

Dynamic Value Shading*

Hugo Hopenhayn[†] Maryam Saeedi[‡]

UCLA

Tepper, CMU

September 27, 2021

Abstract

This paper examines a class of dynamic decision problems with irreversible actions and its embedding in strategic settings. Agents face random opportunities for taking irreversible actions that, together with a private random shock, determine their final payoff. Information is given by a signal process that arrives throughout the decision period. Our main contribution is methodological; we provide a decomposition of the optimal—and equilibrium—solution into a dynamic component and a static one. The solution to the former problem is independent

*We would like to thank John Asker, Matt Backus, Tom Blake, Alessandro Bonati, Yingni Guo, Ken Hendricks, Ali Hortacsu, Moritz Meyer-ter-Vehn, Paul Milgrom, Ariel Pakes, Michael Peters, Robert Porter, Andrea Prat, Debraj Ray, Ali Shourideh, Ennio Stacchetti, Steve Tadelis, Robert Town, Anton Tsoy, as well as seminar and conference participants at Northwestern Economics Department Theory seminar, Einaudi Institute of Economics and Finance, Midwest Economic Theory Meetings, International Industrial Organization Conference, Canadian Economic Theory Conference, The Economics of Information and Communication Technologies Conference, Society of Economic Dynamics, SITE, Dynamic Games, Contracts, and Markets Session, eBay, Econometrics Society World Congress, UCLA mini conference, Workshop in Industrial Organization (Toi 8), Econometrics Society Winter Meetings, Dartmouth Winter IO Conference, Carnegie Mellon University, Tilburg University, NYU, Stanford University, and UC Berkeley. We would also like to thank Xiang Hui for excellent research assistance. Any remaining errors are ours.

[†]hopen@econ.ucla.edu

[‡]msaeedi@andrews.cmu.edu

of the specific payoff function of the agent beyond some general regularity conditions. For games of incomplete information with privately observed actions, this decomposition reduces the problem of finding equilibrium strategies to a solution of a static Bayesian game. The setting applies to a class of strategic problems, such as some tournaments, entry games, intermediation, and dynamic commitments. Two detailed applications are considered. **Keywords:** Dynamic decision problem, Dynamic games, Value shading, Decision under uncertainty. **JEL codes:** C73, D44, D81, L81

1 Introduction

We consider a class of dynamic decision problems that appears in many economic applications, such as search, entry games, dynamic auctions, R&D races, and other contests, with the following features: First, agents make decisions over a period of time, such as bidding or investing, that result in a random final payoff. Second, throughout the decision period, agents observe signals that are correlated with this final payoff. For example, innovators might learn about the value of obtaining a patent over time, employees might learn about outside opportunities or how much they value current employment, and bidders might learn about outside options or alternative uses of their resources. Third, opportunities for taking actions are random. These random times might represent random opportunities for undertaking actions such as R&D investments, or executing trades, or in the case of auctions could be the result of frictions that might impede the precise timing of bids and information acquisition, or simply inattention.¹ Fourth, decisions are irreversible and can only be increased over time, such as sunk investments in R&D or the impossibility of retracting bids in some auctions.

These conditions make solving dynamic decision problems using standard

¹This is a standard assumption in the class of revision games, developed in [Kamada and Kandori \[2011\]](#) and [Kamada and Kandori \[2015\]](#), as well as in models of sticky prices following [Calvo \[1983\]](#).

methods complicated. Particularly, when there is a need to solve dynamic stochastic decision problems repeatedly, standard methods for solving them can pose a computational burden, as in calculating the best responses in strategic settings or estimating a model’s deep parameters.² In this paper, we introduce a new methodology to solve dynamic problems of the above class under a weak separability assumption on payoffs.

Using our methodology, the problem of finding an optimal strategy can be separated into two parts, one involving the solution to a dynamic problem and the other to a static one. The solution to the first depends only on the properties of the stochastic process for values and decision times, but not on the payoff function. As a result, changes in the payoff function do not require solving repeatedly the more complex dynamic problem. This feature is extremely useful in strategic settings, where the payoff function depends on other players’ actions. As an example, in our application to dynamic auctions, equilibrium dynamic bidding strategies can be obtained from the Bayesian Nash equilibria of a static auction. The aforementioned feature is also useful when estimating parameters in the final payoff function, as changes in these parameters do not require solving the dynamic part of the decision problem repeatedly.

The key element of our methodology is defining a correspondence between expected value at a point in time (e.g., the expected value of winning a patent race, given the information acquired) and what we call its equivalent final value. While in a static setting an agent would take the action that maximizes payoffs given its current value, in our dynamic decision problem the agent chooses an action that is optimal for the equivalent final value, had this been the final choice. The latter is usually lower, a property that we call *value shading*. Values and, correspondingly, actions in the dynamic decision problem are shaded because of the option value of future choices and, under mild condi-

²This is known to be a problem with nested fixed point algorithms, as in [Rust \[1987\]](#). While the indirect methods building on [Hotz and Miller \[1993\]](#) and subsequently the application to games by [Bajari et al. \[2007\]](#) can be used for estimation, many of the interesting counterfactuals still require solving the full model.

tions, the incentives for shading decrease as the decision problem progresses. Thus, shading contributes to delay and underinvestment in the early stages of decision problems.

We now provide a more detailed description of our modeling framework and our solution method. The agent or decision maker faces a random sequence of decision times $\{\tau_n\}$ in a time interval $[0, T]$. At these decision nodes, given current information, the agent chooses a non-decreasing sequence of actions $\{a_n\}$ from a totally ordered set, e.g., level of capital in the presence of irreversible investments. Final payoffs are given by a function $U(v_T, a_T)$, where a_T is the final outstanding action and v_T a real valued random variable. The function U is assumed to be linear in v_T and supermodular in (v_T, a_T) but is otherwise unrestricted.

Information and decision times are modeled as a joint Markov process of signals and decision times $\{v_n, \tau_n\}$ that are also sufficient statistics for the expected value of v_T . Our process is very general, allowing, for instance, the arrival of future decision opportunities to depend on the current expected value (e.g., more attentive response when expected values are higher). Moreover, the inclusion of time as a state variable allows for a non-stationary Markov process in values, as is the case in many Bayesian learning environments. An optimal strategy specifies choices $a_n = S(v_n, \tau_n)$ that maximize the expected value $\mathbb{E}_0 U(v_T, a_T)$.

Here we describe our methodology for solving this class of problems in more detail. The first step of our procedure maps each decision node, (v_n, τ_n) , into an *equivalent final value* \tilde{v} , with the property that the optimal action at \tilde{v} , which maximizes $U(\tilde{v}, a)$, is the same as the optimal action at (v_n, τ_n) . This can be illustrated with the aid of Figure 1. Consider a decision node (v_0, t_0) where the agent chooses an optimal action a_0 . The indifference curve represents all points in the value-time space where the agent would choose the same optimal action a_0 . In particular, this action is optimal at the end of the decision time T , when the agent's value is \tilde{v} .

The indifference curves (e.g., the one depicted in Figure 1) have two main




Figure 1: Indifference curves for actions

Note: All points on this curve result in the same optimal action value, a_0 .

properties, given some general assumptions. First, they are downward sloping; second, they are independent of the particular payoff function and depend on the Markov process that determines the uncertainty about valuations and timing of future opportunities to update actions. To understand the intuition of these two findings, it is best to explain how these indifference curves can be used to predict the actions through the decision problem and how they are determined. Consider the choice a_0 at decision node (v_0, t_0) . This choice remains the final decision of the agent if the agent has no other opportunity to revise its choice or chooses not to do so at any future opportunity. The latter happens when at all future opportunities the pair of value and time lies below the indifference curve passing through the initial decision node (v_0, t_0) . This creates a source of *adverse selection* against the agent's future self, as the current choice applies when the agent chooses not to exercise any future revision options, a choice that is correlated with a lower future value.³ This results in a downward sloping indifference curve because these negative outcomes happen more often when agents are far from the end of the decision time and less often when they are closer to the end time. This downward sloping indifference curve results in players' optimal action to be one in which

³An analogue result is found in [Harris and Holmstrom \[1982\]](#), where initially workers' wages are shaded below marginal products, as the wage is effective in the future only if it is less than or equal to the realized marginal product of the worker.

they act as if their value was lower any time that there is a chance of updating their action in the future, a property which we call value shading. Under fairly general conditions discussed in the paper, value shading decreases over time.

Additionally, we show that the indifference curve has the following self-generating property, which is independent of the particular payoff function. Starting at the point (v_0, t_0) , consider all paths where at future decision nodes (v_n, t_n) the value lies below this indifference curve. Those are the paths where the agent will not increase its initial choice a_0 ; therefore, it will remain the agent's final action. The expected final value conditional on this set of paths is precisely \tilde{v} , and the same property holds for any point on this indifference curve. This property, which we call self-generated expectation, depends only on the stochastic process for values and decision times, and is thus independent of the specific payoff function $U(\cdot)$.⁴ While the final payoff function is relevant in determining action a_0 , which is the solution to $\max_a U(\tilde{v}, a)$, it is not relevant for determining the indifference curves and thus for the mapping from decision nodes (v_n, τ_n) to equivalent values \tilde{v} .

Our results extend to dynamic games with incomplete information and privately observed actions, where our main theorem shows that we can find equilibrium strategies by solving a Bayesian Nash equilibrium of an associated static Bayesian game. The corresponding distribution of values for each player in this Bayesian game is a function of only the joint stochastic process of valuations and decision times independently of the payoff function (e.g., it would be exactly the same for different classes of auctions). The equilibrium strategies in the dynamic game are easily derived from those in the associated Bayesian game. This result extends to dynamic games with observable actions when the associated Bayesian game has an equilibrium in weakly dominating strategies.⁵

⁴We borrow this term from [Abreu et al. \[1990\]](#). While related, it is a different concept. Our indifference curves are self-generating as they define the boundaries on future realizations for the calculation of conditional expected values, which in turn are constant along these curves.

