisition in games.

Sequential Sampling. The study of optimal sequential sampling can be traced back to the seminal works of Wald (1947) and Arrow et al. (1949). Sequential sampling has since been used as a modeling device in cognitive psychology and neuroscience to ground a relation between choice and decision time, ${ }^{5}$ and, in particular, to model choice based on memory retrieval (Gold and Shadlen, 2007; Shadlen and Shohamy, 2016; Duncan and Shohamy, 2020; Biderman et al., 2020). Alaoui and Penta (2022) provide an axiomatic foundation of sequential sampling as a representation of iterative of reasoning. Fudenberg et al. (2018) consider a binary-action problem in which a decision-maker sequentially acquires information the payoff difference. They show that at longer stopping times, the agent is closer to being indifferent between the two actions. ${ }^{6}$ Alós-Ferrer et al. (2021) examine the general relation between time-revealed indifference and stochastic choice primitives.

The individual decision-making framework of sequential sampling motivated the experimental study of decision times in games. This led to establishing a number of regularities, such as a positive association between decision times and the strategic sophistication of actions chosen-as given by level- $k$ model (Nagel, 1995; Stahl and Wilson, 1995) or the highest level of $k$-rationalizabilityin dominance-solvable games (see e.g. Agranov et al., 2015; Rubinstein, 2016; Alaoui et al., 2020; Alós-Ferrer and Bruckenmaier, 2021; Gill and Prowse, 2023), and that scaling up incentive levels causally increases decision times and leads to more sophisticated play (Esteban-Casanelles and Gonçalves, 2022). Additionally, as in individual decision-making, response times also reveal

[^3]indifference in global games (Schotter and Trevino, 2021; Frydman and Nunnari, 2023).
Sequential sampling equilibrium adopts sequential sampling to model belief formation in strategic settings, providing a relation between stopping time, beliefs, and choices. It contributes to this theoretical literature with novel results in problems with multiple available actions: comparative statics results on how choices and stopping time relate to payoffs in general decision-problems with arbitrary payoff correlation across actions. It rationalizes the relation between incentives, decision times, and strategic sophistication of choices that has been documented in the experimental literature, Further, in the binary-action case that has been the focus of much of the literature, our model relates the true data generating process to the distribution of choices, stopping times, and posterior beliefs, and obtain a time-revealed indifference prediction in a tractable discrete-time environment, which we show bears out in Friedman and Ward's (2022) data.

Information Acquisition in Games. There is a growing literature on equilibrium solution concepts featuring information acquisition. Osborne and Rubinstein (2003) suppose each player observes a fixed number of samples from their opponents' equilibrium distribution of actions, and the mapping from samples to actions is exogenously specified. Salant and Cherry (2020) study a special case of this solution concept in mean-field games with binary actions, while keeping the sampling procedure exogeneous: players employ unbiased estimators and best-respond to the obtained estimate. ${ }^{7}$ Osborne and Rubinstein (1998) examine a similar notion of equilibrium, where players receive a fixed number of samples from the payoffs of each of their actions and choose the action with the highest average payoff in the sample. More broadly, these correspond to a form of self-confirming equilibrium (Fudenberg and Levine, 1993; Battigalli et al., 1992) in which the feedback function is fixed.

In contrast to these, sequential sampling equilibrium (1) endogenizes the sampling and (2) adds a time dimension via its sequential nature, enabling results regarding the joint distribution of stopping time and choices that cannot be captured with exogenous sampling. If the latter provides the basis for the time-revealed indifference, the former grounds the relationship between payoffs, decision times, and the level of strategic sophistication (in the sense of $k$-rationalizability) of the actions chosen in equilibrium.

The literature also studied solution concepts with costly information acquisition. Yang (2015) examines a coordination game in which players acquire flexible but costly information about an

[^4]exogenous payoff-relevant parameter. As in much of the rational inattention literature (Sims, 2003; Matějka and McKay, 2015), the cost of information is given by prior entropy reduction. Matějka and McKay (2012) and Martin (2017) study pricing games with a similar approach. Denti (2023) allows for players to obtain correlated information under general information cost (as in Caplin and Dean, 2015). Hébert and La'O (2023) study this solution concept in mean-field games.

Our paper provides the first solution concept in which the cost of information acquisition is experimental (Denti et al., 2022), with information acquisition corresponding to costly sequentially sampling from an information structure. While the sequential information acquisition can be studied from a static, ex-ante perspective (Morris and Strack, 2019; Bloedel and Zhong, 2021; Hébert and Woodford, 2023), there are two conceptual features distinguishing sequential sampling equilibrium-beyond, of course, results specific to stopping time. First, in these papers players hold beliefs and can learn about their opponents' action realizations. Second, players' equilibrium beliefs are correct, and so, absent uncertainty about exogenous parameters, equilibria correspond to Nash equilibria of the underlying normal-form game.

In our framework, players are uncertain-neither correct or incorrect-about the prevailing distribution of actions of their opponents. While actions yet to be taken are not learnable, a prevailing stable distribution of opponents' actions is. Further, if the choice of an information structure by our players bears a cost proportional to the expected stopping time, our analysis speaks to the joint distribution of stopping times and choices as determined in equilibrium-noting that the equilibrium stopping time is distributed not by the measure induced by players' prior beliefs, but by that arising from the equilibrium distribution of their opponents' actions (which has measure zero according to their prior).

Finally, we comment on the relation to the work on learning with misspecification, chiefly Esponda and Pouzo's (2016) Berk-Nash equilibrium. This solution concept allows for general forms of misspecification of the players' prior beliefs and is not restricted to either normal-form or complete information games. There, players best-respond to their equilibrium beliefs, those in the support of players' priors that minimize the Kullback-Leibler divergence to equilibrium gameplay, which can be taken as arising as the limit case of Bayesian learning with potentially misspecified priors (see Fudenberg et al., 2021).

## 2. Sequential Sampling Equilibrium

### 2.1. Setup

Preliminaries. Let $\Gamma=\langle I, A, u\rangle$ denote a normal-form game, where $I$ denotes a finite set of players or roles, with generic elements $i, j$ and where $-i$ denotes $I \backslash i ; A:=\times_{i \in I} A_{i}$, where $A_{i}$ is $i$ 's finite set of feasible actions; and $u:=\left(u_{i}\right)_{i \in I}$, with $u_{i}: A_{i} \times \Delta\left(A_{-i}\right) \rightarrow \mathbb{R}$ denoting player $i$ 's payoff function, where $u_{i}$ is continuous, $\Delta\left(A_{-i}\right)$ being endowed with the Euclidean norm. ${ }^{8}$ We extend $u_{i}$ to the space of probability distributions over actions with $u_{i}\left(\sigma_{i}, \sigma_{-i}\right)=\mathbb{E}_{\sigma_{i}}\left[u_{i}\left(a_{i}, \sigma_{-i}\right)\right]$, where $\mathbb{E}_{\sigma_{i}}[\cdot]$ corresponds to the expectation taken with respect to $\sigma_{i}$. While we focus throughout on normal-form games, having payoffs directly depend on opponents' distribution of actions will render proofs readily adaptable to Bayesian games, a setting to which our framework extends naturally as discussed in Section 6.

Beliefs. In contrast to other solution concepts, each player $i$ is uncertain about others' true distribution of actions, $\sigma_{-i}$. While this implies uncertainty about the actions that others ultimately take, $a_{-i}$, the main conceptual difference is that, if others' actions are only observable after these are taken, players may still reason and learn about the prevailing stable action distribution prior to others choosing an action-e.g. the likelihood a restaurant is too crowded, the probability that others abstain in an election, or the distribution of prices of a product across different platforms. Such uncertainty captures the players' limited experience and imperfect memory. Expressing this uncertainty, each player $i$ holds beliefs about $\sigma_{-i}$, given by a Borel probability measure $\mu_{i} \in \Delta\left(\Delta\left(A_{-i}\right)\right)$, where $\Delta\left(\Delta\left(A_{-i}\right)\right)$ is endowed with the topology of weak* convergence, metricized by Lévy-Prokhorov metric $\|\cdot\|_{L P}$. We require player $i$ 's beliefs to have as support, $\operatorname{supp}\left(\mu_{i}\right)$, the set of all distributions, allowing for correlation $-\operatorname{supp}\left(\mu_{i}\right)=\Delta\left(A_{-i}\right)-$ or the set of all distributions assuming independence across opponents, in which case beliefs are given by a product measure $\mu_{i}=\times_{j \epsilon-i} \mu_{i j}$, where each $\mu_{i j}$ is a probability measure on $\Delta\left(A_{j}\right)$ with full support. Results will hold in either unless explicitly mentioned. ${ }^{9}$

Information and Sequential Sampling. We model a player's costly reasoning about others' behavior via a sequential sampling stage that takes place prior to choice, as in sequential sampling models used to describe reasoning in individual decision-making settings (Fudenberg et al., 2018;

[^5]Forstmann et al., 2016). That is, prior to making a choice, player $i$ can acquire signals about the unknown distribution $\sigma_{-i}$ in a sequential and costly manner. Sequential sampling then captures the recollection of past experiences in similar situations (others' behavior, payoffs to tried out actions), asking friends, or more general stepwise reasoning processes-see Alaoui and Penta (2022) for an axiomatic foundation of stepwise reasoning in this guise. Alternatively, it can be taken in more literal fashion, with players acquiring or parsing through existing data.

A key assumption is that this reasoning is informative about others' behavior, that is, that each player $i$ has access to an information structure $\pi_{i}: \Delta\left(A_{-i}\right) \rightarrow \Delta\left(Y_{i}\right)$, where $Y_{i}$ is a finite signal space. Throughout the main text, we restrict attention to the case in which these signals are observations drawn from $\sigma_{-i}$, i.e. $\pi_{i}$ corresponds to the identity. More general signal structures are considered in Appendix B.

As mentioned, information acquisition is sequential. In other words, prior to taking an action, player $i$ can sequentially observe signals $y_{i, t}$ and decide when to stop sampling. Players' stopping time will be interpreted as their decision time, and thus take a prominent role in our endeavor to ground the relation between incentives, choices, and decision time in strategic settings. The sequentiality in sampling marks an important distinction relative to other models of sampling in games (e.g. Osborne and Rubinstein, 1998, 2003), allowing rich joint distributions of actions and decision time to emerge, in which particular actions can be associated with lengthier decision times, and others with shorter.

We will write $y_{i}^{t}=\left(y_{i, \ell}\right)_{\ell \in[1 . . t]}$ to stand for the sample path up to time $t$, where $[n . . n+k]=$ $\{n, n+1, \ldots, n+k\}$ and each realization $y_{i, \ell}$ is distributed according to $\sigma_{-i}$, with the understanding that $y_{i}^{0}=\varnothing$. Formally, we denote $\mathbf{y}_{i}=\left\{y_{i, t}\right\}_{t \in \mathbb{N}}$ as a stochastic process defined on the probability space $(\Omega, \mathscr{F}, \mathbb{P})$ with $\mathbb{F}_{i}$ denoting the natural filtration of $\mathbf{y}_{i}$. The set of sample paths of length $t$ is denoted by $Y_{i}^{t}$ and the set of all finite sample path realizations is denoted by $\mathscr{Y}_{i}:=\bigcup_{t \in \mathbb{N}} Y_{i}^{t}$. Upon observing a given sample path up to time $t, y_{i}^{t}$, player $i$ updates beliefs about $\sigma_{-i}$ according to Bayes' rule, denoted by $\mu_{i} \mid y_{i}^{t} .{ }^{10}$

Sampling Costs. Naturally, sampling is costly, capturing the effort involved in reasoning. For convenience, we will throughout assume that player $i$ 's cost of each observation is given by $c_{i}>0$. It is straightforward to adjust the model in order to accommodate costs that depend on the number

[^6]of observations, insofar as they are eventually bounded away from zero from below, ${ }^{11}$ which, for all purposes, subsumes cases in which there is an upper bound on the number of observations.

Extended Games. An extended game $G$ is then a tuple comprising an underlying normal-form game $\Gamma$, each players' prior beliefs $\mu=\left(\mu_{i}\right)_{i \in I}$, and sampling costs $c=\left(c_{i}\right)_{i \in I}$.

### 2.2. Equilibrium

Having introduced all the primitives of the model, we now turn to defining equilibrium.
Choice. Given a belief $\mu_{i}^{\prime} \in \Delta\left(\Delta\left(A_{-i}\right)\right)$, player $i$ upon stopping sampling chooses an action in order to maximize their expected utility. We denote the player's maximized utility by $v_{i}$ : $\Delta\left(\Delta\left(A_{-i}\right)\right) \rightarrow \mathbb{R}$

$$
v_{i}\left(\mu_{i}^{\prime}\right):=\max _{\sigma_{i} \in \Delta\left(A_{i}\right)} \mathbb{E}_{\sigma_{i}}\left[\mathbb{E}_{\mu_{i}^{\prime}}\left[u_{i}\left(a_{i}, \sigma_{-i}\right)\right]\right],
$$

where $\mathbb{E}_{\sigma_{i}}\left[\mathbb{E}_{\mu_{i}^{\prime}}\left[u_{i}\left(a_{i}, \sigma_{-i}\right)\right]\right]=\int_{A_{i}} \int_{\Delta\left(A_{-i}\right)} u_{i}\left(a_{i}, a_{-i}\right) \mathrm{d} \mu_{i}^{\prime}\left(\sigma_{-i}\right) \mathrm{d} \sigma_{i}\left(a_{i}\right)$. We write $\sigma_{i}^{*}: \Delta\left(A_{-i}\right) \rightarrow \Delta\left(A_{i}\right)$ to denote a selection of optimal choices given beliefs, $\sigma_{i}^{*}\left(\mu_{i}^{\prime}\right) \in \operatorname{argmax}_{\sigma_{i} \in \Delta\left(A_{i}\right)} \mathbb{E}_{\mu_{i}}\left[u_{i}\left(\sigma_{i}, \sigma_{-i}^{\prime}\right)\right]$.

Optimal Stopping. Player $i$ samples optimally in order to maximize expected payoffs. That is, each player $i$ faces an optimal stopping problem: based on the expected value of future reasoning, decide whether to stop and make a choice or obtain another signal. Formally, player $i$ chooses a stopping time $t_{i}$ in the set $\mathbb{T}_{i}$ of all stopping times taking values in $\mathbb{N}_{0} \cup\{\infty\}$ and adapted with respect to natural filtration associated to $\mathbf{y}_{i}$.

Given a prior $\mu_{i} \in \Delta\left(\Delta\left(A_{-i}\right)\right)$, player $i$ 's value function $V_{i}: \Delta\left(\Delta\left(A_{-i}\right)\right) \rightarrow \overline{\mathbb{R}}$ can be written as

$$
V_{i}\left(\mu_{i}\right):=\sup _{t_{i} \in \mathbb{T}_{i}} \mathbb{E}_{\mu_{i}}\left[v_{i}\left(\mu_{i} \mid y_{i}^{t_{i}}\right)-c_{i} \cdot t_{i}\right],
$$

where $\mu_{i} \mid y_{i}^{t_{i}}$ denotes the player's posterior belief when, upon stopping according to stopping time $t_{i}$, the sample $y_{i}^{t_{i}}$ was observed.

It will be useful to consider the dynamic programming formulation of the optimal stopping problem, with $V_{i}$ corresponding to a fixed point of an operator $B_{i}: C_{b}\left(\Delta\left(\Delta\left(A_{-i}\right)\right)\right) \rightarrow C_{b}\left(\Delta\left(\Delta\left(A_{-i}\right)\right)\right)$,

$$
B_{i}\left(\tilde{V}_{i}\right)\left(\mu_{i}^{\prime}\right)=\max \left\{v_{i}\left(\mu_{i}^{\prime}\right), \mathbb{E}_{\mu_{i}^{\prime}}\left[\tilde{V}_{i}\left(\mu_{i}^{\prime} \mid y\right)\right]-c_{i}\right\},
$$

[^7]which will be equivalent for our purposes. This lends our value function a clear interpretation:
$$
\underbrace{V_{i}\left(\mu_{i}^{\prime}\right)}_{\text {value at belief } \mu_{i}^{\prime}}=\max \{\underbrace{v_{i}\left(\mu_{i}^{\prime}\right)}_{\text {value of stopping }}, \quad \underbrace{\mathbb{E}_{\mu_{i}}\left[V_{i}\left(\mu_{i}^{\prime} \mid y_{i}\right)\right]-c_{i}}_{\text {expected value of continuing sampling }}\}
$$

We focus on the earliest optimal stopping time

$$
\tau_{i}(\omega):=\min \left\{t \in \mathbb{N}_{0} \mid V_{i}\left(\mu_{i}^{\prime} \mid y_{i}^{t}(\omega)\right)=v_{i}\left(\mu_{i}^{\prime} \mid y_{i}^{t}(\omega)\right)\right\}
$$

where its optimality follows by standard arguments (Ferguson, 2008, Ch. 3, Theorem 3); while omitted, we note the dependence of $\tau_{i}$ on the prior $\mu_{i}^{\prime}$.

For ease of reference, we summarize properties of optimal sequential sampling in this proposition:
Proposition 1. The following properties hold: (1) $v_{i}$ and $V_{i}$ are bounded, convex, and uniformly continuous. (2) For any prior $\mu_{i}$, player $i$ 's optimal stopping time is finite $\mu_{i}$-a.s. and satisfies $\mathbb{P}_{\mu_{i}}(\tau>$ $T) \leq 2\left\|u_{i}\right\|_{\infty} / c_{i} T$.

This and the remaining omitted proofs are in Appendix A.

Equilibrium Definition. We now close the model by considering equilibrium behavior. Players' strategic uncertainty about others' action distribution motivates their sequential sampling behavior. The randomness inherent to sampling entails randomness in players' optimal actions given their posterior beliefs. Equilibrium is then a fixed point, a consistency condition based on the premise that each player's signals are informative about others' distribution of actions.

Formally, each player acquires information on their opponents' action distribution and their optimal stopping policy, $\tau_{i}$, determines the sequences of signals following which they optimally stop at take an action:

$$
\mathscr{Y}_{i}^{\tau_{i}}:=\left\{y_{i}^{t} \in \mathscr{Y}_{i}: V_{i}\left(\mu_{i} \mid y_{i}^{t}\right)=v_{i}\left(\mu_{i} \mid y_{i}^{t}\right) \text { and } \forall 0 \leq \ell<t, V_{i}\left(\mu_{i} \mid y_{i}^{\ell}\right)>v_{i}\left(\mu_{i} \mid y_{i}^{\ell}\right)\right\} .
$$

Different opponent action distributions $\sigma_{-i}$ induce different distributions over the signals acquired $y_{i}^{\tau_{i}}$. Since different sequences of signals $y_{i}^{\tau_{i}}$ induce different posteriors $\mu_{i} \mid y_{i}^{\tau_{i}}$ at which different actions $a_{i}$ may be optimal, sequential sampling implies a mapping from opponents' action distributions to the player's distribution of actions,

$$
\mathbb{E}_{\sigma_{-i}}\left[\sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right]=\underbrace{\sum_{\begin{array}{c}
t_{i}^{t} \mathscr{Y}_{i}^{\tau_{i}}
\end{array}}^{\prod_{\begin{array}{c}
\text { probability } \\
\text { observing } y_{i}^{t}
\end{array}}^{\prod_{\substack{\text { best response } \\
\text { at posterior } \mu_{i} \mid y_{i}^{t}}} \sigma_{-i}\left(y_{i, \ell}\right)}} \underbrace{\sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{t}\right)}_{i}}_{\begin{array}{c}
\text { set of stopping } \\
\text { sequences }
\end{array}}
$$

That is, the probability of player $i$ taking action $a_{i}$ is given by the probability of taking such
an action once player $i$ after observing $y_{i}^{t}, \sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{t}\right)$, considering every sequence of signals $y_{i}^{t}$ following which player $i$ optimally stops, $y_{i}^{t} \in \mathscr{Y}_{i}^{\tau_{i}}$, and weighting it by the probability of its occurrence. The probability that a sequence of signals $y_{i}^{t}$ is observed - - that is, the true probability of optimally stopping following $y_{i}^{t}, \mathbb{P}_{\sigma_{-i}}\left(y_{i}^{\tau_{i}}=y_{i}^{t}\right)$-is then given by $1_{y_{i}^{t} \in \mathscr{Y}_{i}^{\tau_{i}}} \prod_{\ell \in[1 . . t]} \sigma_{-i}\left(y_{i, \ell}\right)$, as each observation corresponds to an action profile $y_{i, \ell} \in A_{-i}$, sampled independently from $i$ 's opponents' action distribution, $\sigma_{-i}$.

Equilibrium then follows as a consistency condition between players' action distributions:
Definition 1. A sequential sampling equilibrium of an extended game $G=\langle\Gamma, \mu, c\rangle$ is a profile of action distributions $\sigma$ such that, for every $i, \sigma_{i}=\mathbb{E}_{\sigma_{-i}}\left[\sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right]$, where $\tau_{i}$ is player $i$ 's earliest optimal stopping time and $\sigma_{i}^{*}\left(\mu_{i}^{\prime}\right)$ is optimal given belief $\mu_{i}^{\prime}$.

Equilibrium Stopping Time. It is important to emphasize that a sequential sampling equilibrium implies a novel relationship between actions and time. While a sequential sampling equilibrium is defined in the space of action distributions, note that, in a given extended game, the equilibrium distribution of opponents' actions $\sigma_{-i}$ completely pins down the joint distribution of choices and (optimal) stopping time for player $i\left(a_{i}, \tau_{i}\right)$. Such joint distribution is determined in equilibrium, as it crucially depends on the true (equilibrium) action distribution of player $i$ 's opponents about which player $i$ is reasoning. Although one could think about $c \mathbb{E}_{\mu_{i}}\left[\tau_{i}\right]$ as a static cost of information and rephrase our equilibrium notion from a static viewpoint (e.g. Yang, 2015; Denti, 2023; Hébert and La'O, 2023), a major distinctive feature of this model relative to equilibrium models of costly information acquisition in games is that it enables one to speak not of the subjectively expected cost of information acquisition, but of the joint distribution of realized actions and time. As also mentioned earlier, a second difference is that uncertainty here refers to distribution of actions, and therefore it does not require-but can accommodate-exogenous sources of uncertainty.

Interpretation. Sequential sampling equilibrium can be interpreted as positing that, prior to taking an action, players reason through others' behavior to better ground their choices. In this sense, players' sequential sampling reflects an underlying introspective process whereby players reason about how others may act by reaching back in their memory and past experiences. This interpretation is motivated by literature in cognitive science (Shadlen and Shohamy, 2016), which has made use of sequential sampling models to ground the relation between time and choices
not only in association problems and perceptual tasks, ${ }^{12}$ but also in domains where choices are guided by individual preferences (e.g. Clithero, 2018a), with patterns being consistent across these different domains (Smith and Krajbich, 2021). Recent literature provided neurological evidence of memory guiding preference-based choice, conforming with sampling from memory (see Bakkour et al., 2018; Duncan and Shohamy, 2020; Biderman et al., 2020; Biderman and Shohamy, 2021).

A second, complementary and more literal, interpretation of sequential sampling is to see it as acquiring hard information-such as data, experts' opinions, or reviews. For instance, a seller doing market research to better price its product, consumers parsing reviews on a product's quality, voters learning about candidates' platforms through their statements about different issues, or infrequent bidders in online auctions looking at data from other past auctions to reason how to bid.

At the heart of our model is the assumption that sampling is informative about others' behavior, so as to render it valuable. Equilibrium delivers this consistency by positing stationarity of the environment-Section 5 illustrates how it coincides with steady states of a particular dynamic process. Further, it also allows us to obtain comparative statics predictions on how payoffs affect players' choices both by changing the trade-offs in players' information acquisition decisions and the signal realizations. Naturally, stationarity of the environment is a strong assumption that may be unwarranted in situations where non-equilibrium is more apt to describe behavior-see Alaoui and Penta (2016) for one such model.

In order to obtain sharper predictions, we will focus on a simple form of sequential sampling. Players' sampling is then represented in a stylized manner, with signals directly sampled from the prevailing distribution of actions of opponents. In Section 6 we discuss how sequential sampling equilibrium can be easily extended to richer settings and information structures-e.g. noisy recollections, inability to learn about some players', or allowing signals about others' behavior (or types) in one setting to inform players in the situation at hand.

Sequential sampling equilibrium also provides a particular way to relax the implicit epistemic assumption in Nash equilibrium that, in equilibrium, players come to know their opponents' distribution of actions. If players' priors did assign probability one to the same Nash equilibrium of the underlying game, that Nash equilibrium will coincide with a sequential sampling equilibrium of the game. ${ }^{13}$ In our model, however, players are uncertain about the prevailing distribution of

[^8]actions, and it is this uncertainty that drives their sequential sampling behavior. Further, it dispenses with the assumption of mutual knowledge of the game and of others' rationality, since all learning is driven by the procured information and players need not know others' payoff functions.

Existence. We briefly note that a sequential sampling equilibrium exists in all extended games.
Theorem 1. Every extended game has a sequential sampling equilibrium.
The proof proceeds by verifying that, for every player $i, \sigma_{-i} \mapsto b_{i}\left(\sigma_{-i}\right):=\mathbb{E}_{\sigma_{-i}}\left[\sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right]$ maps to a well-defined probability distribution of player $i$ 's actions, and (2) such mapping is continuous. The main difficulty is that, while we know that, by Proposition 1, $\tau_{i}$ is finite with probability 1 with respect to the player's prior $\left(\mathbb{P}_{\mu_{i}}\left(\tau_{i}<\infty\right)=1\right)$, we need player $i$ 's optimal stopping time to be finite with probability 1 with respect to the actual distribution of opponents' actions, $\left(\mathbb{P}_{\sigma_{-i}}\left(\tau_{i}<\infty\right)=1\right)$, as otherwise $b_{i}\left(\sigma_{-i}\right) \notin \Delta\left(A_{i}\right)$. If player $i$ never stops sampling with positive probability (with respect to the true distribution of opponents' actions), then $b_{i}$ does not define a probability distribution over player $i$ 's actions and no equilibrium exists. ${ }^{14}$ The following lemma demonstrates that this condition on stopping time is also sufficient to guarantee the desired properties on $b_{i}$ :

Lemma 1. The following two statements are equivalent: (1) player $i$ 's optimal stopping time is finite with probability 1 with respect to $\sigma_{-i}, \mathbb{P}_{\sigma_{-i}}\left(\tau_{i}<\infty\right)=1$ for any $\sigma_{-i} \in \Delta\left(A_{-i}\right) ;(2) b_{i}\left(\sigma_{-i}\right) \in \Delta\left(A_{i}\right)$ $\forall \sigma_{-i} \in \Delta\left(A_{-i}\right)$. Moreover, if (1) holds, then $b_{i}$ is also continuous.

