

Equivalent Certain Values and Dynamic Irreversibility*

Hugo Hopenhayn[†]

UCLA

Maryam Saeedi[‡]

Tepper, CMU

March 18, 2024

Abstract

We introduce a tractable methodology for analyzing dynamic decision problems and games involving irreversible decisions under uncertainty. By leveraging regularity properties, our approach offers an intuitive method for solving these problems using equivalent certain values, and derives properties and comparative statics of the solution. We show that irreversibility is analogous to information loss, leading agents to act as if they had worse information than with reversible actions. We use our methodology to analyze design features of previously intractable long auctions, establish revenue equivalence, and show that increasing bidding opportunities or allowing bid retraction can harm bidders and benefit the auctioneer.

Keywords: Irreversibility, Dynamic decision problem, Dynamic games, Value shading, Decision under uncertainty, Dynamic Auction Design, Information loss. **JEL codes:** C73, D44, D81

*We would like to thank John Asker, Matt Backus, Tom Blake, Alessandro Bonatti, Dennis Epple, Maryam Farboodi, Yingni Guo, Ken Hendricks, Johannes Hörner, Ali Hortacsu, Moritz Meyer-ter-Vehn, Paul Milgrom, Aniko Oery, Ariel Pakes, Michael Peters, Robert Porter, Andrea Prat, Debraj Ray, Ali Shourideh, Ennio Stacchetti, Andre Sztutman, Steve Tadelis, Robert Town, Anton Tsoy, as well as seminar and conference participants at Northwestern Economics Department Theory seminar, EIEF, Midwest Economic Theory Meetings, IOE, Canadian Economic Theory Conference, The Economics of Information and Communication Technologies Conference, SED, SITE, Dynamic Games, Contracts, and Markets Session, eBay, Econometrics Society World Congress, UCLA mini conference, Workshop in Industrial Organization (Toi 8), Econometrics Society Winter Meetings, Dartmouth Winter IO Conference, CMU, Tilburg University, NYU, Stanford University, and UC Berkeley. Any remaining errors are ours.

[†]hopen@econ.ucla.edu

[‡]msaeedi@andrews.cmu.edu

1 Introduction

Irreversible decisions in the face of uncertainty have been central to much economic research. This class of problems is often difficult to analyze and solve either analytically or numerically. In this paper, we propose a novel and tractable methodology by leveraging a set of regularity properties commonly present in a wide range of dynamic problems and games involving irreversibilities. The proposed approach offers a simple and intuitive method for analyzing and solving this class of problems and provides a range of valuable insights. The decision problems we analyze have the following structure. First, agents make irreversible decisions over a period of time, such as sunk investments in R&D, allocating capacity, or biddings in long auctions. Second, throughout the decision period, agents observe new signals. For example, innovators may learn about the value of obtaining a patent, firms might learn about outside opportunities, and bidders might learn about alternative auctions. Third, the opportunities to take actions are random. These random times represent random opportunities for undertaking actions such as R&D investments, or executing trades.

Our paper has two main contributions. First, we introduce a novel approach to solving and analyzing this class of problems that makes them very intuitive and tractable. Second, our methodology allows us to characterize several properties and comparative statics in this setting, including both dynamic decision problems and a class of games. In particular, we show that irreversibility in these problems is analogous to information loss in the sense of Blackwell, a la Blackwell [1951, 1953]. An agent faced with irreversible decisions acts as if it had worse information than with reversible actions. The implications of irreversibility for expected payoffs can thus be related to the impact of information loss, i.e., a mean-preserving contraction in the distribution of states over which actions are chosen. For example, we show that in the case of dynamic auctions, these irreversibilities result in a distribution of final bids that is a mean-preserving contraction of the distribution of the bids without irreversibility.

Our methodology involves decomposing the general dynamic decision problem into two components: 1) a dynamic option valuation problem (the pricing of an option); and 2) a static decision problem that assigns actions to each possible value of that option. To illustrate this approach, consider the simplest problem in our class given by an irreversible entry decision. A firm can enter a market with a random payoff $v(T)$ accruing in period T . The entry can take place at random decision nodes $\{\tau_n\}$ in the decision period $[0, T]$, at a given cost c_e , and payoff relevant information is received in this interval. The option valuation problem for each decision node is the answer to the following question: What is the highest entry cost (or purchase price) at which this option would be exercised? The assignment of an optimal entry decision is now trivial: enter the first time where this value exceeds the entry

cost c_e . This value is generally lower than the expected value $v(T)$ at this node, as entry gives up the option of waiting. We call this option value the equivalent certain value (ECV). ECV corresponds to the certainty value at which the decision maker would make the same decision in the absence of future opportunities.¹ What is more surprising is that the ECVs are the same *regardless of the decision problem* in our class. The specific choice problem only matters for the second step, the static decision problem of assigning actions to each ECV. This procedure is explained in more detail below. This separation result is particularly useful when considering the design of games, especially when the design involves repeatedly evaluating changes in the payoff functions.

ECVs are constructed by a self-generated expectation method that is independent of the specific payoff function. They are a function of the process of new information and decision-making times. To explain this method, we need to explain our modeling framework a bit further. The agent faces a random sequence of decision times which follow a joint Markov process.² At each decision time, given current information, the agent can make an irreversible decision, e.g., the outstanding bid or current action, which can only be adjusted upward. The final payoff is a function of the final action after the end of the decision period, and the final value, e.g., the expected payoff to a bidder is a function of their outstanding bid and realization of their valuation. The first step of our methodology is to assign an ECV to any decision node, with the property that the optimal action at the decision node is the same as the one at the end node with the valuation ECV, i.e., in a static setting without any further possible actions. This can be illustrated with the aid of Figure 1. In this figure, the x-axis represents the time, and the y-axis represents the expected value at a given point in time. Consider the decision node at (v_0, t_0) , its ECV is going to be a value at the end time, T , where the agent's optimal action is the same at (v_0, t_0) and (ECV, T) , let us call this optimal action a_0 . This figure also illustrates the indifference curve that passes through these two points, which represents all points in the value-time space where the agent would choose the same optimal action a_0 .

These indifference curves have a self-generating property. Starting at the point (v_0, t_0) , where the agent will choose the action a_0 , many paths can occur in the value-time space with various random opportunities for the agent to update the action upward. Among these paths, there is a subset in which the agent *does not* adjust their actions any further. This subset can be further divided into two subsets. First, there is the set of paths with no additional

¹This correspondence satisfies a self-generating property which is explained later on. It has a parallel in the recursive methods in [Abreu et al. \[1990\]](#).

²This is a standard assumption in the class of revision games, developed in [Kamada and Kandori \[2020b\]](#), in models of sticky prices following [Calvo \[1983\]](#), and models of optimal election campaign as in [Kamada and Sugaya \[2020\]](#), among many others.

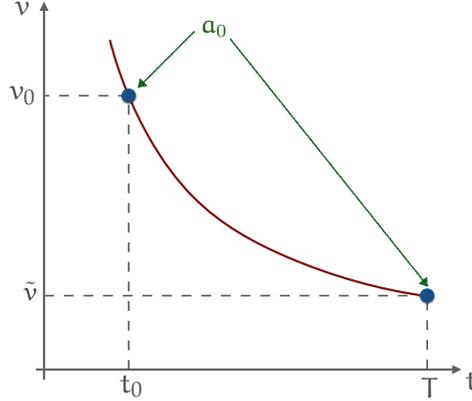


Figure 1: Indifference curves for actions

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

opportunities for adjustment. Second, there is the subset with the paths which have subsequent opportunities, but these opportunities will lie below this indifference curve, thus the agent would want to lower their action but given the irreversibility assumption, they cannot. The expected final value conditional on paths belonging in these two subsets is precisely the ECV. The same property holds for any point on this indifference curve. This property, which we call *self-generated expectation*, depends only on the stochastic process for values and decision times and is thus independent of the specific payoff function.³ Additionally, note that the second subset includes *adverse* future paths, therefore ECVs tend to be lower than unconditional expected value of agents, hence we get value shading.

The indifference curves tend to be downward sloping. This is because these adverse outcomes incorporated into the subset of paths 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. The downward-sloping indifference curve results in larger shading when the agent is further from the end time and less shading when the agent is closer to the end time, therefore it can result in a gradual increase of actions as we get closer to the end time.⁴

We provide a simple recursive method for the computation of ECVs. This can be intuitively illustrated by relating ECVs to a particular option value of a random asset. The asset has a final uncertain payoff $v(T)$. Consider the following security. A buyer can purchase this

³We borrow this term from [Abreu et al. \[1990\]](#). Although related, it is a different concept. Our indifference curves are self-generating as they define the boundaries on future realizations for the calculation of conditional expected values, which in turn are constant along these curves.