⁵Without this assumption, this result can also apply to finding open loop equilibria, which for large games might approximate closed loop equilibria (see [Fudenberg and Levine](#)

Our model applies to many dynamic decision problems that involve optimal stopping, such as job search or exercising an option, or irreversible investments (see [Dixit et al. \[1994\]](#), [Lippman and McCall \[1976\]](#)). Also, the model can be easily extended to allow for flow payoffs that accrue prior to the end of the game, for example, an electricity company selling future contracts to be fulfilled at some future time (see [Ito and Reguant \[2016\]](#)). Our method can be embedded into strategic settings, such as entry games with a deadline, where the number of entrants is known after the deadline. Other possible strategic settings for our model include dynamic tournaments where over time agents might receive information about their private values of winning the tournament, as in the case of a promotion where agents have outside employment options.

As a more thorough application of our methodology, we consider two detailed examples. The first one is the case of dynamic second price auctions, such as eBay and GovDeals. As these auctions take place over a considerable length of time, dynamic considerations can be important for understanding bidding behavior and improving auction design. Our setting is similar to those in [Kamada and Kandori \[2011\]](#), [Kamada and Kandori \[2015\]](#), and [Kapor and Moroni \[2016\]](#), with the difference that we allow the valuation of participants to change throughout the course of the auction. Our second application is on stationary equilibria of anonymous sequential games, based on [Jovanovic and Rosenthal \[1988\]](#).

The rest of the paper is organized as follows. [Section 2](#) provides a simple example that conveys the main intuition and results in the paper. [Section 3](#) describes the general model and provides a set of applications that can fit the general model. [Section 4](#) discusses the intuitive and formal analysis of the model and describes how to embed the results into games. [Section 4.4](#) describes our main results. [Section 5](#) discusses the extension of the paper to random termination time and gives properties for the case where values are independent of Poisson arrivals for bidding times. [Section 6](#) discusses [\[1988\]](#)). It can also apply to solving for equilibria in mean field games.

two applications of the model in more detail. All proofs are deferred to the appendix unless specified.

2 A Simple Example

We start our analysis by considering a two-period contest that illustrates some of the main features of our methodology. There are N players and two periods, $t = \{0, 1\}$. In the first period, after observing a private signal v_{i0} drawn from some distribution $H_i(v_{i0})$, agents choose the level of a private action $a_{i0} \geq 0$, e.g., studying for a test or allocating resources to a project. In the second period, each agent privately observes its value of winning the contest v_{i1} drawn from the conditional distribution $F_i(\cdot|v_{i0})$. With probability p_i the agent has the option of increasing the action to any value $a_{i1} \geq a_{i0}$, e.g., studying more or allocating additional resources to the project. With probability $(1 - p_i)$, the agent is unable to revise its choice; therefore, a_{i0} remains its final action. The agent with the highest final action, a_i , wins the competition and receives payoff $v_{i1} - a_i$. For any other player j , the payoff is equal to $-a_j$. Both the signals and final values are drawn independently across agents. For notational convenience, we suppress the index i unless needed to avoid confusion.

A player's strategy specifies choices $a_0(v_0)$ and $a_1(v_0, v_1)$ for the first and second period, respectively, with the restriction that $a_1 \geq a_0$. The latter choice is only relevant if the agent has an opportunity to increase its action in the second period. Letting G denote the distribution for the highest final action of the other players, an agent's expected utility given final value v and action a is

$$U(v, a) = G(a)v - a. \tag{1}$$

Assume there is a unique action a that maximizes (1) and it is strictly increasing in v . Denote this solution by $S(v)$. This is the optimal action in a static setting.⁶ Given the action a_0 in the first period, there is a unique threshold \tilde{v}

⁶We make these assumptions and others below in the analysis of the example for conve-




Figure 2: Decision tree

such that $S(\tilde{v}) = a_0$. We can use this information to illustrate the tree of the game in Figure 2.

Figure 2 depicts the choices made by the agent in the two-period game. The top branch represents the case where there is no opportunity for revising the first-period choice, so a_0 is the final action. In the second branch, the player would like to choose final action $a < a_0$, but due to the irreversibility condition the final choice is kept at a_0 . Note that the highest value of v_1 that belongs to this branch is equal to \tilde{v} , which, as defined, has the property that $S(\tilde{v}) = a_0$. The bottom and third branches represent values above \tilde{v} where the player will increase the action in the second period to $S(v_1)$.

Considering the best response in the second period, the choice of a_0 maximizes

$$\begin{aligned} \mathbb{E}U_i &= (1-p)U(\mathbb{E}(v_1|v_0), a_0) \\ &+ p \int_{\tilde{v}}^{\tilde{v}} U(v_1, a_0) dF(v_1|v_0) + p \int_{\tilde{v}}^{\infty} U(v_1, S(v_1)) dF(v_1|v_0). \end{aligned}$$

Assuming that G is differentiable, using the envelope theorem, the associated

nience. The set of assumptions that are needed for our main results are given in Sections 3 and 4.

first-order condition simplifies to

$$(1 - p) [G' (a_0) \mathbb{E} (v_1|v_0) - 1] + p \int^{\tilde{v}(a_0)} [G' (a_0) v_1 - 1] dF (v_1|v_0) = 0. \quad (2)$$

Since $S (\tilde{v}) = a_0$, the second-period first-order condition $G' (a_0) \tilde{v} - 1 = 0$ holds too. Substituting for $G' (a_0)$ in (2) gives

$$(1 - p) \left[\frac{\mathbb{E} (v_1|v_0)}{\tilde{v}} - 1 \right] + p \int^{\tilde{v}} \left[\frac{v_1}{\tilde{v}} - 1 \right] dF (v_1|v_0) = 0. \quad (3)$$

This equation defines implicitly \tilde{v} as a function of v_0 only, independently of the other players' strategies. It can be more conveniently rewritten as

$$\tilde{v} = \frac{(1 - p) \mathbb{E} (v_1|v_0) + p \int^{\tilde{v}} v_1 dF (v_1|v_0)}{(1 - p) + pF (\tilde{v}|v_0)}. \quad (4)$$

We call this value \tilde{v} , the equivalent final value of v_0 . The agent makes the same choice in the first period when the agent's value is v_0 as if it were in the final period confronted with value \tilde{v} .

To interpret this relationship, note that the threshold \tilde{v} defines a lottery over final values v under which a_0 will also be the final action of the agent, comprising the following events:

1. The agent does not have an opportunity to revise its first-period choice. This event has probability $(1 - p)$ and expected value $\mathbb{E} (v_1|v_0)$.
2. The agent is able to revise its first-period choice but its final value is less than the threshold \tilde{v} , so the agent would maintain its initial choice. This event has probability $pF (\tilde{v})$ and expected value $\frac{\int^{\tilde{v}} v_1 dF (v_1|v_0)}{F (\tilde{v})}$.

The lottery over these final values has an expected value as given in equation 4, which is equal to \tilde{v} . This is the key property defining the equivalent final value and it holds under a wide class of payoff functions. Let $u (v, a_i, a_{-i})$ denote the final payoff to agent i when its final value is v and vector of final actions

(a_i, a_{-i}) . Assume the expected utility is linear in v and (strictly) supermodular in v, a_i . The former guarantees that $\mathbb{E}_v u(v, a_i, a_{-i}) = u(\mathbb{E}_v v, a_i, a_{-i})$, and the latter guarantees that the optimal choice a_i is a strictly increasing function of v . These two properties are preserved when integrating out the actions of other players with respect to any distribution $G(a_{-i})$. As before, let

$$S(v) = \operatorname{argmax} \mathbb{E}_{a_{-i}} U(v, a_i, a_{-i})$$

denote the optimal strategy for agent i in the final period when faced with a distribution G_{-i} for the strategies of the other players, and let (v_0, a_0) denote the value and optimal strategy of a player in the first period. Given a_0 , the optimal threshold \tilde{v} for increasing this action in the second period will be such that

$$S(\tilde{v}) = a_0. \tag{5}$$

The threshold also defines a lottery over final values v under which a_0 will be the final action of the agent, comprised of the two sets of events defined above, with expected value

$$\frac{(1-p) \mathbb{E}(v_1|v_0) + p \int^{\tilde{v}} v_1 dF(v_1|v_0)}{1-p + pF(\tilde{v})}.$$

Because of the linearity of payoffs in v ,

$$a_0 = S\left(\frac{(1-p) \mathbb{E}(v_1|v_0) + p \int^{\tilde{v}} v_1 dF(v_1|v_0)}{1-p + pF(\tilde{v})}\right). \tag{6}$$

Using (5) and (6) and given that B is strictly increasing, we get to the same relationship as the one in equation 3. Therefore, mapping between equivalent final value \tilde{v} and the initial value v_0 is independent of the specific strategy function S and thus the underlying payoff function U and the distribution of other players' actions.

As suggested in the example, the partition of the space of values into expected-value equivalent pairs (v_0, \tilde{v}) can be used to reduce the dynamic game to

an equivalent static one. Starting with an initial distribution $F_0(v_0)$ and a conditional distribution $F(v_1|v_0)$, we can construct a new distribution of final values as follows. For any initial v_0 , assign a value $\tilde{v}(v_0)$ to the histories where either the corresponding agent does not have a revision opportunity in the second period or gets a value $v_1 \leq \tilde{v}$. In the complement (i.e., histories where the agent can review its choice and $v_1 > \tilde{v}$), set the final value equal to v_1 . Assigning the corresponding probabilities for these histories as determined from F_0 , F , and the review probability p defines a distribution for final values \tilde{F} for each player and thus a static Bayesian game. Letting \tilde{S} denote an equilibrium strategy for the agent in that game, we can now assign $a_0(v_0) = \tilde{S}(\tilde{v}(v_0))$ and $a(v_0, v_1) = \max\{a_0(v_0), \tilde{S}(v_1)\}$ as equilibrium strategies in the dynamic game.

We now consider a key property of equilibria in this class of games. The opportunity of modifying the action in the future introduces an option. From equation (4) it follows immediately that $\tilde{v} < \mathbb{E}(v_1|v_0)$, so the agent in the first period acts as if the final value were lower than the agent's conditional expectation; this is what we call *value shading*. The fact that actions are monotonic in values also results in the shading of actions below the optimal ones. This becomes more severe as the probability p increases, and in the limit when $p \rightarrow 1$, $F(\tilde{v}|v_0) \rightarrow 0$, i.e., the agent acts in the first period as if the value were the lowest in the support. In the other extreme, when $p \rightarrow 0$, $\tilde{v} = \mathbb{E}(v_1|v_0)$, so there is no shading. The intuition for these results goes back to our description of the two sets of events where the action chosen in the first period is the final one. The first event, when the agent has no future opportunity of increasing its initial action, has expected value $\mathbb{E}(v_1|v_0)$. It is the second event, where the agent has this opportunity but chooses not to increase its initial action, that is responsible for shading. Thus, the irreversibility of actions and the opportunity for delay create a negative option value in the first period. This value can also be interpreted as adverse selection against the agent's future self which is responsible for value shading.