The proof can be found in Appendix A. With Lemma 1 in hand, it is then straightforward to show existence of a sequential sampling equilibrium.

Proof. Given Lemma 1, if, for any $\sigma_{-i}, \tau_{i}$ is finite with probability one with respect to $\sigma_{-i}$, then $b_{i}$ is a continuous mapping from $\Delta\left(A_{-i}\right)$ to $\Delta\left(A_{i}\right)$, and existence follows from Brouwer's fixed point theorem. By assumption, $\operatorname{supp}\left(\mu_{i}\right)=\Delta\left(A_{-i}\right)$ and, for any $\sigma_{-i}$,

$$
\mathbb{P}_{\sigma_{-i}}\left(\tau_{i}(\omega) \leq T\right)=\mathbb{P}_{\sigma_{-i}}\left(\left\{\omega: \inf \left\{t \mid \mathbb{E}_{\mu_{i, t}(\omega)}\left[V_{i}\left(\mu_{i, t}(\omega) \mid y_{i, t+1}\right]-V_{i}\left(\mu_{i, t}(\omega)\right) \leq c_{i}\right\} \leq T\right\}\right) .\right.
$$

${ }^{14}$ We provide an example in Online Appendix D to illustrate the potential non-existence of equilibria when priors do not have full support.

As $V_{i}$ is uniformly continuous, there is $\delta>0$ such that, $\forall \mu_{i}, \mu_{i}^{\prime} \in \Delta\left(P_{i}\right)$ satisfying $\left\|\mu_{i}-\mu_{i}^{\prime}\right\|_{L P}<$ $\delta$, $\left|V_{i}\left(\mu_{i}\right)-V_{i}\left(\mu_{i}^{\prime}\right)\right|<c$. Since each observation is subjectively iid, by Berk (1966), $\mu_{i, t}$ weak $^{*}$ converges to a Dirac on $\sigma_{-i}, \sigma_{-i}$-a.s., $\mathbb{P}_{\sigma_{-i}}\left(\lim _{t \rightarrow \infty} \mathbb{E}_{\mu_{i, t}}\left[V_{i}\left(\mu_{i, t} \mid y_{i, t+1}\right]-V_{i}\left(\mu_{i, t}\right)>c_{i}\right)=0\right.$.

In fact, we can obtain an upper bound on the stopping time by combining uniform continuity of $V_{i}$ and the fact that $\mu_{i}$ uniformly accumulates around the empirical frequency:
Remark 1. For every player $i, \exists \bar{T}_{i}<\infty$ such that $\tau_{i} \leq \bar{T}_{i}$, where $\bar{T}_{i}$ depends on $u_{i}$, $\mu_{i}$, and $c_{i}$.
This transforms optimal stopping into a finite horizon problem, a useful result that not only simplifies the analysis, but also makes our solution concept amenable to computational applications.

## 3. Behavioral Implications

In this section, we characterize different behavioral implications of sequential sampling equilibrium. First, we explore the relation between stopping time and action sophistication. Then, we will relate incentives to the joint distribution of choices and stopping time. Finally, we focus on players' beliefs and their relation with stopping time.

### 3.1. Rationality and Sequential Sampling

Existing evidence points toward an association between longer decision times and choices reflecting greater sophistication (e.g. Agranov et al., 2015; Rubinstein, 2016; Alós-Ferrer and Bruckenmaier, 2021), which may express heterogeneity in individual costs of reasoning. However, when facing higher stakes, decision times increase first-order stochastic dominance sense, and choices do reflect higher sophistication, as given by their level of rationalizability (Esteban-Casanelles and Gonçalves, 2022). Existing benchmark models like level-k and cognitive hierarchy (Nagel, 1995; Camerer et al., 2004) are unable to deliver such comparative statics, since choices are invariant with respect to incentive levels. ${ }^{15}$

This section shows how sequential sampling equilibrium can provide a rationale for such an association by relating higher incentives, longer decision times, and a lower bound on the level of rationalizability of action chosen in equilibrium. Further, this establishes a relation between empirical learning (as given by sequential sampling) and introspective learning (as given by rationalizability).

[^9]We first observe that, in our context, higher incentives as given by scaling up a player's payoffs, is equivalent to scaling down the cost to sampling. Optimal sequential sampling naturally predicts an inverse relation between sampling cost and stopping time:

Remark 2. For lower sampling costs $c_{i}$, player $i$ 's optimal stopping time increases in first-order stochastic dominance with respect to the prior $\mu_{i}$ and to any true distribution of opponents' actions $\sigma_{-i} \in \Delta\left(A_{-i}\right)$; that is, both $\mathbb{P}_{\mu_{i}}\left(\tau_{i} \leq t\right)$ and $\mathbb{P}_{\sigma_{-i}}\left(\tau_{i} \leq t\right)$ increase for any $t$.

Our main result of this section goes further in determining a relation between cost and the level of sophistication of actions chosen in equilibrium. Let us recall the definition of rationalizable actions.

Definition 2. An action $a_{i} \in A_{i}$ is 1-rationalizable if there is some $\sigma_{-i} \in \Delta\left(A_{-i}\right)$ such that, $\forall \sigma_{i} \in \Delta\left(A_{i}\right), u_{i}\left(\alpha_{i}, \sigma_{-i}\right) \geq u_{i}\left(\sigma_{i}, \sigma_{-i}\right)$. An action $a_{i} \in A_{i}$ is $k$-rationalizable, for $k \geq 2$, if there is some $\sigma_{-i} \in \Delta\left(A_{-i}^{k-1}\right)$ such that, $\forall \sigma_{i} \in \Delta\left(A_{i}\right), u_{i}\left(a_{i}, \sigma_{-i}\right) \geq u_{i}\left(\sigma_{i}, \sigma_{-i}\right)$, where $A_{-i}^{k-1}:=\times_{j \neq i} A_{j}^{k-1}$ denotes the set of $(k-1)$-rationalizable action profiles of player $i$ 's opponents. An action $a_{i}$ is rationalizable if $a_{i} \in \cap_{k \in \mathbb{N}} A_{i}^{k}$.

For presentation purposes-as implied by the above definition-we focus on a definition of rationalizability allowing for correlation among opponents' actions, and require priors to have full support on $\Delta\left(A_{-i}\right)$. The below result holds as well when considering a definition of rationalizable actions that requires independence across opponents' action distributions, provided beliefs also do not allow for correlation.

We now show that scaling up incentives enough-or, equivalently, for low enough sampling costs-only $k$-rationalizable actions are chosen at sequential sampling equilibria:

Theorem 2. For any normal-form game $\Gamma$, priors $\mu$, and $k \in \mathbb{N}$, there are cost thresholds $\bar{c}_{i}^{k}>0$ such that, for any extended game $G=\langle\Gamma, \mu, c\rangle$ in which $c_{i}<\bar{c}_{i}^{k}$ for all $i$, in any sequential sampling equilibrium $\sigma$ of $G$ only $k$-rationalizable actions are chosen with positive probability.

The result is obtained by combining three observations-the proofs for which can be found in Appendix A. First, if player $i$ believes that, with high enough probability, their opponents only choose ( $k-1$ )-rationalizable actions, then player $i$ will choose a $k$-rationalizable action:

Lemma 2. For any $k \geq 2$, there are $\epsilon, \delta>0$, such that, if $\mu_{i}\left(B_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right)>1-\epsilon\right.$, then $\arg \max _{a_{i} \in A_{i}} \mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, \sigma_{-i}\right)\right] \subseteq A_{i}^{k}$.

Second, that if player $i$ 's opponents do indeed only choose ( $k-1$ )-rationalizable actions, then
player $i$ 's beliefs uniformly accumulate on the event that opponents only choose ( $k-1$ )-rationalizable actions:

Lemma 3. For any $\mu_{i} \in \Delta\left(\Delta\left(A_{-i}\right)\right)$ with full support, and all $\epsilon, \delta>0$, there is $t$ such that, for any sequence of observations $y_{i}^{t}$ for which $y_{i, \ell} \in A_{-i}^{k-1}$ for $\ell \in[1 . . t], \mu_{i} \mid y_{i}^{t}\left(B_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right)\right)>1-\epsilon$.

And third, that, when not all of player $i$ 's actions are rationalizable, it suffices that sampling costs are low enough to ensure that the player acquires a minimum number of signals:

Lemma 4. Suppose that there is no action $a_{i}$ that is a best response to all distribution of opponents' actions $\sigma_{-i} \in \Delta\left(A_{-i}\right)$. Then, for any $T \in \mathbb{N}_{0}$ and any full support prior $\mu_{i} \in \Delta\left(A_{-i}\right)$, there is $\bar{c}_{i}>0$ such that for any sampling cost $c_{i}<\bar{c}_{i}$, the associated earliest optimal stopping time $\tau_{i} \geq T+1$.

The proof of Theorem 2 then proceeds easily:
Proof. The proof follows an induction argument. First, observe that no player will choose actions that are not 1-rationalizable. Now, for $k \geq 1$, assume that players choose only ( $k-1$ )-rationalizable actions with positive probability. From Lemma 3, for any $\delta, \epsilon>0$ there is a $T$ such that, for all $t \geq T$, all $i \in I$, and any $y_{i}^{t} \in A_{-i}^{k-1}, \mu_{i, t}\left(B_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right)\right) \geq \mu_{i, t}\left(B_{\delta}\left(\delta_{\bar{y}_{i}^{t}}\right)\right)>1-\epsilon$. By Lemma 2, this implies that if all players sample for at least $T$ periods, they will only choose $k$-rationalizable actions with positive probability. Lemma 4 ensures that we can find $c^{k}>0$ such that, if $c_{i} \leq c^{k}$ $\forall i$, all players sample at least $T$ periods, i.e. that each player's earliest optimal stopping time is bounded below by $T, \tau_{i} \geq T$. This concludes the proof.

Remark 3. It is possible to generalize the result to classes of priors that satisfy a condition akin to a lower bound on density:

Definition 3 (Diaconis and Freedman 1990). Let $\phi: \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$. The set of $\phi$-positive distributions on $\Delta\left(A_{-i}\right)$ is given by $\mathscr{M}_{i}(\phi):=\left\{\mu_{i} \in \Delta\left(\Delta\left(A_{-i}\right)\right) \mid \inf _{\sigma_{-i} \in \Delta\left(A_{-i}\right)} \mu_{i}\left(B_{\epsilon}\left(\sigma_{-i}\right)\right) \geq \phi(\epsilon), \forall \epsilon>0\right\}$.

Since it is possible to obtain a uniform rate of accumulation around the empirical mean for any prior $\mu_{i} \in \mathscr{M}_{i}(\phi)$ that depends only on $\phi$, we can then extend Theorem 2 so that the same cost thresholds holds for all $\phi$-positive priors $\mu_{i}$. Further, these cost thresholds $c_{i}^{k}$ can be made tight and non-increasing in $k$ by considering all priors within such class.

In short, Theorem 2 uncovers a relationship between sampling costs and degree of sophistication of equilibrium choices. Since scaling up payoffs is tantamount to scaling down costs, providingvia sequential sampling-a rationale for the link from incentive levels to action sophistication,
beliefs about opponents, and decision times documenetd in dominance-solvable games e.g. in Esteban-Casanelles and Gonçalves (2022).

### 3.2. Comparative Statics in Binary Action Games

Binary Action Games. In this section we provide comparative statics results for binary actions games: normal-form games $\Gamma=\langle I, A, u\rangle$ with two players, $|I|=2$, and such that each player $i \in I$ has two actions $A_{i}=\{0,1\}$ and $u_{i}\left(1, \sigma_{-i}\right)-u_{i}\left(0, \sigma_{-i}\right)$ is strictly monotone and continuous in $\sigma_{-i}$, the probability that $i$ 's opponent chooses action $1 .{ }^{16}$ Consequently, we identify $\mu_{i} \in \Delta\left(\Delta\left(A_{-i}\right)\right)$ with a distribution on the unit interval. An extended binary action game $G$ is an extended game for which $\Gamma$ is a binary action game.

Actions and Stopping Time. Our object of interest will be the probability (according to $\sigma_{-i}$ ) that player $i$ stops before time $t$ and, upon stopping, action $a_{i}$ is optimal, that is,

$$
\mathbb{P}_{\sigma_{j}}\left(a_{i} \in A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right) \text { and } \tau_{i} \leq t\right),
$$

where $A_{i}^{*}\left(\mu_{i}^{\prime}\right):=\arg \max _{a_{i} \in A_{i}} \mathbb{E}_{\mu_{i}^{\prime}}\left[u_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right)\right]$ denotes the set of optimal choices at a given belief $\mu_{i}^{\prime}$. Our main result characterizes how the joint distribution of player $i$ 's choices and their stopping times changes along to three dimensions: (1) the player's payoffs, $u_{i}$, (2) their beliefs, $\mu_{i}$, and (3) the true (unknown) distribution of their opponent's actions, $\sigma_{-i}$, taken as exogenous.

Ordering Payoffs and Beliefs. Let us introduce a partial order on player $i$ 's utility functions: Definition 4. Let $u_{i}, u_{i}^{\prime}: A_{i} \times \Delta\left(A_{-i}\right) \rightarrow \mathbb{R} . u_{i}^{\prime}$ is said to has higher incentives to action $a_{i}$ than $u_{i}$, $u_{i}^{\prime} \geq_{a_{i}} u_{i}$, if and only if there is $g: \Delta\left(A_{-i}\right) \rightarrow \mathbb{R}_{+}$such that $u_{i}^{\prime}\left(\alpha_{i}^{\prime}, \sigma_{-i}^{\prime}\right)=u_{i}\left(\alpha_{i}^{\prime}, \sigma_{-i}^{\prime}\right)+1_{a_{i}^{\prime}=a_{i}} g\left(\sigma_{-i}\right)$. Beliefs are ordered according to a generalized version of the monotone likelihood ratio property (cf. Lehrer and Wang, 2020):

Definition 5. Let $\mu_{i}, \mu_{i}^{\prime} \in \Delta([0,1])$. $\mu_{i}^{\prime}$ is said to strongly stochastic dominate $\mu_{i}, \mu_{i}^{\prime} \geq_{S S D} \mu_{i}$, if $\mu_{i}^{\prime} \mid y_{i}^{t}$ first-order stochastically dominates $\mu_{i} \mid y_{i}^{t}$ for any $y_{i}^{t} \in \mathscr{Y}_{i}$.

Note that, when $\mu_{i}$ and $\mu_{i}^{\prime}$ are mutually absolutely continuous, $\geq_{S S D}$ corresponds to the monotone likelihood ratio property, i.e. $d \mu_{i}^{\prime} / d \mu_{i}$ is increasing $\mu_{i}^{\prime}$-a.e.

[^10]Monotone Comparative Statics. The next result characterizes the behavior induced by optimal sequential information acquisition taking $\sigma_{-i}$ are exogenous:

Theorem 3. Let $G$ be an extended binary action game and let $a_{i} \in \operatorname{argmax}_{a_{i}^{\prime} \in A_{i}} u_{i}\left(\alpha_{i}^{\prime}, 1\right)$. Then, $\mathbb{P}_{\sigma_{-i}}\left(a_{i} \in A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right.$ and $\left.\tau_{i} \leq t\right)$ increases (i) in $u_{i}$ with respect to $\geq_{a_{i}}$, (ii) in $\mu_{i}$ with respect $\geq_{S S D}$, and (iii) in $\sigma_{-i}$. Moreover, it is $\mathscr{C}^{\infty}$ in $\sigma_{-i}$.

Let us discuss the intuition behind the theorem (the proof is deferred to the Appendix A).
Claim (i) shows that increasing the payoff associated to action $a_{i}, u_{i}^{\prime} \geq_{a_{i}} u_{i}$, makes the player not only more likely to take that action under the true distribution of actions of the opponent, but to take it faster and to choose the other action less often and slower. While an increase in payoffs does increase the value of sampling at some posterior beliefs-which could lead the player to learn more about the true $\sigma_{-i}$ and find out that perhaps action $a_{i}$ is not optimal after all-this additional information acquisition occurs only when before the player was stopping and taking an action other than $a_{i}$. In other words, player $i$ requires now less information to be convinced to stop and take action $a_{i}$ and more information to stop and choose another action. This result is not particular to binary action games: claim (i) is shown for general settings with arbitrary finitely many actions and general payoff functions. ${ }^{17}$

Claim (ii) can be interpreted as stating that player $i$ is more likely to stop earlier and take action $a_{i}$ the greater the probability their prior assigns to action $a_{i}$ being optimal. The main difficulty is again to show that this seemingly tautological statement holds with respect to the actual, unknown, distribution of the opponent's actions; importantly, note the claim does not depend on whether or how correct player $i$ 's beliefs are. Such monotonicity in beliefs allows one to make predictions on how behavior changes with, for instance, the provision of information that shifts beliefs in the stochastic dominance order (e.g. $\mu_{i}\left|1 \geq_{S S D} \mu_{i}\right| 0$ ).

Finally, the argument for why claim (iii) should hold is straightforward: higher $\sigma_{-i}$ means that player $i$ is more likely to observe higher signals and therefore becoming convinced that action $a_{i}$ is the better alternative. The proof follows from claim (ii) and an induction argument. The fact that the probability of action $a_{i}$ being optimal when stopping before time $t$ is a polynomial with respect to $\sigma_{-i}$ implies the claim on differentiability.

Theorem 3 provides comparative statics on the optimality of a given action, but leaves open the possibility that more than one action is optimal. The next lemma closes this gap by showing that,

[^11]

Figure 1: Generalized Matching Pennies
Note: $\delta_{M}, \gamma_{C}>0$.
in binary action games, a player is never indifferent between the two actions at any belief held upon stopping, provided the player samples at least once or is not indifferent under the prior $\mu_{i}$.

Lemma 5. Let $G$ be an extended binary action game. Then, for any player $i, A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)$ is a singleton or $\tau_{i}=0$. Moreover, if $\tau_{i}>0, A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)=\operatorname{argmax}_{a_{i} \in A_{i}} u_{i}\left(a_{i}, y_{i, \tau_{i}}\right)$.

The reasoning underlying the proof is simple. Without loss of generality, assume that player $i$ 's best-response to $a_{-i}$ is to choose $a_{i}=a_{-i}$. Suppose that player $i$ stops sampling after observing a 0 -valued signal leaving player $i$ indifferent between the two actions (the argument is symmetric if the last signal is 1 -valued). Then, before sampling the last observation, action 1 was already optimal under player $i$ 's prior, as observing a 0 -valued observation induces a lower belief mean. Moreover, if the last observation had instead realized to be 1 -valued, player $i$ would still want to choose action 1 . This implies that if player $i$ stops sampling when indifferent between the two actions, whichever action was optimal before taking the last signal is still optimal regardless of the realization of the signal. Therefore, given that the player will not sample any further, the last signal has no informational value. As the signal is costly, it is suboptimal to take it.

Applications. One immediate implication of Theorem 3 and Lemma 5 is in establishing a strong connection between uniqueness of sequential sampling equilibrium in an extended game and uniqueness of a Nash equilibrium of the underlying binary action game:

Proposition 2. A binary action game $\Gamma$ has a unique Nash equilibrium if and only if any extended game $G=\langle\Gamma, \mu, c\rangle$ in which players with no weakly dominant actions sample at least once there is a unique sequential sampling equilibrium.

An analogous result holds when, in symmetric extended binary action games (same payoff functions, same prior, same sampling cost), one restricts to symmetric Nash equilibria and symmetric sequential sampling equilibria. While uniqueness of a Nash equilibrium implies uniqueness of a sequential sampling equilibrium, it is not the case that the two coincide.

A well-known and counter-intuitive prediction of Nash equilibrium pertains to generalized matching pennies, that is, $2 \times 2$ games with a unique Nash equilibrium in fully mixed strategies, whose structure is illustrated in Figure 1. When the payoffs to action $a_{i}$ of player $i$ increase, Nash equilibrium predicts that the probability with which action $a_{i}$ is chosen remains the same and it is, instead, the opponent's mixed strategy that adjusts to make player $i$ indifferent between choosing any of the two actions-what one could call the opponent-payoff choice effect. However, experimental evidence shows that increasing player $i$ 's payoffs to an action leads that player to choose that action more often, an own-payoff choice effect. ${ }^{18}$ This motivated the emergence of different models, one of the most successful of which quantal response equilibrium (McKelvey and Palfrey, 1995), which directly embeds monotonicity of choices with respect to payoffs in the assumptions for players' behavior (Goeree et al., 2005).

Sequential sampling equilibrium not only rationalizes this empirical regularity via comparative statics pertaining to behavior induced by optimal information acquisition, it delivers novel behavior implications regarding stopping times. Increasing player $i$ 's payoffs to action $a_{i}$,(1) increases the equilibrium probability that player $i$ chooses action $a_{i}$, and (2) leads their opponent, player $j$, in equilibrium, choosing the best response to action $a_{i}$ more often and faster, and their other action less often and slower, in the sense of Theorem 3. If the first observation states sequential sampling equilibrium predicts the own-payoff choice effect, ${ }^{19}$ the second uncovers an entirely novel prediction relating equilibrium choices and stopping time. Both follow directly from combining Proposition 2, Theorem 3, and Lemma 5.

Supporting Evidence. To investigate whether these predictions find support in existing data, we rely on experimental data generously made available by Friedman and Ward (2022) who collected data on choices and decision times for six different generalized matching pennies games. The goal of this exercise is not to fit data or claim that sequential sampling equilibrium perfectly describes subjects' behavior or that it does so better than other existing models, but rather to present suggestive evidence supporting its novel behavioral implications. No feedback or information was provided throughout the experiment; details on the experiment, the data, and further analysis can be found in Appendix C.

[^12]$\mathbb{P}\{$ Player C stops before $t$ and chooses action\}

(a) Action a

(b) Action b
$$
-1 / 20=1 / 10 \quad 1 / 4 \quad 1 / 2=-2=2=4
$$

Figure 2: Opponent-Payoff Time Effect
Notes: The figure compares choices and decision times in generalized matching pennies games as given in Figure 1, for $\gamma_{C}=1$ (and scaled by 20). The data is from Friedman and Ward (2022). The panels exhibit the frequency with which subjects in the player $C$ 's role take a given action ( $a$ in panel (a); $b$ in panel (b)) before time $t$ (in seconds). Different lines correspond to games in which the player $M$ has different payoffs to action $a$. This figure uses only choice data for instances where beliefs were not elicited. The same patterns are present when beliefs are elicited. See Appendix C for further details on the data.

As shown in Figure 2, if one is to interpret stopping time as a proxy for decision time, the data supports our predictions: when increasing $\delta_{M}$ subjects in the Clasher's role do tend to choose action $b$ not only more often but also faster. Moreover, they choose action $a$ less often and slower.

### 3.3. Time-Revealed Preference Intensity

In this section we characterize how stopping time relates to players' posterior beliefs by considering a general family of priors in binary action games. For this section, we restrict attention to games in which payoffs are linear in the opponent's distribution of actions, i.e. $u_{i}\left(a_{i}, \sigma_{-i}\right)=$ $\mathbb{E}_{\sigma_{-i}}\left[u_{i}\left(a_{i}, a_{-i}\right)\right]$.

Beta Beliefs. For tractability, we consider priors that are linear in new information in a manner that mimics Bayesian updating for Gaussian priors:

Definition 6. A prior $\mu_{i}$ is said to be linear in the accumulated information if it is non-degenerate and there are constants $a_{t}, b_{t} \in \mathbb{R}$ such that for any $y_{i}^{t} \in \mathscr{Y}_{i}$ the posterior mean satisfies $\mathbb{E}_{\mu_{i}}\left[\sigma_{-i} \mid\right.$ $\left.y_{i}^{t}\right]=a_{t} \sum_{\ell=1}^{t} y_{i, \ell}^{t}+b_{t}$.

This property, together with the fact that beliefs are a martingale and some algebraic manipulation, allows us to write the posterior mean as a convex combination of the prior mean and the empirical mean of the accumulated information, $\mathbb{E}_{\mu_{i}}\left[\sigma_{-i} \mid y_{i}^{t}\right]=\alpha_{t} / t \cdot \sum_{\ell=1}^{t} y_{i, \ell}^{t}+\left(1-\alpha_{t}\right) \cdot \mathbb{E}_{\mu_{i}}\left[\sigma_{-i}\right]$, where $\alpha_{t} / t=1 /\left(\left(1-\alpha_{1}\right) / \alpha_{1}+t\right) \in(0,1)$. This is extremely convenient as, by linearity of expected utility, one can then analyze optimal stopping just relying on the belief mean and the number of samples. In fact, as shown by Diaconis and Ylvisaker (1979, Theorem 5), identifies a specific parametric class of priors: a prior $\mu_{i}$ is linear in the accumulated information if and only if it is a Beta distribution.

Collapsing Boundaries. When beliefs are linear in the accumulated information, we have the following characterization of the set of beliefs at which player $i$ optimally stops:

Proposition 3. Let $\Gamma$ be a binary action game. Suppose that there is $\tilde{\sigma}_{-i} \in \Delta\left(A_{-i}\right)$ such that $u_{i}\left(1, \tilde{\sigma}_{-i}\right)=u_{i}\left(0, \tilde{\sigma}_{-i}\right)$. For any $c_{i}>0$, there are continuous functions $\bar{\sigma}_{-i}, \underline{\sigma}_{-i}: \mathbb{R}_{++} \rightarrow[0,1]$ such for any Beta distributed prior $\mu_{i}$ with parameters $(\alpha, \beta)$ player $i$ does not optimally stop at $\mu_{i}$ if and only if $\mathbb{E}_{\mu_{i}}\left[\sigma_{-i}\right] \in\left(\underline{\sigma}_{-i}(\alpha+\beta), \bar{\sigma}_{-i}(\alpha+\beta)\right)$. Furthermore, $\bar{\sigma}_{-i}$ is decreasing and $\underline{\sigma}_{-i}$ is increasing, and $\exists T_{i}$ such that $\forall t \geq T_{i} \bar{\sigma}_{-i}(t)=\underline{\sigma}_{-i}(t)=\tilde{\sigma}_{-i}$.

