Dynamic Value Shading[*]

Hugo Hopenhayn[†]   **Maryam Saeedi**[‡]
UCLA   Tepper, CMU

April 17, 2020

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 first is independent 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. A detailed application to dynamic bidding in auctions is 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. In these problems, agents make decisions over a period of time such as bidding or investing, that result in a final payoff, while these payoffs are affected by random events that impact the final values and are revealed throughout the decision period. 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; bidders might learn about outside options or alternative uses of their resources. Additionally, many decisions are irreversible, such as sunk investments in R&D or the impossibility of retracting bids in some auctions. Finally, opportunities for taking actions are random, though possibly correlated with value.[1] These make solving these dynamic problems using standard methods complicated. In this paper, we introduce a new methodology to solve a large class of dynamic problems where we disantangle the problem into a dynamic problem and a static problem. The dynamic problem is independent of the particular payoff function which makes solving the problem more tractable.

Particularly, when there is a need to solve the problems repeatedly, standard methods for solving dynamic stochastic decision problems can pose a computational burden , as in the calculation of best responses in strategic setttings or when estimating a model's deep parameters.[2] 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 a static one. The solution to the first depends only on the properties of the stochastic process for values and decision times while the solution to the latter only depends on the payoff function. As a result, changes in the latter do not require to solve repeatedly the more complex dynamic problem. This 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 bidding strategies can be obtained from the Bayesian Nash equilibria of a static auction. It is also useful in estimating the final payoff function, as changes in parameters do not require to resolve the dynamic part of the decision problem repeatedly.

---

[1]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. This 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].

[2]This 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.
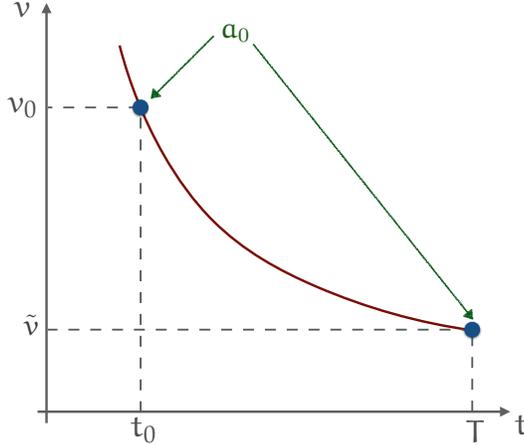
Figure 1: Indifference curves for actions

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

In our model, an 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 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 $\mathbb{E}_0 U(v_T, a_T)$.

We here 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 *eqivalent 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, where the agent's value is $\tilde{v}$.

The indifference curves such as the one depicted in Figure 1 have two main properties, given

3

some general assumptions. First, they are downward slopping, 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). Note that this choice remains the final decision of the agent if the agent has no other opportunity to revise this 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 lie underneath of the indifference curve passing through the initial decision node (v_0, t_0). This creates a source of *adverse selection* againt future self, as the current choice applies when the agent chooses not to exercise any future revision options, and this is correlated with lower future value.[3] This result in a downward slopping 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 curves results in players optimal action to be as if their value was lower anytime that there is a chance of updating their action in future, which we call *value shading.*

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 the initial choice a_0; therefore, it will remain the final action. The expected final value conditional on this set of paths is precisely $\tilde{v}$, and the same property holds for any point in this indifference curve. This property, which we call self-generated expectation, only depends 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 the 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 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

---

[3]An 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 or equal than the realized marginal product of the worker.

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

distribution of values for each player in this Bayesian game is only a function of 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]

Next, we discuss a series of potential applications. The model, for example, 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]). The model can be easily extended to allow for flow payoffs that accrue prior to the end of the game, as an example an electricity company selling future contracts to be fulfilled at some future time (as an example, see Ito and Reguant [2016]). These 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 include dynamic tournaments where over time agent's might receive information about their private values of winning the tournament, as for example in case of a promotion where agent's have outside employment options. The model can be applied to some form of dynamic contracts with incomplete enforcement, as in Harris and Holmstrom [1982], where downward flexibility in wages is a property of the optimal contract. In that setting, equivalent final values have the interpretation of conditional expected productivity.

As a more thorough application of our setting, we consider the case of dynamic second price auctions, such as eBay and GovDeals. These auctions take place over a considerable length of time and as a result dynamic considerations can be important for understanding bidding behavior and considering questions of 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.

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 of imperfect information. Section 5 describes the main value shading results of the paper. Section 6 discusses the extension of the paper to random termination time and gives properties for the

---

[5]Without this assumption, it still applies finding open loop equilibrium that for large games might approximate closed loop equilibria (see Fudenberg and Levine [1988]). It can also apply to solving for equilibria in mean field games.

case where values are independent of Poisson arrivals for bidding times. Section 7 discusses two applications of the model in more details. 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 the actual value of winning the contest v_{i1} drawn from the conditional distribution $F_i (.|v_{i0})$. 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 as 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, payoffs are equal to $-a_j$. Both 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 first and second period, 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 remaining 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 denotes the optimal action in a static setting.[6] Given the action a_0 in the first period, there is a unique threshold $\tilde{v}$ such that $S (\tilde{v}) = a_0$, we can use these 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

---

[6] We make these assumptions and others below in the analysis of the example for convenience. The set of assumptions that are needed for our main results are given in Sections 3 and 4.
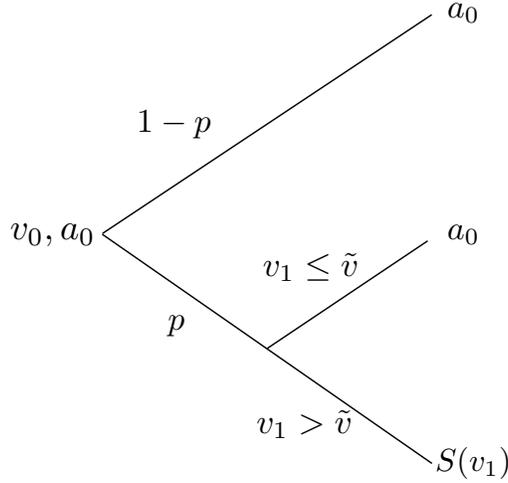
Figure 2: Decision Tree

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 branch represents values above $\tilde{v}$ where the player will increase the action in the second period to $S(v_1)$.

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

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

Assuming that G is differentiable, using the envelope theorem the associated first-order condition simplifies to

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

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:

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

This equation defines implicitly $\tilde{v}$ as a function of v_0 only, independently of the strategies of

other players. It can be more conveniently rewritten as:

$$\tilde{v} = \frac{(1-p)\,\mathbb{E}(v_1|v_0) + p\int^{\tilde{v}} v_1 dF\,(v_1|v_0)}{(1-p) + pF\,(\tilde{v}|v_0)}.$$ (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 their 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 review the first period choice. This has probability $(1-p)$ and expected value $\mathbb{E}\,(v_1|v_0)$.

2. The agent is able to review this choice but has final value less than the threshold $\tilde{v}$ so it would maintain its initial choice . This 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 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 it 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 (v_0, a_0) 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

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

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

Because of 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, it follows that

$$\tilde{v} = \frac{(1-p)\,\mathbb{E}\,(v_1|v_0) + p\int^{\tilde{v}} v_1 dF\,(v_1|v_0)}{1 - p + pF\,(\tilde{v})}, \tag{7}$$

which is identical to equation 3 so the mapping defining equivalent final value $\tilde{v}$ to the initial value v_0 is independent of the specific strategy function S and thus the underlying payoff function U and the distribution of actions of other players.

As suggested in the example above, the partition of the space of values into expected-value equivalent pairs $(v_0, \tilde{v})$ can be used to reduce the dynamic game into 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 (histories where it can review 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 $a\,(v_0, v_1) = \max\left\{a_0\,(v_0)\,, \tilde{S}\,(v_1)\right\}$ 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 (7) 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 its conditional expectation; this is what we call *value shading*. Given that actions are monotonic in values, this 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 $p \to 1$, $F\,(\tilde{v}|v_0) \to 0$, i.e., the agent acts in the first period as if the value was the lowest in the support. In the other extreme, when $p \to 0$, $\tilde{v} = \mathbb{E}\,(v_1|v_0)$ 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 $\mathbb{E}\left(v_1|v_0\right)$. It is the second event, where the agent has this opportunity but chooses not to increase the initial action that is responsible for shading. Thus, the irreversibility of actions and the opportunity for delay creates a negative option value in the first period, that can also be interpreted as adverse selection against the future self, that is responsible for the 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 general class of dynamic games as well. Time is continuous in the interval $[0, T]$. Decision times $\tau_0, \tau_1,...$ 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 captures the irreversible nature of actions. Letting a_T denote the final action, payoffs are given by a function $U\left(v_T, a_T\right)$ where v_T is a bounded real valued random variable in a probability space $(\Omega, \mathcal{F}, \Pi)$.