3 The Decision Problem

We first consider the general structure of a dynamic decision problem. Then we show that it can be embedded in a wide class of dynamic games as well. Time is continuous in the interval $[0, T]$. Decision times τ_0, τ_1, \dots are random according to a process that is detailed below. At these decision times the agent can choose an action a_τ (e.g., capital) from a totally ordered set A , with the restriction that for $\tau' > \tau$, $a_{\tau'} \geq a_\tau$. This restriction captures the irreversible nature of actions. Letting a_T denote the final action, payoffs are given by a function $U(v_T, a_T)$, where v_T is a bounded real valued random variable in a probability space $(\Omega, \mathcal{F}, \Pi)$.

Assumption 1. *The payoff function $U(v, a)$ is linear in a monotone function of v , supermodular in v and a , and admits a maximum with respect to a for all v .*

Information arrival and decision times are modeled as joint stochastic processes on $[0, T]$ as follows. Decision times are given by sequences of stopping times $\{\tau_n(\omega)\}$, where $\tau_{n+1}(\omega) > \tau_n(\omega)$. Information arrival is modeled by a stochastic process $\tilde{v}(t, \omega)$ of signals with the property that $\mathbb{E}(v_T | \tilde{v}(t, \omega) = v) = v$. More formally, let $\{\mathcal{F}_t\}_{0 \leq t \leq T}$ be a filtration representing information available at time t , i.e., increasing σ -algebras on Ω with the property that $\mathcal{F}_t \subset \mathcal{F}_{t+s} \subset \mathcal{F}$. As in the case of a Poisson process, the stopping times $\{\tau_n(\omega)\}$ are modeled as the jumps of a right continuous counting process $\{\eta(t, \omega)\}$. Without loss of generality, we assume that $\{\mathcal{F}_t\}$ is the filtration generated by the pair of stochastic processes $\{\eta(t, \omega), \tilde{v}(t, \omega)\}$ so that the realization of these processes is all the information available at time t and $\mathbb{E}(v_T | \mathcal{F}_t) = \tilde{v}(t, \omega)$.

Since information arrivals are relevant only at decision nodes, we restrict attention to the joint process $\{v_n, \tau_n\}$ where $v_n(\omega) = v(\tau_n(\omega), \omega)$, i.e., the process $v(t, \omega)$ subordinated to the arrival process $\eta(t, \omega)$. We make the following assumption about this process.

Assumption 2. Assume that $\{v_n, \tau_n\}$ follows a joint Markov process, i.e.,

$$P(v_{n+1} = v', \tau_{n+1} = \tau' | \mathcal{F}_{\tau_n}) = P(v_{n+1} = v', \tau_{n+1} = \tau' | v_n, \tau_n).$$

By assuming that values and decision times are Markov, we can identify decision nodes with pairs (v_n, τ_n) corresponding to the realized signal and time in the last arrival. A decision strategy s specifies at each possible decision node a desired action $s(v_n, \tau_n)$, which is the choice the agent would make if unconstrained by past actions. Given that actions can only be increased, $a(s, t) = \max\{s(v_n, \tau_n) | \tau_n \leq t\}$ is the choice that prevails at time t and, in particular, $a(s, T)$ is the final choice. Let \mathbf{S} denote the set of strategies satisfying these conditions.

While decision times are exogenous, our specification is flexible; in particular, it allows for decision times and expected values to be correlated. This specification could capture, for example, a situation where an agent might be more eager to revise its strategy when there is a large information update, or, likewise, the agent might be more attentive when the expected value is high. Moreover, the inclusion of time as a state variable allows for a non-stationary Markov process in values.

Given a strategy $s \in \mathbf{S}$, for each realized path ω we can associate a value $U(v(T, \omega), a(s, T, \omega))$, where $a(s, T, \omega) = \sup\{s(v_n(\omega), \tau_n(\omega)) | \tau_n \leq T\}$. An optimal decision strategy solves

$$\sup_{s \in \mathbf{S}} \mathbb{E}_0 U(v(T), a(s, T)). \quad (7)$$

In Section 4 we provide conditions such that there exists an optimal solution to (7) and develop our method that maps this dynamic problem into an equivalent static one. This method is what makes our structure tractable, facilitating estimation and the analysis of dynamic games. Before getting to the formal analysis, we provide a series of examples that suggest the range of applications of this setup.

3.1 Examples

This general setting embeds various interesting applications of dynamic decision problems. As explained in Section 4.3, under certain conditions, this general setting can be extended to dynamic games as well. Some examples are given below.

Irreversible Investment At random times τ an agent faces an investment opportunity and chooses $i_t \geq 0$ after observing a signal v_t of the final value v_T . The final expected payoff is $v_T R(k_T) - C(k_T)$, where k_T is final cumulative investment, R is total revenue, and C is total cost of investment. This is a direct application of the framework above. It can also be extended to a game where final payoffs depend on total investment k_T of this player and also on the total investment of others.

General Contest and Teamwork The example in Section 2 can be easily generalized. The contest takes place in the interval of time $[0, T]$. Agents can exert effort $e \geq 0$ at random times τ when getting signals v_τ of the final value v_{iT} . Letting a_1, \dots, a_N denote the final cumulative effort of all players, final payoffs have the form $U_i(v_{iT}, a_i, a_{-i})$ satisfying Assumption 1. For example, prizes could depend on the ranking of final efforts as in [Moldovanu and Sela \[2001\]](#). In the case of a team, the functions U_i could be the result of a compensation scheme that depends on a set of signals observed by a principal that are correlated with the vector of final effort choices.

Sequential Trading Commitments At random times, the decision maker is faced with the opportunity of selling at a given price p_τ a quantity of choice q_τ to be delivered at the end of the period. Both arrival time and price are random, following a joint Markov process. Final payoffs are $\sum_\tau p_\tau q_\tau - C(Q)$, where $Q = \sum_\tau q_\tau$ and C is a strictly increasing and convex function. As an example, a utility company might face opportunities to sell forward electricity delivery contracts as in [Ito and Reguant \[2016\]](#). The expected cost

of committing to a larger volume might be convex as more costly energy sources need to be used to fulfill the contracts. Or, a financial trader could sell future contracts that will be fulfilled by resorting to its network of intermediaries, forcing the trader to use more expensive sources in case larger quantities are needed.

While cash flows in the above setting accrue throughout the decision period, the problem can be mapped into the general setting where a payoff $p_t q$ received at time t is equivalent to a random final payoff $p_T q$ when p_t follows a martingale so that $E_t(p_T | p_t) = p_t$. Letting $v = p$, $a = Q$, and $U(v, a) = va - C(a)$ gives the corresponding final payoff function, which is linear in v and supermodular, as required. Alternatively, a monopolist retailer might face random opportunities to buy inventory q_t that will be sold at a final time period at revenue $u(\sum_{\tau} q_{\tau})$. Letting q_t represent negative quantities (interpreted as purchases) and $C(Q) = -u(Q)$ gives the same payoff function as before.⁷

Procrastination in Effort Choice At random times τ an agent chooses effort a_{τ} at cost $c_{\tau} a_{\tau}$. Final payoffs are given by $u(\sum a_{\tau}) - \sum c_{\tau} a_{\tau}$, where u is increasing and concave. Procrastination occurs as an agent might put lower effort in anticipation of the possibility of lower future cost. As a final time T is approached, the incentives for procrastination will decrease. In this example payoffs accrue over time but they can be mapped to final payoffs as in the previous one. This setting can also be embedded in a game where final payoffs depend on the vector of cumulative actions of all players.

Entry Decisions and Search At random times τ , the decision maker gets an opportunity to enter a market and a signal v_{τ} about the expected value of entry. Entry must take place before time T . Final payoffs are $v_T - c$ if the decision maker enters the market and zero otherwise. In this application

⁷These results extend to the case where p_t is not a martingale. This is done by back-loading payoffs as above and redefining payoffs in those histories where there are no further arrivals to pick up the difference $E(v_T | v) - v$. Details of this procedure are available upon request.

the action space is $A = \{0, 1\}$, representing the choice of no entry and entry, respectively. This application can be easily inscribed in an entry game.

Likewise, consider a search environment where an agent must make a decision prior to time T . New offers arrive at random times τ with a value v_τ . The agent may take this offer or continue searching. Assume, as in the standard search environment, that v_τ is sampled from a fixed distribution. Using the method described in the Sequential Trading Commitments example, this problem can be mapped into the structure of our modeling framework. The action space is also $\{0, 1\}$ as in the previous example.

Bidding in Long Auctions In [Hopenhayn and Saeedi \[2020\]](#), we consider a model where a bidder’s value can change over time, capturing the idea that preferences for the object or outside opportunities might change. The bidder can only increase bids over time and there is no retraction of past bids. Examples of these auctions are eBay and GovDeals.⁸ In these auctions, bidders frequently place multiple bids over time and increase them as the auction progresses. To model these auctions in the above class of decision problems, suppose that at random times τ and with an expected final value v , the agent can place (or increase) a bid b_τ . The final expected payoff in the auction will depend on the final value v_T , the final bid b_T of this bidder, and those of others. Integrating over the bids of others, the expected final value has the form $[v_T - \mathbb{E}(b_2 | b_2 \leq b_T)] Prob(b_T \text{ is the highest bid})$. This expected payoff is linear in v_T and supermodular, as required in our general decision problem. This application is examined in more detail in [Section 6](#).