⁴An analogue result is found in [Harris and Holmstrom \[1982\]](#), where initially workers' wages are shaded below marginal products, as the wage is effective in the future only if it is less than or equal to the realized marginal product of the worker.

asset at some price $p(t)$ and is entitled to the final payoff. However, the seller retains the option to buy back the asset *at the same price*, but only *at random times* $\tau' \geq t$. If the seller does not exercise this option or if no repurchase opportunity arises, the buyer retains the final payoffs. At any of these random times, the seller updates the expected value of the asset based on the information received. Letting $v(t)$ be the expected final payoff $v(T)$ given (symmetric) information at that time, the (fair) price of this security $p(t)$ is precisely its ECV. Consequently, the repurchase option will be exercised at the first random trading time τ' whenever the corresponding price of this security (its ECV) exceeds $p(t)$.

Two crucial factors determining the degree of shading are the likelihood of future decision opportunities and the extent of new information received, measured by the conditional variance of expected final values. We show that irreversibility constrains the agent's actions, limiting its ability to condition its choice on the available information. Likewise, we illustrate that with an increase in frequency of arrivals, there is more value shading at the beginning and at the same time more frequent arrivals, which results in more arrivals toward the end time. Consequently, the final ECV is going to be closer to the final value, and hence final decisions are more aligned with the true values. Similar to the removal of irreversibility, more arrival can be bad for agents, for example, in auctions, it can result in increased prices and decreased welfares for the bidders.

Our results apply to dynamic games with incomplete information and privately observed actions or those with a dominant strategy, where our main theorem shows that we can find equilibrium strategies by solving a Bayesian Nash equilibrium of an associated static Bayesian game.⁵ The equilibrium strategies in the dynamic game are easily derived from those in the associated static Bayesian game. This result also extends to dynamic games with *observable* actions when the associated Bayesian game has an equilibrium in weakly dominating strategies.⁶

As a demonstration of the practical value of the model developed, we use our method to analyze optimal dynamic auction design when bidders receive new information about their private value over time. Dynamic auctions that run over longer horizons are becoming increasingly prominent, making it critical to incorporate evolving information. Our setting

⁵Even though these assumptions may appear restrictive, [Gershkov et al. \[2013\]](#) and [Kushnir and Liu \[2019\]](#) demonstrate that, under certain conditions, any allocation supportable through a Bayesian equilibrium can equivalently be supported by an alternative mechanism featuring dominant strategies. Consequently, if a market designer seeks to identify optimal allocations, their attention can be directed solely toward dominant strategies, utilizing our method to address the dynamic aspects of these strategies.

⁶Without this assumption, this result can also apply to finding open loop equilibria, which for large games might approximate closed loop equilibria (see [Fudenberg and Levine \[1988\]](#)). It can also be applied to solving for equilibria in mean-field games.

generalizes those in [Kamada and Kandori \[2020a\]](#), [Kamada and Kandori \[2020b\]](#), and [Kapor and Moroni \[2016\]](#), by allowing the valuation of participants to change throughout the course of the auction.

We first show that a generalization of revenue equivalence holds in our model even with recurring bid opportunities. This result relies on the reduction to a static auction using the distribution of ECVs.⁷ It can also be used to provide insight into optimal reserve prices, similar to those in [Riley and Samuelson \[1981\]](#) for static auctions. Their analysis suggests that if the distribution of values is very diffuse, the seller could benefit from setting higher reserve prices to extract revenue from the highest value bidders without seriously jeopardizing the probability of sale. Our results suggest that the opposite holds in the absence of bid retraction or when bidding opportunities are infrequent.

Additionally, we study the impact of irreversibility and the increase in bidding opportunities in long auctions. As mentioned above, irreversibility in dynamic problems is an analog of loss of information in a static problem. Therefore, allowing bid retraction in long auctions results in a mean-preserving spread of ECVs and the corresponding bids. This increases efficiency as bids become more correlated with true values. Paradoxically, bidders may be worse off as the increase in the spread of bids will typically raise prices and might lead to a reduction in surplus for the winner of the auction.⁸ Similar results are obtained when the frequency of bids is increased (keeping irreversibility in place), higher frequency leads to a mean-preserving spread of bids and therefore to a potential decrease of surplus for bidders. These results point to two design considerations, the irreversibility of bids, and the frequency by which bidders might be reminded about making choices.

Literature Review We propose a new methodology to solve dynamic decision problems involving irreversibility in a tractable way. As discussed, these problems arise in many economic and business contexts (see [Dixit et al. \[1994\]](#), [Lippman and McCall \[1976\]](#)). Some applications that have been studied before include R&D investments, capacity expansion, dynamic pricing under limited capacity, hiring and firing decisions, and dynamic auctions. The challenge posed in these problems is that the future arrival of information creates an option value to delaying decisions, yet actions are irreversible. This tension makes identifying optimal policies complex, [Arrow and Fisher \[1974\]](#), [Henry \[1974\]](#). These papers point out that the optimal solution requires accounting for option value which results in shading, but do not provide a computational solution.

⁷The conditions for symmetry are somewhat stronger, as they require that all agents face identical but independent processes for information and bid opportunities.

⁸This is true for large number of bidders (more than three).

A closely related literature is that of revenue management that looks at optimal dynamic pricing under limited capacity, as explored by [Zhao and Zheng \[2000\]](#) and as surveyed by [Elmaghraby and Keskinocak \[2003\]](#) and [Den Boer \[2015\]](#). This is an extensive literature that studies the problem of monopolies or oligopolists who want to maximize their revenue over a period of time when they are facing uncertainties and irreversibility. In a similar research, [Ito and Reguant \[2016\]](#) examines capacity contracting for power plants. [Betancourt et al. \[2022\]](#) study the problem of two airlines selling seats over time when there is a deadline.

Our methodology has a couple of parallels in the literature. There is a connection between the ECV and the Gittins index [Gittins et al. \[2011\]](#) in the statistical decision literature. Both concepts aim to reduce a dynamic decision problem to an index policy over simpler static problems. The Gittins index identifies the arm with the highest index as the one to play next. ECVs map states into indices to define optimal actions. The self-generating property that identifies the ECVs has a parallel in recursive methods in [Abreu et al. \[1990\]](#).

The rest of the paper is organized as follows. [Section 2](#) provides a simple example that conveys the main intuition and results in the paper. [Section 3](#) describes the general model and provides a set of applications that can fit the general model. [Section 4](#) discusses the intuitive and formal analysis of the model and describes how to embed the results into games. [Section 4.4](#) describes our main results. [Section 5](#) discusses an application of the model to dynamic second-price auctions and explores a few related design problems. In the online appendix, we discuss a few extensions of the baseline model to random termination time and give properties for the case where values are independent of Poisson arrival opportunities and also an application to the identification of anonymous sequential games. All proofs are deferred to the Appendix unless specified.

2 A Simple Example

We start our analysis by considering a two-period contest that illustrates some of the main features of our methodology. There are N players and two periods, $t = \{0, 1\}$. In the first period, after observing a private signal v_{i0} drawn from some distribution $H_i(v_{i0})$, agents choose the level of a private action $a_{i0} \geq 0$, e.g., studying for a test or allocating resources to a project. In the second period, each agent privately observes its value of winning the contest v_{i1} drawn from the conditional distribution $F_i(\cdot|v_{i0})$. With probability p_i , the agent has the option of increasing the action to any value $a_{i1} \geq a_{i0}$, e.g., studying more or allocating additional resources to the project. With probability $(1 - p_i)$, the agent is unable to revise its choice; therefore, a_{i0} remains its final action. The agent with the highest final action,

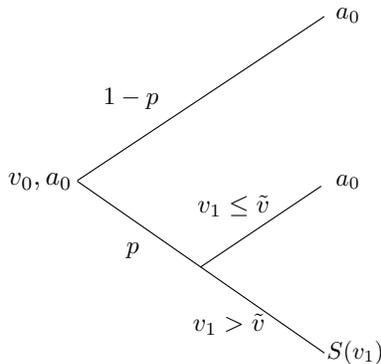


Figure 2: Decision tree

a_i , wins the competition and receives payoff $v_{i1} - a_i$. For any other player j , the payoff is equal to $-a_j$. Both the signals and final values are drawn independently across agents. For notational convenience, we suppress the index i unless needed to avoid confusion.

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

$$U(v, a) = G(a)v - a. \quad (1)$$

Assume there is a unique action a that maximizes (1) and it is strictly increasing in v . Denote this solution by $S(v)$. This is the optimal action in a static setting.⁹ 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 this information to illustrate the tree of the game in Figure 2.

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

Taking into account the best response in the second period, the choice of a_0 maximizes

⁹We make these assumptions and others below in the analysis of the example for convenience. The set of assumptions that are necessary for our main results is given in Sections 3 and 4.

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

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

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

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

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

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

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

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

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

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

The lottery over these final values has an expected value as given in Equation 4, which is equal to \tilde{v} . This is the key property defining the ECV 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 that the expected utility is linear in v and (strictly)

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

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

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

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

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

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

Because of the linearity of payoffs in v ,

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

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

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

that game, we can now assign $a_0(v_0) = \tilde{S}(\tilde{v}(v_0))$ and $a(v_0, v_1) = \max\{a_0(v_0), \tilde{S}(v_1)\}$ as equilibrium strategies in the dynamic game.