Assumption 1. *The payoff function $U\left(v, a\right)$ is linear in 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 $\{\tau_n\left(\omega\right)\}$, where $\tau_{n+1}\left(\omega\right) > \tau_n\left(\omega\right)$. Information arrival is modeled by a stochastic process $\tilde{v}\left(t, \omega\right)$ of signals with the property that $\mathbb{E}\left(v_T|\tilde{v}\left(t, \omega\right) = v\right) = v$. More formally, let $\{\mathcal{F}_t\}_{0 \leq t \leq T}$ be a filtration representing information available at time t, i.e, increasing $\sigma-$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 $\{\tau_n\left(\omega\right)\}$ are modeled as the jumps of a right continuous counting process $\{\eta\left(t, \omega\right)\}$. Without loss of generality, we assume that $\{\mathcal{F}_t\}$ is the filtration generated by the pair of stochastic processes $\{\eta\left(t, \omega\right), \tilde{v}\left(t, \omega\right)\}$ so that the realization of these processes is all the information available at time t and $\mathbb{E}\left(v_T|\mathcal{F}_t\right) = \tilde{v}\left(t, \omega\right)$.

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

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

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

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 $\boldsymbol{S}$ 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 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. Given our assumption of supermodularity, this is precisely when changing actions has more impact. 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 example is given in Section 6.2.1.

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

$$\sup_{s \in \boldsymbol{S}} \mathbb{E}_0 U\left(v\left(T\right), a\left(s, T\right)\right). \tag{8}$$

In Section 4 we provide conditions such that there exists an optimal solution to (8) and provide a simple method that maps this dynamic problem into an equivalent static one. This solution 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 Subsection 4.3, this general setting can be extended to dynamic games as well. Below, we include some examples of applications to games as well.

Irreversible Investment

At random times τ an agent faces an investment opportunity and choose $i_t \geq 0$ 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 the 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 given 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 $a_1, ..., a_N$ 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 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 the 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 larger volume might be convex as more costly energy sources need to be used to fulfill the contracts. A financial trader could sell future contracts that will be fulfilled by resorting to its network of intermediaries forcing 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]

Entry Decisions and Search

At random times τ, the decision maker gets and 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 enter and zero otherwise. In this application the action space is $A = \{0, 1\}$ representing the choice of no entry and entry, respectively. This 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 previous example, this problem can be mapped into the decision problem above. 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 they increase their bid 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 $[v_T - \mathbb{E}(b_2 | b_2 \leq b_T)] Prob(b_T \text{ is highest bid})$. 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 7.

---

[7]These results extend to the case where p_t is not a Martingale. This is done by backloading 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.

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

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. The objective of the decision maker is

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

where $u(v, a)$ is linear in v and supermodular. We will consider here the case where $T = \infty$ but it 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 $(1 - \beta)$ 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. It 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.

Dynamic Insurance with Imperfect Commitment

Starting with Harris and Holmstrom [1982], there is a growing literature that considers the question of optimal dynamic insurance with limited commitment. We show here how this problem can be cast in the context of our model. To simplify the analysis consider first a two-period example. At time zero, a risk averse agent has a random output with a distribution $F(v)$ with support V. The second period output is identical to the one in the first, for sure. First period output is observable by outside insurers and the agent is free to re-contract with a competitive fringe of outsiders. In the first period the agent is offered a long term insurance contract that specifies first period consumption c_0 and

second period consumption $c(v)$ depending on the first period realization. The ability to re-contract and Bertrand competition implies that $c(v) \geq v$. It can be easily shown that in the optimal contract, $c(v) = c_0$ for all $v \leq c_0$ and $c(v) = v$ otherwise, so in practice the action (consumption profile) is non-decreasing, as required in our setting. The insurer's expected payoff is $\mathbb{E}v - c_0 + \beta \int_0^{c_0} (v - c_0) \, dF(v)$, which is driven to zero, by initial competition, thus determining c_0. To map this into our decision problem, let $U(v, c) = -(v - c)^2 / (1 - \beta)$ and $P(dv) = \beta F(dv)$, so as in the previous section we assume the current period is final with probability $(1 - \beta)$. The decision problem is then:

$$\max_{c_0} - (1 - P(V)) \frac{\mathbb{E}(v - c_0)^2}{1 - \beta} - P(V) \int^{c_0} \frac{(v - c_0)^2 \, dF(v)}{1 - \beta}$$

which after substituting gives

$$\min_{c_0} E(v - c_0)^2 + \beta \int^{c_0} (v - c_0)^2 \, dF(v).$$

and first order conditions

$$Ev - c_0 + \beta \int^{c_0} (v - c_0) \, dF(v) = 0.$$

This is identical to the condition determining the equilibrium initial consumption as described above.

For the general case, suppose each period the risk averse agent draws an endowment v distributed according to cdf $F(v|\theta)$ where θ is an unknown parameter. All competing insurers have a common prior $G(\theta)$ and observe all realizations and thus share the same posteriors. An insurance contract can stipulate a sequence of transfers as a function of the history, but the agent is free to re-contract at any time with any other insurer. It is known again that the equilibrium contract has the property that there is no reduction in the agent's consumption, which then follows a non-decreasing sequence. Competition implies that whenever renegotiation binds, insurers make zero profit so expected discounted consumption equals the expected discounted endowment.

Suppose that the posterior distribution for v belongs to a parametric family with sufficient statistic (v_t, t) where $v_t = \mathbb{E}(v|\mathcal{F}_t)$. As an example, this holds in the case of a Kalman-Filter where v is drawn from a normal distribution and v_t is its posterior mean. Let $P(dv', t + 1|v, t) = \beta \phi(dv'|v, t)$ where ϕ is the distribution of next period's posterior given current sufficient statistic (v, t) and assume as in the previous example that $P(A, t'|v, t) = 0$ for all measurable A and $t' > t + 1$. Let $U(v, a) = -(v - a)^2$ where a denotes the level of

15

consumption. This is the candidate decision problem.

4 Main Results

In this section, we first give an intuitive analysis of our main findings. Next, we go through the more formal analysis with stating the main theorems. Finally, we show how our analysis for 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 equilibria of the dynamic game correspond one-to-one to equilibria of an static game with respect to this distribution of values.

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 $\tilde{v}$ that 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 in the introduction.[9] 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 the graphs below.

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 four 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 represent cases where the agent gets a chance of updating the action at a time period in which his 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 $\{(v_0, t_0), (v_1, t_1), ..., (v_n, t_n)\}$ the final