The proof of the result is in Appendix A.
Proposition 3 shows that when beliefs are linear in accumulated information, it is sufficient to consider the posterior mean to characterize the beliefs at which player $i$ continues sampling at any given moment as is illustrated in Figure 3. Note that if $\mu_{i}$ is a Beta distribution with parameters summing to $t$, then $\mu_{i} \mid y_{i}$ has parameters summing to $t+1$. The continuation region is then characterized by an upper and lower threshold that delimit a decreasing interval that "collapses" to a single point: the distribution at which player $i$ is indifferent between either action. This translates to our setting what is commonly known in the neuroscience literature as "collapsing boundaries". ${ }^{20}$

One can then interpret the stopping time as an indicator of the intensity of player $i$ 's preference for one action over another: player $i$ samples for longer if and only if the player is sufficiently close to being indifferent between the two alternatives, a phenomenon that resembles existing experimental evidence in individual decision-making (e.g. Konovalov and Krajbich, 2019). In other words, Proposition 3 entails a behavior marker in the form of time-revealed preference intensity,

[^13]

Figure 3: Stopping Regions for Beta Priors
Notes: The figure exhibits the continuation region (shaded area) and the stopping thresholds (darker blue lines) for posterior means at which player $i$ with a Beta stops. The figure also illustrates the possible realizations of the sampling process for a player with a uniform prior (Beta(1,1)), with the posterior means indicated by circles.
akin to results in Alós-Ferrer et al. (2021).
When the absolute difference in the expected payoffs is known-the case where the prior's support is a doubleton-the stopping region is characterized by fixed bounds in terms of the posterior means as shown by Arrow et al. (1949). In contrast, when there is richer uncertainty about the difference in expected payoffs, as when the prior is given by a Beta distribution, the stopping region is characterized by bounds that collapse to the posterior mean that makes the individual indifferent between the two alternatives. A clear parallel emerges between our setup and that in Fudenberg et al. (2018), where the individual infers the difference in payoffs of two alternatives from the drift of a Brownian motion and a similar contrast between known and unknown payoff differences gives rise to, respectively, fixed and collapsing stopping bounds.

Comparative Statics in Stopping Beliefs. From Proposition 3 and Theorem 3, we obtain that the distribution of beliefs shifts monotonically with respect to the true distribution. Specifically, approximating the stopping posterior mean by the threshold, $\mathbb{E}_{\mu_{i}}\left[\sigma_{-i} \mid y_{i}^{\tau_{i}}\right],{ }^{21}$ and labeling actions so that $u_{i}\left(1, \sigma_{-i}\right)-u_{i}\left(0, \sigma_{-i}\right)$ is increasing in $\sigma_{-i}$, then player $i$ 's stopping (threshold) beliefs increase in a first-order stochastic dominance sense as $\sigma_{-i}$ increases. This is because a higher $\sigma_{-i}$ leads to a higher probability that player $i$ chooses action 1 more often and faster (resp. action 0 less

[^14]|  | Distance to Indifference |  |  |
| :---: | :---: | :---: | :---: |
|  | Player M | Player C | Both |
|  | (1) | (2) | (3) |
| Log Decision Time | $\begin{gathered} -3.682^{* * *} \\ (1.225) \end{gathered}$ | $\begin{gathered} -2.021^{* *} \\ (0.881) \end{gathered}$ | $\begin{gathered} -2.961^{* * *} \\ (0.790) \end{gathered}$ |
| Intercept | $\begin{gathered} 42.314^{* * *} \\ (4.181) \end{gathered}$ | $\begin{gathered} 45.365^{* * *} \\ (2.670) \end{gathered}$ | $\begin{gathered} 48.185^{* * *} \\ (2.435) \end{gathered}$ |
| Fixed Effects | Game | Game | Role $\times$ Game |
| R-Squared | 0.08 | 0.27 | 0.18 |
| Observations | 1620 | 1680 | 3300 |

Heteroskedasticity-robust standard errors in parentheses.
${ }^{*} p<0.1,{ }^{* *} p<0.05,{ }^{* * *} p<0.01$.

## Table 1: Decision Time and Reported Beliefs: Time-Revealed Preference Intensity

Notes: This table presents regression results on the relation between log decision times (in seconds) and the distance between reported beliefs to indifference points. Reported beliefs refer to elicited beliefs about the probability the opponents plays action $a$; the indifference point refers to the probability that makes the player indifferent between taking either action. Columns (1) and (2) only use data for subjects in the roles of player $M$ and $C$, respectively; column (3) uses both. The games are generalized matching pennies games as given in Figure 1, for $\gamma_{C}=1$ (scaled by 20); the data is from Friedman and Ward (2022).

Player C's Mean Beliefs


Figure 4: Opponent Payoff and Beliefs
Notes: This figure exhibits empirical CDF of reported (mean) beliefs about the probability with which subjects in the role of player $C$ believe their opponent (in the role of player $M$ ) will take action $a$; different lines correspond to games with different payoffs to action $a$ for player $M$ as parametrized by $\delta_{M}$. The games are generalized matching pennies games as given in Figure 1, for $\gamma_{C}=1$ (scaled by 20); the data is from Friedman and Ward (2022).
often and slower), implying that the posterior mean has to exceed a higher threshold when the player stops earlier (resp. later), as the upper (resp. lower) bound characterizing the continuation region is decreasing (resp. increasing) in the stopping time.

Supporting Evidence. Relying again on Friedman and Ward's (2022) data, we find support for both these predictions: (1) decision time is significantly negatively related to the distance between the reported mean belief and the indifference point (Table 1), and (2) increasing a player's payoff to an action significantly shifts the opponent's beliefs in the predicted first-order stochastic dominance sense (Figure 4)-see Online Appendix C for additional statistical tests.

## 4. Relation to Nash Equilibrium

One initial interpretation of Nash equilibrium posits that equilibrium beliefs are reached as players "accumulate empirical information" (Nash, 1950, p. 21). In a sequential sampling equilibrium, players accumulate empirical information but at a cost. A natural question is whether, as these costs vanish, sequential sampling equilibria converge to a Nash equilibrium. In this section we show this is the case. Formally,

Theorem 4. Let $\Gamma$ be a normal-form game, $\mu$ a collection of priors, and $\left\{c^{n}\right\}_{n}$ be a sequence of sampling costs such that $c^{n} \rightarrow 0$. For any sequence $\left\{\sigma^{n}\right\}_{n}$ such that each $\sigma_{n}$ is a sequential sampling equilibrium of extended game $G^{n}=\left\langle\Gamma, \mu, c^{n}\right\rangle$, the limit points of $\left\{\sigma^{n}\right\}_{n}$ are Nash equilibria of $\Gamma$.

The claim is conventional in form: players best-respond to their beliefs and their beliefs converge to the true distribution of actions of their opponents. ${ }^{22}$ The main complication comes from the fact that, conditional on stopping, the observations $y_{i}^{\tau_{i}}$ are not independent nor independently distributed according to player $i$ 's opponents' action distribution. To overcome this issue, the proof (see Appendix A) relies on three arguments. First, from Lemma 4 one has that as sampling costs vanish, players acquire a minimum number of observations $T$, and, for that minimum number, each observation $y_{i}^{T}$ is iid according to the opponents' action distribution. Second, we note beliefs accumulate at a uniform rate around the empirical mean of the observed signals. Finally, we use the optional stopping theorem to show that beliefs upon stopping converge to the true underlying distribution in an appropriate manner.

Some comments on which Nash equilibria can be selected in this manner are in order. First, let us define the concept of reachability of a Nash equilibrium:

[^15]Definition 7. A Nash equilibrium $\sigma$ of a normal-form game $\Gamma$ is reachable if there is a collection of priors $\mu$, a sequence of costs $\left\{c^{n}\right\}_{n} \subset \mathbb{R}_{++}$such that $c^{n} \rightarrow 0$, and a sequence $\left\{\sigma^{n}\right\}^{n}$, where for each $n, \sigma^{n}$ is a sequential sampling equilibrium of the extended game $G^{n}=\left\langle\Gamma, \mu, c^{n}\right\rangle$, such that $\sigma^{n} \rightarrow \sigma$. A Nash equilibrium if robustly reachable if it is reachable for any collection of priors $\mu$.

In the remainder of the section, we will restrict player's payoffs to be linear in distributions as usual. In other words, we require that, for every player $i, u_{i}\left(a_{i}, \sigma_{-i}\right)=\mathbb{E}_{\sigma_{-i}}\left[u_{i}\left(a_{i}, a_{-i}\right)\right]$, as conventional. This will be a maintained assumption throughout the rest of this section.

Our first result provides, separately, necessary and sufficient conditions for reachability of a Nash equilibrium.

Proposition 4. Let $\Gamma$ be a normal-form game. (1) If $\sigma$ is a Nash equilibrium of $\Gamma$ involving weakly dominated actions, then $\sigma$ is not reachable. (2) If a is a pure-strategy Nash equilibrium of $\Gamma$ not involving weakly dominated actions, then a is reachable.

Part (1) holds since for any prior, no player will ever choose weakly dominated actions-recall that priors have full support. For (2), note that if $a$ does not involve weakly dominated strategies, then, by Pearce's (1984) Lemma 4, for each player $i$ there is $\sigma_{-i}^{0} \in \operatorname{int} \Delta\left(A_{-i}\right)$ such that $a_{i}$ is a best response to $\sigma_{-i}^{0}$. If we endow each player $i$ with prior $\mu_{i} \in \Delta\left(\Delta\left(A_{-i}\right)\right)$ corresponding to a Dirichlet distribution with mean $\sigma_{-i}^{0}$, then $\alpha_{i}$ is a best response to any posterior belief $\mu_{i} \mid y_{i}^{t}$ when $y_{i, t}=a_{-i}$. Hence, for any costs $c^{n}, \sigma$ is sequential sampling equilibrium of $\left\langle\Gamma, \mu, c^{n}\right\rangle$. Note that we require the Nash equilibrium to be in pure-strategies in order to control posterior beliefs exactly, as otherwise, with some probability, $\sigma_{i}$ may not be a best response to the posterior belief held upon stopping.

For a Nash equilibrium to be reachable with any priors, we obtain a sufficient condition:
Proposition 5. If a is a pure-strategy Nash equilibrium in undominated strategies of the normalform game $\Gamma$ such that, for any player $i, a_{i}$ is a best response to any $\sigma_{-i}^{\prime} \in B_{\epsilon_{i}}\left(\delta_{a_{-i}}\right)$ for some $\epsilon_{i}>0$, then $a$ is robustly reachable.

The intuition for the proof (in Appendix A) is as follows: for any prior, if player $i$ samples enough $a_{-i}$ observations, their posterior mean will lie within $\epsilon_{i}$ of $a_{-i}$ and choosing $a_{i}$ is optimal. Lemma 4 guarantees that players do sample enough. Our requirement that $a_{i}$ is a best response to any distribution of opponents' actions assigning high enough probability to $a_{-i}$ is at the same time more relaxed than strict Nash equilibrium, and more restrictive than trembling hand perfection.

## 5. A Dynamic Formulation

One can view sequential sampling equilibrium as a steady state of a dynamic process in whereby agents sequentially sample from past realizations. This section formalizes that argument.

Dynamic Sequential Sampling. To fix ideas, consider a simple dynamic process, similar to fictitious play. Fix an extended game $G$. Every period, $n=1,2, \ldots$, a unit measure of agents plays the extended game $G$, evenly divided across the different roles $I$. Each agent believes they face a stationary distribution of opponents' actions, matching the empirical frequency of past actions, $\sigma^{n-1} \in \Delta(A)$, not knowing calendar time.

Within period $n$, each agent with role $i$ leans about $\sigma_{-i}^{n-1}$ by optimally sequentially samples according to $\tau_{i}$. Upon stopping, the agent best responds to their posterior beliefs. ${ }^{23}$ This induces a distribution of actions and types in period $n$ given by $b\left(\sigma^{n-1}\right)$, where $b: \Delta(A) \rightarrow \Delta(A)$ is such that $b(\sigma)(\alpha):=\times_{i \in I} b_{i}\left(\sigma_{-i}\right)\left(a_{i}\right)$, with $b_{i}\left(\sigma_{-i}\right):=\mathbb{E}_{\sigma_{-i}}\left[\sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right]$, as before. After taking an action, agents then exit and are replaced by a new population as is standard in evolutionary models of learning in strategic settings. At the start of the following period, the empirical frequency is then $\sigma^{n}=\frac{1}{n+1} b\left(\sigma^{n-1}\right)+\frac{n}{n+1} \sigma^{n-1}$, with $\sigma_{0}$ given. Call any such $\left\{\sigma^{n}\right\}_{n}$ a dynamic sequential sampling process of $G$.

While akin to fictitious play (Brown, 1951), under dynamic sequential sampling, each agent observes but a sample of past play realizations and the sample itself is an endogenous object.

Equilibria and Steady States. We now show an equivalence between sequential sampling equilibria and steady states of dynamic sequential sampling processes.

Theorem 5. Let $G$ be an extended game. $\sigma$ is a sequential sampling equilibrium of $G$ if and only if there is some dynamic sequential process $\left\{\sigma^{n}\right\}_{n}$ of $G$ such that $\sigma^{n} \rightarrow \sigma$.

Proof. We restrict attention to the if part, since the converse is immediate. Let $\bar{\sigma}$ denote the limit of $\sigma^{n}$. Then,

$$
0=\lim _{n \rightarrow \infty}\left\|\sigma^{n}-\bar{\sigma}\right\|_{\infty}=\left\|\lim _{n \rightarrow \infty} \sigma^{n}-\bar{\sigma}\right\|_{\infty}=\left\|\lim _{n \rightarrow \infty} \frac{1}{n+1} \sigma^{0}+\frac{n}{n+1}\left(\frac{1}{n} \sum_{\ell \in[0 \ldots n-1]} b\left(\sigma^{\ell}\right)\right)-\bar{\sigma}\right\|_{\infty} .
$$

As $\sigma^{n} \rightarrow \bar{\sigma}$ and $b$ is continuous, then $b\left(\sigma^{n}\right) \rightarrow b(\bar{\sigma})$. Consequently, the Cesàro mean $\frac{1}{n} \sum_{\ell \in[0 . . n-1]} b\left(\sigma^{\ell}\right)$ also converges to $b(\bar{\sigma})$ and therefore $0=\|b(\bar{\sigma})-\bar{\sigma}\|_{\infty} \Longrightarrow b(\bar{\sigma})=\bar{\sigma}$.

[^16]The steady-state characterization of sequential sampling equilibria in Theorem 5 provides a clear analogue to the characterization of Nash equilibria as steady-states of fictitious play in Fudenberg and Kreps (1993). The main difference between fictitious play and the dynamic process analyzed is that, whereas data is freely observable in fictitious play, sequential sampling players face information acquisition costs. Moreover, as we have seen (Theorem 4), as these costs vanish, limiting sequential sampling equilibria correspond to Nash equilibria. Below we discuss two ways in which the dynamic process can be generalized.

Remark 4. Often it may be the case that information about more recent events is more easily accessible. This can be modeled as a giving a different weight to each period, for instance, exponential discounting past data: $\sigma^{n}=\beta \sigma^{n-1}+(1-\beta) b\left(\sigma^{n-1}\right), \beta \in(0,1)$. Theorem 5 also holds under this alternative definition: as $\sigma^{n} \rightarrow \bar{\sigma} \Longrightarrow b\left(\sigma^{n}\right) \rightarrow b(\bar{\sigma})$ and, for any fixed $\ell, \beta^{n-1-\ell} b\left(\sigma^{\ell}\right) \rightarrow 0$, we have $\sigma^{n}=\beta^{n} \sigma^{0}+(1-\beta) \cdot \sum_{\ell \in[0 . . n-1]} \beta^{n-1-\ell} \cdot b\left(\sigma^{\ell}\right) \rightarrow \bar{\sigma}=b(\bar{\sigma})$.

Remark 5. The assumption that there is a continuum of agents for each role is also not essential: a similar result holds when the populations are finite. Write $a^{n}$ for the realized actions in period $n$ and $\sigma^{n}$ for their empirical frequency (given $a^{0}$ ), with $a^{n} \sim b\left(\sigma^{n-1}\right) .{ }^{24}$ Note that $\sigma^{n} \rightarrow \bar{\sigma}$ still implies that $b\left(\sigma^{n}\right) \rightarrow b(\bar{\sigma})$, and the arguments above remain the same, with $a^{n}$ converging in distribution to a sequential sampling equilibrium.

Convergence. While in general we cannot exclude dynamic sequential sampling from cycling and failing to converge-similarly to what occurs with fictitious play ${ }^{25}$ - in specific classes of games, convergence and asymptotic stability are guaranteed. ${ }^{26}$ This next proposition shows this is the case for binary action games, which we will discuss in more depth in the next section.

Proposition 6. Let $G=\langle\Gamma, \mu, c\rangle$ be a two-player extended game. If $\Gamma$ has a unique Nash equilibrium, the limit of dynamic sequential sampling is a globally asymptotically stable sequential sampling equilibrium.

The proof is deferred to the appendix.

[^17]
## 6. Extensions and Discussion

We conclude with a discussion of possible extensions of sequential sampling equilibrium.
Types and Bayesian Games. Sequential sampling equilibrium can be easily extended to accommodate Bayesian games. This aligns with the idea that sequential sampling equilibria corresponds to the case in which players don't know their opponents' payoffs, as given by their type. Alternatively, one can consider that different types characterize different settings, and players are using information about behavior in similar settings to make inferences about behavior in the particular game they are facing.

In particular, consider games described by $\Gamma:=\langle I, A, \Theta, \rho, u\rangle$, such that $I$ denote the finite set of players, $A:=\times_{i \in I} A_{i}$ the (finite) set of action profiles, $\Theta=\times_{i \in I} \Theta_{i}$ the (finite) set of type profiles, where $\Theta_{i}$ are player $i$ 's possible types, $u_{i}: A \times \Theta \rightarrow \mathbb{R}$ payoff functions, and $\rho \in \Delta(\Theta)$ a distribution over types. Endowing each player $i$ with a prior $\mu_{i} \in \Delta\left(\Delta\left(A_{-i} \times \Theta\right)\right)$ and a sampling cost $c_{i}$ as before we have extended games $G=\langle\Gamma, \mu, c\rangle$. Each player $i$ with type $\theta_{i}$ now learns about the joint distribution of opponents' action profiles and type profiles, $q_{i} \in \Delta\left(A_{-i} \times \Theta\right)$, sequentially sampling from $q_{i}$ at cost $c_{i}$ and stopping according to the earliest optimal stopping time $\tau_{i, \theta_{i}}$.

A sequential sampling equilibrium $\sigma$ would then correspond to a fixed point such that $\sigma_{i, \theta_{i}}=$ $\mathbb{E}_{q_{i}}\left[\sigma_{i, \theta_{i}}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i, \theta_{i}}}\right)\right]$, where $\sigma_{i, \theta_{i}}^{*}\left(\mu_{i}\right)$ is a selection of best responses given belief $\mu_{i}$, and $q_{i, \theta_{i}}\left(a_{-i}, \theta\right)=$ $\rho(\theta) \times{ }_{j \neq i} \sigma_{j, \theta_{j}}\left(a_{j}\right)$ for every $a_{-i}=\left(a_{j}\right)_{j \neq i} \in A_{-i}$ and $\theta=\left(\theta_{j}\right)_{j \in I} \in \Theta$.

Different assumptions on players beliefs will give rise to different equilibria. To apply similar arguments to obtain existence of an equilibrium, we need but to require that players know the distribution of their own types and that $\mu_{i}$ has full support on the set of distributions $q_{i} \in \Delta\left(A_{-i} \times\right.$ $\Theta$ ) satisfying $q_{i}\left(\theta_{i}\right)=\rho\left(\theta_{i}\right)$ for any $\theta_{i} \in \Theta_{i} .{ }^{27}$ Differently, one could assume players know the true distribution of types, or, when types are independent, that they know so.

With the proper adjustments, behavioral implications can also be obtained, now comparing across types. For instance, if the payoff to action $a_{i}$ is higher for type $\theta_{i}$ than for $\theta_{i}^{\prime}$, everything else equal, in every sequential sampling equilibrium type $\theta_{i}$ chooses action $a_{i}$ more often and faster (in the sense of Theorem 3) than type $\theta_{i}^{\prime}$, a result that can be exploited and tested in a number of traditional settings, from global games to voting. Finally, convergence to Bayesian Nash equilibria when sampling costs vanish can be similarly obtained.

[^18]General Information Structures and Analogy Partitions. Throughout, it was assumed that players observe action profiles drawn from a steady state distribution. Often, of course, informationand even memory-is fuzzier, and it is not possible to perfectly distinguish between certain actions taken by others, or even to observe what some other players do at all. In Online Appendix B, we provide sufficient conditions under which is it possible to generalize sequential sampling equilibrium to cases under which players observe not action profiles of their opponents, but a garbling, thereby accommodating situations such as noisy recollections, or missing or misrecorded data. This formalism can also be used to formalize the idea that players' payoffs may depend on the behavior of others about which they are unable to obtain information. For instance, it may be possible to obtain information about behavior within the same firm, but impossible to learn what people do in other firms. For the special case in which players are unable to distinguish between specific action profiles (or types) of their opponents, as sampling costs vanish, sequential sampling equilibria reach not Nash equilibria but analogy-based expectation equilibria (Jehiel, 2005; Jehiel and Koessler, 2008).

Sampling Costs and Discounting. In this paper, we considered a constant additive cost per observation. One could have defined this cost of information in a more general manner, allowing it to depend on the number of observations already acquired, or allowing a finite number of observations at no cost. Alternatively, one could rely on discounting payoffs instead. It is indeed possible to extend the setup to accommodate either, posing no problem for existence of an equilibrium.

Misspecified Priors. Another maintained assumption was that priors are not misspecified, i.e. players are able to learn the true data generating process. This assumption was crucial to obtain existence of an equilibrium: it is the full support of players' prior that guarantees that, as they acquire more and more observations, their beliefs accumulate around a degenerate distribution. When, instead, priors are misspecified, it is possible that players never stop sampling (according to the true data generating process), even though they believe they will (according to their posterior beliefs). We provide one such example of nonexistence in Online Appendix D.

Empirical Estimation. Empirical estimation of the baseline model is facilitated by utilizing Dirichlet beliefs, a flexible class of priors with parameters $\alpha^{i}=\left(\alpha_{1}^{i}, \ldots, \alpha_{n}^{i}\right)$, where $n=\left|A_{-i}\right|$. When $n=2$, it corresponds to the Beta distribution. It can be shown that, for this rich class of conjugate priors, given player $i$ 's payoffs and sampling costs, there is a compact set of param-
eters $\alpha^{i}$ for which player $i$ would sample. This also implies the existence of a tight upper bound for the stopping time $\bar{T}$ for any Dirichlet prior, converting the infinite horizon problem in a finite horizon one and rendering it a computationally tractable problem. Assuming payoffs are known, a particular parameter and sampling cost would map to a given joint distribution of actions and stopping time. Maximum likelihood and other estimation procedures analogous to those relied upon for sequential sampling models in individual decision-making tasks can then be used (see e.g. Myers et al., 2022). ${ }^{28}$

Myopic Sequential Sampling. Finally, a comment on a simpler, alternative, version of sequential sampling equilibrium, in which sequential sampling behavior is myopic. This is closely related to a recent paper by Alaoui and Penta (2022), which provides an axiomatization of the deliberation process as one of myopic sequential information acquisition (cf. their Theorem 4). In other words, the myopic stopping time would be given by $\tau_{i}^{M}(\omega):=\min \left\{t \mid \mathbb{E}_{\mu_{i}}\left[v_{i}\left(\mu_{i} \mid y_{i}^{t+1}\right) \mid\right.\right.$ $\left.\left.y_{i}^{t}(\omega)\right]-v_{i}\left(\mu_{i} \mid y_{i}^{t}(\omega)\right) \leq c_{i}\right\}$, with players stopping whenever the expected value of sampling information is smaller than the cost of one more observation.

One could then define myopic sequential sampling equilibria simply by replacing the optimal stopping time $\tau_{i}$ with the myopic one $\tau_{i}^{M}$. While appealing for its simplicity, one immediate implication is that myopic sequential sampling equilibria do not generically reach Nash equilibria as the sampling costs vanish. This is because the expected value of sampling information is zero whenever the optimal choices under belief $\mu_{i}$ are still optimal under the posterior $\mu_{i} \mid y_{i}$, no matter the realization of the observation. It is therefore immediate that any prior that is sufficiently concentrated around a distribution of opponents' actions for which $a_{i}$ is a strict best response, for any $c_{i}>0$, player $i$ will not see it worthwhile to acquire information. Thus, if in no Nash equilibrium it is optimal to play such an action with probability 1 , myopic sequential sampling will fail to converge to a Nash equilibrium.

## References

Agranov, M., A. Caplin, and C. Tergiman. 2015. "Naive play and the process of choice in guessing games." Journal of the Economic Science Association 1 146-157. 10.1007/s40881-015-0003-5. Alaoui, L., K. A. Janezic, and A. Penta. 2020. "Reasoning about others' reasoning." fournal of Economic Theory 189 (105091): 1-51. 10.1016/j.jet.2020.105091.
Alaoui, L., and A. Penta. 2016. "Endogenous Depth of Reasoning." Review of Economic Studies 83 (4): 1297-1333. 10.1093/restud/rdv052.

[^19]Alaoui, L., and A. Penta. 2022. "Cost-Benefit Analysis in Reasoning." Journal of Political Economy 130 (4): 881-925. 10.1086/718378.
Aliprantis, C. D., and K. C. Border. 2006. Infinite Dimensional Analysis: A Hitchhiker's Guide. Springer.
Alós-Ferrer, C., and J. Bruckenmaier. 2021. "Cognitive sophistication and deliberation times." Experimental Economics 24 558-592. 10.1007/s10683-020-09672-w.
Alós-Ferrer, C., E. Fehr, and N. Netzer. 2021. "Time Will Tell: Recovering Preferences When Choices Are Noisy." Journal of Political Economy 129 (6): 1828-1877. 10.1086/713732.
Alós-Ferrer, C., and M. Garagnani. 2022. "Strength of preference and decisions under risk." Journal of Risk and Uncertainty 64 309-329. 10.1007/s11166-022-09381-0.
Arrow, K., D. Blackwell, and M. Girshick. 1949. "Bayes and Minimax Solutions of Sequential Decision Problems." Econometrica 17 (3/4): 213-244. 10.2307/1905525.
Aumann, R., and A. Brandenburger. 1995. "Epistemic Conditions for Nash Equilibrium." Econometrica 63 (5): 1161-1180. 10.2307/2171725.
Bakkour, A., A. Zylberberg, M. Shadlen, and D. Shohamy. 2018. "Value-based decisions involve sequential sampling from memory." Working Paper. 10.1101/269290.
Battigalli, P., M. Gilli, and M. C. Molinari. 1992. "Learning and Covergence to Equilibrium in Repeated Strategic Interactions: An Introductory Survey." Ricerche Economiche 46 335-378.
Berk, R. 1966. "Limiting Behavior of Posterior Distributions when the Model is Incorrect." Annals of Mathematical Statistics 37 (1): 51-58. 10.1214/aoms/1177699597.
Bhui, R. 2019. "Testing Optimal Timing in Value-Linked Decision Making." Computational Brain and Behavior 2 (2): 85-94. 10.1007/s42113-019-0025-9.
Biderman, N., A. Bakkour, and D. Shohamy. 2020. "What Are Memories For? The Hippocampus Bridges Past Experience with Future Decisions." Trends in Cognitive Sciences 24 (7): 542-556. 10.1016/j.tics.2020.04.004.