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

3 The Decision Problem

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

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

We can relax this assumption by allowing U to be linear in a strictly monotone function of v through a change of variables. For example, we can extend it to cases where the payoff function includes terms such as v^2 . By incorporating this new variable, all results presented in the paper can be extended accordingly.

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

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

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

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

By assuming that values and decision times are Markov, we can identify decision nodes with pairs (v_n, τ_n) corresponding to the realized signal and time in the last arrival. A decision strategy s specifies at each possible decision node a desired action $s(v_n, \tau_n)$, which is the choice the agent would make if unconstrained by previous 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 \mathcal{S} denote the set of strategies satisfying these conditions. These two assumptions will be maintained for the rest of the paper.

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

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

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

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

3.1 Examples

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

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

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

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

contracts. Or, a financial trader could sell future contracts that will be fulfilled by resorting to its network of intermediaries, forcing the trader to use more expensive sources in case larger quantities are needed.

While cash flows in the above setting accrue throughout the decision period, the problem can be mapped into the general setting where a payoff $p_t q$ received at time t is equivalent to a random final payoff $p_T q$ when p_t follows a martingale, therefore, $E_t(p_T | p_t) = p_t$. 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 to achieve 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.¹⁰

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

Entry Decisions At random times τ , the decision maker gets an opportunity to enter a market and a signal v_{τ} about the expected value of entry. The entry must take place before time T or the decision maker receives zero payoff. Final payoffs are $v_T - c$ in the case of entry. In this application, the action space is $A = \{0, 1\}$, representing the choice of no entry and entry, respectively. This application can be easily inscribed in an entry game.

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.¹¹ In these auctions, bidders frequently place multiple bids over time and increase them as the auction progresses. To model these auctions in the above class of decision problems, suppose that at

¹⁰These 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.

¹¹GovDeals is an auction platform used by government agencies to sell used equipment.

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 $[v_T - \mathbb{E}(b_2 | b_2 \leq b_T)] \text{Prob}(b_T \text{ is the highest bid})$. This expected payoff is linear in v_T and supermodular, as required in our general decision problem. This application and the related design questions are examined in more detail in Section 5.

Time Separable Payoffs with Discounting A decision maker has payoffs $u(v, a)$ that are received over time and has a constant discount factor β . Time is discrete. At each point in time payoffs are functions of a random value v_t and an action a_t given by the function $u(v_t, a_t)$. A decision maker chooses a sequence of non-decreasing and contingent actions a_t to maximize

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

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

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

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

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

Therefore, this transformation respects the original payoff structure. While we consider here time-zero payoffs, the same procedure applies to any future period. This formulation easily extends to random arrivals and a structure where arrivals and payoffs follow a general joint Markov process. Linear investment costs of the form $i_t = a_t - a_{t-1}$, as would occur in the

case of irreversible investment, can be easily accommodated in the above payoff function through the rearrangement and collection of the different a_t terms. An application of time separability to an anonymous sequential game is provided in Section C.3.

Building on the general framework and examples laid out above, we now turn to formally analyzing the model to derive our main results on how using ECVs helps reduce the dynamic optimization problem to a more tractable static one.

4 Main Results

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

4.1 Intuitive Analysis

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

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

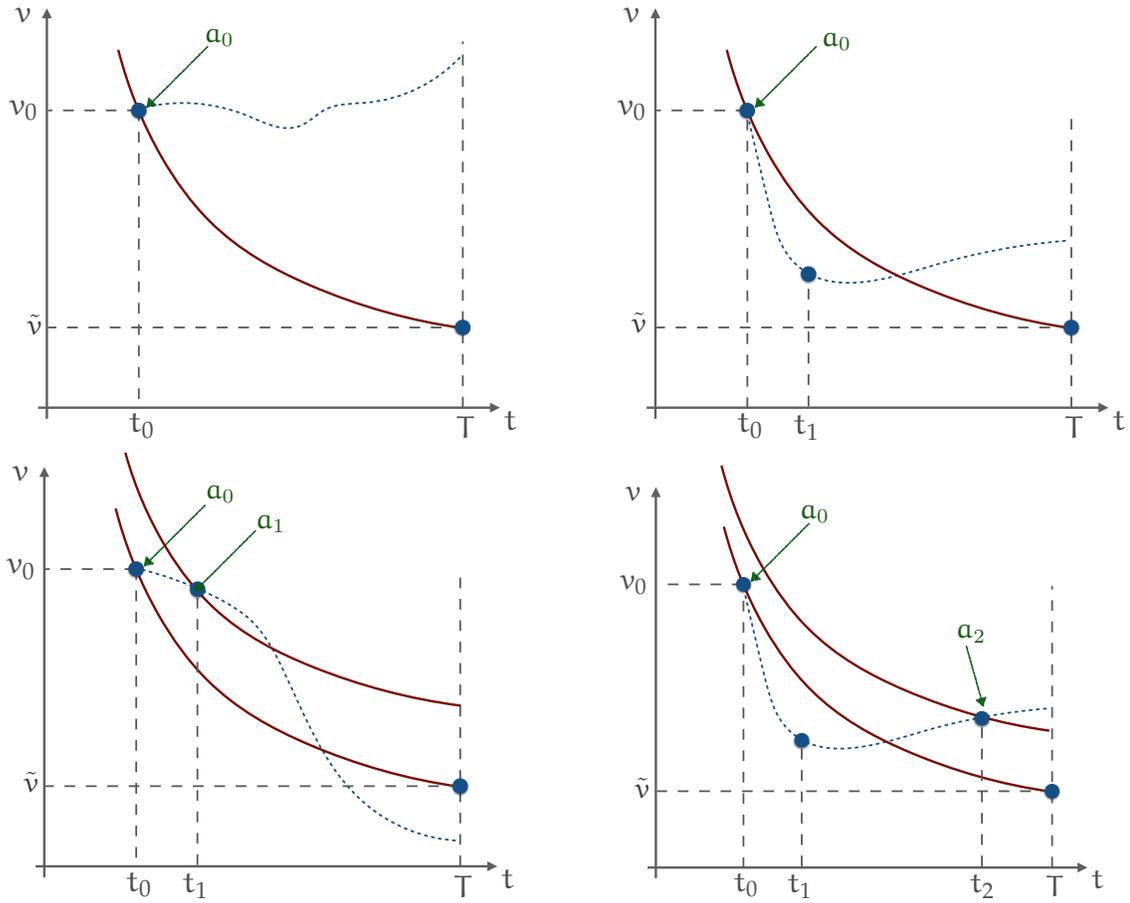


Figure 3: Indifference curves and increasing actions

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

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

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

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

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

An optimal decision strategy is derived as follows. Let

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

denote the solution to the static optimization problem for any value v . Then the optimal decision strategy at decision node (v, t) is given by $\tilde{S}(e(v, t))$.¹²

4.2 Formal Analysis

Here we state our main results that formalize the intuitive arguments given above and provide a sketch of the proof of our main theorem, while the complete proofs are relegated to the appendix. Our analysis in the previous section suggests a general approach to finding the solution to our decision problem. The steps in this proof are as follows.

1. Define a function $e(v, t)$ that partitions the set of states into indifference equivalent certain value (ECV) classes as above. We show that the function is uniquely defined and can be obtained via solving a dynamic programming problem. This is independent of the specific payoff function U .

¹²If the maximum is not unique, \tilde{S} is an increasing selection.

2. Define a candidate-optimal decision strategy at decision node (v, t) as choosing the action $\tilde{S}(e(v, t))$, where $\tilde{S}(\cdot)$ is the function defined by equation 8, and prove that this is an optimal decision strategy for the dynamic problem.

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

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

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

Assumption 3. *The following properties hold:*

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

The first assumption states there is a 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 the case where arrival times and values are independent, we have $\int_{N(v, t)} (v_T(\omega)) d\Pi(\omega|v, t) = v\Pi(N(v, t) | v, t)$, therefore, 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 arrival rate λ .

4.2.1 Self-Generated Expectation

In this section, we define and characterize the ECV function $e(v, t)$ that partitions the decision space into indifference classes. This function is defined implicitly by a recursive problem that satisfies a property that we call self-generated expectation.

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

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

where, as defined earlier, $D(\omega)$ is the set of all decision nodes for a given ω . If the above set is empty, set $\bar{e}(t^+, \omega)$ to an arbitrarily low number.¹³ This will be the maximum equivalent value achieved after time t . For every state (v, t) such that $0 \leq t \leq T$, let

$$H(\varepsilon, v, t) = \{ \omega \mid (v, t) \in D(\omega) \text{ and } \bar{e}(t^+, \omega) \leq \varepsilon \},$$

These are all the histories following (v, t) such that for any subsequent arrival $e(v', t') \leq \varepsilon$. Relating to our previous intuitive analysis, these are all histories where values remain below an indifference class represented by the threshold ε .