---

[9]Under some regularity conditions discussed below, this indifference curve will be downward sloping, as represented.
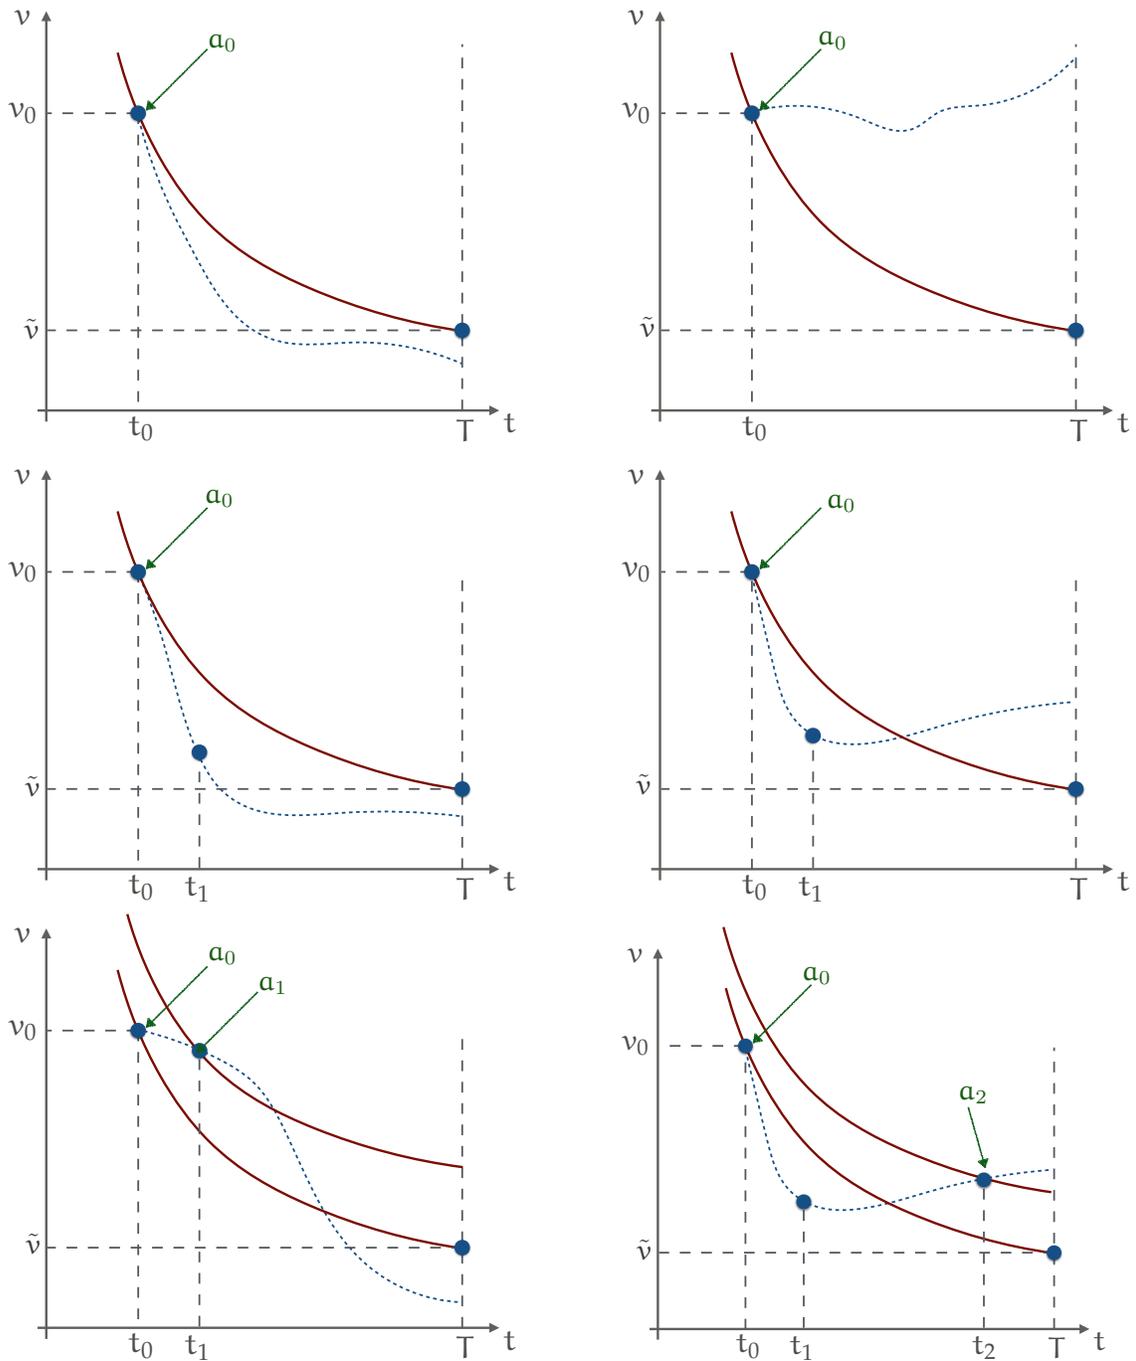
Figure 3: Indifference curves and increasing actions

Note: The dotted line represents changes in valuation of the agent and the red solid lines represent the indifference curves. In the top two graphs, the agent does not get a chance of changing his action. In the middle two graphs, the agent gets a second chance to update his action but does not increase the action. In the bottom two figures, the agent increases his action after getting a chance to do so at a decision note above his original indifference curve. On the left panels, the final value of for the agent is below $\tilde{v}$ and on the left panel they are above it.

action is the one that is optimal for a value equal to $\max\{e(v_0, t_0), e(v_1, t_1), ..., e(v_n, t_n)\}$, i.e., the value associated to the highest indifference curve reached during the decision times $t_0, ..., t_n$.

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 at all subsequent decision nodes lie below the indifference curve corresponding to (v_0, t_0) or there is no subsequent decision node. This includes the first four panels in Figure 3 but not the latter two panels.

An optimal decision strategy is derived as follows. Let

$$\tilde{S}(v) = \operatorname{argmax} U(v, a) \tag{9}$$

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 $\tilde{S}(e(v, t))$.[10]

4.2   Formal Analysis

This sections states our main results that formalizes the intuitive arguments given above and provides 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 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 it 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 9 and prove that this is an optimal decision strategy for the dynamic problem.

---

[10]If the maximum is not unique, $\tilde{S}$ is an increasing selection.

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.,

$$N(v, t) = \{\omega \in \Omega | (v, t) \in D(\omega) \text{ and } \nexists (v', t') \in D(\omega) \ \forall \ t' > t\}. \tag{10}$$

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.

3. $P(v', t' | v, t)$ is continuous in the topology of weak convergence.

The first assumption says 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

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

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. [11] For every state (v, t) such that $0 \leq t \leq T$, let $H(\varepsilon, v, t) = \{\omega | (v, t) \in D(\omega) \text{ and } \bar{e}(t, \omega) \leq \varepsilon\}$.[12]

---

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

[12] Note 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.

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

$$e\left(v,t\right) = \mathbb{E}_{H\left(e\left(v,t\right),v,t\right)}v_T. \tag{11}$$

The above definition is equivalent to the following,

$$\int_{H\left(e\left(v,t\right),v,t\right)} \left(v_T\left(\omega\right) - e\left(v,t\right)\right) d\Pi\left(\omega|v,t\right) = 0. \tag{12}$$

Given a self-generated expectation function $e\left(v,t\right)$, we can define *iso-expectation* level curves $I\left(u\right) = \{\left(v,t\right)|e\left(v,t\right) = u\}$. Intuitively, the level u indicates the conditional expectation of the final value of all paths, starting from a given state in $I\left(u\right)$, 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 $e\left(v,T\right) = v$, since this is a terminal node. Intuitively, working backwards from that point using (11) should give a unique self-generated expectation $e\left(v,t\right)$. While (11) seems like a complicated functional equation, we can find the solution by considering the following auxiliary functional equation:

$$
\begin{aligned}
W\left(\varepsilon,v,t\right) \;=\; & \int_t^T \min\left(W\left(\varepsilon,v',\tau'\right),0\right) dP\left(v',\tau'|v,t\right) \\
& + \int_{N\left(v,t\right)} \left(v_T\left(\omega\right) - \varepsilon\right) d\Pi\left(\omega|v,t\right)
\end{aligned}
\tag{13}
$$

and $W\left(\varepsilon,v,T\right) = v - \varepsilon$. We show below that $W\left(e\left(v,t\right),v,t\right) = 0$ if and only if $e\left(v,t\right)$ is a self-generated expectation.

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

The proof is in the Appendix but here we give a brief overview of the steps involved. First we show that using any self-generated expectation function we can construct a function W that satisfies the functional equation. Then we show that by assumption 3 the functional equation (13) is a contraction mapping so it has a unique solution. It also follows easily that it is strictly decreasing in ε and continuous. Moreover, it is greater 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 11.