Biderman, N., and D. Shohamy. 2021. "Memory and decision making interact to shape the value of unchosen options." Nature Communications 12 1-12. 10.1038/s41467-021-24907-x.
Bloedel, A., and W. Zhong. 2021. "The Cost of Optimally Acquired Information." Working Paper 1-105.
Brown, G. W. 1951. Activity Analysis of Production and Allocation. Chap. XXIV. Iterative Solution of Games by Fictitious Play 374-376, John Wiley and Sons.
Camerer, C., T.-H. Ho, and J.-K. Chong. 2004. "A Cognitive Hierarchy Model of Games." Quarterly fournal of Economics 119 (3): 861-898. 10.1146/annurev-economics-080315-015417.
Caplin, A., and M. Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition." American Economic Review 105 (7): 2183-2203. 10.1257/aer. 20140117.
Che, Y.-K., and K. Mierendorff. 2019. "Optimal Dynamic Allocation of Attention." American Economic Review 109 (8): 2993-3029. 10.1257/aer. 20171000.
Chen, P. N., J. X. He, and H. S. Qin. 2001. "A Proof of the Jacobian Conjecture on Global Asymptotic Stability." Acta Mathematica Sinica 17 (1): 119-132. 10.1007/s101140000098.
Chernoff, H. 1961. "Sequential Tests for the Mean of a Normal Distribution." Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability 179-91.
Chierchia, G., R. Nagel, and G. Coricelli. 2018. "Betting on nature" or "betting on others": anti-coordination induces uniquely high levels of entropy." Scientific Reports 8 (1): 1-11. 10.1038/
s41598-018-21962-1.
Chiong, K., M. Shum, R. Webb, and R. Chen. 2023. "Combining Choice and Response Time Data: A Drift-Diffusion Model of Mobile Advertisements." Management Science 1-20. 10.1287/ mnsc.2023.4738.
Clithero, J. 2018a. "Improving out-of-sample predictions using response times and a model of the decision process." Journal of Economic Behavior and Organization 148 344-375. 10.1016/j.jebo. 2018.02.007.

Clithero, J. 2018b. "Response Times in Economics: Looking Through the Lens of Sequential Sampling Models." Journal of Economic Psychology 69 61-86. 10.1016/j.joep.2018.09.008.
Costa-Gomes, M., and G. Weizsäcker. 2008. "Random belief equilibrium in normal form games." Review of Economic Studies 75 (3): 729-762. 10.1111/j.1467-937X.2008.00498.x.
Denti, T. 2023. "Unrestricted information acquisition." Theoretical Economics (Forthcoming) 18 1101-1140. 10.3982/TE4541.
Denti, T., M. Marinacci, and A. Rustichini. 2022. "Experimental Cost of Information." American Economic Review 112 (9): 3106-3123. 10.1257/aer. 20210879.
Diaconis, P., and D. Freedman. 1990. "On the Uniform Consistency of Bayes Estimates for Multinomial Probabilities." Annals of Statistics 18 (3): 1317-1327. 10.1214/aos/1176347751.
Diaconis, P., and D. Ylvisaker. 1979. "Conjugate Priors for Exponential Families." The Annals of Statistics 7 (2): 269-281. 10.1214/aos/1176344611.
Duncan, K., and D. Shohamy. 2020. "52: Memory, Reward, and Decision-Making." In The Cognitive Neurosciences, edited by Poeppel, D., G. R. Mangun, and M. S. Gazzaniga 617-629, The MIT Press, . 10.7551/mitpress/11442.003.0068.
Esponda, I., and D. Pouzo. 2016. "Berk-Nash Equilibrium: A Framework for Modeling Agents With Misspecified Models." Econometrica 84 (3): 1093-1130. 10.3982/ecta12609.
Esteban-Casanelles, T., and D. Gonçalves. 2022. "The Effect of Incentives on Choices and Beliefs in Games: An Experiment." Working Paper.
Evans, L. C., and R. F. Gariepy. 2015. Measure Theory and Fine Properties of Functions. Chapman and Hall.
Fehr, E., and A. Rangel. 2011. "Neuroeconomic Foundations of Economic Choice-Recent Advances." Journal of Economic Perspectives 25 (4): 3-30. 10.1257/jep.25.4.3.
Ferguson, T. 2008. Optimal Stopping and Applications. UCLA.
Forstmann, B., R. Ratcliff, and E.-J. Wagenmakers. 2016. "Sequential Sampling Models in Cognitive Neuroscience: Advantages, Applications, and Extensions." Annual Review of Psychology 67 (1): 641-666. 10.1146/annurev-psych-122414-033645.
Freedman, D. 1963. "On the Asymptotic Behavior of Bayes' Estimates in the Discrete Case." Annals of Mathematical Statistics 34 (4): 1386-1403. 10.1214/aoms/1177703871.
Friedman, E. 2022. "Stochastic Equilibria: Noise in Actions or Beliefs?" American Economic Journal: Microeconomics 14 (1): 1-65. 10.1257/mic. 20190013.
Friedman, E., and J. Ward. 2022. "Stochastic Choice and Noisy Beliefs in Games: an Experiment." Working Paper 1-78. 10.2139/ssrn. 4190338.
Friedman, J., and C. Mezzetti. 2005. "Random belief equilibrium in normal form games." Games and Economic Behavior 51 (2): 296-323. 10.1016/j.geb.2003.08.004.
Frydman, C., and S. Nunnari. 2023. "Coordination with Cognitive Noise." Working Paper.

Fudenberg, D., and D. Kreps. 1993. "Learning Mixed Equilibria." Games and Economic Behavior 5 (3): 320-367. 10.1006/game.1993.1021.
Fudenberg, D., G. Lanzani, and P. Strack. 2021. "Limit Points of Endogenous Misspecified Learning." Econometrica 89 (3): 1065-1098. 10.3982/ECTA18508.
Fudenberg, D., and D. Levine. 1993. "Self-Confirming Equilibrium." Econometrica 61 (3): 523545. 10.2307/2951716.

Fudenberg, D., and D. Levine. 1998. The Theory of Learning in Games. MIT Press.
Fudenberg, D., P. Strack, and T. Strzalecki. 2018. "Speed, Accuracy, and the Optimal Timing of Choices." American Economic Review 108 (12): 3651-3684. 10.1257/aer. 20150742.
Gill, D., and V. Prowse. 2023. "Strategic Complexity and the Value of Thinking." The Economic Journal 133 (650): 761-786. 10.1093/ej/ueac070.
Goeree, J., and C. Holt. 2001. "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions." American Economic Review 91 (5): 1402-1422. 10.1257/aer.91.5.1402.
Goeree, J., C. Holt, and T. Palfrey. 2005. "Regular Quantal Response Equilibrium." Experimental Economics 8 (4): 347-367. 10.1007/s10683-005-5374-7.
Gold, J., and M. Shadlen. 2007. "The Neural Basis of Decision Making." Annual Review of Neuroscience 30 535-574. 10.1146/annurev.neuro.29.051605.113038.
Hawkins, G., B. Forstmann, E.-J. Wagenmakers, R. Ratcliff, and S. Brown. 2015. "Revisiting the Evidence for Collapsing Boundaries and Urgency Signals in Perceptual Decision-Making." Journal of Neuroscience 35 (6): 2476-2484. 10.1523/jneurosci.2410-14.2015.
Hébert, B., and J. La'O. 2023. "Information Acquisition, Efficiency, and Non-Fundamental Volatility." Journal of Political Economy 131 (10): 000. 10.1086/724575.
Hébert, B., and M. Woodford. 2023. "Rational Inattention when Decisions Take Time." Journal of Economic Theory 208 (105612): 1-27. 10.1016/j.jet.2023.105612.
Hommes, C., and M. Ochea. 2012. "Multiple equilibria and limit cycles in evolutionary games with Logit Dynamics." Games and Economic Behavior 74 (1): 434-441. 10.1016/j.geb.2011.05.014. Jehiel, P. 2005. "Analogy-based expectation equilibrium." Journal of Economic Theory 123 (2): 81-104. 10.1016/j.jet.2003.12.003.
Jehiel, P. 2021. "Analogy-based expectation equilibrium and related concepts: Theory, applications, and beyond." Working Paper 1-107.
Jehiel, P., and F. Koessler. 2008. "Revisiting games of incomplete information with analogybased expectations." Games and Economic Behavior 62 (2): 533-557. 10.1016/j.geb.2007.06.006.
Jordan, J. S. 1993. "Three Problems in Learning Mixed-Strategy Nash Equilibria." Games and Economic Behavior 5 (3): 368-386. 10.1006/game.1993.1022.
Konovalov, A., and I. Krajbich. 2019. "Revealed strength of preference: Inference from response times." Judgment and Decision Making 14 (4): 381-394.
Krajbich, I., C. Armel, and A. Rangel. 2010. "Visual fixations and the computation and comparison of value in simple choice." Nature Neuroscience 13 (10): 1292-1298. 10.1038/nn.2635.
Krajbich, I., D. Lu, C. Camerer, and A. Rangel. 2012. "The attentional drift-diffusion model extends to simple purchasing decisions." Frontiers in Psychology 31-18. 10.3389/fpsyg.2012.00193. Krajbich, I., B. Oud, and E. Fehr. 2014. "Benefits of Neuroeconomics Modeling: New Policy Interventions and Predictors of Preference." American Economic Review: Papers \& Proceedings 104 (5): 501-506. 10.1257/aer.104.5.501.

Lehrer, E., and T. Wang. 2020. "Strong Stochastic Dominance." Working Paper 1-31.
Liang, A., X. Mu, and V. Syrgkanis. 2022. "Dynamically Aggregating Diverse Information." Econometrica 90 (1): 47-80. 10.3982/ECTA18324.
Luce, R. D. 1986. Response Times: Their Role in Inferring Elementary Mental Organization. Oxford University Press.
Martin, D. 2017. "Strategic pricing with rational inattention to quality." Games and Economic Behavior 104 131-145. 10.1016/j.geb.2017.03.007.
Massart, P. 1990. "The Tight Constant in the Dvoretzky-Kiefer-Wolfowitz Inequality." Annals of Probability 18 (3): 1269-1283. 10.1214/aop/1176990746.
Matějka, F., and A. McKay. 2015. "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model." American Economic Review 105 (1): 272-298. 10.1257/aer. 20130047.

Matějka, F., and A. McKay. 2012. "Simple Market Equilibria with Rationally Inattentive Consumers." American Economic Review: Papers \& Proceedings 102 (3): 24-29. 10.1257/aer.102.3.24.
McKelvey, R., and T. Palfrey. 1995. "Quantal Response Equilibrium." Games and Economic Behavior 10 (1): 6-38. 10.1006/game.1995.1023.
McKelvey, R., T. Palfrey, and R. Weber. 2000. "The effects of payoff magnitude and heterogeneity on behavior in $2 \times 2$ games with unique mixed strategy equilibria." Journal of Economic Behavior and Organization 42 (4): 523-548. 10.1016/S0167-2681(00)00102-5.
Morris, S., and P. Strack. 2019. "The Wald Problem and the Equivalence of Sequential Sampling and Static Information Costs." Working Paper. 10.2139/ssrn. 2991567.
Moscarini, G., and L. Smith. 1963. "The Optimal Level of Experimentation." Econometrica 69 (6): 1629-1644. 10.1111/1468-0262.00259.

Mosteller, F., and P. Nogee. 1951. "An Experimental Measurement of Utility." Journal of Political Economy 59 (5): 371-404.
Myers, C. E., A. Interian, and A. A. Moustafa. 2022. "A practical introduction to using the drift diffusion model of decision-making in cognitive psychology, neuroscience, and health sciences." Frontiers in Psychology 13 1039172. 10.3389/fpsyg.2022.1039172.
Nagel, R. 1995. "Unraveling in Guessing Games: An Experimental Study." American Economic Review 85 (5): 1313-1326.
Nash, J. 1950. Non-Cooperative Games. Ph.D. dissertation, Princeton University.
Ochs, J. 1995. "Games with Unique, Mixed Strategy Equilibria: An Experimental Study." Games and Economic Behavior 10 (1): 202-217. 10.1006/game.1995.1030.
Osborne, M., and A. Rubinstein. 1998. "Games with Procedurally Rational Players." American Economic Review 88 (4): 834-847.
Osborne, M., and A. Rubinstein. 2003. "Sampling equilibrium, with an application to strategic voting." Games and Economic Behavior 45 (2): 434-441. 10.1016/s0899-8256(03)00147-7.
Pearce, D. 1984. "Rationalizable Strategic Behavior and the Problem of Perfection." Econometrica 52 (4): 1029-1050. 10.2307/1911197.
Ratcliff, R. 1978. "A theory of memory retrieval." Psychological Review 85 (2): 59-108. 10.1037/ 0033-295X.85.2.59.
Ratcliff, R., P. Smith, S. Brown, and G. McKoon. 2016. "Diffusion Decision Model: Current Issues and History." Trends in Cognitive Sciences 20 (4): 260-281. 10.1016/j.tics.2016.01.007.

Recalde, M. P., A. Riedl, and L. Vesterlund. 2018. "Error-prone inference from response time: The case of intuitive generosity in public-good games." Journal of Public Economics 160 132-147. 10.1016/j.jpubeco.2018.02.010.

Rubinstein, A. 2016. "A Typology of Players: Between Instinctive and Contemplative." Quarterly fournal of Economics 117 (523): 859-890. 10.1093/qje/qjw008.
Rudin, W. 1973. Functional Analysis. McGraw-Hill.
Salant, Y., and J. Cherry. 2020. "Statistical Inference in Games." Econometrica 88 (4): 1725-1752. 10.3982/ECTA17105.

Schotter, A., and I. Trevino. 2021. "Is response time predictive of choice? An experimental study of threshold strategies." Experimental Economics 24 87-117. 10.1007/s10683-020-09651-1. Shadlen, M., and D. Shohamy. 2016. "Decision Making and Sequential Sampling from Memory." Neuron 90 (5): 927-939. 10.1016/j.neuron.2016.04.036.
Shapley, L. S. 1964. Some Topics in Two-Person Games. 1-28, Princeton University Press, . 10. 1515/9781400882014-002.
Sims, C. 2003. "Implications of rational inattention." Journal of Monetary Economics 50 (3): 665690. 10.1016/S0304-3932(03)00029-1.

Smith, S. M., and I. Krajbich. 2021. "Mental representations distinguish value-based decisions from perceptual decisions." Psychonomic Bulletin \& Review 28 1413-1422. 10.3758/ s13423-021-01911-2.
Spiliopoulos, L., and A. Ortmann. 2018. "The BCD of response time analysis in experimental economics." Experimental Economics 21 383-433. 10.1007/s10683-017-9528-1.
Stahl, D. O., and P. W. Wilson. 1995. "On Players' Models of Other Players: Theory and Experimental Evidence." Games and Economic Behavior 10 (1): 218-254. 10.1006/game.1995.1031.
Steiner, J., C. Stewart, and F. Matějka. 2017. "Rational Inattention Dynamics: Inertia and Delay in Decision-Making." Econometrica 85 (2): 521-553. 10.3982/ECTA13636.
Wald, A. 1947. "Foundations of a General Theory of Sequential Decision Functions." Econometrica 15 (4): 279-313. 10.2307/1905331.
Yang, M. 2015. "Coordination with flexible information acquisition." Journal of Economic Theory 158 (Part B): 721-738. 10.1016/j.jet.2014.11.017.
Zhong, W. 2022. "Optimal Dynamic Information Acquisition." Econometrica 90 (4): 1537-1582. 10.3982/ECTA17787.

## Appendix A. Omitted Proofs

Proof of Proposition 1. For ease of notation, I'll write $Y_{i}:=A_{-i}, P_{i}:=\Delta\left(Y_{i}\right)$, and $\mu_{i} \in \Delta\left(P_{i}\right)$. We consider general continuous utility functions $u_{i}: A_{i} \times P_{i} \rightarrow \mathbb{R}$. Boundedness of $V_{i}$ and $v_{i}$ follows immediately from boundedness of $u_{i}$. Below we prove the remaining properties.

Claim 1. For any $\mu_{i} \in \Delta\left(P_{i}\right)$, and any optimal stopping time $\tau_{i}, \mathbb{P}_{\mu_{i}}\left(\tau_{i}>T\right) \leq 2\left\|u_{i}\right\|_{\infty} / c_{i} T$.
Proof. Let $\left\|u_{i}\right\|_{\infty}=\max _{\left(a_{i}, p_{i}\right) \in A_{i} \times P_{i}}\left|u\left(a_{i}, p_{i}\right)\right|<\infty$. We then have that, for any $\mu_{i} \in \Delta\left(P_{i}\right)$ and
$T \in \mathbb{N}$,

$$
-\left\|u_{i}\right\|_{\infty} \leq V_{i}\left(\mu_{i}\right) \leq \mathbb{P}_{\mu_{i}}\left(\tau_{i} \leq T\right)\left\|u_{i}\right\|_{\infty}+\mathbb{P}_{\mu_{i}}\left(\tau_{i}>T\right)\left(\left\|u_{i}\right\|_{\infty}-c_{i} T\right) \Longrightarrow \mathbb{P}_{\mu_{i}}\left(\tau_{i}>T\right) \leq 2\left\|u_{i}\right\|_{\infty} / c_{i} T .
$$

Claim 2. $v_{i}$ is uniformly continuous.
Proof. $\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, p_{i}\right)\right]$ is jointly continuous in $\left(a_{i}, \mu_{i}\right)$ with respect to product topology. Let $\mu_{i, n} \rightarrow$ $\mu_{i}, a_{i, n} \rightarrow a_{i}$. Note that $\forall \epsilon>0, \exists N$ such that $\forall n \geq N, \forall p_{i},\left|u_{i}\left(a_{i, n}, p_{i}\right)-u_{i}\left(a_{i}, p_{i}\right)\right|<\epsilon / 2$ and $\left|\mathbb{E}_{\mu_{i, n}}\left[u_{i}\left(a_{i}, p_{i}\right)\right]-\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, p_{i}\right)\right]\right|<\epsilon / 2$. Hence,
$\left|\mathbb{E}_{\mu_{i, n}}\left[u_{i}\left(a_{i, n}, p_{i}\right)\right]-\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, p_{i}\right)\right]\right| \leq\left|\mathbb{E}_{\mu_{i, n}}\left[u_{i}\left(a_{i, n}, p\right)-u_{i}\left(a_{i}, p_{i}\right)\right]\right|+\left|\mathbb{E}_{\mu_{i, n}}\left[u_{i}\left(a_{i}, p_{i}\right)\right]-\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, p_{i}\right)\right]\right|<\epsilon$.
Continuity of $v_{i}$ follows from Berge's maximum theorem and uniform continuity from HeineCantor theorem.

Claim 3. $V_{i}$ is uniformly continuous.
Proof. Let $\mathbb{T}_{i, T}$ denote the set of stopping times $\tau^{\prime} \in \mathbb{T}_{i}$ that are bounded above by $T$ and, for every $T \in \mathbb{N}, V_{i, T}: \Delta\left(P_{i}\right) \rightarrow \mathbb{R}$ be given by

$$
V_{i, T}(\mu):=\sup _{\tau^{\prime} \in \mathbb{T}_{i, T}} \mathbb{E}_{\mu_{i}}\left[v_{i}\left(\mu_{i} \mid y_{i}^{\tau^{\prime}}\right)-c_{i} \cdot \tau^{\prime}\right]
$$

Note that, as $\mathbb{T}_{i, T}$ is finite, it is compact with respect to the discrete topology, and an application of Berge's maximum theorem implies $V_{i, T}$ is continuous.

Note that, for any $\mu_{i} \in \Delta\left(P_{i}\right), T \in \mathbb{N}, 0 \leq V_{i}\left(\mu_{i}\right)-V_{i, T}\left(\mu_{i}\right) \leq \mathbb{P}_{\mu_{i}}\left(\tau_{i}>T\right)\left\|u_{i}\right\|_{\infty} \leq 2\left\|u_{i}\right\|_{\infty}^{2} / c_{i} T$. Hence, $\left\|V_{i}-V_{i, T}\right\| \leq 2\left\|u_{i}\right\|_{\infty}^{2} / c_{i} T$, and $V_{i, T}$ converges uniformly to $V_{i}$. Since for any $T, V_{i, T}$ is in the space of bounded continuous functions $C_{b}^{0}\left(\Delta\left(P_{i}\right)\right)$, which, endowed with the sup-norm is a Banach space, $V_{i}$ is continuous; by the Heine-Cantor theorem, it is uniformly continuous.
Claim 4. $v_{i}, V_{i}$, and $V_{i, T}$ are convex, for any $T \in \mathbb{N}$.
Proof. This follows since each of these can be seen as the pointwise supremum over a family of convex functions over $\Delta\left(P_{i}\right)$, which is compact with respect to $\|\cdot\|_{L P}$.

Proof of Lemma 1. By contrapositive, that (2) implies (1) is straightforward. We prove a more general claim that implies the converse. Let $\Sigma_{-i}^{\tau_{i}}:=\left\{\sigma_{-i} \in \Delta\left(A_{-i}\right) \mid \mathbb{P}_{\sigma_{-i}}\left(\tau_{i}<\infty\right)=1\right\}$ denote the opponents' distribution actions with respect to which player $i$ 's optimal stopping time is finite with probability 1.

Lemma 6. On $\Sigma_{-i}^{\tau_{i}}, b_{i}$ is a continuous mapping to $\Delta\left(A_{i}\right)$.

Proof. Fix a selection of optimal choices $\sigma_{i}^{*}\left(\mu_{i}^{\prime}\right) \in \arg \max _{\sigma_{i} \in \Delta\left(A_{i}\right)} \mathbb{E}_{\mu_{i}^{\prime}}\left[u_{i}\left(\sigma_{i}, \sigma_{-i}^{\prime}\right)\right]$. For $t \in \mathbb{N}$, let $b_{i, t}: \Sigma_{-i}^{\tau_{i}} \rightarrow[0,1]^{\left|A_{i}\right|}$ be given by $b_{i, t}\left(\sigma_{-i}\right):=\mathbb{E}_{\sigma_{-i}}\left[1_{\tau_{i}<t} \sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right]$, and $p_{i, t}: \Sigma_{-i}^{\tau_{i}} \rightarrow[0,1]$ denote $p_{i, t}\left(\sigma_{-i}\right):=\mathbb{P}_{\sigma_{-i}}\left(\tau_{i}<t\right)$. Note that, for every $t^{\prime} \geq t, b_{i, t} \leq b_{i, t^{\prime}}$ and $p_{i, t} \leq p_{i, t^{\prime}}$, ensuring pointwise convergence to $b_{i}$ and 1 , respectively. Further, both $b_{i, t}$ and $p_{i, t}$ are continuous.

Since $p_{i, t} \rightarrow 1$, by Dino's theorem $p_{i, t}$ converges uniformly on $\Sigma_{-i}^{\tau_{i}}$. We then have $\| b_{i}\left(\sigma_{-i}\right)-$ $b_{i, t}\left(\sigma_{-i}\right) \|_{1} \leq 1-i, t\left(\sigma_{-i}\right)$, and $b_{i, t}$ also converges uniformly to $\left.b_{i}\right|_{\Sigma_{-i}^{\tau_{i}}}: \Sigma_{-i}^{\tau_{i}} \rightarrow[0,1]^{\left|A_{i}\right|}$, ensuring that, on $\Sigma_{-i}^{\tau_{i}}, b_{i}$ is continuous. To see that $b_{i}\left(\sigma_{-i}\right) \in \Delta\left(A_{i}\right)$ for any $\sigma_{-i} \in \Sigma_{-i}^{\tau_{i}}$, note that $1 \geq$ $\left\|b_{i}\left(\sigma_{-i}\right)\right\|_{1} \geq\left\|b_{i, t}\left(\sigma_{-i}\right)\right\|_{1} \geq p_{i, t}\left(\sigma_{-i}\right) \rightarrow 1$.

Proof of Remark 1. Let $\sum_{\ell \in[1 . . t]} \delta_{y_{i, \ell}^{t}} / t=:{\overline{y_{i}}}^{t} \in \Delta\left(A_{-i}\right)$ denote the empirical frequency. When $\mu_{i}$ has full support on $\Delta\left(A_{-i}\right)$, then there is $\phi: \mathbb{R}_{++} \rightarrow \mathbb{R}_{++}$such that $\inf _{\sigma_{-i} \in \Delta\left(A_{-i}\right)} \mu\left(B_{\epsilon}\left(\sigma_{-i}\right)\right) \geq$ $\phi(\epsilon)$. Then, by Diaconis and Freedman (1990), for every $\epsilon>0$, there is $T$ such that for all $t \geq T$, $\mu_{i, t}\left(B_{\epsilon / 2}\left({\overline{y_{i}}}^{t}\right)\right) /\left(1-\mu_{i, t}\left(B_{\epsilon / 2}\left({\overline{y_{i}}}^{t}\right)\right) \geq \frac{2-\epsilon}{\epsilon}\right.$, which implies $\left\|\mu_{i, t}-\delta_{\bar{y}_{i}}\right\|_{L P} \leq \epsilon$. Immediately, for every $\epsilon>0$, there is $T$ such that for all $t \geq T$,

$$
\left\|\mu_{i, t}-\mu_{i, t+1}\right\|_{L P} \leq\left\|\mu_{i, t}-\delta_{{\overline{y_{i}}}^{t}}\right\|_{L P}+\left\|\mu_{i, t+1}-\delta_{{\overline{y_{i}}}^{t+1}}\right\|_{L P}+\left\|\delta_{{\overline{y_{i}}}^{t}}-\delta_{{\overline{y_{i}}}^{t+1}}\right\|_{L P} \leq 2 / 3 \epsilon+2 / t \leq \epsilon .
$$

As $V_{i}$ is uniformly continuous (cf. Proposition 1), this implies $\exists \bar{T}_{i}$ such that $\forall t \geq T, \forall y_{i}^{t+1}$, $V_{i}\left(\mu_{i} \mid y_{i}^{t+1}\right)-V_{i}\left(\mu_{i} \mid y_{i}^{t}\right) \leq c_{i} \Longrightarrow \mathbb{E}_{\mu_{i}}\left[V_{i}\left(\mu_{i} \mid y_{i}^{t+1}\right)-V_{i}\left(\mu_{i} \mid y_{i}^{t}\right) \mid y_{i}^{t}\right] \leq c_{i} \Longrightarrow \tau_{i} \leq \bar{T}_{i}$.