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

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

The above definition is equivalent to the following:

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

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

The derivation of a self-generated expectation follows a recursive structure. First, note that $e(v, T) = v$, since this is a terminal node. Intuitively, working backward from that point using (10) should give a unique self-generated expectation $e(v, t)$. While (10) seems like a complicated functional equation, we can find the solution by considering the following auxiliary functional equation:

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

and $W(\varepsilon, v, T) = v - \varepsilon$. This functional equation is a contraction mapping. Using additional assumptions on the arrival time process, one can transform this functional equation to a

¹³For example, set it equal to $\inf e(v, t)$.

simple PED that is easy to compute.¹⁴

Proposition 1. *Given Assumption 3, the unique function satisfying equation (12) is given by:*

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

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

Notice that this is a standard dynamic programming problem that can be easily solved by backward induction, starting at the terminal node T where $W(\varepsilon, v, T) = v - \varepsilon$.¹⁵ Here we give an overview of the steps involved in the proof. The first step consists of showing that using any self-generated expectation function we can construct a function W that satisfies the functional equation (12) and the condition that $W(e(v, t), v, t) = 0$. Then we show that by Assumption 3 the functional equation (12) is a contraction mapping, so it has a unique solution. It also follows easily that this function is strictly decreasing in ε and continuous. Moreover, it is greater than or equal to zero when $\varepsilon = 0$ and negative for large ε . It follows by the intermediate value theorem that there is a unique value $e(v, t)$ such that $W(e(v, t), v, t) = 0$. It immediately follows that this solution satisfies (10).

4.2.2 The Optimal Solution

Consider the problem

$$\max_a U(v, a). \tag{13}$$

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 deciding at the final node T with a value $v_T = v$.

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

¹⁴In a companion empirical paper, we apply this method to the case of online auctions, in that paper we assume that arrival times follow a poisson distribution and values change using a brownian motion, [Hopenhayn and Saeedi \[2020\]](#). Given these assumptions, equation 12 becomes a simple PDE, similar to the heat transfer PDF that is easy to compute in any programming language.

¹⁵As we show in Section C, this problem is further simplified when the Markov process has independent increments, a condition that is satisfied in many applications, where the value function does not depend on v .

This theorem allows us to break down the problem of finding the optimal strategy $S(v, t)$ in two steps: (i) solve for the equivalent certain values $e(v, t)$ as indicated in Section 4.2.1 and (ii) Solve the static problem (13) and use its solution to assign strategies $S(v, t) = \tilde{S}(e(v, t))$. Furthermore, the solution $e(v, t)$ for the first step is *independent* of the specific function $U(v, a)$. This is particularly useful when estimating this function or in games, where the function might depend on strategies of other players, as the first step must be performed only once.

In this section, we formally showed how ECVs are calculated and how they can be used to find optimal actions in a dynamic problem. The following two examples can help further with intuition. Example 1 describes an easy way to interpret the ECVs as the option value for an asset with random repurchase. Example 2 shows how ECVs can be used to solve a dynamic entry problem.

Example 1: ECVs can be related to an intuitive asset pricing problem. Take an asset with random payoff $v(T)$ that will accrue in period T . Consider the following contract: at a price p , a trader can buy this asset, but the original owner maintains the option to repurchase the asset at the same price, but only at random arrival times. In turn, the trader has the right to the final payoff only in the event that the repurchase does not take place. The joint distribution of arrival times and information is identical to that in the above setting. The (break-even) price of this asset is precisely its ECV, as demonstrated in Proposition 2. The key step in the proof is to show that if the price of the asset $p = e(v, t)$, the original owner will exercise the repurchase option in any future trading window if and only if $e(v', t') \geq e(v, t)$.

Proposition 2. *The equilibrium price of the asset with a random repurchase option is equal to the ECV.*

Example 2: This example is related to the optimal entry or purchase option. Our procedure separates the dynamic option value problem, which partitions states (v, t) into equivalent classes of ECVs, from the assignment of strategies. A firm can enter a market (or a trader buys an asset) with a random payoff $v(T)$ that accrues in period T . Information and opportunities to exercise this option are given by the process $\{(v, t)\}_{t \leq T}$ as described in our setting. ECVs are the answer to the following question: What is the highest entry cost (or purchase price) at which this option would be exercised? The assignment of strategy for the optimal decision rule is now trivial: enter in the first opportunity where $e(v, t) \geq c_e$.

Having shown how ECVs allow simplifying the dynamic decision problem, we now expand

the approach to apply to strategic environments and a class of games with incomplete information.

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 ECVs, as defined above. In contrast to above decision problem, the vector of strategies has the additional restriction that the strategies must conform an equilibrium, i.e., be mutually best responses. We define a static Bayesian game where the distribution for each player's type is the distribution of the ECVs for that player, and strategies map these values into their corresponding action sets. Finally, we establish that any equilibrium of this static Bayesian game defines equilibrium strategies for all players in the original game.

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

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

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

Consider player i . For every history ω , we can identify a unique value corresponding to the highest final equivalent value reached, $v_i(\omega) = \max\{e_i(v_n(\omega), \tau_n(\omega))\}$, for the corresponding path. This procedure uniquely determines a distribution Ψ_i of ECVs for this player that depends only on the corresponding Markov process P_i for decision nodes (v, t) . Define the

(static) Bayesian game as follows: set of players $I = \{1, \dots, N\}$, distribution of values for each player Ψ_1, \dots, Ψ_N , strategy sets A_1, \dots, A_N , and payoff function $u_{iT}(v_i, a_i, a_{-i})$.

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

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

Theorem 2. Consider a game Γ that satisfies Assumption 4 and its associated Bayesian game Γ_B . Additionally Markov processes $P_i(v, t)$ satisfy Assumption 3, for all i . For any vector of equilibrium strategies $\{\tilde{S}_i\}_{i \in N}$ of Γ_B the strategies defined by $S_i(v, t) = \tilde{S}_i(e_i(v, t))$ are an equilibrium for Γ , where the function $e_i(v, t)$ is the self-generating expectation for player i .

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 ECVs—and (2) a static equilibrium determination—the Bayesian game. This decomposition is possible because the dynamic decision problem depends only on the stochastic process for values and decision times, not opponent strategies. The actual choices made in the dynamic decision problem depend on the strategies of others, but not the determination of ECVs. Loosely speaking, the dynamic problem determines the indifference maps (the function $e(v, t)$), while the solution to the Bayesian game labels them with the actions.

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

The decomposition described above is possible in part from our assumption that no information from other players' actions or values is revealed throughout the game. For some special cases, this assumption can be relaxed. In particular, if for all v_i there is a weakly dominant action choice a_i that maximizes $u(v_i, a_i, a_{-i})$ for all a_{-i} , the equilibrium derived above remains an equilibrium of the dynamic game with any added information about opponents' values and strategies. This assumption holds, for example, in the case of a dynamic second-price auction as described in Section 5. This method can also be useful when considering the question of implementation in dominant strategies. Another potential setting for our method is similar to that of [Bonatti et al. \[2017\]](#), where in a dynamic Cournot setting sellers have incomplete information about their rivals: sellers observe the revenue process but not individual rivals' actions.

4.4 ECVs and Shading

Our leading example suggests that agents will shade values, and consequently actions, as a result of the adverse selection problem introduced by the option of future decision opportunities and the irreversibility of decisions. Value shading is defined by the property that early in the decision process agents make choices as if the final values were lower than the conditional expectation.¹⁶ Value shading arises because of the irreversibility of actions, the opportunity of making future choices, and the arrival of new information concerning the final value. Two of the key forces determining the extent of shading are the likelihood of future decision nodes and the extent of new information received 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 C.2. Our first result in this section concerns the general existence of value shading. We next discuss the conditions under which value shading decreases over time as the end period approaches. Finally, we show that as the frequency of future bidding increases, so does the shading.

Assumption 5. *We assume the following assumptions hold.*

1. For all $t < T$, $\int_{N(v,t)} v_T(\omega) d\Pi(\omega|v, t)$ is strictly increasing in v .
2. For all $t < T$ and v , $\int_{v' > v, t < t' < T} dP(v', t'|v, t) > 0$.

The first part of the assumption would immediately hold given the martingale property if expected values are independent of arrival times. The second part of the assumption holds when the process is independent of the arrival time t' , for example, if the probability of a future arrival is strictly positive and the conditional distribution for value is nondegenerate, i.e., $P(v' = v|v, t) < 1$.

Proposition 3. *The equivalent value is less than or equal to the expected value: $e(v, t) \leq \mathbb{E}[v_T|v, t]$. If Assumptions 5 hold, then the inequality will be strict.*

In terms of decisions, shading of values implies shading of actions given the assumptions on utility function, so that $S(v, t) = \tilde{S}(e(v, t)) \leq \tilde{S}(\mathbb{E}[v_T|v, t])$.