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 8, where the function $e(v, t)$ is the self-generated expectation corresponding to the process P.*

We provide an intuitive argument here, while the formal proof is in the Appendix. 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, this 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.[13] 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

$$\frac{\partial}{\partial a} U\left(\mathbb{E}_{\bar{H}(a,v,t)} v(T), a\right) \Pi\left(\bar{H}(a, v, t)|v, t\right) = 0, \tag{14}$$

---

[13]In the formal proof we do not assume differentiability.

so the optimal action satisfies $a = \tilde{S}\left(\mathbb{E}_{\bar{H}(a,v,t)}v\left(T\right)\right)$. It only remains to show that $\mathbb{E}_{\bar{H}(a,v,t)}v\left(T\right) = e\left(v,t\right)$. When $a = \tilde{S}\left(e\left(v,t\right)\right)$, this follows easily; the set $\bar{H}\left(a,v,t\right)$, is exactly the set such that for all future decision nodes $e\left(v_m, \tau_m\right) \leq e\left(v,t\right)$. On the other hand, suppose some other a satisfies equation (14), then for all future values where $S\left(v', \tau'\right) = a$, it must be the case that $e\left(v', \tau'\right) = \mathbb{E}_{\bar{H}(a,v,t)}v\left(T\right)$. It then follows from the definition of a self-generating expectation that $e\left(v,t\right) = \mathbb{E}_{\bar{H}(a,v,t)}v\left(T\right)$. 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 they 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 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 $\Gamma = \left(I, \{A_i\}_{i\in I}, \{Z_i\}_{i\in I}, \{P_i\}_{i\in I}, \{u_{iT}\}_{i\in I}\right)$ as follows. There is a fixed set of players $I = \{1, ..., N\}$. Each player faces a process for values $v \in Z_i$ and decision times in $[0, T]$ with Markov transition $P_i\left(v, t\right)$ that are independent across players. Final payoffs are given by utility functions $u_{iT}\left(v_{iT}, a_{iT}, a_{-iT}\right)$ where v_{iT} is the vector of final values for player i and $\left(a_{iT}, a_{-iT}\right)$ the vector of final actions coming from totally ordered sets $A_1, ..., A_N$. We assume that information sets for each player contain only their own histories, and as a result strategies $S_i : Z_i \times [0, T] \to A_i$ for each player specify choices of actions as a function of these histories and without loss of generality we can restrict to Markov strategies $S_i\left(v, t\right)$. Let $\mathbf{S_i}$ denote the set of these strategies. Let $u_i\left(S_i, S_{-i}\right) = \mathbb{E}_0 u_{iT}\left(v_{iT}, a_{iT}, a_{-iT} | S_i, S_{-i}\right)$.

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

Finding the Nash equilibria of this game seems a quite 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 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 only depends on the corresponding Markov process P_i for decision nodes (v, t). Define the (static) Bayesian game as follows: Set of players $I = \{1, ..., N\}$, distribution of values for each player $\Psi_1, ..., \Psi_N$, strategy sets $A_1, ..., A_N$ and payoff function $u_{iT}(v_i, a_i, a_{-i})$.

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

Assumption 4. *Assume the functions $u_{iT}(v_i, a_i, a_{-i})$ are linear in 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 $\left\{\tilde{S}_i\right\}_{i\in N}$ 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

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

for all $a\epsilon A_i$. 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 can be achieved as 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 is possible in part from our assumption that no information from other players' actions or values are 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 holds, for example, in the case of a dynamic second price auction as described in Section 7. This can be also 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 rival, they observe revenue process but not individual rival's actions.

5 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 earlier in the decision process agents make choices as if the final values were lower than the conditional expectation.[14] This 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 as 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 6.2.1. Our first result in this section concerns the general existence of value shading. The following discuss conditions under which value shading decreases over time as the end period is approached.

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

---

[14]Value 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.

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

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

This 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 t', this holds if the probability of a future arrival is strictly positive and the conditional distribution for value is not-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' are only a function of 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.[15]

While the assumptions of this proposition cover many relevant cases, some other useful ones are excluded. In particular, this proposition assumes that the new arrival τ' and the new value v' are independent. This will not be the case when information is accumulated over time prior to the next decision node, as in the case considered in Section 7. The following proposition provides an alternative set of sufficient conditions that apply to those cases.

---

[15]This last assumption rules out 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.

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' only depend on v and the time elapsed until this next opportunity.

6   Extension and Special Cases

This Section first considers an extension to the case where the ending time T is random and re-examines the properties of shading over time discussed before. It subsequently analyzes 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 only a function of time t.

6.1   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 (13) as repeated below, which is the key equation used to find the self-generated expectation.

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

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 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.*

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

The conditions of the Proposition 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.

6.2 Independent Increments

In this section, we consider two cases where optimal actions are simplified: 1) increments in value independent of the current value v and 2) increments in value proportional to v. In particular, these 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 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 (13) is a contraction mapping. Assume that W has the property stated above. It follows that:

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

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. $\square$

This proposition simplifies shading to a simple one-dimensional function $s(t) = v - e(v, t)$ that is independent of v.

Consider now the case where $P(\gamma v', \tau' | \gamma v, t) = P(v', \tau' | v, t)$. As the next Proposition shows, this 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 Proposition. Assume

the function W has this property. Then evaluate

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

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. $\qquad\square$

6.2.1   Shading over Time with Independent Increments

While the propositions derived in Section 6.2.1 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 $v\left(T\right)=v$ with probability one, there is 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 $e\left(v,t\right)$.*

Note that while the assumptions require that the next arrival τ' is independent of current value v, they do not require that 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 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 $x\left(t\right)$. 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

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

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

7   A More Detailed Application

To show that our method can be numerically applied to dynamic games, this section considers an application that is developed in Hopenhayn and Saeedi [2020]. 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 considering questions of 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.[16] Changes in value could come here from several sources. This could be due to 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 is the case in electricity markets, see Ito and Reguant [2016]), information regarding complementary goods (as in the case of spectrum auctions, see Börgers 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 $v\left(t\right)$ follows a Brownian motion with zero drift and volatility σ, and the process for rebidding is Poisson with arrival ρ. The functional equation

---

[16]There is also a literature strand on modeling and estimating dynamics across auctions. The classic paper is Jofre-Bonet and Pesendorfer [2003], who estimate 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 option.

(13) can be rewritten as

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

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 $x = v - \varepsilon$,

$$\tilde{W}(x,t) = \rho\int_t^T \exp\left(-\rho\left(\tau'-t\right)\right)\int \min\left(0, \tilde{W}\left(x+\sqrt{\tau'-t}\sigma z, \tau'\right)\right)d\Phi(z)\,d\tau' \qquad (16)$$
$$+ \exp\left(-\rho\left(T-t\right)\right)x.$$

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

$$0 = \rho\int_t^{t+\Delta} \exp\left(-\rho\left(\tau'-t\right)\right)\int\left[\min\left(0, \tilde{W}\left(x+\sqrt{\tau'-t}\sigma z, \tau'\right)\right) - \tilde{W}(x,t)\right]d\Phi(z)\,d\tau' \qquad (17)$$
$$+ \exp\left(-\rho\Delta\right)\left[\int \tilde{W}\left(x+\sqrt{\Delta}\sigma z, t+\Delta\right) - \tilde{W}(x,t)\right]d\Phi(z).$$

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 GovDeals 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 no revenues and total value of alternative designs.

8   Final Remarks

This paper 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.

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 contributes to delay and under-investments in 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 tradeoffs in research and corners to cut. Our paper is no exception 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 that of 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. 16

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. 16

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. 7

Jeremy Bulow, Jonathan Levin, and Paul Milgrom. Winning play in spectrum auctions. Technical report, National Bureau of Economic Research, 2009. 7

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. 16

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. 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, 1, 3.1

Kenneth Hendricks and Alan Sorensen. The role of intermediaries in dynamic auction markets. Technical report, Working Paper, 2015. 16

Hugo Hopenhayn and Maryam Saeedi. Bidding dynamics in auctions. 2020. 3.1, 7, 7

Ali Hortaçsu, Jakub Kastl, and Allen Zhang. Bid shading and bidder surplus in the u.s. treasury auctionsystem. 2015. URL https://goo.gl/PC9crD. 14

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, 7

Mireia Jofre-Bonet and Martin Pesendorfer. Estimation of a dynamic auction game. *Econometrica*, 71(5):1443–1489, 2003. 7, 16

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. 16

Robert Zeithammer. Forward-looking bidding in online auctions. *Journal of Marketing Research*, 43(3):462–476, 2006. 16

A   Appendix. Proofs

A.1   Proof of Proposition 1

Recall $e(v, t)$ was defined as a threshold function with the property that

$$E\left[v_T(\omega) \mid (v, t) \in D(\omega) \text{ and } \bar{e}(t, \omega) \leq e(v, t)\right] = e(v, t). \tag{18}$$

where $\bar{e}(t, \omega) = \sup\left\{e(v', t') \mid (v', t') \,\epsilon D(\omega) \text{ and } t' > t\right\}$ or set to an arbitrarily low number if this set is empty. It is convenient to expand the set $H(\varepsilon, v, t)$ dynamically as follows:

1. It contains all ω such that there is no arrival following (v, t), i.e., $(v, t) \in D(\omega)$ and $(v', t') \notin D(\omega)$ for all v' and $t' > t$. Call this set $N(v, t)$.