Time Separable Payoffs with Discounting A decision maker has payoffs $u(v, a)$ that are received over time and has a constant discount factor β . Time is discrete, and at each point in time payoffs are functions of a random value v_t and an action a_t as given by function $u(v_t, a_t)$. A decision maker chooses a

⁸GovDeals is an auction platform used by government agencies to sell used equipment.

sequence of non-decreasing and contingent actions a_t to maximize

$$\max_{\{a_t \text{ increasing}\}} \sum_{t=0}^T \beta^t \mathbb{E}u(v_t, a_t),$$

where we assume $u(v, a)$ is linear in v and supermodular. The restriction to increasing actions could capture, for instance, returns from irreversible past investments or cumulative R&D. While this problem does not fit directly in our setting, we exploit time separability of payoffs to provide an equivalent formulation that does so. This is done by treating all payoffs as final with appropriately defined weights. We consider here the case where $T = \infty$, but this case is easily extended to finite or even random T . Let $B = \frac{1}{1-\beta}$ and define $U(v, a) = Bu(v, a)$. Let $P(t+1|t) = \beta$ and $P(t'|t) = 0$ for all $t' > t+1$. As of time zero this implies that the probability of no arrivals is $(1-\beta)$, and that of only n arrivals, $(1-\beta)\beta^n$. Expected final value at time zero is

$$(1-\beta)v_0 + \beta(1-\beta)\mathbb{E}v_1 + (1-\beta)\beta^2\mathbb{E}v_2 + \dots$$

The corresponding final actions are a_0, a_1, \dots and

$$\begin{aligned} \mathbb{E}(U(v, a)) &= (1-\beta)\mathbb{E}U(v_0, a_0) + (1-\beta)\beta\mathbb{E}U(v_1, a_1) + \dots + (1-\beta)\beta^t\mathbb{E}U(v_t, a_t) \\ &= \sum_{t=0}^T \beta^t \mathbb{E}u(v_t, a_t), \end{aligned}$$

so this transformation respects the original payoff structure. While we consider here time zero payoffs, the same procedure applies to any future period. This formulation easily extends to random arrivals and a structure where arrivals and payoffs follow a general joint Markov process. Linear investment costs of the form $i_t = a_t - a_{t-1}$, as would occur in the case of irreversible investment, can be easily accommodated in the above payoff function through the rearrangement and collection of the different a_t terms. An application of time separability to an anonymous sequential game is provided in Section 6.2.

4 Main Results

In this section, we first provide an intuitive analysis of our main findings. Next, we go through the more formal analysis with stating the main theorems. We also show how our analysis for the optimal decision problem can be extended to a large class of dynamic games. The key insight is that we identify a distribution of valuations for each player that is independent of the game and opponents' valuations and strategies. Then we show that the equilibria of the dynamic game correspond one-to-one to the equilibria of a static game with respect to this distribution of values. Finally, we consider the dynamic properties of value shading.

4.1 Intuitive Analysis

The example in Section 2 identified initial values v_0 with a threshold \tilde{v} with the property that for *any* game or decision problem with payoffs that satisfy the given assumptions, the initial action chosen at v_0 equals the optimal final choice at this threshold. This defined a partition of initial and final values into equivalent classes. In the general model where $t \in [0, T]$, a similar representation can be obtained. We can partition the set of value and time pairs (v, t) into indifference classes that can be identified by a final value \tilde{v} which we denote by $e(v, t)$. These have the property that optimal actions are identical for all pairs in an indifference class, as depicted in Figure 1. Moreover, our assumption of supermodularity of payoffs ensures that optimal actions are increasing in the final equivalent value, i.e., in the northeast direction in Figure 1.

These indifference curves can be used to define an agent's optimal strategy over time and in particular the final action chosen. This is illustrated in Figure 3. In the paths shown, the first decision node is (v_0, t_0) where the agent chooses an action a_0 . This is also the final action in the first two panels, where either the agent has no opportunity for future actions or is faced with this opportunity at a decision node (v_1, t_1) in a lower indifference curve. The last two panels of

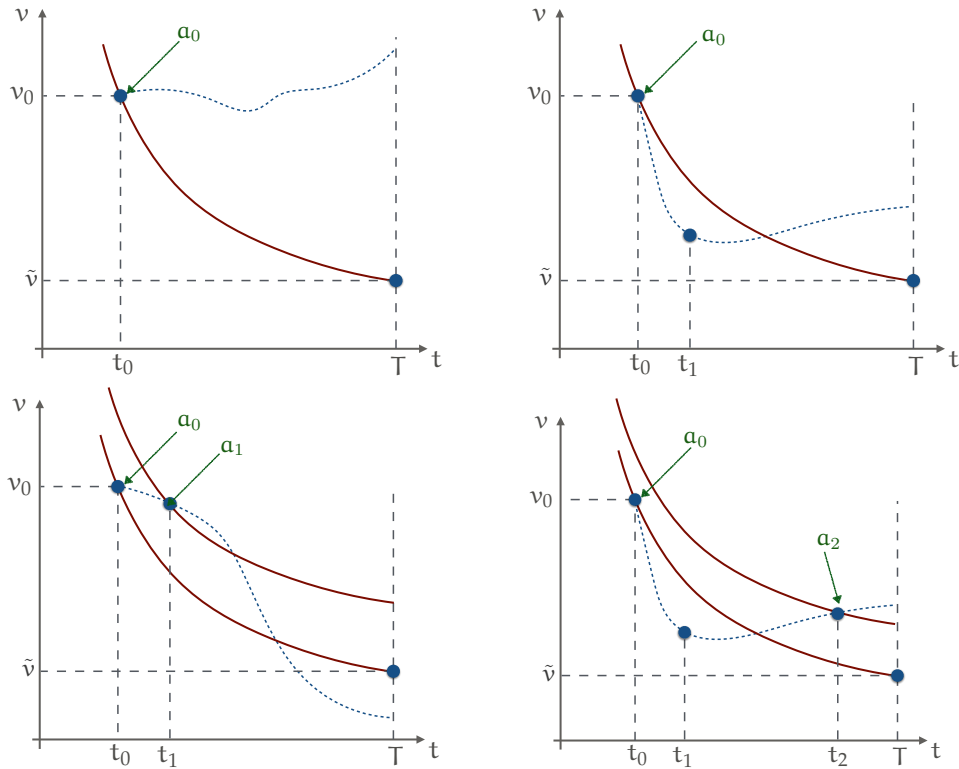


Figure 3: Indifference curves and increasing actions

Note: The dotted line represents changes in the valuation of the agent, and the red solid lines represent the indifference curves. In the top left graph, the agent does not get a chance of changing its action. In the top right graph, the agent gets a second chance to update its action but chooses not to increase it. In the bottom two graphs, the agent increases its action after getting a chance to do so at a decision node above the agent's original indifference curve.

Figure 3 represent cases where the agent gets a chance of updating its action at a time period in which the agent's valuation is above the original indifference curve. In these two cases, the agent chooses a higher action, so the choice made at decision node (v_0, t_0) no longer binds. More generally, for any path with decision nodes $\{(v_0, t_0), (v_1, t_1), \dots, (v_n, t_n)\}$ the final action is the one that is optimal for a value equal to $\max\{e(v_0, t_0), e(v_1, t_1), \dots, e(v_n, t_n)\}$, i.e., the value associated to the highest indifference curve reached during the decision times t_0, \dots, t_n .

This procedure can be formalized as follows. In our example, the threshold \tilde{v} was defined by the following property:

$$\mathbb{E}(v_1 | \text{no decision opportunity with } v_1 > \tilde{v}) = \tilde{v},$$

i.e., the expected value for all realizations where the action chosen in the first period remains the final one. Similarly, $e(v_0, t_0)$ is the expected final value v_T on the set of all paths following (v_0, t_0) such that all subsequent decision nodes lie below the indifference curve corresponding to (v_0, t_0) or there is no subsequent decision node. This is the case in the top two panels in Figure 3 but not the last two panels.

An optimal decision strategy is derived as follows. Let

$$\tilde{S}(v) = \operatorname{argmax} U(v, a) \tag{8}$$

denote the solution to the static optimization problem for any value v . Then the optimal decision strategy at decision node (v, t) is given by $\tilde{S}(e(v, t))$.⁹

4.2 Formal Analysis

Here we state our main results that formalize the intuitive arguments given above and provide a sketch of the proof of our main theorem, while the complete proofs can be found in the appendix. Our analysis in the previous section

⁹If the maximum is not unique, \tilde{S} is an increasing selection.

suggests a general approach to finding the solution to our decision problem. The steps to this proof are as follows:

1. Define a function $e(v, t)$ that partitions the set of states into indifference equivalent final value classes as above. We show that the function is uniquely defined and can be obtained via solving a dynamic programming problem.
2. Define a candidate-optimal decision strategy at decision node (v, t) as choosing the action $\tilde{S}(e(v, t))$, where $\tilde{S}(\cdot)$ is the function defined by equation 8, and prove that this is an optimal decision strategy for the dynamic problem.

Let $D(\omega)$ denote the decision nodes (v, t) for path $\omega \in \Omega$ and $N(v, t)$ denote the set of paths $\omega \in \Omega$ such that there are no arrivals after (v, t) , i.e.,

$$N(v, t) = \{\omega \in \Omega \mid (v, t) \in D(\omega) \text{ and } \nexists (v', t') \in D(\omega) \forall t' > t\}. \quad (9)$$

Let $\Pi(\omega|v, t)$ denote the conditional probability of ω given $(v, t) \in D(\omega)$. The following assumptions are used throughout the paper.

Assumption 3. *The following properties hold:*

1. $\exists \delta > 0$ such that $\Pi(N(v, t) | v, t) > \delta$ for all (v, t) ,
2. $\int_{N(v, t)} (v_T(\omega)) d\Pi(\omega|v, t)$ is continuous in v, t , and
3. $P(v', t'|v, t)$ is continuous in the topology of weak convergence.

The first assumption states there is positive probability bounded away from zero that the current decision node is the last one. The last two assumptions are standard continuity requirements on the stochastic process for signals. In particular, in case arrival times and values are independent $\int_{N(v, t)} (v_T(\omega)) d\Pi(\omega|v, t) = v\Pi(N(v, t) | v, t)$, so the second condition states that the probability of a next arrival before T is a continuous function of t . This is satisfied for example in the Poisson arrival case where $\Pi(N(v, t)) = \exp(-\lambda(T - t))$, with the arrival rate λ .