When the prior does not allow for correlation, $\mu_{i}=\times_{j \neq i} \mu_{i j}$ and each marginal $\mu_{i j}$ uniformly accumulates around the empirical frequency projected on $\Delta\left(A_{j}\right)$. As $I$ is finite, one can similarly obtain a uniform rate of convergence that depends only on $t$. It is then straightforward to adjust the proof to obtain the result.

Proof of Proposition 6. Let $A_{i}=\{0,1\}, i \in I$; denote the probability that player $i$ chooses action 1 by $\sigma_{i}$. By manner of a continuous-time approximation as in Fudenberg and Levine (1998, Ch. 2), the dynamic system can be written as $\dot{\sigma}_{i}=b_{i}\left(\sigma_{j}\right)-\sigma_{i}, i, j \in I, i \neq j$. The Jacobian of the dynamic system is given by

$$
\left(\begin{array}{cc}
-1 & b_{i}^{\prime}\left(\sigma_{j}\right) \\
b_{j}^{\prime}\left(\sigma_{i}\right) & -1
\end{array}\right)
$$

and its eigenvalues are given by $\lambda=-1 \pm \sqrt{b_{i}^{\prime}\left(\sigma_{j}\right) b_{j}^{\prime}\left(\sigma_{i}\right)}$, where differentiability of $b_{i}, b_{j}$ is ensured by Theorem 3. If $b_{i}^{\prime}\left(\sigma_{j}\right) b_{j}^{\prime}\left(\sigma_{i}\right) \leq 0$, there is a unique $\sigma$ such that $b_{i}\left(\sigma_{j}\right)=\sigma_{i}$ and, since the real parts of the eigenvalues of the Jacobian matrix are strictly negative, by the Jacobian conjec-
ture on global asymptotic stability-proved to hold on the plane (Chen et al., 2001) $-\sigma$ is globally asymptotically stable.

In particular, if there is a unique Nash equilibrium, either one player has a dominant strategy (and then $b_{i}^{\prime}=0$ for some player $i$ ) or neither does. If some player does not sample, i.e. $\tau_{i}=0$ for some player $i$, then again $b_{i}^{\prime}=0$. If both players sample, we must then have $b_{i}^{\prime}\left(\sigma_{j}\right) b_{j}^{\prime}\left(\sigma_{i}\right)<0$. In any of these cases, $b_{i}^{\prime}\left(\sigma_{j}\right) b_{j}^{\prime}\left(\sigma_{i}\right) \leq 0$.

Proof of Remark 2. Let $c_{i}^{\prime} \geq c_{i}$ and denote $V_{i}^{\prime}$ and $V_{i}$ the value functions associated with $c_{i}^{\prime}$ and $c_{i}$, respectively. Since $v_{i}\left(\mu_{i}\right) \leq V_{i}^{\prime}\left(\mu_{i}\right) \leq V_{i}\left(\mu_{i}\right)$, it is immediate that $V_{i}\left(\mu_{i}\right)=v_{i}\left(\mu_{i}\right) \Longrightarrow V_{i}^{\prime}\left(\mu_{i}\right)=$ $v_{i}\left(\mu_{i}\right)$. Let $\tau_{i}^{\prime}$ and $\tau_{i}$ the earliest optimal stopping times associated with $c_{i}^{\prime}$ and $c_{i}$, respectively. Then

$$
\begin{aligned}
\left\{\omega \in \Omega \mid \tau_{i}(\omega) \leq t\right\} & =\left\{\omega \in \Omega \mid \exists t^{\prime} \leq t: \begin{array}{l}
V_{i}\left(\mu_{i} \mid y_{i}^{t^{\prime}}(\omega)\right)=v_{i}\left(\mu_{i} \mid y_{i}^{t^{\prime}}(\omega)\right) \\
\\
V_{i}\left(\mu_{i} \mid y_{i}^{\ell}(\omega)\right)>v_{i}\left(\mu_{i} \mid y_{i}^{\ell}(\omega)\right), \forall \ell<t^{\prime}
\end{array}\right\} \\
& \subseteq\left\{\omega \in \Omega \mid \exists t^{\prime} \leq t: \begin{array}{l}
V_{i}^{\prime}\left(\mu_{i} \mid y_{i}^{t^{\prime}}(\omega)\right)=v_{i}\left(\mu_{i} \mid y_{i}^{t^{\prime}}(\omega)\right) \\
V_{i}^{\prime}\left(\mu_{i} \mid y_{i}^{\ell}(\omega)\right)>v_{i}\left(\mu_{i} \mid y_{i}^{\ell}(\omega)\right), \forall \ell<t^{\prime}
\end{array}\right\}=\left\{\omega \in \Omega: \tau_{i}^{\prime}(\omega) \leq t\right\} .
\end{aligned}
$$

Proof of Lemma 2. Take any $a_{i} \in A_{i} \backslash A_{i}^{k}$. For $\delta \geq 0$ define

$$
\begin{gathered}
\bar{B}_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right):=\left\{\sigma_{-i} \in \Delta\left(A_{-i}\right) \mid \exists \sigma_{-i}^{\prime} \in \Delta\left(A_{-i}^{k-1}\right):\left\|\sigma_{-i}-\sigma_{-i}^{\prime}\right\| \leq \delta\right\} \\
h_{i}^{k}(\delta):=\max _{\sigma_{-i} \in \bar{B}_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right)}\left(u_{i}\left(a_{i}, \sigma_{-i}\right)-\max _{\sigma_{i}^{\prime} \in \Delta\left(A_{i}\right)} u_{i}\left(\sigma_{i}^{\prime}, \sigma_{-i}\right)\right) .
\end{gathered}
$$

Since $u_{i}$ is continuous and $\delta \mapsto \bar{B}_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right)$ is a continuous, convex-, compact-, and nonemptyvalued correspondence, $h_{i}^{k}$ is continuous by Berge's maximum theorem. As $\bar{B}_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right)$ increases in subset order with $\delta, h_{i}^{k}$ is nondecreasing.
By definition of $k$-rationalizability, $h_{i}^{k}(0)<0$. This implies that there is $\delta^{\prime}>0$ such that, $\forall \delta \leq \delta^{\prime}$, $h_{i}^{k}(\delta) \leq h_{i}^{k}(0) / 2<0$. Then, for any $\delta<\delta^{\prime}$, and any $\sigma_{-i} \in B_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right), u_{i}\left(a_{i}, \sigma_{-i}\right)-\max _{\sigma_{i}^{\prime} \in \Delta\left(A_{i}\right)} u_{i}\left(\sigma_{i}^{\prime}, \sigma_{-i}\right) \leq$ $h_{i}^{k}(0) / 2<0$. Finally, observe that $\max _{\sigma_{-i} \in \Delta\left(A_{-i}\right) \backslash \Delta\left(A_{-i}^{k-1}\right)} \max _{\sigma_{i} \in \Delta\left(A_{i}\right)} u_{i}\left(a_{i}, \sigma_{-i}\right)-u_{i}\left(\sigma_{i}, \sigma_{-i}\right)<$ $2\left\|u_{i}\right\|_{\infty}$. Let $\epsilon<-h_{i}^{k}(0) /\left(4\left\|u_{i}\right\|_{\infty}-h_{i}^{k}(0)\right)$ and $\mu_{i}\left(B_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right)\right)>1-\epsilon$. It then follows that

$$
\begin{aligned}
\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, \sigma_{-i}\right)\right]-\max _{\sigma_{i} \in \Delta\left(A_{i}\right)} \mathbb{E}_{\mu_{i}}\left[u_{i}\left(\sigma_{i}, \sigma_{-i}\right)\right] & \leq \mu_{i}\left(B_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right)\right) h_{i}^{k}(0) / 2+\left(1-\mu_{i}\left(B_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right)\right)\right) 2\left\|u_{i}\right\|_{\infty} \\
& <(1-\epsilon) h_{i}^{k}(0) / 2+\epsilon 2\left\|u_{i}\right\|_{\infty}<0 .
\end{aligned}
$$

Proof of Lemma 3. By assumption, $\bar{y}_{i}^{t} \in \Delta\left(A_{-i}^{k-1}\right)$. Since every player's prior has full support, by Diaconis and Freedman (1990), each player's prior concentrates on an $\delta$-ball around the empirical
frequency of at a uniform rate. This implies that, for any $\delta, \epsilon>0$ there is a $T$ such that, for all $t \geq T$ and any $y_{i}^{t} \in A_{-i}^{k-1}, \mu_{i, t}\left(B_{\delta}\left(\Delta\left(A_{-i}^{k-1}\right)\right)\right) \geq \mu_{i, t}\left(B_{\delta}\left(\delta_{\bar{y}_{i}^{t}}\right)\right)>1-\epsilon$.

## Auxiliary Results.

Lemma 7. For any $\epsilon>0, \exists \delta>0$ such that for any $T \in \mathbb{N}, c_{i}>0, \mathbb{E}_{\mu_{i}}\left[v_{i}\left(\delta_{\sigma_{-i}}\right)\right]-V_{i}\left(\mu_{i}\right) \leq(1-$ $\left.2 \exp \left(-2 T \delta^{2}\right)\right) \epsilon / 4+8 \exp \left(-2 T \delta^{2}\right)\left\|u_{i}\right\|_{\infty}+c_{i} T$.
Proof. Let $\hat{T}_{i}$ be the stopping time such that player $i$ stops after $T$ periods and let $\bar{y}_{i}^{T}$ denote the empirical frequency of $y_{i}^{T}$, i.e. $\bar{y}_{i}^{T}:=\sum_{t \in[1 . . T]} \delta_{y_{i, t}} / T$. Let $\hat{\alpha}_{i}: \Delta\left(A_{-i}\right) \rightarrow A_{i}$ be such that $\hat{a}_{i}\left(\mu_{i}^{\prime}\right) \in$ $\arg \max \mathbb{E}_{\mu_{i}^{\prime}}\left[u_{i}\left(a_{i}, \sigma_{-i}\right)\right]$. Since $u_{i}\left(a_{i}\left(\sigma_{-i}\right), \sigma_{-i}\right)=\max _{a_{i} \in A_{i}} u_{i}\left(a_{i}, \sigma_{-i}\right)$ is continuous by Berge's maximum theorem and uniformly so by Heine-Cantor theorem, let $\delta>0$ be such that for any $\left\|\sigma_{-i}-\sigma_{-i}^{\prime}\right\|<\delta \Longrightarrow\left|u_{i}\left(a_{i}\left(\sigma_{-i}\right), \sigma_{-i}\right)-u_{i}\left(a_{i}\left(\sigma_{-i}^{\prime}\right), \sigma_{-i}^{\prime}\right)\right|+\left|u_{i}\left(a_{i}\left(\sigma_{-i}^{\prime}\right), \sigma_{-i}\right)-u_{i}\left(a_{i}\left(\sigma_{-i}^{\prime}\right), \sigma_{-i}^{\prime}\right)\right|<\epsilon / 4$. Since $\hat{T}_{i}$ and $\hat{a}_{i}\left(\bar{y}_{i}^{\hat{T}_{i}}\right)$ are potentially suboptimal stopping time and choices for player $i, V_{i}\left(\mu_{i}\right) \geq$ $\mathbb{E}_{\mu_{i}}\left[u_{i}\left(\hat{\alpha}_{i}\left(\bar{y}_{i}^{\hat{T}_{i}}\right), \sigma_{-i}\right)-c_{i} \hat{T}_{i}\right]$. For $c_{i}=k /(4 T)>0$,
$\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\delta_{\sigma_{-i}}\right)\right]-V_{i}\left(\mu_{i}\right) \leq \mathbb{E}_{\mu_{i}}\left[u_{i}\left(\hat{a}_{i}\left(\sigma_{-i}\right), \sigma_{-i}\right)-u_{i}\left(\hat{a}_{i}\left(\bar{y}^{\hat{T}_{i}}\right), \sigma_{-i}\right)+c_{i} \hat{T}_{i}\right]$
$=\mathbb{E}_{\mu_{i}}\left[u_{i}\left(\hat{a}_{i}\left(\sigma_{-i}\right), \sigma_{-i}\right)-u_{i}\left(\hat{a}_{i}\left(\bar{y}^{\hat{T}_{i}}\right), \bar{y}_{i}^{\hat{T}_{i}}\right)+u_{i}\left(\hat{a}_{i}\left(\bar{y}_{i}^{\hat{T}_{i}}\right), \bar{y}_{i}^{\hat{T}_{i}}\right)-u_{i}\left(\hat{a}_{i}\left(\bar{y}_{i}^{\hat{T}_{i}}\right), \sigma_{-i}\right)\right]+c_{i} T$ $\leq\left(1-2 \exp \left(-2 T \delta^{2}\right)\right) \epsilon / 4+8 \exp \left(-2 T \delta^{2}\right)\left\|u_{i}\right\|_{\infty}+c_{i} T$,
where the last inequality follows by the Dvoretzky-Kiefer-Wolfowitz-Massart inequality (Massart, 1990), which delivers that, for any $\delta>0$ and $\sigma_{-i}, \mathbb{P}_{\sigma_{-i}}\left(\left\|\sigma_{-i}-\bar{y}_{i}^{T}\right\|>\delta\right)<2 \exp \left(-2 T \delta^{2}\right)$.

Proof of Lemma 4. We first prove the weaker statement:
Claim 5. Suppose that there is no action $a_{i}$ that is a best response to all distribution of opponents' actions $\sigma_{-i} \in \Delta\left(A_{-i}\right)$. Then, for any full support prior $\mu_{i} \in \Delta\left(A_{-i}\right)$, there is a sampling cost $c_{i}>0$ such that the associated earliest optimal stopping time $\tau_{i} \geq 1$.

Proof. Let us start by showing that $\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\delta_{\sigma_{-i}}\right)\right]-v_{i}\left(\mu_{i}\right)>0$. Since there is no action $a_{i}$ that is a best response to all distribution of opponents' actions $\sigma_{-i} \in \Delta\left(A_{-i}\right)$, for any $\mu_{i}$ with full support and any $a_{i}\left(\mu_{i}\right) \in \arg \max _{a_{i} \in A_{i}} \mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, \sigma_{-i}\right)\right]$, there is $\sigma_{-i}^{\prime} \in \Delta\left(A_{-i}\right)$ such that $v_{i}\left(\delta_{\sigma_{-i}^{\prime}}\right)=$ $\max _{a_{i} \in A_{i}} u_{i}\left(\alpha_{i}, \sigma_{-i}^{\prime}\right)>u_{i}\left(a_{i}\left(\mu_{i}\right), \sigma_{-i}^{\prime}\right)$. By continuity, there is an $\epsilon>0$ such that for all $\sigma_{-i}^{\prime \prime} \in$ $B_{\epsilon}\left(\sigma_{-i}^{\prime}\right), v_{i}\left(\delta_{\sigma_{-i}^{\prime \prime}}\right)-u_{i}\left(a_{i}\left(\mu_{i}\right), \sigma_{-i}^{\prime \prime}\right) \geq\left(v_{i}\left(\delta_{\sigma_{-i}}^{\prime}\right)-u_{i}\left(a_{i}\left(\mu_{i}\right), \sigma_{-i}^{\prime}\right)\right) / 2>0$. Hence, $\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\delta_{\sigma_{-i}}\right)\right]-v_{i}\left(\mu_{i}\right) \geq$ $\mu_{i}\left(B_{\epsilon}\left(\delta_{\sigma_{-i}^{\prime}}\right)\right)\left(v_{i}\left(\delta_{\sigma_{-i}}^{\prime}\right)-u_{i}\left(a_{i}\left(\mu_{i}\right), \sigma_{-i}^{\prime}\right)\right) / 2>0$.

Next, we show that, if $c_{i}$ is low enough, $V_{i}\left(\mu_{i}\right)>v_{i}\left(\mu_{i}\right)$ (implying $\tau_{i} \geq 1$ ) by proving that $\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\delta_{\sigma_{-i}}\right)\right]-$ $V_{i}\left(\mu_{i}\right)<\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\delta_{\sigma_{-i}}\right)\right]-v_{i}\left(\mu_{i}\right)=: k$. By Lemma 7, for any $\epsilon>0, \exists \delta>0$ such that for any $T \in \mathbb{N}$,
$c_{i}>0, \mathbb{E}_{\mu_{i}}\left[v_{i}\left(\delta_{\sigma_{-i}}\right)\right]-V_{i}\left(\mu_{i}\right) \leq\left(1-2 \exp \left(-2 T \delta^{2}\right)\right) \epsilon / 4+8 \exp \left(-2 T \delta^{2}\right)\left\|u_{i}\right\|_{\infty}+c_{i} T$. Letting $\epsilon=k$, $c_{i}=k /(4 T)$, we have $\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\delta_{\sigma_{-i}}\right)\right]-V_{i}\left(\mu_{i}\right) \leq k / 2+8 \exp \left(-2 T \delta^{2}\right)\left\|u_{i}\right\|_{\infty}$. It is then straightforward to see that, for $T$ large enough (and $c_{i}>0$ small enough), $\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\delta_{\sigma_{-i}}\right)\right]-V_{i}\left(\mu_{i}\right)<k=$ $\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\delta_{\sigma_{-i}}\right)\right]-v_{i}\left(\mu_{i}\right)$, proving the result.

Now consider all possible posteriors following any possible $t$ realizations, $\left\{\mu_{i} \mid y_{i}^{t}, y_{i}^{t} \in \cup_{t \in[1 . . T]} Y^{t}\right\}$, observing that this is a finite set. Given the nature of the information process, since the prior $\mu_{i}$ has full support, so do the posterior beliefs. By the above claim, for each $\mu_{i} \mid y_{i}^{t}$, there is a cost $c_{i}>0$ such that player $i$ would find it optimal to acquire at least one more signal. Taking the lowest of all such costs implies that under such cost, player $i$ would deem it optimal to acquire at least $T$ signals, concluding the proof.

Proof of Theorem 3. Label player $i$ 's actions so that $u_{i}\left(1, \sigma_{-i}\right)-u_{i}\left(0, \sigma_{-i}\right)$ is increasing in $\sigma_{-i}$.
Proof of Theorem 3(i). We first prove a more general comparative statics result:
Proposition 7. $\mathbb{P}_{\sigma_{-i}}\left(a_{i} \in A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right.$ and $\left.\tau_{i} \leq t\right)$ is increasing with respect to $\geq_{a_{i}}$, and $\mathbb{P}_{\sigma_{-i}}\left(a_{i} \notin\right.$ $A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)$ and $\left.\tau_{i} \leq t\right)$ is decreasing with respect to $\geq_{a_{i}}$, for any $t \in \mathbb{N}$ and $\sigma_{-i} \in \Delta\left(A_{-i}\right)$.

Proof. Let $\tilde{u}_{i} \geq_{a_{i}} u_{i}$ and denote the respective (i) value functions, (ii) earliest optimal stopping times, (iii) optimal choices at given beliefs, and (iv) selections of optimal choices by (i) $\tilde{V}_{i}$ and $V_{i}$, (ii) $\tilde{\tau}_{i}$ and $\tau_{i}$, (iii) $\tilde{A}_{i}^{*}$ and $A_{i}^{*}$, and (iv) $\tilde{\sigma}_{i}^{*}$ and $\sigma_{i}^{*}$, respectively. With $g:=\tilde{u}_{i}\left(a_{i}, \cdot\right)-u_{i}\left(\alpha_{i}, \cdot\right)$, by definition we obtain

$$
V_{i}\left(\mu_{i}\right)+\mathbb{E}_{\mu_{i}}\left[\tilde{\sigma}_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tilde{\tau}_{i}}\right)\left(a_{i}\right) g\left(\sigma_{-i}^{\prime}\right)\right] \geq \tilde{V}_{i}\left(\mu_{i}\right) \geq V_{i}\left(\mu_{i}\right)+\mathbb{E}_{\mu_{i}}\left[\sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\left(a_{i}\right) g\left(\sigma_{-i}^{\prime}\right)\right] \geq V_{i}\left(\mu_{i}\right)
$$

Lemma 8. $V_{i}\left(\mu_{i}\right)=\mathbb{E}_{\mu_{i}}\left[u_{i}\left(\alpha_{i}, \sigma_{-i}^{\prime}\right)\right] \Longrightarrow \tilde{V}_{i}\left(\mu_{i}\right)=\mathbb{E}_{\mu_{i}}\left[\tilde{u}_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right)\right]$ and, for $\alpha_{i}^{\prime} \neq a_{i} \tilde{V}_{i}\left(\mu_{i}\right)=\mathbb{E}_{\mu_{i}}\left[\tilde{u}_{i}\left(a_{i}^{\prime}, \sigma_{-i}^{\prime}\right)\right] \Longrightarrow$ $V_{i}\left(\mu_{i}\right)=\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}^{\prime}, \sigma_{-i}^{\prime}\right)\right]$.
Proof. Since $\mathbb{E}_{\mu_{i}}\left[\tilde{\sigma}_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tilde{\tau}_{i}}\right)\left(a_{i}\right) g\left(\sigma_{-i}^{\prime}\right)\right] \leq \mathbb{E}_{\mu_{i}}\left[g\left(\sigma_{-i}^{\prime}\right)\right]$, if $V_{i}\left(\mu_{i}\right)=\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right)\right]$, then $\mathbb{E}_{\mu_{i}}\left[\tilde{u}_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right)\right]=$ $\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right)\right]+\mathbb{E}_{\mu_{i}}\left[g\left(\sigma_{-i}^{\prime}\right)\right] \geq V_{i}\left(\mu_{i}\right)+\mathbb{E}_{\mu_{i}}\left[\tilde{\sigma}_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tilde{\tau}_{i}}\right)\left(\alpha_{i}\right) g\left(\sigma_{-i}^{\prime}\right)\right] \geq \tilde{V}_{i}\left(\mu_{i}\right) \geq \mathbb{E}_{\mu_{i}}\left[\tilde{u}_{i}\left(\alpha_{i}, \sigma_{-i}^{\prime}\right)\right]$. Moreover, if $\tilde{V}_{i}\left(\mu_{i}\right)=\mathbb{E}_{\mu_{i}}\left[\tilde{u}_{i}\left(a_{i}^{\prime}, \sigma_{-i}^{\prime}\right)\right]=\mathbb{E}_{\mu_{i}}\left[u_{i}\left(\alpha_{i}^{\prime}, \sigma_{-i}^{\prime}\right)\right]$, then $\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}^{\prime}, \sigma_{-i}^{\prime}\right)\right]=\tilde{V}_{i}\left(\mu_{i}\right) \geq V_{i}\left(\mu_{i}\right) \geq \mathbb{E}_{\mu_{i}}\left[u_{i}\left(\alpha_{i}^{\prime}, \sigma_{-i}^{\prime}\right)\right]$.

Note that, by the contrapositive of Lemma 8, $\tilde{V}_{i}\left(\mu_{i}\right)>\mathbb{E}_{\mu_{i}}\left[\tilde{u}_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right)\right] \Longrightarrow V_{i}\left(\mu_{i}\right)>\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right)\right]$ and $V_{i}\left(\mu_{i}\right)>\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}^{\prime}, \sigma_{-i}^{\prime}\right)\right] \Longrightarrow \tilde{V}_{i}\left(\mu_{i}\right)>\mathbb{E}_{\mu_{i}}\left[\tilde{u}_{i}\left(\alpha_{i}^{\prime}, \sigma_{-i}^{\prime}\right)\right]$ for $a_{i}^{\prime} \neq a_{i}$. This implies

$$
\left\{\omega \in \Omega \mid \tau_{i}(\omega) \leq t \text { and } a_{i} \in A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}(\omega)\right)\right\}
$$

$$
\begin{aligned}
& =\left\{\omega \in \Omega \mid \exists t^{\prime} \leq t: \begin{array}{l}
V_{i}\left(\mu_{i} \mid y_{i}^{t^{\prime}}(\omega)\right)=\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right) \mid y_{i}^{t^{\prime}}(\omega)\right] \\
V_{i}\left(\mu_{i} \mid y_{i}^{\ell}(\omega)\right)>\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}^{\prime}, \sigma_{-i}^{\prime}\right) \mid y_{i}^{t^{\prime}}(\omega)\right], \forall \ell<t^{\prime}, \forall a_{i}^{\prime}
\end{array}\right\} \\
& \subseteq\left\{\omega \in \Omega \mid \exists t^{\prime} \leq t: \begin{array}{l}
\tilde{V}_{i}\left(\mu_{i} \mid y_{i}^{t^{\prime}}(\omega)\right)=\mathbb{E}_{\mu_{i}}\left[\tilde{u}_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right) \left\lvert\, y_{i}^{\left.t^{\prime}(\omega)\right]} \begin{array}{r}
\tilde{V}_{i}\left(\mu_{i} \mid y_{i}^{\ell}(\omega)\right)>\mathbb{E}_{\mu_{i}}\left[u_{i}\left(\alpha_{i}^{\prime}, \sigma_{-i}^{\prime}\right) \mid y_{i}^{t^{\prime}}(\omega)\right], \forall \ell<t^{\prime}, \forall a_{i}^{\prime} \neq a_{i}
\end{array}\right.\right\} \\
=\left\{\omega \in \Omega \mid \tilde{\tau}_{i}(\omega) \leq t \text { and } a_{i} \in \tilde{A}_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tilde{\tau}_{i}}(\omega)\right)\right\}
\end{array}\right.
\end{aligned}
$$

and

$$
\begin{aligned}
& \left\{\omega \in \Omega \mid \tilde{\tau}_{i}(\omega) \leq t \text { and } a_{i} \notin \tilde{A}_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tilde{\tau}_{i}}(\omega)\right)\right\} \\
& =\left\{\omega \in \Omega \mid \exists t^{\prime} \leq t: \begin{array}{l}
\tilde{V}_{i}\left(\mu_{i} \mid y_{i}^{t^{\prime}}(\omega)\right)=\tilde{v}_{i}\left(\mu_{i} \mid y_{i}^{t^{\prime}}(\omega)\right)>\mathbb{E}_{\mu_{i}}\left[\tilde{u}_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right) \mid y_{i}^{t^{\prime}}(\omega)\right] \\
\tilde{V}_{i}\left(\mu_{i} \mid y_{i}^{\ell}(\omega)\right)>\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}^{\prime}, \sigma_{-i}^{\prime}\right) \mid y_{i}^{t^{\prime}}(\omega)\right], \forall \ell<t^{\prime}, \forall a_{i}^{\prime}
\end{array}\right\} \\
& \subseteq\left\{\omega \in \Omega \mid \exists t^{\prime} \leq t: \begin{array}{l}
V_{i}\left(\mu_{i} \mid y_{i}^{t^{\prime}}(\omega)\right)=v_{i}\left(\mu_{i} \mid y_{i}^{t^{\prime}}(\omega)\right)>\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right) \mid y_{i}^{t^{\prime}}(\omega)\right] \\
V_{i}\left(\mu_{i} \mid y_{i}^{\ell}(\omega)\right)>\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right) \mid y_{i}^{\left.t^{\prime}(\omega)\right], \forall \ell<t^{\prime}}\right.
\end{array}\right\} \\
& =\left\{\omega \in \Omega \mid \tau_{i}(\omega) \leq t \text { and } a_{i} \notin A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}(\omega)\right)\right\} .
\end{aligned}
$$

The above, together with Lemma 5—which is proved independently from Theorem 3(i)-delivers the result. Note that Proposition 7 further implies

Corollary 1. $\mathbb{P}_{\sigma_{-i}}\left(a_{i} \in(\right.$ resp. $\left.\notin) A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right) \mid \tau_{i} \leq t\right)$ is increasing (resp. decreasing) with respect to $\geq_{a_{i}}, \forall \sigma_{-i} \in \Delta\left(A_{-i}\right)$ and $t \in \mathbb{N}$.