2. It contains all ω such that for the next arrival $(v', t') \in D(\omega)$ has $e(v', t') \leq \varepsilon$ and $\omega \in H(\varepsilon, v', t')$. Call this set $A(\varepsilon, v, t)$. Note that $N(v, t)$ and $A(\varepsilon, v, t)$ are a partition of the set $H(\varepsilon, v, t)$.

We will prove that there is a unique threshold function satisfying property (18) and that it solves $W(e(v, t), v, t) = 0$, where $W(\varepsilon, v, t)$ is the unique solution to the following Bellman equation:

$$
\begin{aligned}
W(\varepsilon, v, t) &= \int_t^T \min\left(W(\varepsilon, v', \tau'), 0\right) dP(v', \tau' \mid v, t) \\
&\quad + \int_{N(v,t)} (v_T(\omega) - \varepsilon) \, d\Pi(\omega \mid v, t)
\end{aligned} \tag{19}
$$

with terminal value $W(\varepsilon, v, T) = v - \varepsilon$, where $\Pi(\omega \mid v, t)$ is the distribution over ω conditional on $(v, t) \in D(\omega)$. The first step is to establish *necessity*, that is, any function $e(v, t)$ with property (18) corresponds to a function $W(\varepsilon, v, t)$ satisfying this functional equation and $W(e(v, t), v, t) = 0$. The second step is to show sufficiency and this is established by showing that the Bellman equation is a contraction mapping and that the unique solution $W(\varepsilon, v, t)$ is strictly decreasing in ε. Consequently, there is a unique function $e(v, t)$ for which $W(e(v, t), v, t) = 0$. Finally we show that this $e(v, t)$ satisfies property 18.

Step 1. Necessity   Take the candidate value function:

$$W(\varepsilon, v, t) = \int_{H(\varepsilon, v, t)} (v_T(\omega) - \varepsilon) \, d\Pi(\omega \mid v, t).$$

We now show that $W(\varepsilon, v, t)$ is a solution to functional equation (19). Substituting in the right hand side of (19) results in

$$W(\varepsilon, v, t) = \int \min \left\{ \int_{H(\varepsilon, v', t')} (v_T(\omega) - \varepsilon) \, d\Pi(\omega | v', t'), 0 \right\} dP(v', t' | v, t)$$
$$+ \int_{N(v, t)} (v_T(\omega) - \varepsilon) \, d\Pi(\omega | v, t).$$

By Lemma 1, the term in brackets will be zero iff $e(v', t') \geq \varepsilon$ since $H(e(v', t'), v', t')$ is the integration set in definition (18) for $t = t'$. So the first integral is over paths $\{\omega : e(v', t') \leq \varepsilon\} \cap H(\varepsilon, v', t')$ where (v', t') is the next arrival following (v, t). This is precisely the set $A(\varepsilon, v, t)$ defined above. Moreover, since $\int \Pi(\omega | v', t') \, dP(v', t' | v, t) = \Pi(\omega | v, t)$ it follows that:

$$W(\varepsilon, v, t) = \int_{A(v, t)} (v_T(\omega) - \varepsilon) \, d\Pi(\omega | v, t)$$
$$+ \int_{N(v, t)} (v_T(\omega) - \varepsilon) \, d\Pi(\omega | v, t).$$
$$= \int_{H(\varepsilon, v, t)} (v_T(\omega) - \varepsilon) \, d\Pi(\omega | v, t)$$

For $\varepsilon = e(v, t)$, $W(e(v, t), v, t) = \int_{H(e(v, t), v, t)} (v_T(\omega) - e(v, t)) \, d\Pi(\omega | v, t)$ and by property (12) this is equal to zero, which completes this step of the proof.

Step 2. Sufficiency and Uniqueness We first show that there is a unique function $W(\varepsilon, v, t)$ satisfying (19) by establishing it is a contraction mapping in the space of continuous and bounded functions endowed with the sup norm. Given that v_T is bounded, it is immediate that the Bellman equation (19) preserves boundedness, provided ε belongs to a bounded set. By Assumption 3, it also preserves continuity and thus maps the space of continuous and bounded into itself. To prove that it is a contraction mapping, we verify Blackwell sufficient conditions. Monotonicity is trivially satisfied. To check discounting, consider the function $W(\varepsilon, v, t) + a$ for $a \geq 0$ on the right hand side of the Bellman equation

(19).

$$
\int_t^T \min\left(W\left(\varepsilon, v', \tau'\right) + a, 0\right) dP\left(v', \tau'|v, t\right)
$$

$$
+ \int_{N(v,t)} \left(v_T\left(\omega\right) - \varepsilon\right) d\Pi\left(\omega|v, t\right)
$$

$$
\leq \int_t^T \min\left(W\left(\varepsilon, v', \tau'\right), 0\right) dP\left(v', \tau'|v, t\right) + a\left(1 - \Pi\left(N\left(v, t\right)|v, t\right)\right) \tag{20}
$$

$$
+ \int_{N(v,t)} \left(v_T\left(\omega\right) - \varepsilon\right) d\Pi\left(\omega|v, t\right)
$$

$$
= W\left(\varepsilon, v, t\right) + a\left(1 - \Pi\left(N\left(v, t\right)|v, t\right)\right).
$$

By Assumption 3 $\Pi\left(N\left(v, t\right)|v, t\right) > \delta$ for some $0 < \delta < 1$, proving the second Blackwell sufficient condition.

We now show that there is a unique function $e\left(v, t\right)$ satisfying $W\left(e\left(v, t\right), v, t\right) = 0$ for all $\left(v, t\right)$. We show this recursively by establishing that W is strictly decreasing in the first argument, for any $\left(v, t\right)$. The proof is by induction showing that the Bellman equation (19) maps weakly decreasing functions into strictly decreasing ones. So, assume that the W function in the right hand side of the Bellman equation (19) is weakly decreasing. Letting $\varepsilon' > \varepsilon$,

$$
TW\left(\varepsilon', v, t\right) = \int_t^T \min\left(W\left(\varepsilon', v', \tau'\right), 0\right) dP\left(v', \tau'|v, t\right)
$$

$$
+ \int_{N(v,t)} \left(v_T\left(\omega\right) - \varepsilon'\right) d\Pi\left(\omega|v, t\right)
$$

$$
\leq \int_t^T \min\left(W\left(\varepsilon, v', \tau'\right), 0\right) dP\left(v', \tau'|v, t\right)
$$

$$
+ \int_{N(v,t)} \left(v_T\left(\omega\right) - \varepsilon'\right) d\Pi\left(\omega|v, t\right),
$$

which is strictly less than $TW\left(\varepsilon, v, t\right)$ since by Assumption 3 $\Pi\left(N\left(v, t\right)|v, t\right) > 0$. The function W is continuous so to prove that there exists and ε such that $W\left(\varepsilon, v, t\right) = 0$, it suffices to show that it will be negative for large ε and positive for small one. Looking at the above Bellman equation, the first term is non-positive and the second term is strictly decreasing in ε, so will also be arbitrarily negative for large ε. Furthermore, by Assumption 3 $\Pi\left(N\left(v, t\right)|v, t\right) > \delta$, so for large enough ε this term will dominate. The same argument can be used by making ε small enough (negative if needed) to make the second term become positive enough to dominate.