4.2.1 Self-Generated Expectation

Consider a real-valued (Borel) measurable function $e(v, t)$. For $\omega \in \Omega$, an element of the underlying probability space and $0 \leq t \leq T$, if there is any arrival after t , define

$$\bar{e}(t, \omega) = \max \{e(v', t') \mid (v', t') \in D(\omega) \text{ and } t' > t\},$$

where, as defined earlier, $D(\omega)$ is the set of all decision nodes for a given ω . If the above set is empty, set $\bar{e}(t, \omega)$ to an arbitrarily low number.¹⁰ For every state (v, t) such that $0 \leq t \leq T$, let $H(\varepsilon, v, t) = \{\omega \mid (v, t) \in D(\omega) \text{ and } \bar{e}(t, \omega) \leq \varepsilon\}$.¹¹

Definition. The function $e(v, t)$ is a *self-generated expectation (SGE)* for the process defined by transition function P if it satisfies the following property for all (v, t) :

$$e(v, t) = \mathbb{E}_{H(e(v,t), v, t)} v_T. \quad (10)$$

The above definition is equivalent to the following:

$$\int_{H(e(v,t), v, t)} (v_T(\omega) - e(v, t)) d\Pi(\omega \mid v, t) = 0. \quad (11)$$

Given a self-generated expectation function $e(v, t)$, we can define *iso-expectation* level curves $I(u) = \{(v, t) \mid e(v, t) = u\}$. Intuitively, the level u indicates the conditional expectation of the final value of all paths, starting from a given state in $I(u)$, that never cross above this iso-expectation curve at a future decision node. These are the indifference curves described in the previous section.

The derivation of a self-generated expectation follows a recursive structure. First note that $e(v, T) = v$, since this is a terminal node. Intuitively, working

¹⁰For example, set it equal to $\inf e(v, t)$.

¹¹Note that this definition is a given $e(v, t)$ function. To simplify the notation, we have suppressed this term from an argument of the function H .

backwards from that point using (10) should give a unique self-generated expectation $e(v, t)$. While (10) seems like a complicated functional equation, we can find the solution by considering the following auxiliary functional equation:

$$W(\varepsilon, v, t) = \int_t^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau' | v, t) + \int_{N(v, t)} (v_T(\omega) - \varepsilon) d\Pi(\omega | v, t) \quad (12)$$

and $W(\varepsilon, v, T) = v - \varepsilon$.

Proposition 1. *Given Assumption 3, the function $e(v, t)$ defined implicitly by $W(e(v, t), v, t) = 0$ exists and is the unique self-generated expectation for the process defined by transition function P .*

Here we give an overview of the steps involved. The first step consists in showing that using any self-generated expectation function we can construct a function W that satisfies functional equation (12) and the condition that $W(e(v, t), v, t) = 0$. Then we show that by Assumption 3 the functional equation (12) is a contraction mapping, so it has a unique solution. It also follows easily that this function is strictly decreasing in ε and continuous. Moreover, it is greater than or equal to zero when $\varepsilon = 0$ and negative for large ε . It follows by the intermediate value theorem that there is a unique value $e(v, t)$ such that $W(e(v, t), v, t) = 0$. We finally show that this solution satisfies 10.

4.2.2 The Optimal Solution

Consider the problem

$$\max_a U(v, a).$$

Let $\tilde{S}(v)$ be a (weakly) increasing selection of the set of maximizers, which is guaranteed to exist by Assumption 1. This strategy gives an optimal action for the agent if it were choosing at time T with a value $v_T = v$.

Theorem 1. For any payoff function $U(v, a)$ satisfying Assumption 1 and Markov process $P(v, t)$ satisfying Assumption 3, the strategy defined by $S(v, t) = \tilde{S}(e(v, t))$ is an optimal strategy for the dynamic decision problem 7, where the function $e(v, t)$ is the self-generated expectation corresponding to process P .

We provide an intuitive argument here. Take this strategy S for all $t' > t$ and consider the choice of action a in state (v, t) . Let $\bar{H}(a, v, t)$ denote all paths $\omega \in \Omega$ following (v, t) such that either there are no more future decision nodes or $S(v_m, \tau_m) \leq a$ for all decision nodes that follow. Using the candidate strategy function for future states, $S(v_m, \tau_m) \leq a$ corresponds to states where $\tilde{S}(e(v_m, \tau_m)) \leq a$. Let $a(T, \omega)$ denote the final action when applying this strategy in the future, so $a(T, \omega) = a$ for $\omega \in \bar{H}(a, v, t)$ and is greater than a in the complement. It follows that the expected value of choosing a in (v, t) followed by this strategy function equals

$$V(v, t, a, S) = \int_{\bar{H}(a, v, t)} U(v(T), a) d\Pi(\omega|v, t) + \int_{\bar{H}(a, v, t)^c} U(v(T), a(T, \omega)) d\Pi(\omega|v, t).$$

First, note that the boundary of the set $\bar{H}(a, v, t)$ (and $\bar{H}(a, v, t)^c$) consists of all those paths starting from (v, t) for which the final action $a(T)$ is equal to a . For the purpose of providing a heuristic argument, assume $U(v, a)$ is differentiable in a .¹² So, given the envelope condition, when considering the derivative of the above, we can ignore the effect of the change in the supports of the two integrals. The first-order condition is then $\partial V / \partial a = \frac{\partial}{\partial a} \int_{\bar{H}(a, v, t)} U(v(T), a) d\Pi(\omega|v, t) = 0$. By Assumption 1 (linearity in v), this is equivalent to the condition

$$\frac{\partial}{\partial a} U(\mathbb{E}_{\bar{H}(a, v, t)} v(T), a) \Pi(\bar{H}(a, v, t) | v, t) = 0, \quad (13)$$

so the optimal action satisfies $a = \tilde{S}(\mathbb{E}_{\bar{H}(a, v, t)} v(T))$. It only remains to show that $\mathbb{E}_{\bar{H}(a, v, t)} v(T) = e(v, t)$. When $a = \tilde{S}(e(v, t))$, this follows easily; the set $\bar{H}(a, v, t)$ is exactly the set such that for all future decision nodes $e(v_m, \tau_m) \leq$

¹²In the formal proof we do not assume differentiability.

$e(v, t)$. Otherwise, suppose some other a satisfies equation (13); then for all future values where $S(v', \tau') = a$, it must be the case that $e(v', \tau') = \mathbb{E}_{\bar{H}(a, v, t)} v(T)$. It then follows from the definition of a self-generating expectation that $e(v, t) = \mathbb{E}_{\bar{H}(a, v, t)} v(T)$. Since by Theorem 1, there is a unique self-generating expectation, the optimal action is uniquely determined by the condition given in Proposition 1.

4.3 Embedding in Games

As our leading example suggests, our results for decision problems can be extended to a class of games of incomplete information. Fixing the strategies of the other players, the choice of a best response is a decision problem that falls within the class discussed above. This best response can be found by maximizing expected payoffs at equivalent final values, as defined above. In contrast to our above decision problem, the vector of strategies has the additional restriction that the strategies must conform an equilibrium, i.e., be mutually best responses. We define a static Bayesian game where the distribution for each player's type is the distribution of the equivalent final values for that player, and strategies map these values into their corresponding action sets. Finally, we establish that any equilibrium of this static Bayesian game defines equilibrium strategies for all players in the original game.

Define a game $\Gamma = (I, \{A_i\}_{i \in I}, \{Z_i\}_{i \in I}, \{P_i\}_{i \in I}, \{u_{iT}\}_{i \in I})$ as follows. There is a fixed set of players $I = \{1, \dots, N\}$. Each player faces a process for values $v \in Z_i$ and decision times in $[0, T]$ with Markov transition $P_i(v, t)$ that are independent across players. Final payoffs are given by utility functions $u_{iT}(v_{iT}, a_{iT}, a_{-iT})$, where v_{iT} is the vector of final values for player i and (a_{iT}, a_{-iT}) is the vector of final actions coming from totally ordered sets A_1, \dots, A_N . We assume that information sets for each player contain only their own histories, and as a result strategies $S_i : Z_i \times [0, T] \rightarrow A_i$ for each player specify choices of actions as a function of these histories, and without loss of generality we can restrict to Markov strategies $S_i(v, t)$. Let \mathbf{S}_i denote the set of strategies. Let $u_i(S_i, S_{-i}) = \mathbb{E}_0 u_{iT}(v_{iT}, a_{iT}, a_{-iT} | S_i, S_{-i})$.

Definition 1. An equilibrium for game $\Gamma = (I, \{A_i\}_{i \in I}, \{Z_i\}_{i \in I}, \{P_i\}_{i \in I}, \{u_i\}_{i \in I})$ is a vector of functions $S_i : Z_i \times [0, T] \rightarrow A_i$ such that for all i $u_i(S_i, S_{-i}) \geq u_i(S'_i, S_{-i})$ for all $S'_i \in \mathbf{S}_i$.

Finding the Nash equilibria of this game seems a formidable task, given the high dimensionality of the strategy space. In what follows we show how this problem can be reduced to solving for the one-dimensional strategies that specify the actions for each player in a static Bayesian game.

Consider player i . For every history ω , we can identify a unique value corresponding to the highest final equivalent value reached, $v_i(\omega) = \max\{e_i(v_n(\omega), \tau_n(\omega))\}$, for the corresponding path. This procedure determines uniquely a distribution Ψ_i of equivalent final values for this player that depends only on the corresponding Markov process P_i for decision nodes (v, t) . Define the (static) Bayesian game as follows: set of players $I = \{1, \dots, N\}$, distribution of values for each player Ψ_1, \dots, Ψ_N , strategy sets A_1, \dots, A_N , and payoff function $u_{iT}(v_i, a_i, a_{-i})$.

Definition 2. $\Gamma_B = (I, \{\Psi_i\}_{i \in I}, \{A_i\}_{i \in I}, \{u_{iT}\}_{i \in I})$ is the static Bayesian game associated to dynamic Bayesian game $\Gamma = (I, \{A_i\}_{i \in I}, \{Z_i\}_{i \in I}, \{P_i\}_{i \in I}, \{u_i\}_{i \in I})$.

Assumption 4. Assume the functions $u_{iT}(v_i, a_i, a_{-i})$ are linear in an increasing function of v_i and supermodular in (v_i, a_i) .