Proof of Theorem 3(ii). For $a_{i} \in\{0,1\}$, define (i) $u_{i}^{a_{i}}\left(a_{i}^{\prime}, \sigma_{-i}\right):=u_{i}\left(a_{i}^{\prime}, \sigma_{-i}\right)-u_{i}\left(1-a_{i}, \sigma_{-i}\right)$, (ii) $v_{i}^{a_{i}}\left(\mu_{i}\right):=\max _{a_{i}^{\prime} \in A_{i}} \mathbb{E}_{\mu_{i}}\left[u_{i}^{a_{i}}\left(a_{i}^{\prime}, \sigma_{-i}\right)\right]$, and (iii) $V_{i}^{a_{i}}\left(\mu_{i}\right):=\sup _{t_{i} \in \mathbb{T}_{i}} \mathbb{E}_{\mu_{i}}\left[v_{i}^{a_{i}}\left(\mu_{i} \mid y_{i}^{t_{i}}\right)-c_{i} \cdot t_{i}\right]$. Note that, by definition, $u_{i}^{a_{i}}\left(1-a_{i}, \sigma_{-i}\right)=0$. Moreover, $v_{i}^{a_{i}}\left(\mu_{i}\right)=v_{i}\left(\mu_{i}\right)-\mathbb{E}_{\mu_{i}}\left[u_{i}\left(1-a_{i}, \sigma_{-i}\right)\right]$, which also implies that $V_{i}^{a_{i}}\left(\mu_{i}\right)=V_{i}\left(\mu_{i}\right)-\mathbb{E}_{\mu_{i}}\left[u_{i}\left(1-a_{i}, \sigma_{-i}\right)\right]$. A useful property of $V_{i}^{a_{i}}$ is as follows:
Lemma 9. For any $\mu_{i}^{\prime} \geq_{S S D} \mu_{i} V_{i}^{1}\left(\mu_{i}^{\prime}\right) \geq V_{i}^{1}\left(\mu_{i}\right)$ and $V_{i}^{0}\left(\mu_{i}^{\prime}\right) \leq V_{i}^{0}\left(\mu_{i}\right)$.
Proof. Let $B_{i}^{a_{i}}: \mathscr{C}^{0}\left(\Delta\left(\Delta\left(A_{-i}\right)\right)\right) \rightarrow \mathscr{C}^{0}\left(\Delta\left(\Delta\left(A_{-i}\right)\right)\right)$ be such that $B_{i}^{a_{i}}(w)\left(\mu_{i}\right):=\max \left\{v_{i}^{a_{i}}\left(\mu_{i}\right), \mathbb{E}_{\mu_{i}}\left[w\left(\mu_{i} \mid\right.\right.\right.$ $\left.\left.\left.y_{i}\right)\right]-c_{i}\right\}$. As argued in Section 2, $V_{i}^{a_{i}}$ is a fixed-point of $B_{i}^{a_{i}}$. Moreover, by Remark 1, there is a finite $n \in \mathbb{N}$, such that $V_{i}^{a_{i}}=B_{i}^{a_{i}(n)}\left(v_{i}^{a_{i}}\right)$, where $B_{i}^{a_{i}(1)}=B_{i}^{a_{i}}$ and, for $n \geq 1, B_{i}^{a_{i}(n+1)}=B_{i}^{a_{i}} \circ B_{i}^{a_{i}(n)}$. Note that $v_{i}^{1}$ (resp. $v_{i}^{0}$ is increasing (resp. decreasing) in $\geq_{S S D}$. If $w \in \mathscr{C}^{0}\left(\Delta\left(\Delta\left(A_{-i}\right)\right)\right.$ ) is increasing in $\geq_{S S D}$, then so is $B_{i}^{1}(w)-$ a symmetric argument applies to $B_{i}^{0}$. To see this, note that

$$
\begin{aligned}
B_{i}^{1}(w)\left(\mu_{i}^{\prime}\right) & =\max \left\{v_{i}^{1}\left(\mu_{i}^{\prime}\right), \mathbb{E}_{\mu_{i}^{\prime}}\left[\sigma_{-i}\right] w\left(\mu_{i}^{\prime} \mid 1\right)+\mathbb{E}_{\mu_{i}^{\prime}}\left[1-\sigma_{-i}\right] w\left(\mu_{i}^{\prime} \mid 0\right)-c_{i}\right\} \\
& \geq \max \left\{v_{i}^{1}\left(\mu_{i}\right), \mathbb{E}_{\mu_{i}^{\prime}}\left[\sigma_{-i}\right] w\left(\mu_{i} \mid 1\right)+\mathbb{E}_{\mu_{i}^{\prime}}\left[1-\sigma_{-i}\right] w\left(\mu_{i} \mid 0\right)-c_{i}\right\}
\end{aligned}
$$

$$
\geq \max \left\{v_{i}^{1}\left(\mu_{i}\right), \mathbb{E}_{\mu_{i}}\left[\sigma_{-i}\right] w\left(\mu_{i} \mid 1\right)+\mathbb{E}_{\mu_{i}}\left[1-\sigma_{-i}\right] w\left(\mu_{i} \mid 0\right)-c_{i}\right\}=B_{i}^{1}(w)\left(\mu_{i}\right),
$$

where the first inequality follows from monotonicity of $v_{i}^{1}$ with respect to $\geq_{S S D}$, by monotonicity of $w$ and the fact that $\mu_{i}^{\prime} \geq_{S S D} \mu_{i} \Longrightarrow \mu_{i}^{\prime}\left|y_{i} \geq_{S S D} \mu_{i}\right| y_{i}$ for $y_{i} \in\{0,1\}$, and the second because $\mathbb{E}_{\mu_{i}^{\prime}}\left[\sigma_{-i}\right] \geq \mathbb{E}_{\mu_{i}}\left[\sigma_{-i}\right]$ (by FOSD) and, as can be shown, $\mu_{i}\left|1 \geq_{S S D} \mu_{i}\right| 0$, implying $w\left(\mu_{i} \mid 1\right) \geq w\left(\mu_{i} \mid 0\right)$.

Lemma 9 implies:
Corollary 2. Let $\mu_{i}^{\prime} \geq_{S S D} \mu_{i}$. $V_{i}\left(\mu_{i}^{\prime} \mid y_{i}^{t}\right)=\mathbb{E}_{\mu_{i}^{\prime}}\left[u_{i}\left(0, \sigma_{-i}\right) \mid y_{i}^{t}\right] \Longrightarrow 0=V_{i}^{1}\left(\mu_{i}^{\prime} \mid y_{i}^{t}\right) \geq V_{i}^{1}\left(\mu_{i} \mid y_{i}^{t}\right) \geq$ $0 \Longrightarrow V_{i}\left(\mu_{i} \mid y_{i}^{t}\right)=\mathbb{E}_{\mu_{i}}\left[u_{i}\left(0, \sigma_{-i}\right) \mid y_{i}^{t}\right]$ and $V_{i}\left(\mu_{i} \mid y_{i}^{t}\right)=\mathbb{E}_{\mu_{i}}\left[u_{i}\left(1, \sigma_{-i}\right) \mid y_{i}^{t}\right] \Longrightarrow 0=V_{i}^{0}\left(\mu_{i} \mid y_{i}^{t}\right) \geq$ $V_{i}^{0}\left(\mu_{i}^{\prime} \mid y_{i}^{t}\right) \geq 0 \Longrightarrow V_{i}\left(\mu_{i}^{\prime} \mid y_{i}^{t}\right)=\mathbb{E}_{\mu_{i}^{\prime}}\left[u_{i}\left(1, \sigma_{-i}\right) \mid y_{i}^{t}\right]$.

In order to conclude the proof of Theorem 3(ii), let $\tau_{i}$ and $\tau_{i}^{\prime}$ denote the earliest optimal stopping times associated with $\mu_{i}$ and $\mu_{i}^{\prime}$. Then, by Corollary $2,\left\{\omega \in \Omega \mid \tau_{i}(\omega) \leq t\right.$ and $1 \in($ resp. $\left.\notin) A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}(\omega)\right)\right\} \subseteq$ (resp. $\supseteq)\left\{\omega \in \Omega \mid \tau_{i}^{\prime}(\omega) \leq t\right.$ and $1 \in($ resp. $\left.\notin) A_{i}^{*}\left(\mu_{i}^{\prime} \mid y_{i}^{\tau_{i}^{\prime}}(\omega)\right)\right\}$.

Proof of Theorem 3(iii). Let
$\mathscr{N}:=\left\{n \in \mathbb{N}_{0}^{2} \mid \exists y_{i}^{t}:(\mathrm{i}) t=n_{0}+n_{1}\right.$, (ii) $\sum_{\ell \in[1 . . t]} y_{i, \ell}=n_{1}$, and (iii) $\left.\forall \ell \leq t, V_{i}\left(\mu_{i} \mid y_{i}^{\ell}\right)>v_{i}\left(\mu_{i} \mid y_{i}^{\ell}\right)\right\}$, and, for $j \in\{0,1\}$, let $\mathscr{N}_{j}:=\left\{n \in \mathbb{N}_{0}^{2} \mid n-(j, 1-j) \in \mathscr{N}\right\}$. Note that, if $n \in \mathscr{N}_{j}$, then there is some sequence $y_{i}^{t}$ satisfying $t=n_{0}+n_{1}, \sum_{\ell \in[1 . . t]} y_{i, \ell}=n_{1}$, and along which player $i$ decides to keep sampling every period (according to $\tau_{i}$ ), i.e. $V_{i}\left(\mu_{i} \mid y_{i}^{\ell}\right)>v_{i}\left(\mu_{i} \mid y_{i}^{\ell}\right)$ for all $\ell<t$, and decides to stop at $y_{i}^{t}$ and take action $j$-a consequence of Lemma 5. Let $T_{i}:=\sup \operatorname{supp}\left\{\tau_{i}\right\}$ (where supp is defined with respect to $\mu_{i}$ ). By Remark $1, T_{i}<\infty$ and thus, $\forall n \in \mathscr{N}, n_{0}+n_{1}<T_{i}$ and $\mathscr{N}$ is finite. Below we implicitly rely on the fact that, if $\left(n_{0}, n_{1}\right),\left(n_{0}^{\prime}, n_{1}^{\prime}\right) \in \mathscr{N}_{j}$ and $n_{1-j}^{\prime}>n_{1-j}$, then $n_{j}^{\prime} \geq n_{j}$, which is implied by Corollary 2.

We recursively define the probability of stopping and choosing action 1 . Define the asymmetric part of a linear order on $\mathscr{N}$ given by $n \triangleright n^{\prime}$ if and only if $n_{1}^{\prime}>n_{1}$ or $n_{1}^{\prime}=n_{1}$ and $n_{0}^{\prime}>n_{0}$. Let $p: \mathscr{N} \times[0,1] \rightarrow[0,1]$ be given by $p\left(n ; \sigma_{-i}\right):=\sigma_{-i}$ if $n+(0,1) \in \mathscr{N}_{1}$ and $p\left(n ; \sigma_{-i}\right):=\sigma_{-i} p(n+$ $\left.(0,1) ; \sigma_{-i}\right)+\left(1-\sigma_{-i}\right) p\left(n+(1,0) ; \sigma_{-i}\right)$ if otherwise; $p$ can be recursively defined on $n \in \mathscr{N}$ increasing with respect to $\triangleright$. Extend $p$ to $n \in \mathscr{N}_{j}$ by letting $p\left(n ; \sigma_{-i}\right)=j$ if $n \in \mathscr{N}_{j}$. Note that $p\left((0,0) ; \sigma_{-i}\right)=\mathbb{P}_{\sigma_{-i}}\left(1=A_{i}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right)=\mathbb{E}_{\sigma_{-i}}\left[\sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right]$.

We now show by induction that, if $(0,0) \in \mathscr{N}, p((0,0) ; \bullet)$ is $\mathscr{C}^{\infty}$ and strictly increasing. Note that for $n: n^{\prime} \triangleright n$ for all $n^{\prime} \neq n$ in $\mathscr{N}, n_{0}+n_{1}=T_{i}-1$ and $p\left(n ; \sigma_{-i}\right)=\sigma_{-i}$ is $\mathscr{C}^{\infty}$ and strictly increasing
in $\sigma_{-i}$ and $0=p\left(n+(0,1) ; \sigma_{-i}\right)<p\left(n+(0,1) ; \sigma_{-i}\right)=1, \forall \sigma_{-i} \in[0,1)$.
Suppose that, for all $n^{\prime} \in \mathscr{N}: n \triangleright n^{\prime}, p\left(n^{\prime} ; \bullet\right)$ is $\mathscr{C}^{\infty}$ and strictly increasing, and that $p\left(n^{\prime}+\right.$ $(1,0) ; \bullet)<p\left(n^{\prime}+(0,1) ; \bullet\right)$. As $p\left(n ; \sigma_{-i}\right)=\sigma_{-i} p\left(n+(0,1) ; \sigma_{-i}\right)+\left(1-\sigma_{-i}\right) p\left(n+(1,0) ; \sigma_{-i}\right)$ and $n \triangleright n+(1,0), n+(0,1)$, then $p(n+(0,1) ; \bullet)$ and $p(n+(1,0) ; \bullet)$ are $\mathscr{C}^{\infty}$, strictly increasing, and satisfy $p(n+(1,0) ; \bullet)<p(n+(0,1) ; \bullet)$. Then $p(n ; \bullet) \in \mathscr{C}^{\infty}$ and $\frac{\partial}{\partial \sigma_{-i}} p\left(n ; \sigma_{-i}\right)=p\left(n+(0,1) ; \sigma_{-i}\right)-$ $p\left(n+(1,0) ; \sigma_{-i}\right)+\sigma_{-i} \frac{\partial}{\partial \sigma_{-i}} p\left(n+(0,1) ; \sigma_{-i}\right)+\left(1-\sigma_{-i}\right) \frac{\partial}{\partial \sigma_{-i}} p\left(n+(1,0) ; \sigma_{-i}\right)>0$ for all $\sigma_{-i} \in[0,1)$. To obtain that $\mathbb{P}_{\sigma_{-i}}\left(1=A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right.$ and $\left.\tau_{i} \leq t\right)=\mathbb{E}_{\sigma_{-i}}\left[\sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right) 1_{\tau_{i} \leq t}\right]$ is also $\mathscr{C}^{\infty}$ and strictly increasing in $\sigma_{-i} \in[0,1)$ for $t \geq n_{1}(0)$, it is necessary to restrict $\mathscr{N}$. Define $\underline{n}_{1}(m):=n_{1}$ if $\left(m, n_{1}\right) \in$ $\mathscr{N}_{1}$ and let $\mathscr{N}^{t}:=\left\{n \in \mathscr{N} \mid n_{0}+\underline{n}_{1}\left(n_{0}\right) \leq t\right\}, \mathscr{N}_{j}^{t}:=\{n \mid n-(1-j, j) \in \mathscr{N}\}$. Let $p^{t}: \mathscr{N} \times[0,1] \rightarrow[0,1]$ be given by $p^{t}\left(n ; \sigma_{-i}\right):=\sigma_{-i}$ if $n+(0,1) \in \mathscr{N}_{1}^{t}$ and $p^{t}\left(n ; \sigma_{-i}\right):=\sigma_{-i} p\left(n+(0,1) ; \sigma_{-i}\right)+\left(1-\sigma_{-i}\right) p(n+$ (1,0); $\sigma_{-i}$ ) if otherwise. Extend $p^{t}$ to $n \in \mathscr{N}_{j}^{t}$ by letting $p^{t}\left(n ; \sigma_{-i}\right)=j$ if $n \in \mathscr{N}_{j}^{t}$. An analogous inductive argument applied to $p^{t}$ delivers the result. For $t<n_{1}(0), \mathbb{P}_{\sigma_{-i}}\left(1=A_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right.$ and $\tau_{i} \leq$ $t)=0$, for all $\sigma_{-i} \in[0,1]$.

Proof of Lemma 5. Label player $i$ 's actions so that $u_{i}\left(1, \sigma_{-i}\right)-u_{i}\left(0, \sigma_{-i}\right)$ is increasing in $\sigma_{-i}$. If, (i) $u_{i}\left(1, \sigma_{-i}\right)-u_{i}\left(0, \sigma_{-i}\right) \leq 0 \forall \sigma_{-i}$ or (ii) $u_{i}\left(1, \sigma_{-i}\right)-u_{i}\left(0, \sigma_{-i}\right) \geq 0 \forall \sigma_{-i}$, then $\tau_{i}=0$. Suppose then that $\tau_{i}>0$ and note $a_{-i}=\arg \max _{a_{i} \in A_{i}} u_{i}\left(a_{i}, a_{-i}\right)$. Let $V_{i}\left(\mu_{i}\right)>v_{i}\left(\mu_{i}\right)$ and $V_{i}\left(\mu_{i} \mid a_{-i}\right)=$ $\mathbb{E}_{\mu_{i}}\left[u_{i}\left(1-a_{-i}, \sigma_{-i}\right) \mid a_{-i}\right]$ and observe that, from Lemma 9, $0=V_{i}^{a_{-i}}\left(\mu_{i} \mid a_{-i}\right) \geq V_{i}^{a_{-i}}\left(\mu_{i}\right) \geq 0$, which implies $V_{i}\left(\mu_{i}\right)=\mathbb{E}_{\mu_{i}}\left[u_{i}\left(1-a_{-i}, \sigma_{-i}\right)\right] \leq v_{i}\left(\mu_{i}\right)$, a contradiction.

Proof of Proposition 2. Note that in a binary action game there are multiple unique Nash equilibria if and only if for any $i \in I, u_{i}(1,1)-u_{i}(0,1), u_{i}(0,0)-u_{i}(1,0)>0$. This implies that, for both players, both actions are undominated and so, by Lemma 4, players sample at least once whenever the sampling costs are sufficiently low. By Lemma 5, if $\sigma_{-i}=1(=0)$ and $\tau_{i}>0$, then $\mathbb{E}_{\sigma_{-i}}\left[\sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}}\right)\right]=1(=0)$. Hence, $(0,0)$ and $(1,1)$ are both Nash equilibria and sequential sampling equilibria.

Recall that, by assumption, $u_{i}\left(1, \sigma_{-i}\right)-u_{i}\left(0, \sigma_{-i}\right)$ is strictly monotone. If there is a unique Nash equilibrium, either (i) both players have a weakly dominant action; (ii) one player has both actions undominated, and the other has a weakly dominant action; or (iii) both players have undominated actions. In (i), uniqueness of a sequential sampling equilibrium follows as both players always choose their weakly dominant action. In (ii), the player with the weakly dominant action chooses it with probability 1 , and the opponent, whenever they sample at least once, by Lemma 5, will choose the best response to the weakly dominant action with probability 1 , thus
entailing a unique sequential sampling equilibrium. In (iii), uniqueness of a Nash equilibrium implies a payoff structure akin to a matching pennies game: one player, $i$, wants to match the other, i.e. $u_{i}\left(1, \sigma_{j}\right)-u_{i}\left(0, \sigma_{j}\right)$ is strictly increasing, and their opponent $j$ seeks to mismatch, i.e. $u_{j}\left(1, \sigma_{i}\right)-u_{j}\left(0, \sigma_{i}\right)$ is strictly decreasing. Consequently, by Theorem 3, $\sigma_{j} \mapsto b_{i}\left(\sigma_{j}\right)=\mathbb{E}_{\sigma_{j}}\left[\sigma_{i}^{*}\left(\mu_{i} \mid\right.\right.$ $\left.\left.y_{i}^{\tau_{i}}\right)\right] \in[0,1]$ is increasing, and $\sigma_{i} \mapsto b_{j}\left(\sigma_{i}\right)=\mathbb{E}_{\sigma_{i}}\left[\sigma_{j}^{*}\left(\mu_{j} \mid y_{j}^{\tau_{j}}\right)\right] \in[0,1]$ is decreasing, both are continuous, and therefore their graph has a unique intersection. Lemma 5 clarifies that the sequential sampling equilibrium is unaffected by the selection of tie-breaking $\sigma_{i}^{*}, \sigma_{j}^{*}$.

Proof of Proposition 3. Let $B\left(\sigma_{-i}, t\right)$ to denote a Beta distribution with parameters $\alpha, \beta \geq 0$, such that $t=\alpha+\beta>0$ and $\sigma_{-i}=\alpha / t$, with the convention that $B(1, t)$ and $B(0, t)$ correspond to Dirac measures on 1 and 0 , respectively. Label player $i$ 's actions so that $u_{i}\left(1, \sigma_{-i}\right)-u_{i}\left(0, \sigma_{-i}\right)$ is increasing in $\sigma_{-i}$.

Let $V_{i}^{a_{i}}$ be as defined in the proof of Lemma 9. We show the following properties of $V_{i}^{a_{i}}$ :
Lemma 10. $V_{i}^{1}\left(B\left(\sigma_{-i}, t\right)\right)\left(r e s p . V_{i}^{0}\left(B\left(\sigma_{-i}, t\right)\right)\right)$ is (1) increasing (resp. decreasing) in $\sigma_{-i} ;$ (2) convex in $\sigma_{-i}$; (3) decreasing in $t$; (4) continuous in ( $\sigma_{-i}, t$ ).
Proof. (1) follows immediately from Lemma 9, since for $\sigma_{-i}>\sigma_{-i}^{\prime}, B\left(\sigma_{-i}, t\right) \geq_{S S D} B\left(\sigma_{-i}^{\prime}, t\right)$. For (2), note that $V_{i}^{1}\left(B\left(\sigma_{-i}, t\right)\right)=V_{i}^{0}\left(B\left(\sigma_{-i}, t\right)\right)+u_{i}\left(1, \sigma_{-i}\right)-u_{i}\left(0, \sigma_{-i}\right)$, and thus it suffices to show convexity of $V_{i}^{1}\left(B\left(\sigma_{-i}, t\right)\right)$ in $\sigma_{-i}$. Let $z\left(\sigma_{-i}, t\right)$ be a random variable that delivers $\frac{1}{t+1}\left(t \sigma_{-i}+1\right)$ with probability $\sigma_{-i}$, and $\frac{1}{t+1}\left(t \sigma_{-i}+0\right)$ with probability $1-\sigma_{-i}$. Note that, for any $w:[0,1] \times \mathbb{R}_{++}$ that is increasing and convex in the first argument, for any $\sigma_{-i} \geq \sigma_{-i}^{\prime}, \lambda \in(0,1)$, and $\sigma_{-i}^{\prime \prime}:=$ $\lambda \sigma_{-i}+(1-\lambda) \sigma_{-i}^{\prime}$, straightforward algebra shows that $\mathbb{E}\left[w\left(z\left(\sigma_{-i}^{\prime \prime}, t\right)\right)\right] \leq \lambda \mathbb{E}\left[w\left(z\left(\sigma_{-i}, t\right)\right)\right]+(1-$ $\lambda) \mathbb{E}\left[w\left(z\left(\sigma_{-i}^{\prime}, t\right)\right)\right]$.

Observe that, by definition, $v_{i}^{1}\left(B\left(\sigma_{-i}, t\right)\right)=\max _{a_{i} \in A_{i}} u_{i}\left(a_{i}, \sigma_{-i}\right)-u_{i}\left(0, \sigma_{-i}\right)$, which is increasing and convex in $\sigma_{-i}$ and invariant with respect to $t$.

Now we show that if $w\left(B\left(\sigma_{-i}, t\right)\right)$ is increasing and convex in $\sigma_{-i}$, so is $B_{i}^{1}(w)\left(B\left(\sigma_{-i}, t\right)\right)$ :

$$
\begin{aligned}
& B_{i}^{1}(w)\left(B\left(\sigma_{-i}^{\prime \prime}, t\right)\right)=\max \left\{v_{i}^{1}\left(B\left(\sigma_{-i}^{\prime \prime}, t\right)\right), \mathbb{E}\left[w\left(z\left(\sigma_{-i}^{\prime \prime}, t+1\right)\right)\right]-c_{i}\right\} \\
& \leq \max \left\{\lambda v_{i}^{1}\left(B\left(\sigma_{-i}, t\right)\right)+(1-\lambda) v_{i}^{1}\left(B\left(\sigma_{-i}^{\prime}, t\right)\right), \lambda \mathbb{E}\left[w\left(z\left(\sigma_{-i}, t+1\right)\right)\right]+(1-\lambda) \mathbb{E}\left[w\left(z\left(\sigma_{-i}^{\prime}, t+1\right)\right)\right]-c_{i}\right\} \\
& \leq \lambda B_{i}^{1}(w)\left(B\left(\sigma_{-i}, t\right)\right)+(1-\lambda) B_{i}^{1}(w)\left(B\left(\sigma_{-i}^{\prime}, t\right)\right) .
\end{aligned}
$$

By similar arguments as in Lemma 9- $V_{i}^{1}$ is a fixed point of $B_{i}^{1}$ which can be obtained by applying the $n$-th composition of $B_{i}^{1}$ with itself to $v_{i}^{1}$ - we have that $V_{i}^{1}\left(B\left(\sigma_{-i}, t\right)\right)$ is convex in $\sigma_{-i}$, and thus so is $V_{i}\left(B\left(\sigma_{-i}, t\right)\right)$.