This proposition implies that at any time before the endpoint, values and the corresponding choices are shaded. One might wrongly conclude that distribution of ECVs, and the corresponding choices are downward biased. In fact, ECVs are unbiased, as established in the

¹⁶Value shading can occur for strategic reasons; for example, in Hortaçsu et al. [2018], it is a result of dealer market power in uniform-price Treasury bill auctions. Our source of shading is distinct and fundamentally nonstrategic.

following proposition. There is an intuitive reason: the value $e(v, t)$ and the corresponding action $\tilde{S}(e(v, t))$ will remain final only when no higher equivalent value is reached. While values are shaded according to Proposition 3, they remain unbiased for the set of paths where these values are the terminal ECVs. This follows directly from the definition of self-generated expectation given by Equation 10. The following proposition formally proves this intuition.

Proposition 4. *The terminal ECV is unbiased, namely: $\mathbb{E}[\max\{\bar{e}(t^+, \omega), e(v, t)\} | \omega] = \mathbb{E}[v_T | v, t]$.*

In words, this proposition says that, starting at any decision node (v, t) , the distribution of terminal ECVs is unbiased, so its mean equals the expected final value at that point.

Interestingly, this proposition implies that the stochastic process defined by $e(v, t)$ can be used to recover the values at all relevant decision nodes (v, t) and the corresponding shading $E(v_T | v, t) - e(v, t)$. Equivalently, given the strategies $\tilde{S}(\cdot)$ corresponding to optimal choices, shading can be identified from the stochastic process for observed choices $S(t)$ without any knowledge of the corresponding values by inverting $\tilde{S}(e(v, t)) = S(t)$.

Corollary 1. *(Identification) For a given strategy $\tilde{S}(\cdot)$, shading at all relevant decision nodes can be identified from the observed stochastic process for choices $S(t)$.*

Shading Over Time Shading arises from the value associated with the option of a future choice. This is affected by the likelihood of having this opportunity and the arrival of new information. For example, if after a decision node (v, t) there is a certainty of another arrival before T , then $e(v, t)$ should be equal to the minimum value in support of v_T . It then seems natural that as the time limit T is approached, the degree of shading should decrease. In our general formulation, this might not be true for a couple of reasons. First, opportunities for making decisions might not be independent of each other.¹⁷ Second, the process for information arrival does not need to be stationary and could depend on decision nodes.¹⁸ To examine the evolution of shading over time, we consider two different cases where the degree of shading, given by $v - e(v, t)$, decreases. The first case is when decision times

¹⁷As 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.

¹⁸As an example, suppose there are three periods. If the variance of the distribution of arrivals in the third period increases when there is an arrival in the second period, this could lead to an increase in shading between decision node 1 and decision node 2. In contrast, if the event occurs, $e(v, t_0)$

are independent of values, and the second is when the next arrival time is independent of calendar time.

Proposition 5. *ECVs, $e(v, t)$, are increasing in t if either of the following assumptions holds.*

1. Values and arrival times are independent of each other and $F(\tau'|t)$ is (weakly) decreasing in t . Poisson process with arrival rate $\lambda(t)$. $F(\tau|t) = 1 - e^{-\int_t^\tau \lambda(s)ds}$
2. $P(v', \tau'|v, t) = P_v(v'|v, \tau' - t) F(\tau' - t)$.

It follows immediately that holding fixed values, v , optimal actions will be increasing over time. For the first set of conditions to hold, it is sufficient that hazard rates for arrival $h(t)$ are independent of previous arrivals, for example in the case of the Poisson process with arrival rates $h(t)$, regardless of the properties of this function. This can accommodate the case where agents are more attentive as end time approaches, so $h(t)$ would increase. While the first assumption covers many relevant cases, some other useful ones are excluded. In particular, assuming that the new arrival τ' and the new value v' are independent excludes learning environments where information is accumulated over time prior to the next decision node, as considered in Section 5. The second assumption is applicable to a learning environment. It requires that the time to the next decision node be independent of the current calendar time t , and that the next value v' depend only on v and the time elapsed up to this decision node. The former would hold, for instance, if arrivals follow a standard Poisson process with constant hazard rate.

Implications of Irreversibility Proposition 4 shows that the final equivalent value gives an unbiased estimate of the final value. In the absence of irreversibility, i.e., when the action a is not restricted to be positive, we can follow a similar method to solve for the optimal (or equilibrium) policy. For any path ω , define $\bar{e}(\omega)$ as the maximum equivalent value achieved on that path and define *final value* for the path ω by

$$\bar{v}(\omega) = E(v(T) | v(t, \omega), t \text{ is the last arrival for path } \omega).$$

If arrival times and valuation follow independent values, then $\bar{v}(\omega)$ will simply be $v(T)$. This is the only relevant value for the final action that will be taken on this path, which then solves $a(w) = \max_a U(a, \bar{v}(\omega))$. To assess the implications of irreversibility, the following proposition compares the distribution of $\bar{e}(\omega)$ with that of $\bar{v}(\omega)$.

Proposition 6. *The distribution of final values $\{\bar{v}(\omega)\}$ is a mean preserving spread of the distribution of final equivalent values $\{\bar{e}(\omega)\}$.*

The intuition behind this result, is that irreversibility constrains the actions of the agent, limiting its ability to condition choice on information. The specific implications for the choice of actions depend on more specific details of the utility function $U(a, v)$. For example, if the optimal policy $a(v) = \max_a U(a, v)$ is concave (convex), the expected action will be higher (lower) in the presence of irreversibility. The comparison is more complicated in the case of games, as the utility function itself is a function of the equilibrium strategies of the other players. In Section 5 we consider a specific application to a dynamic second-price auction such as the one corresponding to eBay’s setting with proxy bidding.

The increase in spread has an additional intuitive interpretation related to how informed the agent’s decisions are. At one extreme, with full information, the agent observes the true value $v_T(\omega)$ when choosing the final action. In our setting, the distribution of $\bar{e}(\omega)$ represents the implicit information in the agent’s decisions. Being a mean-preserving contraction of $v_T(\omega)$, it indicates limited information. Furthermore, compared to the case with reversible actions, the distribution of terminal values $\bar{e}(\omega)$ under irreversibility is itself a mean-preserving contraction of the one with reversible actions, and this implies that it has a lower covariance with the true values $v_T(\omega)$. This means that the agent faced with irreversible decisions acts as if they had worse information than with reversible actions. The mean-preserving contraction in the distribution of $\bar{e}(\omega)$ captures the loss of information due to irreversibility constraints.

Frequency of Arrival Times What is the effect of increasing the frequency of arrival times? This can be a design feature; agents are often reminded of making a choice or the time left before the end of an auction. This can be a design feature available to the market designer to either make it easier to change actions or make it harder. Intuition suggests that this should increase the option value, therefore ECVs should decrease as well. To obtain analytical result, we assume that arrival times are independent of the process for the values.

Assumption 6. $P(v', t'|v, t) = P_v(v'|v, t) F(t'|t)$

We assume also that v follows a martingale, so $E(v'|v) = v$. We will say that the distribution $\tilde{F}(t)$ has more frequent arrivals than $F(t)$ is $F(t'|t) \leq \tilde{F}(t'|t)$, so $F(t'|t)$ first order stochastically dominates $\tilde{F}(t'|t)$.

Proposition 7. *Suppose $F(t'|t) \leq \tilde{F}(t'|t)$ and $F(t', t)$ is non-increasing in t , for all t, t' . Then the corresponding ECVs $e(v, t) \geq \tilde{e}(v, t)$. Furthermore, this implies that the distribu-*

tion of terminal ECVs $\bar{v}(\omega)$ under \tilde{F} is a mean preserving spread of the one corresponding to F .

This proposition shows that more frequent opportunities for taking actions will result in a mean preserving spread of the distribution of over $\bar{v}(\omega)$. As in the case of reversible actions, this represents improved information and a higher covariance between $\bar{e}(\omega)$ and $v_T(\omega)$.

5 Application to Dynamic Second-Price Auction Design

We now demonstrate an important practical application of the model to optimal dynamic auction design. As discussed in Section 3.1, many auctions take place over extended periods of time. The tools we have developed help provide insights into bidding behavior and auction performance based on various aspects important in auction design.

By using our methodology, analyzing problems that fit our set up becomes more tractable. This tractability can help with optimal design questions related to dynamics, as well as numerical estimation of dynamic problems. In this section, we show how our methodology can be applied to dynamic second-price auctions and explore a few related design problems. In the appendix, we show how this method can help with estimations.