Theorem 2. Consider a game Γ that satisfies Assumption 4 and its associated Bayesian game Γ_B . For any vector of equilibrium strategies $\{\tilde{S}_i\}_{i \in N}$ of Γ_B the strategies defined by $S_i(v, t) = \tilde{S}_i(e_i(v, t))$ are an equilibrium for Γ , where the function $e_i(v, t)$ is the self-generating expectation for player i .

Proof. Let $U_i(v_{iT}, a_{iT}) = \mathbb{E}_{a_{-iT}} u(v_{iT}, a_{iT}, a_{-iT} | S_{-i})$, that is, the expected final payoff given v_{iT}, a_{iT} after integrating out the strategies of the other players. Assumption 4 implies that U_i is linear in v and supermodular. This payoff function and the stochastic process P_i for values and decision times define a dynamic decision problem that satisfies the assumptions of Theorem 1. Since

\tilde{S}_i is a best response for agent i in the Bayesian game, it follows that almost surely for \tilde{v} in the support of Ψ_i

$$\begin{aligned} U_{iT}(\tilde{v}, \tilde{S}_i(\tilde{v})) &= \mathbb{E}_{a_{-i}} u_{iT}(\tilde{v}, \tilde{S}_i(\tilde{v}), a_{-i} | S_{-i}) \\ &\geq \mathbb{E}_{a_{-i}} u_{iT}(\tilde{v}, a, a_{-i} | S_{-i}) = U_{iT}(\tilde{v}, a) \end{aligned}$$

for all $a \in A_i$. So, $\tilde{S}_i(\tilde{v})$ is an optimal solution for any equivalent final value \tilde{v} and thus the corresponding S_i as defined is an optimal strategy for the dynamic decision problem defined by the best response. As a result, the strategy vector $\{S_i\}_{i \in N}$ is a Nash equilibrium for game Γ . \square

Our result decomposes the problem of finding an equilibrium to the dynamic game Γ into two steps: (1) a dynamic decision problem—that of finding the equivalent final values—and (2) a static equilibrium determination—the Bayesian game. This decomposition can be achieved because this dynamic decision problem is determined by the stochastic process for values and decision times independently of the strategies of other players. The actual choices made in the dynamic decision problem do depend on the strategies of others, but not the determination of equivalent final values. Loosely speaking, the dynamic problem determines the indifference maps (the function $e(v, t)$), while the solution to the Bayesian game labels them with the actions.

This decomposition can be very useful in applications. As an example, consider the dynamic contest described earlier. An optimal payoff structure can be designed simply by considering the static Bayesian game defined by the corresponding distributions of final equivalent values.

The decomposition described above is possible in part from our assumption that no information from other players' actions or values is revealed throughout the game. For some special cases, this assumption can be relaxed. In particular, if for all v_i there is a weakly dominant action choice a_i that maximizes $u(v_i, a_i, a_{-i})$ for all a_{-i} , the equilibrium derived above remains an equilibrium of the dynamic game with any added information about opponents' values and strategies. This assumption holds, for example, in the case of a dynamic

second price auction as described in Section 6. This method can also be useful when considering the question of implementation in dominant strategies. Another potential setting for our method is similar to that of [Bonatti et al. \[2017\]](#), where in a dynamic Cournot setting sellers have incomplete information about their rivals: sellers observe the revenue process but not individual rivals' actions.

4.4 Value Shading

Our leading example suggests that agents will shade values, and consequently actions, as a result of the adverse selection problem introduced by the option of future decision opportunities and the irreversibility of decisions. Value shading is defined by the property that early in the decision process agents make choices as if the final values were lower than the conditional expectation.¹³ Value shading arises because of the irreversibility of actions, the opportunity of making future choices, and the arrival of new information concerning the final value. Two of the key forces determining the extent of shading are the likelihood of future decision nodes and the extent of new information received as measured by the conditional variance of expected final values. The results in this section deal with conditions relating to the former, while the role of the variability of values is discussed in Section 5.2. Our first result in this section concerns the general existence of value shading. We next discuss the conditions under which value shading decreases over time as the end period approaches.

Proposition 2. $e(v, t) \leq \mathbb{E}[v_T | v, t]$.

A strong inequality can be obtained under very weak conditions that are given in the following assumptions.

¹³Value shading can occur for strategic reasons; for example, in [Hortaçsu et al. \[2015\]](#), it is a result of dealer market power in uniform-price Treasury bill auctions. Our source of shading is distinct and fundamentally non-strategic.

Assumption 5. For all $t < T$, $\int_{N(v,t)} v_T(\omega) d\Pi(\omega|v, t)$ is strictly increasing in v .

This assumption would immediately hold given the martingale property if expected values are independent of arrival times.

Assumption 6. For all $t < T$ and v , $\int_{v' > v, t < t' < T} dP(v', t'|v, t) > 0$.

When the process is independent of the arrival time t' , this assumption holds if the probability of a future arrival is strictly positive and the conditional distribution for value is non-degenerate, i.e., $P(v' = v|v, t) < 1$.

Proposition 3. If Assumptions 5 and 6 hold, then $e(v, t) < \mathbb{E}(v_T|v, t)$.

In terms of decisions, shading of values implies shading of actions so that $S(v, t) \leq \tilde{S}(\mathbb{E}[v_T|v, t])$.

To examine the evolution of shading over time, we first consider the case where decision times are independent of values. Consistent with Assumption 2, the distribution of future decision times τ' is a function of only the last arrival time t ; denote this by the cdf $F(\tau'|t)$. The following proposition gives conditions such that $e(v, t)$ is increasing in t ; therefore, the amount of shading will go down over time.

Proposition 4. Assume that values and time arrivals are independent of each other and that $F(\tau'|t)$ is decreasing in t . Then $e(v, t)$ is increasing in t .

It follows immediately that holding fixed values v , actions will be increasing over time. The condition for the proposition follows immediately when the distribution of arrivals is independent of the initial time t , so $F(\tau'|t) = G(\tau' - t)$ for some cdf G .¹⁴ While the assumptions of this proposition cover many relevant cases, some other useful ones are excluded. In particular, this proposition

¹⁴This last assumption rules out the cases where a recent arrival might accelerate the onset of future ones. As an example, suppose there are two possible states of nature, one where arrivals never occur and one where there is a Poisson arrival rate λ of this happening at any time. In addition, assume that in both states of nature there is an arrival for sure at time zero. If a second arrival occurs at some time $t > 0$, then $F(\tau'|t) = 1 - \exp(-\lambda(\tau' - t))$ can be greater than $F(\tau'|0)$ if the initial prior is sufficiently pessimistic.

assumes that the new arrival τ' and the new value v' are independent, excluding learning environments where information is accumulated over time prior to the next decision node, as considered in Section 6. The following proposition provides an alternative set of sufficient conditions that apply to those situations.

Proposition 5. *Suppose that $P(v', \tau'|v, t) = P_v(v'|v, \tau' - t) F(\tau' - t)$. Then $e(v, t)$ is increasing in t .*

The assumptions in this proposition require that the time to the next decision node be independent of current calendar time t , and that the next value v' depend only on v and the time elapsed until this next opportunity.

5 Extension and Special Cases

We first consider an extension to the case where the ending time T is random and re-examine the aforementioned properties of shading over time. We subsequently analyze two special cases of practical importance where the determination of equivalent final values is greatly simplified and shading is either independent or proportional to value, so it is only a function of time t .

5.1 Random Termination

We have assumed that the decision problem lasts for a fixed time $[0, T]$. Our formulation allows for random termination without modification. Consider equation (12), repeated below, which is the key equation used to find the self-generated expectation:

$$W(\varepsilon, v, t) = \int_t^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau'|v, t) + \int_{N(v, t)} (v_T(\omega) - \varepsilon) d\Pi(\omega|v, t).$$

We can consider T in this equation as a random termination without any changes, with a slightly different interpretation: the term $P(v', \tau' | v, t)$ can be interpreted as the probability of the event that the next decision node is τ' and that $\tau' < T$ (i.e., the decision problem has not ended by then). Similarly, the second term, $N(v, t)$, can be interpreted as the set of paths following (v, t) where the random termination occurs before the next arrival. With this change of interpretation, the same equation applies and so are all the results that follow.

It is useful to examine the conditions of Proposition 4 in light of this reinterpretation. Rewriting the assumption as $F(\tau' - T, t)$ stochastically *increasing* in t (which in the case of deterministic T is equivalent to the condition given in Proposition 4), the result follows. This is now an assumption regarding the difference between the two random variables, τ' and T . The following corollary gives sufficient conditions for this assumption to hold.

Corollary 1. *Let $H(y|t) = P(T - t \leq y|t)$ denote the cdf for the remaining time of the decision problem conditional on $T \geq t$. Let $G(x|t) = P(\tau' - t \leq x|t)$ denote the conditional cdf of the time to the next arrival. Assume that τ' and T are conditionally independent given t and*

1. $H(y, t)$ is weakly increasing in t and
2. $G(x|t)$ is weakly decreasing in t .

Then $F(\tau' - T|t)$ is (weakly) decreasing and $e(v, t)$ is (weakly) increasing in t .

Proof. Note that $\tau' - T = \tau' - t - (T - t)$. Let $(T - t) = x$ so $P(\tau' - T \leq z|t, x) = P(\tau' - t \leq z + x) = G(z + x|t)$. Integrating over x results in $F(z|t) = \int G(z + x|t) dH(x|t)$. By the second assumption, the integrand is point-wise decreasing in t . By the first assumption, the distribution H is stochastically decreasing in t and, since G is an increasing function, it also implies that the integral is decreasing in t . This proves that $F(z|t)$ is decreasing in t . The second conclusion follows directly from Proposition 4. □

The assumptions of this corollary have an intuitive interpretation. The second one is the analogue of the assumption made in Proposition 4. The first assumption simply states that the hazard rate for termination of the decision time increases with duration, which seems a natural assumption in the case of random termination. These assumptions imply that the level curves for self-generated expectations are decreasing, as depicted in Figure 1. In the special—time stationary—case where both conditional distributions G and H are independent of t , $W(\varepsilon, v, t)$ and $\varepsilon(v, t)$ will also be independent of t , so the level curves will be flat. Shading will still occur but will not change over time.