For (3), note that $\mathbb{E}\left[z\left(\sigma_{-i}, t\right)\right]=\sigma_{-i}$ and that, for $t<t^{\prime}, z\left(\sigma_{-i}, t\right)$ is a mean-preserving spread of $z\left(\sigma_{-i}, t^{\prime}\right)$; hence for any convex function $f, \mathbb{E}\left[f\left(z\left(\sigma_{-i}, t\right)\right)\right] \geq \mathbb{E}\left[f\left(z\left(\sigma_{-i}, t^{\prime}\right)\right)\right]$. Take any $w$ such that $w\left(B\left(\sigma_{-i}, t\right)\right)$ is convex in $\sigma_{-i}$ and decreasing in $t$. Then, for any $t^{\prime}>t$,

$$
\begin{aligned}
& B_{i}(w)\left(B\left(\sigma_{-i}, t^{\prime}\right)\right)=\max \left\{v_{i}\left(B\left(\sigma_{-i}, t^{\prime}\right)\right), \mathbb{E}\left[w\left(B\left(z\left(\sigma_{-i}, t^{\prime}\right), t^{\prime}\right)\right)\right]-c_{i}\right\} \\
& =\max \left\{v_{i}\left(B\left(\sigma_{-i}, t\right)\right), \mathbb{E}\left[w\left(B\left(z\left(\sigma_{-i}, t^{\prime}\right), t^{\prime}\right)\right)\right]-c_{i}\right\} \leq \max \left\{v_{i}\left(B\left(\sigma_{-i}, t\right)\right), \mathbb{E}\left[w\left(B\left(z\left(\sigma_{-i}, t^{\prime}\right), t\right)\right)\right]-c_{i}\right\} \\
& \leq \max \left\{v_{i}\left(B\left(\sigma_{-i}, t\right)\right), \mathbb{E}\left[w\left(B\left(z\left(\sigma_{-i}, t\right), t\right)\right)\right]-c_{i}\right\}=B_{i}(w)\left(B\left(\sigma_{-i}, t\right)\right) .
\end{aligned}
$$

By the same argument as before, $V_{i}\left(B\left(\sigma_{-i}, t\right)\right)$ is then decreasing in $t$.
Finally, (4) follows immediately from the fact that for any $\left(\sigma_{-i}^{n}, t^{n}\right) \rightarrow\left(\sigma_{-i}, t\right), B\left(\sigma_{-i}^{n}, t^{n}\right) \rightarrow B\left(\sigma_{-i}, t\right)$ (with respect to $\left.\|\cdot\|_{L P}\right)$, and therefore by Proposition $1, V_{i}\left(B\left(\sigma_{-i}, t\right)\right.$ ) is continuous in $\left(\sigma_{-i}, t\right)$.

By Corollary 2, if $V_{i}\left(B\left(\sigma_{-i}, t\right)\right)=u_{i}\left(a_{i}, \sigma_{-i}\right)$ for $a_{i}=1$ (resp. $\left.a_{i}=0\right)$, then $V_{i}\left(B\left(\sigma_{-i}^{\prime}, t\right)\right)=u_{i}\left(a_{i}, \sigma_{-i}^{\prime}\right)$ for any $\sigma_{-i}^{\prime} \geq \sigma_{-i}$ (resp. $\leq$ ), since $B\left(\sigma_{-i}^{\prime}, t\right) \geq_{S S D} B\left(\sigma_{-i}, t\right)$. Define $\bar{\sigma}_{-i}(t):=\min \left\{\sigma_{-i} \in[0,1] \mid\right.$ $\left.V_{i}\left(B\left(\sigma_{-i}, t\right)\right)=u_{i}\left(1, \sigma_{-i}\right)\right\}$ and $\underline{\sigma}_{-i}(t):=\max \left\{\sigma_{-i} \in[0,1] \mid V_{i}\left(B\left(\sigma_{-i}, t\right)\right)=u_{i}\left(0, \sigma_{-i}\right)\right\}$.

That $\bar{\sigma}_{-i}(t)$ is continuous and decreasing in $t$ follows from continuity of $V_{i}\left(B\left(\sigma_{-i}, t\right)\right)$ in $\left(\sigma_{-i}, t\right)$, continuity of $u_{i}\left(1, \sigma_{-i}\right)$ in $\sigma_{-i}$, and the fact that $V_{i}\left(B\left(\sigma_{-i}, t\right)\right)$ is decreasing in $t$ and $u_{i}\left(1, \sigma_{-i}\right)$ $u_{i}\left(0, \sigma_{-i}\right)$ is increasing in $\sigma_{-i}$. An analogous argument applies to show that $\underline{\sigma}_{-i}(t)$ is continuous and increasing in $t$.

Finally, we prove that both these functions converge to $\tilde{\sigma}_{-i}$. To see this, note that for $\mu_{i}$ given by $B\left(\sigma_{-i}, t\right)$, simple algebra shows that $\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\mu_{i} \mid y_{i}\right)\right]-v_{i}\left(\mu_{i}\right)=\mathbb{E}\left[v_{i}\left(z\left(\sigma_{-i}, t\right)\right)\right]-v_{i}\left(\sigma_{-i}\right)$ is maximized for any $t$ at $\sigma_{-i}=\tilde{\sigma}_{-i}: \tilde{\sigma}_{-i} u_{i}(1,1)+\left(1-\tilde{\sigma}_{-i}\right) u_{i}(1,0)=\tilde{\sigma}_{-i} u_{i}(0,1)+\left(1-\tilde{\sigma}_{-i}\right) u_{i}(0,0)$, with maximum value $\frac{1}{t+1}\left(u_{i}(1,1)-u_{i}(0,1)+u_{i}(0,0)-u_{i}(1,0)\right)\left(1-\tilde{\sigma}_{-i}\right) \tilde{\sigma}_{-i}$. Therefore, it is always optimal to keep sampling at belief $B\left(\tilde{\sigma}_{-i}, t\right)$ if $t<T:=\left(u_{i}(1,1)-u_{i}(0,1)+u_{i}(0,0)-u_{i}(1,0)\right)\left(1-\tilde{\sigma}_{-i}\right) \tilde{\sigma}_{-i} / c_{i}-1$. From here, one can deduce that $\tau_{i} \leq T$ for any Beta prior and that $\bar{\sigma}_{-i}(t)>\tilde{\sigma}_{-i}>\underline{\sigma}_{-i}(t)$ for all $t<T$ and $\bar{\sigma}_{-i}(t)=\tilde{\sigma}_{-i}=\underline{\sigma}_{-i}(t)$, for all $t \geq T$.

Proof of Theorem 4. We prove the result when priors allow for correlation; adjusting the proof to accommodate the case in which they do not is tedious but straightforward.

Let $\Sigma_{i}^{*}\left(\mu_{i}\right):=\arg \max _{\sigma_{i} \in \Delta\left(A_{i}\right)} \mathbb{E}_{\mu_{i}}\left[u_{i}\left(\sigma_{i}, \sigma_{-i}\right)\right]$ denote the set of maximizers at belief $\mu_{i}$.
We first prove that if a sequence of probability measures $\mu_{i}^{m} \in \Delta\left(\Delta\left(A_{-i}\right)\right)$ weak ${ }^{*}$ converges to $\delta_{\sigma_{-i}}$ and $a_{i} \in A_{i}$ is not a best response to $\sigma_{-i}$, then for any sequence of distributions $\sigma_{i}^{m} \in \Sigma_{i}^{*}\left(\mu_{i}^{m}\right)$, $\sigma_{i}^{m}\left(a_{i}\right) \rightarrow 0$. Note that $\mathbb{E}_{\mu_{i}}\left[u_{i}\left(\sigma_{i}, \sigma_{-i}\right)\right]$ is jointly continuous in ( $\sigma_{i}, \mu_{i}$ ) with respect to the product metric, where $\Delta\left(A_{i}\right)$ is endowed with the standard Euclidean metric and $\Delta\left(\Delta\left(A_{-i}\right)\right)$ with the

Lévy-Prokhorov metric. Then, by Berge's maximum theorem, $\Sigma_{i}^{*}$ is upper-hemicontinuous and compact-valued. Supposing that $\sigma_{i}^{m}\left(a_{i}\right)$ does not converge to zero, implies that for any convergent subsequence of $\sigma_{i}^{m}$, its limit assigns strictly positive probability to $a_{i}$ being chosen, while also belonging to $\Sigma_{i}^{*}\left(\delta_{\sigma_{-i}}\right)$, a contradiction.

Now take any sequence of profiles of action distributions $\left\{\sigma^{n}\right\}_{n}$ as in the statement of the theorem and let $\left\{\sigma^{m}\right\}_{m}$ be a convergent subsequence of $\left\{\sigma^{n}\right\}_{n}$ with limit $\sigma^{k}$. Suppose that $\sigma_{i}^{k}$ is not a best response to $\sigma_{-i}^{k}$. This implies that $\exists a_{i} \in A_{i}$ such that $\sigma_{i}^{k}\left(a_{i}\right) \geq \delta>0$ for some $\delta>0$ and $a_{i}$ is not a best response to $\sigma_{-i}^{k}$. By continuity of $u_{i}$, note that if $a_{i}$ is not a best response to $\sigma_{-i}^{k}$, then it is not a best response to any $\sigma_{-i} \in B_{\epsilon}\left(\sigma_{-i}^{k}\right)$ for small enough $\epsilon>0$. By convergence of $\sigma^{m}$, we then obtain that for all large enough $m, a_{i}$ is not a best response to $\sigma_{-i}^{m}$ and $\sigma_{i}^{m}\left(a_{i}\right) \geq$ $\delta / 2$. That is, $\mathbb{E}_{\sigma_{-i}^{m}}\left[\sigma_{i}^{*}\left(\mu_{i} \mid y_{i}^{\tau_{i}^{m}}\right)\left(a_{i}\right)\right] \geq \delta / 2 \forall$ large enough $m$, from which we deduce $\mathbb{P}_{\sigma_{-i}^{m}}\left(\sigma_{i}^{*}\left(\mu_{i} \mid\right.\right.$ $\left.\left.y_{i}^{\tau_{i}^{m}}\right)\left(a_{i}\right) \geq \delta / 4\right) \geq \delta / 4$. In turn, from the above, this implies that there must be $\epsilon>0$ such that $\mathbb{P}_{\sigma_{-i}^{m}}\left(\left\|\mu_{i} \mid y_{i}^{\tau_{i}^{m}}-\delta_{\sigma_{-i}^{m}}\right\|_{L P} \geq \epsilon\right) \geq \delta / 4$. We now prove that this cannot be the case, that is, we show that $\lim _{m \rightarrow \infty} \mathbb{P}_{\sigma_{-i}^{k}}\left(\left\|\mu_{i} \mid y_{i}^{\tau_{i}^{m}}-\delta_{\sigma_{-i}^{k}}\right\|_{L P}>\epsilon\right)=0$.
It would be natural to expect that, with sampling costs going to zero, optimal stopping time grows unboundedly and, by the law of large numbers, players learn the true distribution of actions of their opponents, best respond to it, and sequential sampling equilibrium converges to a Nash equilibrium. But, conditional on stopping, the set of signals are neither independent or identically distributed, so we cannot apply the law of large number directly. We then need to take a detour.

Denote player $i$ 's the associated earliest optimal stopping time by $\tau_{i}^{m}$ and their value function (which depends on $c_{i}^{m}$ ) as $V_{i}^{m}$. From Lemma 4, there is $\left\{T^{m}\right\}_{m}$ such that $\tau_{i}^{m} \geq T^{m}$ and $T^{m} \uparrow \infty$.

By Diaconis and Freedman (1990), there is $\epsilon(t)$ nonincreasing and such that $\epsilon(t) \rightarrow 0$ as $t \rightarrow \infty$ such that $\left\|\mu_{i} \mid y_{i}^{t}-\delta_{\bar{y}_{i}}\right\|_{L P} \leq \epsilon(t)$ uniformly over sequences of $t$ observations, $y_{i}^{t} \in \mathscr{Y}_{i}$. Since, taking $y_{i, \ell} \sim \sigma_{-i},\left\|\delta_{\bar{y}_{i}^{t}}-\delta_{\sigma_{-i}}\right\|_{L P}=\left\|{\overline{y_{i}}}^{t}-\sigma_{-i}\right\|$ is a (bounded) supermartingale with respect to $\sigma_{-i}$, by the optional stopping theorem, for $\tau_{i} \geq t, \mathbb{E}_{\sigma_{-i}}\left[\left\|\delta_{\bar{y}_{i} \tau_{i}}-\delta_{\sigma_{-i}}\right\|_{L P}\right] \leq \mathbb{E}_{\sigma_{-i}}\left[\left\|\delta_{\bar{y}_{i}} t-\delta_{\sigma_{-i}}\right\|_{L P}\right]$. Hence, for any $\sigma_{-i}, \mathbb{E}_{\sigma_{-i}}\left[\left\|\mu_{i} \mid y_{i}^{\tau_{i}^{m}}-\delta_{\sigma_{-i}}\right\|_{L P}\right] \leq \mathbb{E}_{\sigma_{-i}}\left[\left\|\delta_{\bar{y}_{i} \tau_{i}^{m}}-\delta_{\sigma_{-i}}\right\| \|_{L P} \mid y_{i}^{T^{m}}\right]+\epsilon\left(T^{m}\right) \leq \mathbb{E}_{\sigma_{-i}}\left[\left\|\delta_{\bar{y}_{i} T^{m}}-\delta_{\sigma_{-i}}\right\|_{L P}\right]+$ $\epsilon\left(T^{m}\right) \leq \mathbb{E}_{\sigma_{-i}}\left[\left\|{\overline{y_{i}}}^{T^{m}}-\sigma_{-i}\right\|\right]+\epsilon\left(T^{m}\right)$.
Let $x_{m}:=\mathbb{E}_{\sigma_{-i}^{m}}\left[\left\|\bar{y}_{i}^{T^{m}}-\sigma_{-i}^{m}\right\|\right] \in[0,2]$. Suppose that $\left\{x_{m}\right\}_{m}$ does not converge to 0 . Take any convergent subsequence $x_{\ell} \rightarrow \gamma>0$. For all large enough $\ell, x_{\ell} \geq \gamma / 2$. That is, $\mathbb{E}_{\sigma_{-i}^{\ell}}\left[\left\|{\overline{y_{i}}}^{T^{\ell}}-\sigma_{-i}^{\ell}\right\|\right] \geq \gamma / 2$, implying that $\mathbb{P}_{\sigma_{-i}^{\ell}}\left(\left\|{\overline{y_{i}}}^{T^{\ell}}-\sigma_{-i}^{\ell}\right\| \geq \gamma / 4\right) \geq \gamma / 4$, as otherwise $\mathbb{E}_{\sigma_{-i}^{\ell}}\left[\left\|{\overline{y_{i}}}^{T^{\ell}}-\sigma_{-i}^{\ell}\right\|\right] \leq(1-\gamma / 4) \gamma / 4+\gamma / 4 \leq$ $\gamma / 4<\gamma / 2$. However, by the Dvoretzky-Kiefer-Wolfowitz-Massart inequality (Massart, 1990)
$\mathbb{P}_{\sigma_{-i}^{\ell}}\left(\left\|{\overline{y_{i}}}^{T^{\ell}}-\sigma_{-i}^{\ell}\right\| \geq \gamma / 4\right) \leq 2 \exp \left(-T^{\ell} \gamma^{2} / 8\right) \rightarrow 0$, a contradiction.
We conclude that $\lim _{m \rightarrow \infty} \mathbb{E}_{\sigma_{-i}^{m}}\left[\left\|\mu_{i} \mid y_{i}^{\tau_{i}^{m}}-\delta_{\sigma_{-i}^{m}}\right\|_{L P}\right] \leq \mathbb{E}_{\sigma_{-i}^{m}}\left[\left\|{\overline{y_{i}}}^{T^{m}}-\sigma_{-i}^{m}\right\|\right]+\epsilon\left(T^{m}\right) \rightarrow 0$.
Proof of Proposition 5. By Diaconis and Freedman (1990), for any $\mu_{i}$, there is $T_{i}<\infty$ such that $\forall t \geq T_{i}, \mathbb{E}_{\mu_{i}}\left[\sigma_{-i}^{\prime} \mid a_{-i}^{t}\right] \in B_{\varepsilon_{i}}\left(\delta_{a_{-i}}\right)$. Take $T:=\max _{i \in I} T_{i}$. By Lemma 4, there is $\bar{c}$ such that, $\tau_{i} \geq T$ for every player $i$ for which $a_{i}$ is not always a best response (where $\bar{c}$ may depend on $\mu$ ). Hence, for any $\mu$ there is an $N$ such that $\forall n \geq N, a$ is a sequential sampling equilibrium of $\left\langle\Gamma, \mu, c^{n}\right\rangle$.

## Online Appendix

## Appendix B. General Information Structures

In this section, we extend sequential sampling equilibrium to accommodate analogy partitions and more general information structures. For brevity, let us consider the case in which players' beliefs allows for correlation.

## B.1. Existence of a Sequential Sampling Equilibrium

Let us simplify the notation in our baseline setup: $Y_{i}:=A_{-i}, P_{i}:=\Delta\left(Y_{i}\right), \mu_{i} \in \Delta\left(P_{i}\right)$ with full support, and $u_{i}: A_{i} \times P_{i} \rightarrow \mathbb{R}$, a continuous function.

An analogy partition for player $i$ can be represented by a surjective function $f_{i}: Y_{i} \rightarrow Z_{i}$, where $\left|Z_{i}\right|<\left|Y_{i}\right|<\infty$. Naturally, it defines a garbling: instead of observing $y_{i}$, the player has access to coarser information $f_{i}\left(y_{i}\right)$. More generally, one can consider $\left|Z_{i}\right| \times\left|Y_{i}\right|$ stochastic matrices $B_{i}$, where $B_{i}(z, y) \in[0,1]$ and $\sum_{z \in Z} B_{i}(z, y)=1$, and such that $B_{i}$ has rank $\left|Z_{i}\right|<\left|Y_{i}\right| .{ }^{29}$

Let $Q_{i}:=\Delta\left(Z_{i}\right)$ and $v_{i}$ be the pushforward measure on $Q_{i}$ given $\mu_{i}$ and $B_{i}$, where for every measurable set $S \subseteq Q_{i}, v_{i}(S):=\mu_{i}\left(\left\{p_{i} \in P_{i} \mid B_{i} p_{i}=\epsilon S\right\}\right)$. We assume that the player now has access to iid draws from a fixed $B_{i} p_{i, 0} \in Q_{i}$. It is straightforward to adjust the definition of the optimal stopping problem, expand the definition of an extended game with the additional primitives $\left\{\pi_{i}\right\}_{i \in I}$, where $\pi_{i}$ denotes the information structure defined by $\left(B_{i}, Z_{i}\right)$, and have sequential sampling equilibrium accommodate such more general information structures.

We provide the following sufficient condition for existence of a sequential sampling equilibrium:
Theorem 6. Let $G:=\langle\Gamma, \mu, c, \pi\rangle$ be an extended game such that for every player $i, \mu_{i}$ admits a continuous density, and $\operatorname{rank}\left(B_{i}\right)=\left|Z_{i}\right|<\left|A_{-i}\right|$. Then $G$ admits a sequential sampling equilibrium. Proof. Let $J_{B_{i}}:=\operatorname{det}\left(B_{i} B_{i}^{T}\right)$, which is strictly positive, as $\operatorname{rank}\left(B_{i}\right)=\left|Z_{i}\right|$. For convenience, define $\pi_{i}: P_{i} \rightarrow Q_{i}$ as $\pi_{i}\left(p_{i}\right):=B_{i} p_{i}$. Denote by $g_{\mu_{i}}$ the density of $\mu_{i}$. Denoting $\lambda^{n}$ the $n$-dimensional Lebesgue measure, by the coarea formula (see Evans and Gariepy, 2015, Theorem 3.10) we have

$$
\begin{aligned}
\mathbb{E}_{\mu_{i}}\left[u_{i}\left(a_{i}, p_{i}\right) \mid z_{i}^{t}\right] & =\int_{P_{i}} \prod_{\ell \in[1 . . t]}\left(B_{i} p_{i}\right)\left(z_{i, \ell}^{t}\right) u_{i}\left(a_{i}, p_{i}\right) g_{\mu_{i}}\left(p_{i}\right) \mathrm{d} \lambda^{\left|Y_{i}\right|-1}\left(p_{i}\right) \\
& =\int_{Q_{i}} \prod_{\ell \in[1 . . t]} q_{i}\left(z_{i, \ell}^{t}\right) \int_{\pi_{i}^{-1}\left(q_{i}\right)} u_{i}\left(a_{i}, p_{i}\right) g_{\mu_{i}}\left(p_{i}\right) J_{B_{i}}^{-1 / 2} \mathrm{~d} \lambda^{\left|Y_{i}\right|-\left|Z_{i}\right|}\left(p_{i}\right) \mathrm{d} \lambda^{\left|Z_{i}\right|-1}\left(q_{i}\right) .
\end{aligned}
$$

[^20]Define

$$
u_{i}\left(a_{i}, q_{i}\right):=\int_{\pi_{i}^{-1}\left(q_{i}\right)} u_{i}\left(a_{i}, p_{i}\right) g_{\mu_{i}}\left(p_{i}\right) \mathrm{d} \lambda^{\left|Y_{i}\right|-\left|Z_{i}\right|}\left(p_{i}\right) J_{B_{i}}^{-1 / 2} .
$$

We will now prove the following:
Lemma 11. $u_{i}\left(a_{i}, q_{i}\right)$ is continuous in $q_{i}$.
Proof. We first show that $\pi_{i}^{-1}$ is a continuous correspondence.
Let $K\left(P_{i}\right)$ denote the set of nonempty, compact, and convex subsets of $P_{i}$. Take any $q_{i, n} \rightarrow q_{i}$ and $p_{i, n} \in \pi_{i}^{-1}\left(q_{i, n}\right)$ converging to $p_{i}$ and note that $B_{i} p_{i}=\lim _{n} B_{i} p_{i, n}=\lim _{n} q_{i, n}=q_{i}$, and thus $\pi_{i}^{-1}$ is upper-hemicontinuous (uhc).

To show that it is lower-hemicontinuous (lhc) take any open set $U \subseteq P_{i}$ such that $U \cap \pi_{i}^{-1}\left(q_{i}\right) \neq \varnothing$. This implies that there is $p_{i} \in \operatorname{int}\left(U \cap \pi_{i}^{-1}\left(q_{i}\right)\right)$ and $\epsilon>0$ such that $B_{\epsilon}\left(p_{i}\right) \subset U \cap \pi_{i}^{-1}\left(q_{i}\right)$. As $\pi_{i}$ is a linear mapping, then by the open mapping theorem (Rudin, 1973, Theorem 2.11), $\pi_{i}\left(B_{\epsilon}\left(p_{i}\right)\right)$ is open and, therefore, $\exists \delta>0$ such that $B_{\delta}\left(q_{i}\right) \subseteq \pi_{i}\left(B_{\epsilon}\left(p_{i}\right)\right)$. Consequently, $\forall q_{i}^{\prime} \in B_{\delta}\left(q_{i}\right), \varnothing \neq$ $\pi_{i}^{-1}\left(q_{i}^{\prime}\right) \cap B_{\epsilon}\left(p_{i}\right) \subseteq \pi_{i}^{-1}\left(q_{i}^{\prime}\right) \cap U$.
We highlight that $\pi_{i}^{-1}$ is not only continuous, but also compact- and convex-valued correspondence, from $Q_{i}$ to $P_{i}$. When restricted to $\pi_{i}\left(P_{i}\right)$, it is also nonempty, and thus then $\pi_{i}^{-1}: \pi_{i}\left(P_{i}\right) \rightarrow$ $K\left(P_{i}\right)$ is continuous with respect to the Hausdorff metric (see Aliprantis and Border, 2006, Theorem 7.15).

Let $h: P_{i} \rightarrow \pi_{i}^{-1}\left(q_{i}\right)$ be such that $h\left(p_{i}\right):=\operatorname{argmin}_{p_{i}^{\prime} \in \pi_{i}^{-1}\left(q_{i}\right)}\left\|p_{i}-p_{i}^{\prime}\right\|_{\infty}$. By continuity of $\pi_{i}^{-1}$, for any $\epsilon$, there is an $N$ such that $\forall n \geq N, \pi_{i}^{-1}\left(q_{i, n}\right) \subset B_{\epsilon}\left(\pi_{i}^{-1}\left(q_{i}\right)\right)$. Hence, for large $n$, for any point in $\pi_{i}^{-1}\left(q_{i, n}\right)$, there is a point in $\pi_{i}^{-1}\left(q_{i}\right)$ that is at most $\epsilon$ away. As $u\left(a_{i}, p_{i}\right) g_{\mu_{i}}\left(p_{i}\right)$ is continuous in $p_{i}$ and, by Heine-Cantor theorem, uniformly so ( $P_{i}$ is compact). Hence, for any $q_{i}^{\prime}$ close enough to $q_{i}$, the difference in the payoff function will be well-approximated by the the difference in measure (up to a constant scaling factor), $\left|u_{i}\left(a_{i}, q_{i}^{\prime}\right)-u_{i}\left(a_{i}, q_{i}\right)\right| \approx \mid \lambda^{\left|Y_{i}\right|-\left|Z_{i}\right|}\left(\pi_{i}^{-1}\left(q_{i}^{\prime}\right)\right)-$ $\lambda^{\left|Y_{i}\right|-\left|Z_{i}\right|}\left(\pi_{i}^{-1}\left(q_{i}\right)\right) \mid$.

In the sequel, we show that the measure is continuous in $q_{i}$. Take any sequence $\left(q_{i, n}\right)_{n} \subseteq \pi_{i}\left(P_{i}\right)$ (a compact set) satisfying $q_{i, n} \rightarrow q_{i}$. As $\pi_{i}^{-1}$ is continuous, and, in particular, upper hemicontinuous,

$$
\limsup _{n} \lambda^{\left|Y_{i}\right|-\left|Z_{i}\right|}\left(\pi_{i}^{-1}\left(q_{i, n}\right)\right) \leq \lambda^{\left|Y_{i}\right|-\left|Z_{i}\right|}\left(\pi_{i}^{-1}\left(q_{i}\right)\right),
$$

since for any open set containing $\pi_{i}^{-1}\left(q_{i}\right)$, it will contain $\pi_{i}^{-1}\left(q_{i, n}\right)$ for $q_{i, n}$ sufficiently close to $q_{i}$.