Many auctions take place over a considerable length of time; such is the case of trading platforms (e.g., eBay, GovDeals) and other settings (e.g., procurement, spectrum). As a result, dynamic considerations can be important for understanding bidding behavior and improving auction design. In particular, during these long auctions, bidders' valuations and strategies are likely to be affected by information that arrives during the auction; however, most of the literature has abstracted from this feature.¹⁹ Here, changes in value could come from several sources: the existence of alternatives that change the outside value, preference shocks (e.g., change of plans when buying event tickets), cost or capacity shocks (as in electricity markets, see Ito and Reguant [2016]), information regarding complementary goods

¹⁹There is also a literature strand on modeling and estimating dynamics across auctions. The classic paper is Jofre-Bonet and Pesendorfer [2003], which estimates dynamic auctions in procurement by controlling for the utilized capacity of participants. More recent papers that consider the option value faced by bidders in sequential auctions include Zeithammer [2006], Said [2011], Board and Skrzypacz [2016], Hendricks and Sorensen [2018], Backus and Lewis [2012], Bodoh-Creed et al. [2021], and Coey et al. [2020]. As a result of this option value, changes in the alternative items can alter the reservation price for bidders over time. While these papers focus on dynamic bidding across auctions, they assume that bidding within each auction happens instantaneously. Nevertheless, these papers motivate our reduced form approach toward the change in valuation to be a result of changes in these outside options.

(as in the case of spectrum auctions, see [Börger and Dustmann \[2005\]](#) and [Bulow et al. \[2009\]](#)), and alternative demands for use of resources in the face of capacity constraints (as in procurement auctions, see [Jofre-Bonet and Pesendorfer \[2003\]](#)).

For simplicity, we consider the case of sealed-bid second-price auctions. Using our approach, the optimal bid at time (v, t) in a dynamic auction, equals the optimal bid in a static auction for the corresponding ECV; since this is a second price auction, $b(v, t) = e(v, t)$. Furthermore, the optimal bid of a static sealed-bid second-price auction is simply bidding one’s value or the ECV in this case. Consequently, at any point in time, the standing bid of a bidder is the maximum of ECVs of the bidder during their decision nodes.

We also assume that Assumption 5 holds and that the changes in values and arrival times are independent of each other. Therefore, Proposition 3 and Proposition 5 hold. The propositions tell us that the bid is strictly less than the expected value at the time and increases over time. The exact shading function will be determined by solving the functional equation 12, given the process for the valuation and bidding time opportunities.²⁰

Using the results so far, we can act as a market designer and explore the impact of changing the environment on auctioneer revenue and also expected revenue of bidders, and in particular the winner of the auction. We will first show that revenue equivalence holds, then argue how optimal reserve price can be found in this environment, third, we study the impact of allowing for bid retraction on bidders, and finally we explore the impact of increasing arrival rates.

5.1 Revenue Equivalence

How does the type of auction chosen by the auctioneer affect expected revenues in the dynamic setting? Like in the case of static auctions, the answer depends on details. Only in the case of independent private values with symmetric distributions and risk neutrality, revenue equivalence holds. This result translates to our setting also in the special case where these conditions hold for the corresponding distributions $\bar{e}_i(\omega)$ for all bidders. The assumption of risk aversion holds in our setting. The assumption of independence requires that the processes for arrival times and information be independent across bidders. The assumption of symmetry requires these processes to be identical. Under these two conditions, the expected revenue for the seller and the bidders’ values are independent of the auction format. Next, we will discuss different design choices.

²⁰In [Hopenhayn and Saeedi \[2020\]](#), this model is estimated with the data from eBay and GovDeal. The estimates show considerable shading and explain a considerable amount of skewness in bidding times. The model is then used to perform a series of counterfactuals and to assess the implications of alternative designs on bidders’ welfare and sellers’ revenue.

5.2 Optimal Reserve Price

In our environment, bidders are indifferent between winning the object at the bid they submit and losing the auction at this bid. Therefore, having a dynamic reserve bid cannot increase bidders' bid they will submit throughout the auction.²¹ On the one hand, if the reserve price is higher than the second highest bid but lower than the highest bid, it can increase the revenue that the auctioneer receives. On the other hand, if the reserve price is higher than the highest bid, it will result in no transaction. These forces are exactly the same as those present when studying the optimal reserve price for static auctions. Our methodology shows that we can apply the same techniques as in the static case and find the reserve price to be equal to the value with zero virtual value. The difference is that the distribution to consider is the distribution of $\bar{e}(\omega)$ rather than the distribution of final value of the bidders. This is affected by the dynamics and will change by change of rebidding opportunities or how much uncertainty exists about the value of the object.

5.3 Allowing Bid Retraction

Another design feature worth exploring is allowing bidders to *retract* their bids or simply adjust their bids downward if they wish at any future opportunity. This will remove the incentive to shade one's bid in earlier opportunities. Because in the future, bidders are not bound by their earlier bids. Therefore, when allowing bidders to retract their bids, they would submit their expected value of the object at any given point in time. This expected value might fluctuate due to new information that arises, but it does not result in the shading we discussed earlier.

We can use the result of Proposition 6 for this case. The distribution of bids under bid retraction will equate $\{\bar{v}(\omega)\}$, while the distribution of bids under our baseline model equates $\{\bar{e}(\omega)\}$. This proposition indicates that the bids under retraction are mean-preserving spread of the baseline's bid distribution. This implies that the expected auction price, which corresponds to the second-order stochastic of bids, is likely to change. If the number of bidders is small, permitting bid retraction could lead to a reduction in expected prices. Furthermore, this has the potential to enhance allocation efficiency, as bidders would have the option to decrease their bid upon receiving unfavorable news. In such scenarios, allowing bid retraction results in an increase in bidders' expected revenue, while decreasing the expected revenue for the auctioneer.

²¹Liu et al. [2019] study the problem of dynamic reserve price, in their study dynamics arises across sequential auctions. They find that without commitment, the auctioneer benefits from changing the reserve price over time.

However, when the number of bidders is large, an increase in the spread will lead to a higher expected price in the equilibrium, which is favorable for the auctioneer. The effect on the bidders in this scenario is somewhat more nuanced: there will be an enhancement in allocation efficiencies and, simultaneously, an increase in prices paid when a bidder secures the highest bid. In our numerical simulations, the bidders were on expectation generally worse off.

We find that the ability to retract one's bid can harm the bidders, which is counter-intuitive at first glance. One should consider the price and value of the winner to interpret this result. This is similar to the prisoner's dilemma – each bidder would benefit from having the ability to retract their bid if they wish. However, they are negatively affected by their peers' ability, as it leads to higher expected prices. Another way to interpret these results is similar to the interpretation of common-value auctions and the winner's curse. In presence of common values, bidders will shade their bids, therefore, in equilibrium, instances of regret are not so frequent. This is also the case for the dynamic auctions. When retracting bids are not possible, bidders are more cautious and shade their earlier bids, and therefore, the instances of regret are not very frequent, especially as number of bidders increase. Consequently, the harm of higher prices outweighs the benefits of improved allocation efficiencies.

5.4 Increasing Arrival Rates

Assuming that the process for arrival rate and the process for change in values are independent, we can use Proposition 7 to show that more frequent arrivals results in higher shading. As a result, an increase in arrival time will result in lower bids at the beginning but more frequent bids as well; given that the average of all bids received are equal to expected value of the object, this result in a mean preserving spread of bids. This mean-preserving spread is formally shown in Proposition 7.

The implications of bids being mean-preserving spread is similar to those in bid retraction section, if number of bidders is low that will result in lower expected price (second-order stochastic), and if the number of bidders are large then the expected price rises.

As a designer who wants to choose the arrival rate, if the number of bidders is deterministic, setting them to zero or infinity will be the best option depending on the number of bidders, but if there is any uncertainty on the number of active bidders in an auction with a good chance of having a low number of bidders, then the optimal policy will be interior and determined based on the parameters of the market.

6 Final Remarks

In this paper, we propose a novel methodology to analyze dynamic decision-making problems with irreversibility and its extensions to a class of dynamic games. The theory relies on simple yet rich regularity conditions with potentially broad applicability. Our methodology decomposes the problem of finding the optimal solution into two steps; the first consists of solving a dynamic option valuation problem which is independent of the specific payoff function. The second step assigns strategies to decision nodes by solving a static decision problem, where the specifics of the payoff function are used.

Our approach has several advantages over existing methods. First, it provides a more tractable and intuitive solution method. This is particularly useful in the class of dynamic games considered, where solving for best responses reduces to a static Bayesian game. Similarly, when estimating parameters of the payoff functions, the dynamic option value problem needs to be solved only once, as opposed to the standard approach, where this must be done every time parameters are changed.

This separation result is also particularly useful when considering the design of the game, as it involves evaluating changes in the payoff functions. As an application, we consider design features of previously intractable long auctions, establish revenue equivalence, and show that increasing bidding opportunities or allowing bid retraction can harm bidders and benefit the auctioneer.

Second, our approach simplifies the analysis of this class of problems, characterizing properties of the solution, and evaluating the comparative statics. In particular, we provide an intuitive interpretation of the effect of irreversibility in dynamic problems. The irreversibilities are akin to information loss in static problems; the agent will act as if they have less accurate information about their payoff relevant state.