We need to show that $e\left(v, t\right)$ satisfies (18). First note that we can rewrite $W\left(e\left(v, t\right), v, t\right)$ as following:

$$W\left(e\left(v,t\right),v,t\right) \;=\; \int \min\left(W\left(e\left(v,t\right),v',t'\right),0\right) dP\left(v',t'|v,t\right) \qquad (21)$$

$$+\int_{N(v,t)} \left(v_T\left(\omega\right)-e\left(v,t\right)\right) d\Pi\left(w|v,t\right)$$

$$=\;\int_{e(v',t')\le e(v,t)} W\left(e\left(v,t\right),v',t'\right) dP\left(v',t'|v,t\right)$$

$$+\int_{N(v,t)} \left(v_T\left(\omega\right)-e\left(v,t\right)\right) d\Pi\left(w|v,t\right)$$

where the equality follows from the property that $W\left(e\left(v',t'\right),v',t'\right)=0$ and decreasing in the first term, therefore, $W\left(e\left(v,t\right),v',t'\right)\le 0$ iff $e\left(v',t'\right)\le e\left(v,t\right)$.

Let $N_k\left(v,t\right)$ denote the set of paths such that $(v,t)\in D\left(\omega\right)$ and there no more than k arrivals following t and for all these arrivals $e\left(v_l,t_l\right)\le e\left(v,t\right)$. Note that $N_k\left(v,t\right)\subset N_{k+1}\left(v,t\right)$ and that $\cup_k N_k\left(v,t\right)=H\left(e\left(v,t\right),v,t\right)$. By repeated substitution in (21) it can be shown that

$$\int_{N_k(v,t)} \left(v_T\left(\omega\right)-e\left(v,t\right)\right) d\Pi\left(w|v,t\right) \;\ge\; W\left(e\left(v,t\right),v,t\right)$$

$$\ge\; \left(1-\delta\right)^k \min\left\{\inf_{v',t'} W\left(e\left(v,t\right),v',t'\right),0\right\}$$

$$+\int_{N_k(v,t)} \left(v_T\left(\omega\right)-e\left(v,t\right)\right) d\Pi\left(w|v,t\right)$$

Taking limits ($\liminf$ on the left hand side and $\limsup$ on the right hand side), this implies that:

$$W\left(e\left(v,t\right),v,t\right) \;=\; \lim_{k\to\infty}\int_{N_k(v,t)} \left(v_T\left(\omega\right)-e\left(v,t\right)\right) d\Pi\left(w|v,t\right)$$

$$=\;\int_{\cup_k N_k(v,t)} \left(v_T\left(\omega\right)-e\left(v,t\right)\right) d\Pi\left(w|v,t\right)$$

$$=\;\int_{H(e(v,t),v,t)} \left(v_T\left(\omega\right)-e\left(v,t\right)\right) d\Pi\left(w|v,t\right)$$

Further, since $W\left(e\left(v,t\right),v,t\right)=0$, this implies property (12).

Supporting Lemmas

Lemma 1. $\int_{H(\varepsilon,v,t)}\left(v_T\left(\omega\right)-\varepsilon\right) d\Pi\left(\omega|v,t\right)$ *is strictly decreasing in* ε *and equal to zero when* $e\left(v,t\right)=\varepsilon$.

Proof. The last part follows from the definition of self-generated expectation. To show that it is strictly decreasing in ε, consider $\delta > 0$. Then

$$\int_{H(\varepsilon+\delta,v,t)} (v_T(\omega) - (\varepsilon+\delta))\, d\Pi(\omega|v,t) = \int_{H(\varepsilon,v,t)} (v_T(\omega) - (\varepsilon+\delta))\, d\Pi(\omega|v,t)$$

$$+ \int_{\varepsilon < \bar{e}(t,\omega) \leq \varepsilon+\delta} (v_T(\omega) - (\varepsilon+\delta))\, d\Pi(\omega|v,t)$$

$$\leq \int_{H(\varepsilon,v,t)} (v_T(\omega) - (\varepsilon+\delta))\, d\Pi(\omega|v,t)$$

$$+ \int_{\varepsilon < \bar{e}(t,\omega) \leq \varepsilon+\delta} (v_T(\omega) - \bar{e}(t,\omega))\, d\Pi(\omega|v,t)$$

The last term is zero by Lemma 2, thus completing the proof. $\square$

Lemma 2. $\int_{\bar{e}(t,\omega)\epsilon B} (v_T(\omega) - (\bar{e}(t,\omega)))\, d\Pi(\omega|v,t) = 0$ *for any (Borel set) B.*

The proof of this Lemma uses the following property of conditional expectation in Ash (Theorem 5.3.3. pg. 210):

Theorem. *Let Y be an extended random variable on $(\Omega, \mathcal{F}, P)$, and $X : (\Omega, \mathcal{F}) \to (\Omega', \mathcal{F}')$ a random object. If $E(Y)$ exists, there is a function $g : (\Omega', \mathcal{F}') \to (\bar{R}, \mathscr{B})$ such that for each $A\epsilon\mathcal{F}'$,*

$$\int_{\{X\epsilon A\}} Y\, dP = \int_A g(x)\, dP_x(x), \tag{22}$$

where $P_x(A) = P(\omega|X(\omega)\epsilon A)$. The function $g(x)$ is interpreted as $\mathbb{E}(Y|X = x)$.

Proof of Lemma   Consider the following random variables:

$\bar{e}(t,\omega)$ as defined above

$\tau(t,\omega)$: time at which it was reached

$v_T(\omega)$: as defined above.

We will suppress the index t for notational convenience.

Let $X(\omega) = (\tau(\omega), \bar{e}(\omega))$. The random variable that we consider in applying the above theorem is $Y(\omega) = v_T(\omega) - \bar{e}(\omega)$. For any measurable subset $A \subset \{t < \tau \leq T\}$ and Borel subset B of $\mathbb{R}$, let $P_x(A \times B) = \Pi(\tau(\omega) \in A, \bar{e}(\omega) \in B|v,t)$.

For all (τ, e),

$$\mathbb{E}\left(Y|\tau\left(\omega\right) = \tau, \bar{e}\left(\omega\right) = e, \left(v, t\right)\right) = \mathbb{E}\left(Y|\tau\left(\omega\right) = \tau, \bar{e}\left(\omega\right) = e\right)$$

$$= \mathbb{E}\left(v_T\left(\omega\right) - e|\left(v', \tau\right) \in D\left(\omega\right), e\left(v', t\right) = e, \bar{e}\left(v', \tau\right) \leq e\right) = 0,$$

by the definition of self-generated expectation. Substituting $x = (\tau, \varepsilon)$ and using (22),

$$\int_{\{\tau(\omega) > \tau_0, \bar{e}(\omega) \epsilon B\}} \left(v_T\left(\omega\right) - \bar{e}\left(\omega\right)\right) dP\left(\omega\right) = \int_{\{\tau > \tau_0, \varepsilon \epsilon B\}} E\left(Y|\tau, \varepsilon\right) dP_x\left(\tau, \varepsilon\right) = 0.$$

A.2   Proof of Theorem 1

We need to show that the strategy defined in Theorem 1 is a solution to the dynamic decision problem.

Consider a node (v, t) and some alternative action $a_2 \neq a_1 \equiv S\left(v, t\right)$. We will show that this one period deviation is not an improvement. Let $e_1 = e\left(v, t\right)$ so $a_1 = \tilde{S}\left(e_1\right)$. Consider first the case where $a_2 > a_1$ and let $e_2 = \sup\left\{v|\tilde{S}\left(v\right) \leq a_2\right\}$.

Let $V\left(v, t, a\right)$ denote the expected utility of choosing a at this state and following the (candidate) optimal policy for the future. We need to prove that $V\left(v, t, a_1\right) \geq V\left(v, t, a_2\right)$. Let $\bar{s}\left(\omega, v, t\right) = \max\left\{S\left(v', t'\right) | \left(v', t'\right) \epsilon D\left(\omega\right) \text{ and } \left(v', t'\right) \neq \left(v, t\right)\right\}$. For a path ω such that $(v, t) \in D\left(\omega\right)$ this is the maximal action excluding the choice at node (v, t) and it is also the final action if it is greater or equal than the choice at this node. Define $H\left(e, v, t\right)$ as in Section 4.2.1 and $H\left(e, v, t\right)^c$ its complement in the set of paths following $(v, t) : \left\{\omega| \left(v, t\right) \in \mathrm{D}(\omega)\right\}$. To simplify notation, from now on we drop the argument (v, t) from these functions, unless needed to avoid confusion. We can decompose the histories following $a_i, i\epsilon\left\{1, 2\right\}$ into these two sets. In the set of histories $H\left(e_i, v, t\right)$, a_i will be the final choice as no higher equivalent value than e_i is reached. In H^c, the choice a_i does not bind as an equivalent value higher than e_i is reached and the final corresponding action is $\bar{s}\left(\omega\right)$. It follows that

$$V\left(v, t, a_i\right) = \int_{H(e_i)} U\left(v\left(T\right), a_i\right) dP\left(\omega\right)$$

$$+ \int_{H(e_i)^c} U\left(v\left(T\right), \bar{s}\left(\omega\right)\right) dP\left(\omega\right),$$

for $i = \{1, 2\}$.