We now argue that the above inequality holds with equality. To see this, fix an arbitrary $\epsilon>0$ and take a collection of points on the boundary of $\pi_{i}^{-1}(q)$ such that any two points are not closer than $\epsilon / 2$ and not farther away than $\epsilon>0$. This implies that we have a finite collection of such points. As $\pi_{i}^{-1}$ is lhc, there is a $\delta>0$ such that for every $q_{i}^{\prime} \in B_{\delta}\left(q_{i}\right), \pi_{i}^{-1}\left(q_{i}^{\prime}\right)$ contains a point in an $\epsilon / 4$ neighborhood of every point in our collection, and by convexity, their convex hull. This implies that we can approximate arbitrarily well the interior of $\pi^{-1}\left(q_{i}\right)$ taking any $n$ sufficiently large; i.e. there is a $\gamma(\epsilon)>0$ such that $\lambda^{\left|Y_{i}\right|-\left|Z_{i}\right|}\left(\pi_{i}^{-1}\left(q_{i}\right)\right)-\gamma(\epsilon) \leq \lambda^{\left|Y_{i}\right|-\left|Z_{i}\right|}\left(\pi_{i}^{-1}\left(q_{i}^{\prime}\right)\right)$, with $\gamma(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$.

Hence, as $q_{i, n} \rightarrow q_{i}$, for any $q_{i, n} \neq q_{i}, \lambda^{\left|Y_{i}\right|-\left|Z_{i}\right|}\left(\pi_{i}^{-1}\left(q_{i, n}\right)\right) \rightarrow \lambda^{\left|Y_{i}\right|-\left|Z_{i}\right|}\left(\pi_{i}^{-1}\left(q_{i}\right)\right)$.
We can then redefine the problem by considering $v_{i}$ to be uniform on $Q_{i}$ and take $u_{i}\left(a_{i}, q_{i}\right)$ as the utility function.

Redefining

- $v_{i}: \Delta\left(Q_{i}\right) \rightarrow \mathbb{R}$ as $v_{i}\left(v_{i}^{\prime}\right):=\max _{a_{i} \in A_{i}} \mathbb{E}_{v_{i}}\left[u_{i}\left(a_{i}, q_{i}\right)\right] ;$
- $V_{i}: \Delta\left(Q_{i}\right) \rightarrow \mathbb{R}$ as $V_{i}\left(v_{i}\right):=\sup _{\tau^{\prime} \in \mathbb{T}_{i}} \mathbb{E}_{v_{i}}\left[v_{i}\left(v_{i} \mid z_{i}^{\tau^{\prime}}-c_{i} \tau_{i}\right] ;\right.$
- $\tau_{i}(\omega):=\inf \left\{t \mid V_{i}\left(v_{i} \mid z_{i}^{t}(\omega)\right)=v_{i}\left(v_{i} \mid z_{i}^{t}(\omega)\right)\right\} ;$
- $b_{i}\left(p_{i}\right):=\mathbb{E}_{\pi_{i}\left(p_{i}\right)}\left[\sigma_{i}^{*}\left(z_{i}^{\tau_{i}}\right)\right]$, for some fixed $z_{i}^{t} \mapsto \sigma_{i}^{*}\left(z_{i}^{t}\right) \in \operatorname{argmax}_{\sigma_{i} \in \Delta\left(A_{i}\right)} \mathbb{E}_{v_{i}}\left[u_{i}\left(a_{i}, q_{i}\right) \mid z_{i}^{t}\right]$.
we obtain-by analogous arguments to Proposition 1-that $V_{i}$ is continuous and, by Berk (1966), that $v_{i, t}$ weak $^{*}$ converges to $B_{i} p_{i}, p_{i}$-a.s., when $z_{i, t} \sim \pi_{i}\left(p_{i}\right)$, for all $t$. Therefore, $\tau_{i}$ is finite $p_{i}$-a.s., for any $p_{i} \in P_{i}$. Finally, an analogous version of Lemma 1 applies and $b_{i}$ is continuous in $p_{i}$ and maps to $\Delta\left(A_{i}\right)$. Hence, by essentially the same arguments as in Theorem 1, a sequential sampling equilibrium exists.

In the above, we restricted to the case in which $\left|Z_{i}\right|<\left|Y_{i}\right|$. If instead $\operatorname{rank}\left(B_{i}\right)=\left|Z_{i}\right|=\left|Y_{i}\right|$, we have the following:

Proposition 8. Let $G:=\langle\Gamma, \mu, c, B, Z\rangle$ be an extended game such that for every player $i, \operatorname{rank}\left(B_{i}\right)=$ $\left|Z_{i}\right|=\left|A_{-i}\right|$. Then $G$ admits a sequential sampling equilibrium.

Proof. Note that now $B_{i}$ is invertible and $\pi_{i}\left(p_{i}\right):=B_{i} p_{i}$ is bijective when restricting its domain to $\pi_{i}\left(P_{i}\right)$. Hence, $\pi_{i}$ admits a continuous inverse (note it is a linear mapping). Since for any $p_{i} \in \Delta\left(A_{-i}\right)$, and $\mu_{i}$ has full support, $\pi_{i}\left(p_{i}\right)$ is in the support of the pushforward measure $v_{i}:=$ $\pi_{i} \# \mu_{i} \in \Delta\left(Q_{i}\right)$. Thus, for any $p_{i}$, by Berk (1966), $v_{i, t}$ weak $^{*}$ converges to a Dirac on $\pi_{i}\left(p_{i}\right)$. Moreover, if $v_{i} \in \Delta\left(Q_{i}\right)$ is the pushforward measure given $\mu_{i} \in \Delta\left(P_{i}\right)$ and $\pi_{i}: P_{i} \rightarrow Q_{i}$, we now
also have that $\mu_{i}$ is $v_{i}$ 's pushforward measure given $\pi_{i}^{-1}: Q_{i} \rightarrow P_{i}$. Then, weak ${ }^{*}$ convergence of $v_{i, t}$ to $\delta_{q_{i}}$ implies weak* convergence of the $\mu_{i, t}$ to $\delta_{\pi_{i}^{-1}\left(q_{i}\right)}$. Uniform continuity of $V_{i}: \Delta\left(P_{i}\right) \rightarrow \mathbb{R}$ (as originally defined on the main text) delivers existence of an equilibrium as in Theorem 1.

Naturally, all the above can also be extended to Bayesian games.

## B.2. Relation to Analogy-Based Expectation Equilibrium

Finally, we discuss convergence to analogy-based expectation equilibria-see Jehiel (2021) for a survey. In line with the literature, we consider payoffs that are linear in the distribution of actions. This solution concept when applied to normal-form games (including Bayesian games) can be readily translate to our setup: $\sigma$ is an analogy-based expectation equilibrium if, for each player $i \in I$,
(1) $q_{i}\left(z_{i}\right)=\sum_{a_{-i} \in f_{i}^{-1}\left(z_{i}\right)} \sigma_{-i}\left(a_{-i}\right)$ for every $z_{i} \in Z_{i}$;
(2) $u_{i}\left(a_{i}, z_{i}\right):=\sum_{a_{-i} \in f_{i}^{-1}\left(z_{i}\right)} \frac{1}{\left|f_{i}^{-1}\left(z_{i}\right)\right|} u_{i}\left(a_{i}, a_{-i}\right)$; and
(3) $\sigma_{i} \in \arg \max _{\sigma_{i}^{\prime} \in \Delta\left(A_{i}\right)} \mathbb{E}_{q_{i}} u_{i}\left(a_{i}, z_{i}\right)$.

Recalling that $f_{i}$ is a surjective mapping from $A_{-i}$ to $Z_{i}$, condition (1) states that each player $i$ bundles difference action profiles (or players, or types, contingencies) $a_{-i}$ into the same 'analogy class' $z_{i}$. Condition (2) can be read as a simplification device by player $i$ : the player cannot distinguish across the different $a_{-i}$ within the same analogy class $z_{i}$, they consider the average behavior, as if the probability of each $a_{-i}$ within the same analogy class were the same. Then (condition (3)), they best respond to the expected payoff given the actual distribution over analogy classes, but assuming that, within the analogy class, distribution over contingencies is uniform.

In the above setup, this is achieved whenever $\mu_{i}$ is uniform. The result follows from arguments analogous to those in the proof of Theorem 4, but with a crucial simplification: the linearity of payoffs in distributions, the posterior means pin-down the set of best responses. And, upon stopping, $\mathbb{E}_{\mu_{i}}\left[q_{i}^{\prime} \mid z_{i}^{\tau_{i}}\right]=\frac{\tau_{i}}{\tau_{i}+\left|A_{-i}\right|} \bar{z}_{i}^{\tau_{i}}+\frac{1}{\tau_{i}+\left|A_{-i}\right|}$. Since, for any player $i$ for which no action is always a best response, there is a lower bound to the stopping time $\tau_{i} \geq T^{n}$, that grows unboundedly as sampling costs vanish, $c_{i}^{n} \rightarrow 0$ (Lemma 4), by a similar application of the optional stopping theorem as in the proof of Theorem $4, \mathbb{E}_{q_{i}}\left[\left\|\mathbb{E}_{\mu_{i}}\left[q_{i}^{\prime} \mid z_{i}^{\tau_{i}^{m}}\right]-q_{i}\right\|\right] \leq \mathbb{E}_{q_{i}}\left[\left\|\mathbb{E}_{\mu_{i}}\left[q_{i}^{\prime} \mid z_{i}^{T^{m}}\right]-q_{i}\right\|\right] \leq$ $\frac{T^{m}}{T^{m}+\left|A_{-i}\right|} \mathbb{E}_{q_{i}}\left[\left\|\overline{z_{i}} T^{m}-q_{i}\right\|\right]+2 \frac{1}{T^{m}+\left|A_{-i}\right|}$. Then, by the law of large numbers, $\mathbb{E}_{q_{i}}\left[\left\|\overline{z_{i}} \bar{T}^{m}-q_{i}\right\|\right] \rightarrow 0$. This provides a shorter route to show that players' posterior means converge to the underlying true distribution.

## Appendix C. Experimental Data and Analysis Details

In this appendix, we provide details on the experimental data used along with additional analysis. A total of 164 subjects were recruited for sessions run in the Columbia Experimental Laboratory in the Social Sciences (CELSS) to play matching pennies games as the one depicted in Figure 1. Subjects are randomly and anonymously matched, but their roles are fixed throughout. Player $M$ 's payoff to action $a, \delta_{M}$, took one of six values (here rescaled by a factor of 20 for convenience): $4,2,1 / 2,1 / 4,1 / 10$, and $1 / 20$.

The experiment consisted of two stages. In the first stage, actions elicited and each game is played twice. In the second stage, each game is played 5 times and either both actions and beliefs about the probability that the opponent chooses action $a$ are elicited or only actions are elicited. Beliefs here refer to point estimates reported by the subjects, neglecting any strategic uncertainty. In other words, belief reports would correspond to posterior means in our framework. Elicitation of actions and beliefs is incentive-compatible and robust to risk attitudes and game payoffs correspond to probability points towards prizes of $\$ 10$. Throughout, no feedback was provided, game order was randomized and, importantly for our purposes, decision times are recorded. Other details on the experimental design can be found in Friedman and Ward (2022).

There are some important caveats to note. First, beliefs elicited in the second stage refer to opponent's actions from the first stage. This, together with the fact that elicitation of actions and beliefs is sequential instead of simultaneous, with beliefs being elicited first, may raise concerns of whether reported beliefs are a good proxy for the beliefs that subjects hold when taking an action. Second, while decision time was recorded, subjects were forced to wait a minimum of 10 seconds before reporting their beliefs. As the subjects' decision times will be used as a proxy to test sequential sampling equilibrium's predictions for stopping times, the forced minimum decision time may undermine the exercise. Finally, the authors highlight there being evidence of "no-feedback learning" as the same subject plays the same game multiple times. This is especially worrying when comparing instances where only actions are elicited with those where both actions and beliefs are. In order to avoid issues due to experience or learning, and focus on initial response as much as possible, we restrict attention to choice data when beliefs are not elicited.

Table 2 documents the own- and opponent-payoff choice effects mentioned on Section 3.2: as player $M$ 's payoffs to action $a$ increase, subjects in that tend to choose the action more often and action $b$ less often, while the opposite is true for subjects in the opponent's role, player $C$. In

|  | Player $M$ chooses $a$ |  |  | Player $C$ chooses $a$ |  |
| :--- | :---: | :---: | :---: | :---: | :---: |
|  | OLS | Logit |  | OLS | Logit |
|  | $(1)$ | $(2)$ |  | $(3)$ | $(4)$ |
| $\delta_{M}$ | $0.230^{* * *}$ | $0.949^{* * *}$ |  | $-0.772^{* * *}$ | $-3.430^{* * *}$ |
|  | $(0.041)$ | $(0.169)$ |  | $(0.036)$ | $(0.197)$ |
| Intercept | $0.329^{* * *}$ | $-0.702^{* * *}$ |  | $0.842^{* * *}$ | $1.522^{* * *}$ |
|  | $(0.018)$ | $(0.079)$ |  | $(0.017)$ | $(0.090)$ |
| (Pseudo) R | 0.02 | 0.01 |  | 0.20 | 0.15 |
| Observations | 1782 | 1782 |  | 1806 | 1806 |

Heteroskedasticity-robust standard errors in parentheses.
${ }^{*} p<0.1,{ }^{* *} p<0.05,{ }^{* * *} p<0.01$.
Table 2: Payoffs and Choices: Own- and Opponent-Payoff Choice Effect
Notes: The table exhibits the association between player $M$ 's payoff to action $a$ and the frequency with which subjects in each role choose action $a . \delta_{M}$ parametrizes player $M$ 's payoff to action $a$. The games in question are generalized matching pennies games as given in Figure 1, for $\gamma_{C}=1$ (and scaled by 20). The sample includes data from rounds in which there is no belief elicitation. The data is from Friedman and Ward (2022).

| Player $C$ Beliefs |  |  |
| :---: | :---: | :---: |
| $\sigma_{M}^{\tau_{C}}\left\|\delta_{M}^{\text {High }} \geq_{F O S D} \sigma_{M}^{\tau_{C}}\right\| \delta_{M}^{\text {Low }}$ |  |  |
| High | Low | KS-Statistic |
| $(1)$ | $(2)$ | $(3)$ |
| 4 | 2 | $0.33^{* * *}$ |
| 2 | $1 / 2$ | $0.76^{* * *}$ |
| $1 / 2$ | $1 / 4$ | $0.40^{* * *}$ |
| $1 / 4$ | $1 / 10$ | $0.23^{* * *}$ |
| $1 / 10$ | $1 / 20$ | $0.23^{* * *}$ |
| ${ }^{*} p<0.1,{ }^{* *} p<0.05,{ }^{* * *} p<0.01$. |  |  |

Table 3: Opponent Payoff and Beliefs: FOSD Tests
Notes: The table exhibits the results of two-sample Kolmogorov-Smirnov tests for first-order stochastic dominance of the distribution of reported beliefs by subjects in the role of player $C$ in games with different values of $\delta_{M}$, High (col (1)) and Low (col (2)). Column (3) presents the test statistic, with number of observations $(n, m)=(280,280) . \delta_{M}$ parametrizes player $M$ 's payoff to action $a$. The games in question are generalized matching pennies games as given in Figure 1, for $\gamma_{C}=1$ (and scaled by 20). The data is from Friedman and Ward (2022).
the main text, Table 1 provides support for the collapsing boundaries result presented in Section 3.3: a negative association between the distance of subjects' reported beliefs to their indifference point and the decision time. Related to this result, the main text discussed a first-order stochastic
dominance shift of beliefs for player $C$ as $\delta_{M}$ (player $M$ 's payoffs to action $a$ ) increases. Figure 4 exhibits such lawful relation; appropriate statistical tests are now given in Table 3.

## Appendix D. Misspecified Priors

As in finite dimensional spaces, the Bayesian learning is consistent for any distribution if and only if the prior has full support (Freedman, 1963), Proposition 1 uncovers an important consequence of Bayesian learning for optimal stopping: Not only is the decision-makers' optimal stopping time finite with probability one, for any true distribution of their samples, it is also bounded uniformly across all distributions of samples. This effectively transforms the optimal stopping problem from infinite to finite horizon, allowing for a solution to be obtained by backward recursion, simplifying the problem significantly.

The intuition underlying the result is that if the prior has full support, the posterior accumulates around the empirical mean. Then, one can guarantee a bound on the rate at which the posterior accumulates around the empirical mean, depending on the number of observations but not on the sample path itself (Diaconis and Freedman, 1990). With this, it is possible to bound the gains in expected payoff of sampling further regardless of the realized sample path and show that there is a number of observations after which the cost of an additional observation dwarfs the expected gain, regardless of realizations. Hence, one concludes that the decision-maker necessarily stops after such number of samples and we can find an explicit upper bound for the stopping time that depends only on the prior $\mu_{i}$, payoffs $u_{i}$, and sampling cost $c_{i}$.

This stands in contrast to the canonical problem in Arrow et al. (1949) where the prior has finite support, and optimal stopping time is not bounded. ${ }^{30}$ Further, it stands in sharp contrast to the case in which beliefs are misspecified.

We now provide an example in which misspecification leads to a player never stopping with probability 1 (with respect to the true distribution of opponents' actions), and sequential sampling equilibrium fails to exist.

Let $\Gamma$ be a two-player game in which player $i$ 's opponent has three possible actions, $a, b$, and $c$, and always chooses $c$ (e.g. because $c$ is dominant, or because their sampling cost is too high

[^21]and $c$ is uniquely optimal under their prior). Denote $\sigma_{-i}=\left(\sigma_{-i}(a), \sigma_{-i}(b), \sigma_{-i}(c)\right) \in \Delta(\{a, b, c\})$. Suppose player $i$ 's prior beliefs about $\sigma_{-i}, \mu_{i}$, are such that player $i$ assigns probability $1 / 2$ to $(1 / 2,1 / 6,1 / 3)$ and probability $1 / 2$ to $(1 / 6,1 / 2,1 / 3)$. Player $i$ can choose either $a$ or $b$ and player $i$ 's payoffs are given by $u_{i}\left(a_{i}, a_{-i}\right)=1$ if $a_{i}=a_{-i}$, and 0 if otherwise. Then, if $y_{i}^{t}$ is such that $y_{i, \ell}=c$ for all $\ell \in[1 . . t], \mu_{i} \mid y_{i}^{t}=\mu_{i}$. Under their prior, $v_{i}\left(\mu_{i}\right)=1 / 4, v_{i}\left(\mu_{i} \mid a\right)=v_{i}\left(\mu_{i} \mid b\right)=3 / 4$, and $v_{i}\left(\mu_{i} \mid c\right)=v_{i}\left(\mu_{i}\right)=1 / 4$, hence $\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\mu_{i} \mid y_{i}\right)\right]=\frac{2}{3} \frac{3}{4}+\frac{1}{3} \frac{1}{4}=\frac{7}{12}$. Note that, a necessary condition for player $i$ to stop is that $\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\mu_{i} \mid y_{i}\right)\right]-v_{i}\left(\mu_{i}\right) \leq c_{i}$. But, since at any sequential sampling equilibrium player $i$ 's opponent chooses $c$ with probability 1 , we have that $\mu_{i} \mid y_{i}^{t}=\mu_{i}$ and, for any $c_{i}<1 / 3$, we always obtain $\mathbb{E}_{\mu_{i}}\left[v_{i}\left(\mu_{i} \mid y_{i}\right)\right]-v_{i}\left(\mu_{i}\right)=\frac{1}{3}>c_{i}$. Therefore, since $\sigma_{-i}(c)=1$, $\mathbb{P}_{\sigma_{-i}}\left(\tau_{i}=\infty\right)=1$.


[^0]:    * Department of Economics, University College London; duarte.goncalves@ucl.ac.uk. I am very grateful to Yeon-Koo Che, Mark Dean, and Navin Kartik for the continued encouragement and advice. I also thank Larbi Alaoui, César Barilla, Martin Cripps, Tommaso Denti, Teresa Esteban-Casanelles, Evan Friedman, Drew Fudenberg, Philippe Jehiel, Elliot Lipnowski, Qingmin Liu, Antonio Penta, Jacopo Perego, Philip Reny, Ariel Rubinstein, Evan Sadler, Ran Spiegler, Jakub Steiner, Yu Fu Wong, and the audiences at various seminars and conferences for valuable feedback.
    First posted draft: 25 November 2020. This draft: 2 November 2023.

[^1]:    ${ }^{1}$ See e.g. Shadlen and Shohamy (2016), Duncan and Shohamy (2020), and Biderman et al. (2020).
    ${ }^{2}$ Fehr and Rangel (2011), Krajbich et al. (2014), Clithero (2018b), and Spiliopoulos and Ortmann (2018) provide reviews of the existing literature in economics. Forstmann et al. (2016) surveys the literature in cognitive sciences.
    ${ }^{3}$ Early evidence can be found in Mosteller and Nogee (1951); Krajbich et al. (2010), Clithero (2018a), and Alós-Ferrer and Garagnani (2022) provide more recent experimental evidence.

[^2]:    ${ }^{4}$ See Alaoui and Penta (2022) for an axiomatization of the value of stepwise reasoning as the value of sampling information, analogous to our model.

[^3]:    ${ }^{5}$ The classic reference is Ratcliff (1978). See Ratcliff et al. (2016) and Forstmann et al. (2016) for a review of the literature and Krajbich et al. (2012); Spiliopoulos and Ortmann (2018); Clithero (2018b); Chiong et al. (2023) for economic applications.
    ${ }^{6}$ A related literature on optimal sequential information acquisition studies the dynamic choice of information, be it deciding about its intensity (Moscarini and Smith, 1963), selecting across sources (Che and Mierendorff, 2019; Liang et al., 2022), or choosing it in a fully flexible manner (Steiner et al., 2017; Zhong, 2022).

[^4]:    ${ }^{7}$ Related are solution concepts with noisy but unbiased beliefs, e.g. Friedman and Mezzetti (2005); Friedman (2022).

[^5]:    ${ }^{8}$ We will use $\|\cdot\|_{p}$ to denote the $p$-norm and $\|\cdot\|_{\infty}$ for the sup-norm.
    ${ }^{9}$ It is also possible to extend this framework to accommodate other cases potentially of interest, e.g. ruling out opponents play strictly dominated actions; we omit these cases to simplify the presentation.

[^6]:    ${ }^{10}$ Note that $\mu_{i}$ induces a measure on $\Delta\left(A_{-i}\right) \times \mathscr{Y}_{i}$.

[^7]:    ${ }^{11}$ Formally: there is some $N$ and $\underline{c}_{i}>0$ such that player $i$ 's cost for any observation following the $N$-th is greater than $\underline{c}_{i}$.

[^8]:    ${ }^{12}$ See Ratcliff (1978) for a pioneering study of the use of sequential sampling models in cognitive sciences and Forstmann et al. (2016); Ratcliff et al. (2016) for recent review articles.
    ${ }^{13}$ As it is implicit in this statement, even though beliefs are degenerate and coincide on the same Nash equilibrium,

[^9]:    ${ }^{15}$ A notable exception is the non-equilibrium model by Alaoui and Penta (2016), which endogenizes level- $k$ via a cost-benefit analysis. Differently from our model, though, it predicts degenerate stopping times.

[^10]:    ${ }^{16}$ Monotonicity in $\sigma_{-i}$ is automatically satisfied when $u_{i}\left(a_{i}, \sigma_{-i}\right)=\mathbb{E}_{\sigma_{-i}}\left[u_{i}\left(a_{i}, a_{-i}\right)\right]$. We require strict monotonicity to prevent the case in which players are always indifferent between both actions $\left(u_{i}\left(1, \sigma_{-i}\right)=u_{i}\left(0, \sigma_{-i}\right), \forall \sigma_{-i}\right)$, which is a trivial case. Since we will allude to extensions that may require non-linearity in $\sigma_{-i}$, we impose only minimal conditions on payoffs.

[^11]:    ${ }^{17}$ See Proposition 7 in the Appendix A.

[^12]:    ${ }^{18}$ This finding has been replicated several times, namely by Ochs (1995), McKelvey et al. (2000) and Goeree and Holt (2001).
    ${ }^{19} \mathrm{~A}$ similar result holds in my model with respect to symmetric anti-coordination (extended) games. In such case, the unique symmetric sequential sampling equilibrium exhibits the own-payoff effect under the same conditions as in generalized matching pennies. This matches gameplay patterns documented in experimental settings by Chierchia et al. (2018) in the context of symmetric two-player anti-coordination games.

[^13]:    ${ }^{20}$ See Hawkins et al. (2015) for a discussion on the evidence of collapsing boundaries and Bhui (2019) for supporting experimental evidence in an environment in which, as in our model, there is uncertainty about the difference in the binary actions' expected payoffs.

[^14]:    ${ }^{21}$ This is so as to avoid discreteness issues inherent to the sampling procedure.

[^15]:    ${ }^{22}$ This will hold regardless of whether players' prior beliefs allow or not for opponent action correlation.

[^16]:    ${ }^{23}$ We keep fixed a selection of best responses $\sigma_{i}^{*}$ used to break-ties.

[^17]:    ${ }^{24}$ If agents directly sample data with past actions, $\left\{a^{\ell}\right\}_{\ell<n}$, one may worry that about whether sampling without replacement affects the result; this is not the case-provided, of course, the starting dataset large enough (but still finite; cf. Remark 1) so that sequential sampling without replacement is well defined.
    ${ }^{25}$ Classical references are Shapley (1964) and Jordan (1993). Cycling can occur even with stochastic fictitious play: see Hommes and Ochea (2012).
    ${ }^{26} \mathrm{An}$ equilibrium $\sigma$ is asymptotically stable if for all $\epsilon>0$, there is a $\delta>0$ such that for any $\sigma^{0}:\left\|\sigma^{0}-\sigma\right\|_{\infty}<\delta$, $\left\|\sigma^{n}-\sigma\right\|_{\infty}<\epsilon$ for all $n$. That is, if the dynamic sequential sampling process starting close enough to the equilibrium remains closeby thereafter.

[^18]:    ${ }^{27}$ This renders their expected payoff given their type, $\mathbb{E}_{q_{i}}\left[u_{i}\left(a_{i}, a_{-i}, \theta\right) \mid \theta_{i}\right]$, to be continuous in $q_{i} \in \operatorname{supp} \mu_{i}$.

[^19]:    ${ }^{28}$ In order to map theoretical stopping times and empirical decision times, one can use quantiles instead of the values.

[^20]:    ${ }^{29}$ Allowing for $\left|Z_{i}\right|=\left|Y_{i}\right|$ is possible, but makes the proofs more cumbersome.

[^21]:    ${ }^{30}$ Similarly, optimal stopping time is also not bounded in the continuous-time version of the canonical problem, with Gaussian noise, be it with (Moscarini and Smith, 1963) or without experimentation concerns (Chernoff, 1961). In some cases with finite support prior, however, stopping time can be bounded, as in the case with Poisson arrival of conclusive information, but not when the decision-maker can choose from different information sources (Che and Mierendorff, 2019).