5.2 Independent Increments

We consider two special cases that come up frequently in applications, where the derivation of optimal strategies is considerably simplified: (1) increments in value independent of the current value v and (2) increments in value proportional to v . In particular, these conditions apply to the cases where v follows an arithmetic and geometric Brownian motion, respectively. In both cases, we assume that decision times are given by a homogeneous Poisson process that is independent of the past realized signals $\{v_n\}$. These assumptions considerably simplify the derivation of shading that becomes either independent from or proportional to v . In addition, we provide a new result connecting shading to the variance of innovations.

Proposition 6. *Assume $P(v' = v + \delta | v, t)$ is independent of v for all δ and all t , and decision times are independent of v . Then*

$$W(\varepsilon + \delta, v + \delta, t) = W(\varepsilon, v, t), \forall \varepsilon, \delta, v \in \mathbb{R}, t \in \mathbb{R}_+$$

and consequently $e(v + \delta, t) = e(v, t) + \delta$.

Proof. We have previously shown that the functional equation (12) is a contraction mapping. Assume that W has the property stated above. It follows

that

$$\begin{aligned}
TW(\varepsilon + \delta, v + \delta, t) &= \int_t^T \min(W(\varepsilon + \delta, v' + \delta, \tau'), 0) dP(v' + \delta, \tau' | v + \delta, t) \\
&\quad + \int_{N(v, t)} (v_T(\omega) + \delta - (\varepsilon + \delta)) d\Pi(\omega + \delta | v + \delta, t) \\
&= \int_t^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau' | v, t) + \int_{N(v, t)} (v_T(\omega) - \varepsilon) d\Pi(\omega | v, t) \\
&= TW(\varepsilon, v, t).
\end{aligned}$$

This property is thus preserved under the functional equation and is clearly closed in the space of continuous and bounded functions under the sup norm. Therefore, it must hold for the unique fixed point. The second property stated in the proposition follows immediately from the definition of a self-generated expectation. \square

Letting $\delta = -\varepsilon$, the above proposition implies that $W(\varepsilon, v, t) = W(0, v - \varepsilon, t)$. Letting $s = v - \varepsilon$, functional equation (12) can be written as

$$W(s, t) = \int_t^T \min(W(s + z, \tau'), 0) dF(z) + \int_{N(s, t)} (s_T(\omega)) d\Pi(\omega | s, t),$$

where F is the distribution of the increments. Defining $s(t)$ implicitly by $W(s(t), t) = 0$, equivalent final values $e(v, t) = v - s(t)$, so the shading factor $s(t)$ thus defined is independent of t .

Consider now the case where $P(\gamma v', \tau' | \gamma v, t) = P(v', \tau' | v, t)$, which, as the next proposition shows, implies that $W(\gamma \varepsilon, \gamma v, t) = \gamma W(\varepsilon, v, t)$.

Proposition 7. *Assume $P(\gamma v', \tau' | \gamma v, t) = P(v', \tau' | v, t)$ for all $\gamma, v, v' \in \mathbb{R}, t, \tau' \in \mathbb{R}_+$. Then $W(\gamma \varepsilon, \gamma v, t) = \gamma W(\varepsilon, v, t)$ for all $\varepsilon, \gamma, v \in \mathbb{R}, t \in \mathbb{R}_+$ and consequently $e(\gamma v, t) = \gamma e(v, t)$.*

Proof. The proof follows a similar inductive argument as in the previous propo-

sition. Assume the function W has this property. Then evaluate

$$\begin{aligned}
TW(\gamma\varepsilon, \gamma v, t) &= \int_t^T \min(W(\gamma\varepsilon, \gamma v', \tau'), 0) dP(\gamma v', \tau' | \gamma v, t) \\
&\quad + \int_{N(v, t)} (\gamma v_T(\omega) - \gamma\varepsilon) d\Pi(\gamma\omega | \gamma v, t) \\
&= \int_t^T \min(\gamma W(\varepsilon, v', \tau'), 0) dP(v, \tau' | v, t) + \int_{N(v, t)} \gamma (v_T(\omega) - \varepsilon) d\Pi(\omega | v, t) \\
&= \gamma TW(\varepsilon, v, t).
\end{aligned}$$

This property is thus preserved under the functional equation and is clearly closed in the space of continuous and bounded functions under the sup norm. Therefore, it must hold for the unique fixed point. The second property stated in the proposition follows immediately from the definition of a self-generated expectation. \square

Shading over Time with Independent Increments While the propositions derived in Section 5.2 apply to this special case, an additional intuitive and useful result can be proved. A natural question is how the variance of new values affects shading, as it affects the option value of future actions. In the extreme, if variance were zero so $v(T) = v$ with probability one, there would be no shading. We prove a monotonicity result for the case of independent increments considered in Proposition 6.

Proposition 8. *Under the assumptions of Proposition 6, W is concave in v . A mean preserving increase in spread of the distribution of increments decreases $e(v, t)$.*

Shading Over Time and Learning Note that while the assumptions require that the next arrival τ' be independent of current value v , they do not require that the next value v' be independent from either t or τ' .

In a Bayesian learning environment, the weight of new information decreases over time and so does the variance of the change in the posterior, which gives

another reason for decreasing the level of shading over time. As an example, consider an environment where the signals for the value v_T are given by a Brownian motion with drift v_T , where v_T is itself drawn from a normal distribution with known mean and variance. The history at time t is the state of the Brownian motion $x(t)$. Letting v_0 be the mean of the distribution of v_T , σ_0 its variance, and σ the volatility of the Brownian motion, the posterior mean at time t is

$$v(t) = x_0 \frac{1/\sigma_0}{1/\sigma_0 + t/\sigma} + \frac{x(t)}{t} \frac{t/\sigma}{1/\sigma_0 + t/\sigma},$$

and the variance is $1/(1/\sigma_0 + t/\sigma)$. Together with an independent arrival process for decision nodes, this formula can be used recursively to define the Markov process $P(v', t'|v, t)$ that satisfies the assumption of independent increments in Proposition 6. As the variance of the increments decreases over time, Proposition 8 implies that shading decreases over time.

6 Two Detailed Applications

To show that our method can be numerically applied to dynamic games, we now consider two applications. The first one concerns long auctions, developed in [Hopenhayn and Saeedi \[2020\]](#). This application also derives an easily solvable partial differential equation to derive the shading function when values follow a Brownian motion. The second application considers a special case of anonymous sequential games, which were introduced by [Jovanovic and Rosenthal \[1988\]](#).

6.1 Long Auctions

Many auctions take place over a considerable length of time; such is the case of trading platforms (e.g., eBay, GovDeals) and other settings (e.g., procurement, spectrum). As a result, dynamic considerations can be important for understanding bidding behavior and improving auction design. In particular, during

these long auctions, bidders' valuations and strategies are likely to be affected by information that arrives during the auction; yet most of the literature has abstracted from this feature.¹⁵ Here, changes in value could come from several sources: the existence of alternatives that change the outside value (modeled here exogenously), preference shocks (e.g., change of plans when buying event tickets), cost or capacity shocks (as in electricity markets, see [Ito and Reguant \[2016\]](#)), information regarding complementary goods (as in the case of spectrum auctions, see [Börger and Dustmann \[2005\]](#) and [Bulow et al. \[2009\]](#)), and alternative demands for use of resources in the face of capacity constraints (as in procurement auctions, see [Jofre-Bonet and Pesendorfer \[2003\]](#)). We estimate a model that fits in the class of games described above.

The specification is as follows. The value $v(t)$ follows a Brownian motion with zero drift and volatility σ , and the process for rebidding is Poisson with arrival ρ . The functional equation (12) can be rewritten as

$$W(\varepsilon, v, t) = \rho \int_t^T \exp(-\rho(\tau' - t)) \int \min\left(0, W\left(\varepsilon, v + \sqrt{\tau' - t}\sigma z, \tau'\right)\right) d\Phi(z) d\tau' + \exp(-\rho(T - t))(v - \varepsilon). \quad (14)$$

Note that by Proposition 4, $e(v, t)$ is increasing in t .

Our specification satisfies the condition in Proposition 6 so $W(\varepsilon, v, t) = W(0, v - \varepsilon, t)$. In consequence, we can write the value function $\tilde{W}(x, t)$ where $x = v - \varepsilon$,

$$\tilde{W}(x, t) = \rho \int_t^T \exp(-\rho(\tau' - t)) \int \min\left(0, \tilde{W}\left(x + \sqrt{\tau' - t}\sigma z, \tau'\right)\right) d\Phi(z) d\tau' + \exp(-\rho(T - t))x. \quad (15)$$

To numerically compute the shading function, it is more convenient to solve

¹⁵There is also a literature strand on modeling and estimating dynamics across auctions. The classic paper is [Jofre-Bonet and Pesendorfer \[2003\]](#), which estimates dynamic auctions in procurement by controlling for the utilized capacity of participants. More recent papers that consider the option value faced by bidders in sequential auctions include [Zeithammer \[2006\]](#), [Said \[2011\]](#), [Hendricks and Sorensen \[2015\]](#), [Backus and Lewis \[2012\]](#), [Bodoh-Creed et al. \[2016\]](#), and [Coey et al. \[2015\]](#). As a result of this option value, changes in the available alternative items can alter the reservation price for bidders over time. While these papers focus on dynamic bidding across auctions, they assume that bidding within each auction happens instantaneously. Nevertheless, these papers motivate our reduced-form approach toward the change in valuation to be a result of changes in these outside options.

for the value function \tilde{W} using the PDE corresponding to the above Bellman equation (the Hamilton-Jacobi equation). To do so, consider a small interval $[t, t + \Delta]$ and subtract $\tilde{W}(x, t)$ to rewrite the Bellman equation as follows:

$$0 = \rho \int_t^{t+\Delta} \exp(-\rho(\tau' - t)) \int \left[\min\left(0, \tilde{W}\left(x + \sqrt{\tau' - t}\sigma z, \tau'\right)\right) - \tilde{W}(x, t) \right] d\Phi(z) dt \\ + \exp(-\rho\Delta) \left[\int \tilde{W}\left(x + \sqrt{\Delta}\sigma z, t + \Delta\right) - \tilde{W}(x, t) \right] d\Phi(z).$$

Taking derivative with respect to Δ and evaluating at $\Delta = 0$ results in the following PDE

$$0 = \rho \left[\min\left(0, -\tilde{W}(x, t)\right) \right] + \frac{1}{2}\sigma^2 \frac{\partial^2}{\partial x^2} \tilde{W}(x, t) + \frac{\partial}{\partial t} \tilde{W}(x, t).$$

This PDE belongs to a standard class and can be easily solved with numerical methods. Once this function is derived, optimal shading can be obtained by finding the root $\tilde{W}(s, t) = 0$.