Note that $H(e_2) = H(e_1) \cup \{\omega | e_1 < \bar{e}(t, \omega) \le e_2\}$ so

$$
\begin{aligned}
V(v, t, a_2) &= \int_{H(e_1)} U(v(T), a_2) \, dP(\omega) \\
&+ \int_{e_1 < \bar{e}(\omega) \le e_2} U(v(T), a_2) \, dP(\omega) \quad\quad (23) \\
&+ \int_{H(e_2)^c} U(v(T), \bar{s}(\omega)) \, dP(\omega).
\end{aligned}
$$

Consider a more relaxed problem where this agent is allowed to follow the original strategy, i.e., if he arrives at histories where $e_1 < \bar{e}(\omega) \le e_2$, he is unconstrained by the pre-existing choice a_2, so his final action is $\bar{s}(\omega)$. As a result,

$$
\begin{aligned}
V(v, t, a_2) &\le \int_{H(e_1)} U(v(T), a_2) \, dP(\omega) \\
&+ \int_{H(e_1)^c} U(v(T), \bar{s}(\omega)) \, dP(\omega).
\end{aligned}
$$

where the right hand side value is the optimal for the relaxed problem. It follows that

$$
\begin{aligned}
V(v, t, a_2) - V(v, t, a_1) &\le \int_{H(e_1)} [U(v(T), a_2) - U(v(T), a_1)] \, dP(\omega) \quad\quad (24) \\
&= U\left(E_{H(e_1)} v(T), a_2\right) - U\left(E_{H(e_1)}, a_1\right) \le 0.
\end{aligned}
$$

Where the equality follows the linearity of U in v, and the last inequality follows from the $e(v, t) = \mathbb{E}_{H(e_1)} v(T)$ and $\tilde{S}(e(v, t)) = a_1$.

Now suppose instead that $a_2 < a_1$.

Define e_1 as before and let $e_2 = \inf\left\{v | \tilde{S}(v) > a_2\right\}$. For the alternative action $a_2 < a_1$, it easily follows that

$$
\begin{aligned}
V(v, t, a_2) &= \int_{H(e_2)} U(v(T), a_2) \, dP(\omega) + \int_{H(e_2)^c} U(v(T), \bar{s}(\omega)) \, dP(\omega) \quad\quad (25) \\
&= \int_{H(e_1)} U(v(T), a_2) \, dP(\omega) + \int_{H(e_1)^c} U(v(T), \bar{s}(\omega)) \, dP(\omega) \\
&+ \int_{e_2 < \bar{e}(w) \le e_1} (U(v(T), \bar{s}(\omega)) - U(v(T), a_2)) \, dP(\omega).
\end{aligned}
$$

Subtracting the above from $V(v, t, a_1)$ we can write

$$\begin{aligned} V\left(v,t,a_{1}\right)-V\left(v,t,a_{2}\right) & = \int_{H(e_{1})} U\left(v\left(T\right),a_{1}\right)dP\left(\omega\right)-\int_{H(e_{1})} U\left(v\left(T\right),a_{2}\right)dP\left(\omega\right)\\ & \quad -\int_{e_{2}<\bar{e}(w)\leq e_{1}}\left(U\left(v\left(T\right),\bar{s}\left(\omega\right)\right)-U\left(v\left(T\right),a_{2}\right)\right)dP\left(\omega\right) \qquad (26)\\ & = \left(U\left(e_{1},a_{1}\right)-U\left(e_{1},a_{2}\right)\right)P\left(H\left(e_{1}\right)\right)\\ & \quad -\int_{e_{2}<\bar{e}(w)\leq e_{1}}\left(U\left(v\left(T\right),\bar{s}\left(\omega\right)\right)-U\left(v\left(T\right),a_{2}\right)\right)dP\left(\omega\right)\\ & = \left(U\left(e_{1},a_{1}\right)-U\left(e_{1},a_{2}\right)\right)P\left(H\left(e_{1}\right)\right) \qquad (27)\\ & \quad -\int_{e_{2}<\bar{e}(\omega)\leq e_{1}}\left(U\left(\bar{e}\left(\omega\right),\tilde{S}\left(\bar{e}\left(\omega\right)\right)\right)-U\left(\bar{e}\left(\omega\right),a_{2}\right)\right)dP\left(\omega\right) (28) \end{aligned}$$

where as before the second equality follows from linearity of U in v and that $e_{1}=e\left(v,t\right)=E_{H(e_{1})}\left(v\left(T\right)\right)$ and the third equality from Lemma 3

Consider the last integral. Supermodularity of the U function implies that $\tilde{S}\left(v\right)$ must be an increasing function, so for ω such that $e_{2}<\bar{e}\left(w\right)\leq e_{1}$ it follows that

$$a_{1}\geq\tilde{S}\left(\bar{e}\left(\omega\right)\right)=s\left(\omega\right)\geq\tilde{S}\left(e_{2}\right)\geq a_{2}.$$

It also follows from supermodularity and $e_{1}\geq\bar{e}\left(\omega\right)$ that

$$U\left(\bar{e}\left(\omega\right),\tilde{S}\left(\bar{e}\left(\omega\right)\right)\right)-U\left(\bar{e}\left(\omega\right),a_{2}\right) \leq U\left(e_{1},\tilde{S}\left(\bar{e}\left(\omega\right)\right)\right)-U\left(e_{1},a_{2}\right). \qquad (29)$$

Finally, since $a_{1}=\tilde{S}\left(e_{1}\right)$ it follows that

$$U\left(e_{1},\tilde{S}\left(\bar{e}\left(\omega\right)\right)\right)-U\left(e_{1},a_{2}\right) \leq U\left(e_{1},a_{1}\right)-U\left(e_{1},a_{2}\right). \qquad (30)$$

Combining (26), (29) and (30), it follows that

$$\begin{aligned} V\left(v,t,a_{1}\right)-V\left(t,v,a_{2}\right) & \geq \left(U\left(e_{1},a_{1}\right)-U\left(e_{1},a_{2}\right)\right)P\left(H\left(e_{1}\right)\right) \qquad (31)\\ & \quad -\int_{e_{2}<\bar{e}(\omega)\leq e_{1}} U\left(e_{1},a_{1}\right)-U\left(e_{1},a_{2}\right)dP\left(\omega\right), \end{aligned}$$

and since the set of paths $\left\{e_{2}\leq\bar{e}\left(\omega\right)\leq e_{1}\right\}$ is a subset of $H\left(a_{1}\right)$, then $V\left(v,t,a_{1}\right)-V\left(t,v,a_{2}\right)\geq 0$, so the proof is complete.

Lemma 3. *For $e_1 > e_2$,*

$$\int_{e_2 < \bar{e}(\omega) \leq e_1} U\left(v\left(T\right), \bar{s}\left(\omega\right)\right) dP\left(\omega\right)$$
$$= \int_{e_2 < \bar{e}(\omega) \leq e_1} U\left(\bar{e}\left(\omega\right), \tilde{S}\left(\bar{e}\left(\omega\right)\right)\right) dP\left(\omega\right).$$

Proof. Follows from the linearity of U and application of Lemma 2.