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

The analysis of dynamic games has proven to be a difficult problem. There are obvious trade-offs in research and corners to cut. Our paper is no exception, and we have our share of strong assumptions. In particular, we have chosen to represent the impact of information on values and the existence of decision time frictions in a reduced form, given by the stochastic process for values and decision time opportunities. There are obvious shortcomings, but the

payoff is a parsimonious representation of equilibria and a very tractable general structure that could be easily used in answering many design questions or other 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. **1, 3, 1**
- Kenneth J Arrow and Anthony C Fisher. Environmental preservation, uncertainty, and irreversibility. *The Quarterly Journal of Economics*, 88(2):312–319, 1974. **1**
- Robert B Ash. *Real Analysis and Probability: Probability and Mathematical Statistics: a Series of Monographs and Textbooks*. Academic press, 1972. **B**
- 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>. **19**
- Jose M Betancourt, Ali Hortaçsu, Aniko Oery, and Kevin R Williams. Dynamic price competition: Theory and evidence from airline markets. Technical report, National Bureau of Economic Research, 2022. **1**
- David Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102. The Regents of the University of California, 1951. **1**
- David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953. **1**
- Simon Board and Andrzej Skrzypacz. Revenue management with forward-looking buyers. *Journal of Political Economy*, 124(4):1046–1087, 2016. **19**
- Aaron L Bodoh-Creed, Joern Boehnke, and Brent Hickman. How efficient are decentralized auction platforms? *The Review of Economic Studies*, 88(1):91–125, 2021. **19**
- 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. **5**

- Jeremy Bulow, Jonathan Levin, and Paul Milgrom. Winning play in spectrum auctions. Technical report, National Bureau of Economic Research, 2009. 5
- Guillermo A Calvo. Staggered prices in a utility-maximizing framework. *Journal of monetary Economics*, 12(3):383–398, 1983. 2
- Dominic Coey, Bradley J Larsen, and Brennan C Platt. Discounts and deadlines in consumer search. *American Economic Review*, 110(12):3748–3785, 2020. 19
- Arnoud V Den Boer. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science*, 20(1):1–18, 2015. 1
- Avinash K Dixit, Robert K Dixit, Robert S Pindyck, and Robert Pindyck. *Investment under uncertainty*. Princeton university press, 1994. 1
- Wedad Elmaghraby and Pinar Keskinocak. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management science*, 49(10):1287–1309, 2003. 1
- Drew Fudenberg and David K Levine. Open-loop and closed-loop equilibria in dynamic games with many players. *Journal of Economic Theory*, 44(1):1 – 18, 1988. ISSN 0022-0531. doi: [https://doi.org/10.1016/0022-0531\(88\)90093-2](https://doi.org/10.1016/0022-0531(88)90093-2). URL <http://www.sciencedirect.com/science/article/pii/0022053188900932>. 6
- Alex Gershkov, Jacob K Goeree, Alexey Kushnir, Benny Moldovanu, and Xianwen Shi. On the equivalence of bayesian and dominant strategy implementation. *Econometrica*, 81(1): 197–220, 2013. 5
- John Gittins, Kevin Glazebrook, and Richard Weber. Multi-armed bandit allocation indices. *John Wiley & Sons*, 2011. 1
- Milton Harris and Bengt Holmstrom. A theory of wage dynamics. *The Review of Economic Studies*, 49(3):315–333, 1982. 4
- Kenneth Hendricks and Alan Sorensen. Dynamics and efficiency in decentralized online auction markets. Technical report, National Bureau of Economic Research, 2018. 19
- Claude Henry. Investment decisions under uncertainty: the” irreversibility effect”. *The American Economic Review*, 64(6):1006–1012, 1974. 1
- Hugo Hopenhayn and Maryam Saeedi. Bidding dynamics in auctions. 2020. 3.1, 14, 20

- Ali Hortaçsu, Jakub Kastl, and Allen Zhang. Bid shading and bidder surplus in the us treasury auction system. *American Economic Review*, 108(1):147–169, 2018. 16
- Koichiro Ito and Mar Reguant. Sequential markets, market power and arbitrage. *American Economic Review*, 106(7):1921–1957, 2016. 1, 3.1, 5
- Mireia Jofre-Bonet and Martin Pesendorfer. Estimation of a dynamic auction game. *Econometrica*, 71(5):1443–1489, 2003. 19, 5
- Yuichiro Kamada and Michihiro Kandori. Cooperation in revision games and some applications. *Global Economic Review*, 49(4):329–348, 2020a. 1
- Yuichiro Kamada and Michihiro Kandori. Revision games. *Econometrica*, 88(4):1599–1630, 2020b. 2, 1
- Yuichiro Kamada and Takuo Sugaya. Optimal timing of policy announcements in dynamic election campaigns. *The Quarterly Journal of Economics*, 135(3):1725–1797, 2020. 2
- Adam Kapor and Sofia Moroni. Sniping in proxy auctions with deadlines. 2016. 1
- Alexey Kushnir and Shuo Liu. On the equivalence of bayesian and dominant strategy implementation for environments with nonlinear utilities. *Economic Theory*, 67:617–644, 2019. 5
- Steven A Lippman and John J McCall. The economics of job search: A survey. *Economic inquiry*, 14(2):155–189, 1976. 1
- Qingmin Liu, Konrad Mierendorff, Xianwen Shi, and Weijie Zhong. Auctions with limited commitment. *American Economic Review*, 109(3):876–910, 2019. 21
- Benny Moldovanu and Aner Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91(3):542–558, 2001. 3.1
- John G Riley and William F Samuelson. Optimal auctions. *The American Economic Review*, 71(3):381–392, 1981. 1
- Maher Said. Sequential auctions with randomly arriving buyers. *Games and Economic Behavior*, 73(1):236–243, 2011. 19
- Robert Zeithammer. Forward-looking bidding in online auctions. *Journal of Marketing Research*, 43(3):462–476, 2006. 19
- Wen Zhao and Yu-Sheng Zheng. Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management science*, 46(3):375–388, 2000. 1

A Dictionary of Notation

The following list can be helpful when referring to the proofs.

$v_T(\omega)$ random final value

$D(\omega)$: Set of decision nodes (v, t) for path ω .

$\Pi(\omega|v, t)$: conditional probability over ω given that $(v, t) \in D(\omega)$

$\bar{e}(t^+, \omega) = \sup \{e(v', t') \mid (v', t') \in D(\omega) \text{ and } t' > t\}$ and if this set is empty it is equal to an arbitrary low number.

$\bar{e}(\omega)$: maximum equivalent value achieved on path ω .

$H(\varepsilon, v, t) = \{\omega \mid (v, t) \in D(\omega) \text{ and } \bar{e}(t^+, \omega) \leq \varepsilon\}$, paths subsequent to (v, t) where the highest ECV is less than or equal to ε . This represents the set of all ω where there is no arrival with $e(v', t') > \varepsilon$ exist for $\forall t' > t$.

$A(\varepsilon, v, t) = \{\omega \mid (v, t) \in D(\omega) \text{ and } \bar{e}(t^+, \omega) \leq \varepsilon, \exists (v', t') \in D(\omega), t' > t\}$ It contains all $\omega \in H(\varepsilon, v', t')$ with at least one arrival.

$N(v, t)$: set of paths following (v, t) with no subsequent decision nodes

$e(v, t) = E(v_T \mid \omega \in H(e(v, t), v, t))$. This is the self-generating expectation property of ECVs.

$I(u) = \{(v, t) \mid e(v, t) = u\}$ are the iso-expectation sets.

Functional equation defining $W(\varepsilon, v, t)$:

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

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

$\tilde{S}(v)$ is the correspondence defined by the set of maximizers of $U(v, a)$

$S(v, t) = \tilde{S}(e(v, t))$ is the candidate optimal strategy for the dynamic decision problem.

$\tilde{S}(\omega) = \tilde{S}(\bar{e}(\omega))$

Ψ_i : Distribution of ECVs for player i

Γ : Dynamic Bayesian game

Γ_B : Static Bayesian game associated to Γ

B Proofs

The following results are used in the proofs.

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

$$\int_{\{X \in A\}} Y dP = \int_A g(x) dP_x(x), \quad (14)$$

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

Proof. See Ash [1972] (Theorem 5.3.3. pg. 210). \square

Lemma 1. $\int_{\bar{e}(t^+, \omega) \in B} (v_T(\omega) - (\bar{e}(t^+, \omega))) d\Pi(\omega | v, t) = 0$ for any (Borel set) B .