In [Hopenhayn and Saeedi \[2020\]](#), this model is estimated with eBay and Gov-Deals data. The estimates show considerable shading and explain a considerable amount of skewness in bidding times. The model is then used to perform a series of counterfactuals and to assess the implications of alternative designs on bidders' welfare and sellers' revenue.

6.2 Anonymous Sequential Game

Consider a stationary population of agents. Time is discrete. Each period an agent continues in the game with probability δ and a value v that follows a Markov process with conditional distribution $F(v'|v)$. Exiting agents are replaced by new ones with values v drawn from some initial distribution $G(v)$. At each of these decision nodes, the agent chooses whether to increase its capital k by Δk at unit cost c and is then faced with a random match to a subset of other players in the population. Profits in the period (gross of investment costs) are given by $\pi(v, k, \tilde{k})$, where \tilde{k} is the vector of capital of other competing agents. Assume this is linear in v and supermodular in

v, k . For simplicity, suppose the Markov process has independent increments so that $v' = v + z$, where z has cdf $\Phi(z)$. Define the payoff function as follows: $u(v, k, \tilde{k}) = \frac{1}{1-\delta} (\pi(v, k, \tilde{k}) - ck)$.

We consider a stationary equilibrium where the measure over firm capital stocks $\mu(k)$ is time invariant. Each period, competing firms are drawn randomly from the corresponding distribution. A stationary equilibrium is given by investment strategies $k' = g(v, k, \mu)$ that solve the firm's dynamic problem of capital accumulation and such that μ is an invariant measure generated by these decision rules.

We explain now how to derive the stationary equilibrium using our approach. Given the assumption of independent increments and noting that as a consequence of stationarity there is no time argument, shading is given by a shift s independent of v so that $e(v) = v - s$. The shading factor s satisfies $\tilde{W}(s) = 0$, where the function \tilde{W} is the solution to functional equation

$$\tilde{W}(x) = \delta \int \min(0, \tilde{W}(x+z)) d\Phi(z) + (1-\delta)x.$$

Having solved for s , we can define the distribution of equivalent final values for a player. Letting v_t denote the random value at time t , then the distribution of equivalent final values for a given player is the mixture of the random variable $v_t - s$ for $t = 1, \dots$ with weights $(1-\delta)\delta^{t-1}$, as explained in Section 3.1. In the symmetric case this can be interpreted as the distribution from which all competitors in a period draw their values. The steps for finding the equilibrium and estimating parameters to match moments in the data are explained below. For comparison, we describe first the standard nested fixed point algorithm that is used in practice.

Solving this through a nested fixed point algorithm would require the following steps:

1. Derive a stationary distribution of values $F(v)$.
2. Outer fixed point:
 - (a) Choose the estimated parameters θ of the payoff function.

- (b) Inner fixed point:
 - i. Guess a strategy $k(v)$ for all players.
 - ii. Solve the dynamic problem for an agent to get the best response strategy. Define this as a new guess for strategy.
 - iii. Get new strategies for all players.
- (c) Adjust parameters θ until a good match to the data is obtained.

Using our method, the steps would be:

1. Find equivalent values (in this case the scalar s).
2. Derive the stationary distribution of equivalent values (this is the same distribution as in step 1 of the previous procedure, shifted by the shading factor s).
3. Outer fixed point:
 - (a) Choose the estimated parameters θ of the payoff function.
 - (b) Inner fixed point:
 - i. Guess the strategies $k(v)$ for all players.
 - ii. Calculate the static best responses solving the static problem

$$k(v) = \arg \max_k \int u(k, \tilde{k}(\tilde{v}), v - s) dF(\tilde{v}).$$
 - iii. Get new strategies for all players.
 - (c) Adjust parameters θ until a good match to the data is obtained.

While in both cases a dynamic programming problem needs to be solved, the nested fixed point algorithm requires this to be done in the most inner loop (for each parameter vector θ and each strategy of other players), while in our setting it is done only once.

A simple solution to our first step can be found for the following stochastic process. Assume that with probability $(1 - p)$ the value continues to be the

same, while with probability p it is drawn again from distribution $F(v)$.¹⁶ It easily follows that

$$e(v) = \frac{(1 - \delta(1 - p(1 - F(v))))v + \delta p \int^v y dF(y)}{1 - \delta(1 - p)}.$$

This is a weighted average of v and $E(y|y \leq v)$, so it is clearly lower than v . Accordingly, the agent behaves as if the value v were lower.

7 Final Remarks

In this paper we considered a general theory for dynamic decision problems with an extension to dynamic games of imperfect information. The theory relies on a simple yet rich structure with potentially broad applicability. The problem of finding optimal actions or equilibria in the dynamic setting is reduced to solving for the corresponding optimal actions or equilibria in an equivalent static setting with respect to a distribution of values that is independent of the decision problem, making this general class very tractable.

The values and actions in this static problem are shaded because of the option of future actions, and under mild conditions, the incentives for shading decrease as the decision problem progresses. This shading contributes to delay and underinvestment in the early stages of decision problems.

The theory could be extended in several directions. It seems that risk aversion can be introduced relatively easily by defining the self-generated expectations in terms of certainty equivalent values. Our methods might also extend to the case of some information revelation during the game, with the obvious complication that self-generated expectations would require to be solved jointly across all players.

The analysis of dynamic games has proven to be a difficult problem. There are obvious trade-offs in research and corners to cut. Our paper is no exception

¹⁶Note that this stochastic process does not have independent increments.

and we have our share of strong assumptions. In particular, we have chosen to represent the impact of information on values, alternatives and opportunities, and the existence of decision time frictions in a reduced form, given by the stochastic process for values and decision time opportunities. There are obvious shortcomings, but the payoff is a parsimonious representation of equilibria and a very tractable general structure that could be easily used in further applications.

References

- Dilip Abreu, David G. Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58: 1041–63, 02 1990. doi: 10.2307/2938299. 4
- Matthew Backus and Gregory Lewis. A demand system for a dynamic auction market with directed search. *Harvard University, October*, 2012. URL <http://goo.gl/zdxan3>. 15
- P. Bajari, C.L. Benkard, and J. Levin. Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370, 2007. 2
- Aaron Bodoh-Creed, Joern Boehnke, and Brent Richard Hickman. How efficient are decentralized auction platforms? 2016. 15
- Alessandro Bonatti, Gonzalo Cisternas, and Juuso Toikka. Dynamic oligopoly with incomplete information. *The Review of Economic Studies*, 84(2):503–546, 2017. 4.3
- Tilman Börgers and Christian Dustmann. Strange bids: bidding behaviour in the united kingdom’s third generation spectrum auction. *The Economic Journal*, 115(505):551–578, 2005. 6.1
- Jeremy Bulow, Jonathan Levin, and Paul Milgrom. Winning play in spectrum auctions. Technical report, National Bureau of Economic Research, 2009. 6.1

- Guillermo A Calvo. Staggered prices in a utility-maximizing framework. *Journal of monetary Economics*, 12(3):383–398, 1983. 1
- Dominic Coey, Bradley Larsen, and Brennan C Platt. A theory of bidding dynamics and deadlines in online retail. 2015. 15
- Avinash K Dixit, Robert K Dixit, Robert S Pindyck, and Robert Pindyck. *Investment under uncertainty*. Princeton university press, 1994. 1
- Drew Fudenberg and David K Levine. Open-loop and closed-loop equilibria in dynamic games with many players. *Journal of Economic Theory*, 44(1): 1 – 18, 1988. ISSN 0022-0531. doi: [https://doi.org/10.1016/0022-0531\(88\)90093-2](https://doi.org/10.1016/0022-0531(88)90093-2). URL <http://www.sciencedirect.com/science/article/pii/0022053188900932>. 5
- Milton Harris and Bengt Holmstrom. A theory of wage dynamics. *The Review of Economic Studies*, 49(3):315–333, 1982. 3
- Kenneth Hendricks and Alan Sorensen. The role of intermediaries in dynamic auction markets. Technical report, Working Paper, 2015. 15
- Hugo Hopenhayn and Maryam Saeedi. Bidding dynamics in auctions. 2020. 3.1, 6, 6.1
- Ali Hortaçsu, Jakub Kastl, and Allen Zhang. Bid shading and bidder surplus in the u.s. treasury auctionsystem. 2015. URL <https://goo.gl/PC9crD>. 13
- V Joseph Hotz and Robert A Miller. Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3): 497–529, 1993. 2
- Koichiro Ito and Mar Reguant. Sequential markets, market power and arbitrage. *American Economic Review*, 106(7):1921–1957, 2016. 1, 3.1, 6.1
- Mireia Jofre-Bonet and Martin Pesendorfer. Estimation of a dynamic auction game. *Econometrica*, 71(5):1443–1489, 2003. 6.1, 15

- Boyan Jovanovic and Robert W Rosenthal. Anonymous sequential games. *Journal of Mathematical Economics*, 17(1):77–87, 1988. [1](#), [6](#)
- Yuichiro Kamada and Michihiro Kandori. Asynchronous revision games. Technical report, mimeo, 2011. [1](#), [1](#)
- Yuichiro Kamada and Michihiro Kandori. Revision games part i: Theory. *Working paper*, December 2015. [1](#), [1](#)
- Adam Kapor and Sofia Moroni. Sniping in proxy auctions with deadlines. 2016. [1](#)
- Steven A Lippman and John J McCall. The economics of job search: A survey. *Economic inquiry*, 14(2):155–189, 1976. [1](#)
- Benny Moldovanu and Aner Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91(3):542–558, 2001. [3.1](#)
- John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033, 1987. [2](#)
- Maher Said. Sequential auctions with randomly arriving buyers. *Games and Economic Behavior*, 73(1):236–243, 2011. [15](#)
- Robert Zeithammer. Forward-looking bidding in online auctions. *Journal of Marketing Research*, 43(3):462–476, 2006. [15](#)