□

A.3   Proof of Propositions 2 and 3

We first prove by the property $W\left(\varepsilon, v, t\right) \leq E\left(v_T | v, t\right) - \varepsilon$ is preserved under the Bellman equation. Suppose that $W\left(\varepsilon, v', \tau'\right) \leq E\left(v\left(T\right) | v', \tau'\right) - \varepsilon$. Then

$$
\begin{aligned}
W\left(\varepsilon, v, t\right) &= \int_t^T \min\left(W\left(\varepsilon, v', \tau'\right), 0\right) dP\left(v', \tau' | v, t\right) + \int_{N(v,t)} \left(v_T\left(\omega\right) - \varepsilon\right) d\Pi\left(\omega | v, t\right) \\
&\leq \int_t^T W\left(\varepsilon, v', \tau'\right) dP\left(v', \tau' | v, t\right) + \int_{N(v,t)} \left(v_T\left(\omega\right) - \varepsilon\right) d\Pi\left(\omega | v, t\right) \\
&\leq \int_t^T \mathbb{E}\left(\left[v_T | v', \tau'\right] - \varepsilon | v', \tau'\right) dP\left(v', \tau' | v, t\right) + \int_{N(v,t)} \left(v_T\left(\omega\right) - \varepsilon\right) d\Pi\left(\omega | v, t\right). \\
&= \mathbb{E}\left[v_T | v, t\right] - \varepsilon,
\end{aligned}
$$

where the last equality follows from the law of iterated expectation. It immediately follows from $W\left(e\left(v, t\right), v, t\right) = 0$ that $e\left(v, t\right) \leq \mathbb{E}\left(v_T | v, t\right)$. By Assumption 5 it follows that $W\left(\varepsilon, v, t\right)$ is strictly increasing in v. Together with Assumption 6 implies that the first inequality is strict. In particular, for $\varepsilon = e\left(v, t\right)$ it follows that the first inequality above is strict, so $\mathbb{E}\left(v_T | v, t\right) > e\left(v, t\right)$.

A.4   Proof of Proposition 4

Consider the Bellman equation:

$$
W\left(\varepsilon, v, t\right) \;=\; \int_t^T \left[\int \min\left(W\left(\varepsilon, v', \tau'\right), 0\right) dP\left(v'|v\right)\right] dF\left(\tau'|t\right) \tag{32}
$$
$$
+ \int_{N(v,t)} \left(v_T\left(\omega\right) - \varepsilon\right) d\Pi\left(\omega|v, t\right)
$$
$$
=\; \int_t^T \left[\int \min\left(W\left(\varepsilon, v', \tau'\right), 0\right) dP\left(v'|v\right)\right] dF\left(\tau'|t\right) \tag{33}
$$
$$
+ \left(1 - F\left(T|t\right)\right)\left[\mathbb{E}\left(v_T|v\right) - \varepsilon\right]. \tag{34}
$$

We show that the condition is preserved under the Bellman equation. By way of induction, assume that the right hand side W is increasing in τ'. Given the assumption of stochastic dominance, $F\left(\tau'|t\right)$ is stochastically increasing in t and so is the integral. The other effect of increasing t involves shifting mass from the first to the second term. By Proposition 2, $W\left(\varepsilon, v, t\right) \leq E\left(v_T|v, t\right) - \varepsilon$, implying that the second term in the right hand side of (32) is greater in expectation that the first term. So the shift in mass also contributes to increasing the overall expectation.

A.5   Proof of Proposition 5

We show that monotonicity is preserved under the Bellman equation. Assume by way of induction that $W\left(\varepsilon, v', \tau'\right)$ is increasing in its last argument. Using the first assumption in the proposition, equation (13) can be rewritten as

$$
W\left(\varepsilon, v, t\right) \;=\; \int_0^{T-t} \int \min\left(W\left(\varepsilon, v', t+x\right), 0\right) dP_v\left(v'|v, x\right) dF\left(x\right)
$$
$$
+ \int_{x>T-t} \int \left(v' - \varepsilon\right) dP_v\left(v'|v, x\right) dF\left(x\right).
$$

Take $\tau_2 > \tau_1$.

$$
\begin{aligned}
W\left(\varepsilon, v, \tau_2\right) &= \int_0^{T-\tau_2} \int \min\left(W\left(\varepsilon, v', \tau_2 + x\right), 0\right) dP_v\left(v'|v, x\right) dF\left(x\right) \\
&\quad + \int_{x>T-\tau_2} \int \left(v' - \varepsilon\right) dP_v\left(v'|v, x\right) dP_\tau\left(x\right) \\
&\geq \int_0^{T-\tau_2} \int \min\left(W\left(\varepsilon, v', \tau_1 + x\right), 0\right) dP_v\left(v'|v, x\right) dF\left(x\right) \\
&\quad + \int_{x>T-\tau_2} \int \left(v' - \varepsilon\right) dP_v\left(v'|v, x\right) dF\left(x\right) \\
&= \int_0^{T-\tau_1} \int \min\left(W\left(\varepsilon, v', \tau_1 + x\right), 0\right) dP_v\left(v'|v, x\right) dF\left(x\right) \\
&\quad + \int_{x>T-\tau_1} \int \left(v' - \varepsilon\right) dP_v\left(v'|v, x\right) dF\left(x\right) \\
&\quad - \int_{T-\tau_2}^{T-\tau_1} \int \min\left(W\left(\varepsilon, v', \tau_1 + x\right), 0\right) dP_v\left(v'|v, x\right) dF\left(x\right) \\
&\quad + \int_{T-\tau_2}^{T-\tau_1} \int \left(v' - \varepsilon\right) dP_v\left(v'|v, x\right) dF\left(x\right) \\
&\geq W\left(\varepsilon, v, \tau_1\right) - \int_{T-\tau_2}^{T-\tau_1} \int \left[W\left(\varepsilon, v', \tau_1 + x\right) - \left(v' - \varepsilon\right)\right] dP_v\left(v'|v, x\right) dF\left(x\right)
\end{aligned}
$$

where the first inequality follows from the induction hypothesis. As in the proof of Proposition 2, $W\left(\varepsilon, v', \tau_1 + x\right) \leq \mathbb{E}\left(v_T|v', \tau_1 + x\right) - \varepsilon = v' - \varepsilon$ by the Martingale assumption. As a consequence, the term subtracted in the last line above is negative. It follows that $W\left(\varepsilon, v, \tau_2\right) \geq W\left(\varepsilon, v, \tau_1\right)$.

A.6 Proof of Proposition 8

Proposition 9. $P\left(v' = v + \delta|v, t\right)$ *is independent of v for all δ and all t, and decision times are independent of v.*

Consider the dynamic programming equation for W,

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

Given the assumption of scale invariance and that arrival times and values are independent, it follows that $\int_{N(v,t)} \left(v_T\left(\omega\right) - \varepsilon\right) d\Pi\left(\omega|v, t\right) = \mathbb{E}\left[v_T|v, t, \tau' > T\right] \Pi(N(v,t)|v, t) = \left(v + \mathbb{E}\left[v_T|0, t, \tau' > T\right]\right) \Pi(N$
where $\Pi(N(v,t)|v, t)$ is the probability of no arrivals which is independent of v given the as-

45

sumptions. This implies that the second term is linear in v, hence, we just need to prove that the first term is concave in v. We show that concavity is preserved under the Bellman equation. Assume W is concave in v'. Then

$$
\int_t^T \min\left(W\left(\varepsilon, v', \tau'\right), 0\right) dP\left(v', \tau' | \left(\alpha v_2 + \left(1 - \alpha\right) v_1\right), t\right)
$$

$$
= \int_t^T \min\left(W\left(\varepsilon, \alpha\left(v' + v_2\right) + \left(1 - \alpha\right)\left(v' + v_1\right), \tau'\right), 0\right) dP\left(v', \tau' | 0, t\right)
$$

$$
\geq \int_t^T \left(\alpha \min\left\{W\left(\varepsilon, \left(v' + v_2\right), \tau'\right), 0\right\} + \left(1 - \alpha\right) W\left(\varepsilon, \left(v' + v_1\right), \tau'\right)\right) dP\left(v', \tau' | 0, t\right)
$$

$$
= \alpha \int_t^T \min\left\{W\left(\varepsilon, v', \tau'\right), 0\right\} dP\left(v', \tau' | v_2, t\right)
$$

$$
+ \left(1 - \alpha\right) \int_t^T \min\left\{W\left(\varepsilon, v', \tau'\right), 0\right\} dP\left(v', \tau' | v_1, t\right).
$$

The second result follows immediately from concavity and the definition of a mean preserving increase in spread.