Proof. Consider the following random variables: $\bar{e}(t^+, \omega)$ as defined above, $\tau(t, \omega)$: time at which it was reached, and $v_T(\omega)$: as defined above. Let $X(\omega) = (\tau(t, \omega), \bar{e}(t^+, \omega))$. The random variable that we consider in applying Theorem B is $Y(\omega) = v_T(\omega) - \bar{e}(t^+, \omega)$. For any measurable subset $A \subset \{t < \tau \leq T\}$ and Borel subset B of \mathbb{R} , let $P_x(A \times B) = \Pi(\tau(\omega) \in A, \bar{e}(t^+, \omega) \in B | v, t)$. For all (τ, e) ,

$$\begin{aligned} \mathbb{E}(Y | \tau(t, \omega) = \tau, \bar{e}(\tau^+, \omega) = e, (v, t)) &= \mathbb{E}(Y | \tau(t, \omega) = \tau, \bar{e}(\tau^+, \omega) = e) \\ &= \mathbb{E}(v_T(\omega) - e | (v', \tau) \in D(\omega), e(v', \tau) = e, \bar{e}(\tau^+, \omega) \leq e) \\ &= 0, \end{aligned}$$

by the definition of self-generated expectation. Substituting $x = (\tau, \varepsilon)$ and using (14),

$$\int_{\{\tau(\omega) > \tau_0, (t^+, \omega) \in B\}} (v_T(\omega) - (t^+, \omega)) dP(\omega) = \int_{\{\tau > \tau_0, \varepsilon \in B\}} E(Y | \tau, \varepsilon) dP_x(\tau, \varepsilon) = 0.$$

\square

Lemma 2. $\int_{H(\varepsilon, v, t)} (v_T(\omega) - \varepsilon) d\Pi(\omega | v, t)$ is strictly decreasing in ε and equal to zero when $e(v, t) = \varepsilon$.

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

$$\begin{aligned}
\int_{H(\varepsilon+\delta, v, t)} (v_T(\omega) - (\varepsilon + \delta)) d\Pi(\omega|v, t) &= \int_{H(\varepsilon, v, t)} (v_T(\omega) - (\varepsilon + \delta)) d\Pi(\omega|v, t) \\
&+ \int_{\varepsilon < (t^+, \omega) \leq \varepsilon + \delta} (v_T(\omega) - (\varepsilon + \delta)) d\Pi(\omega|v, t) \\
&\leq \int_{H(\varepsilon, v, t)} (v_T(\omega) - (\varepsilon + \delta)) d\Pi(\omega|v, t) \\
&+ \int_{\varepsilon < (t^+, \omega) \leq \varepsilon + \delta} (v_T(\omega) - (t^+, \omega)) d\Pi(\omega|v, t).
\end{aligned}$$

The last term is zero by Lemma 1, thus completing the proof. \square

Proof of Proposition 1 Recall $e(v, t)$ was defined as a threshold function with the property that

$$E[v_T(\omega) | (v, t) \in D(\omega) \text{ and } \bar{e}(t^+, \omega) \leq e(v, t)] = e(v, t), \quad (15)$$

where $\bar{e}(t^+, \omega) = \sup\{e(v', t') | (v', t') \in D(\omega) \text{ and } t' > t\}$ 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 (15) 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:

$$\begin{aligned}
W(\varepsilon, v, t) &= \int_t^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau'|v, t) \\
&+ \int_{N(v, t)} (v_T(\omega) - \varepsilon) d\Pi(\omega|v, t)
\end{aligned} \quad (16)$$

with terminal value $W(\varepsilon, v, T) = v - \varepsilon$, where $\Pi(\omega|v, t)$ 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 (15) 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 15.

Step 1. Necessity Take the candidate value function:

$$W(\varepsilon, v, t) = \int_{H(\varepsilon, v, t)} (v_T(\omega) - \varepsilon) d\Pi(\omega|v, t). \quad (17)$$

We show that $W(\varepsilon, v, t)$ is a solution to (16). Substituting into (16) results in

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

By Lemma 2, 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 (15) 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 $A(\varepsilon, v, t)$. Moreover, since $\int \Pi(\omega|v', t') dP(v', t'|v, t) = \Pi(\omega|v, t)$ we have

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

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 (11) 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 (16) 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 then the Bellman equation (16) 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 (16):

$$\begin{aligned} & \int_t^T \min(W(\varepsilon, v', \tau') + a, 0) dP(v', \tau'|v, t) + \int_{N(v,t)} (v_T(\omega) - \varepsilon) d\Pi(\omega|v, t) \\ & \leq \int_t^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau'|v, t) + a(1 - \Pi(N(v, t)|v, t)) + \int_{N(v,t)} (v_T(\omega) - \varepsilon) d\Pi(\omega|v, t) \\ & = W(\varepsilon, v, t) + a(1 - \Pi(N(v, t)|v, t)). \end{aligned} \tag{18}$$

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

It follows that the function $W(\varepsilon, v, t)$ is uniquely defined and from the necessity part of the proof it satisfies:

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

For $\varepsilon = e(v, t)$, it follows from the definition of $e(v, t)$ given in equation (11) that $W(e(v, t), v, t) = 0$. We now show that there is a unique function satisfying this property, by establishing that $W(\varepsilon, v, t)$ is strictly increasing in ε .

The proof is by induction, showing that the Bellman equation (16) maps weakly decreasing functions into strictly decreasing ones. So, assume that the W function on the right hand side of the Bellman equation (16) is weakly decreasing. Letting $\varepsilon' > \varepsilon$,

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

which is strictly less than $TW(\varepsilon, v, t)$ since by Assumption 3 $\Pi(N(v, t)|v, t) > 0$. The function W is continuous, so to prove that there exists an ε such that $W(\varepsilon, v, t) = 0$, it suffices to show that it will be negative for large values of ε and positive for small ones. Looking at the above Bellman equation, the first term is non-positive and the second term is strictly decreasing in ε , so it will also be arbitrarily negative for large ε . Furthermore, by Assumption 3 $\Pi(N(v, t)|v, t) > \delta$, 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.

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 $a_2 \neq a_1 \equiv S(v, t)$. We will show that this one-period deviation is not an improvement. Let $e_1 =$

$e(v, t)$ so $a_1 = \tilde{S}(e_1)$. Consider first the case where $a_2 > a_1$ and let $e_2 = \sup \left\{ v \mid \tilde{S}(v) \leq a_2 \right\}$.

Let $V(v, t, a)$ denote the expected utility of choosing a at this state and following the (candidate) optimal policy for the future. We need to prove that $V(v, t, a_1) \geq V(v, t, a_2)$. Let $\bar{s}(\omega, v, t) = \max \{S(v', t') \mid (v', t') \in D(\omega) \text{ and } (v', t') \neq (v, t)\}$. For a path ω such that $(v, t) \in D(\omega)$ this is the maximal action excluding the choice at node (v, t) and it is also the final action if it is greater than or equal to the choice at this node. Define $H(e, v, t)$ as in Section 4.2.1 and $H(e, v, t)^c$ its complement in the set of paths following (v, t) : $\{\omega \mid (v, t) \in D(\omega)\}$. To simplify notation, from now on we drop the argument (v, t) from these functions. We can decompose the histories following $a_i, i \in \{1, 2\}$ into these two sets. In the set of histories $H(e_i, v, t)$, 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 because an equivalent value higher than e_i is reached and the final corresponding action is $\bar{s}(\omega)$. It follows that

$$V(v, t, a_i) = \int_{H(e_i)} U(v(T), a_i) dP(\omega) + \int_{H(e_i)^c} U(v(T), \bar{s}(\omega)) dP(\omega),$$

for $i = \{1, 2\}$. Note that $H(e_2) = H(e_1) \cup \{\omega \mid e_1 < \bar{e}(t^+, \omega) \leq e_2\}$ so

$$\begin{aligned} V(v, t, a_2) &= \int_{H(e_1)} U(v(T), a_2) dP(\omega) + \int_{e_1 < \bar{e}(t^+, \omega) \leq e_2} U(v(T), a_2) dP(\omega) \\ &+ \int_{H(e_2)^c} U(v(T), \bar{s}(\omega)) dP(\omega). \end{aligned} \quad (19)$$

Consider a more relaxed problem where this agent is allowed to follow the original strategy, i.e., if the agent arrives at histories where $e_1 < \bar{e}(t^+, \omega) \leq e_2$, the agent is unconstrained by the preexisting choice a_2 , so its final action is $\bar{s}(\omega)$. As a result,

$$V(v, t, a_2) \leq \int_{H(e_1)} U(v(T), a_2) dP(\omega) + \int_{H(e_1)^c} U(v(T), \bar{s}(\omega)) dP(\omega),$$

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

$$\begin{aligned} V(v, t, a_2) - V(v, t, a_1) &\leq \int_{H(e_1)} [U(v(T), a_2) - U(v(T), a_1)] dP(\omega) \\ &= U(E_{H(e_1)} v(T), a_2) - U(E_{H(e_1)}, a_1) \leq 0, \end{aligned} \quad (20)$$

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 $e_2 = \inf \left\{ v \mid \tilde{S}(v) > a_2 \right\}$. For

the alternative action $a_2 < a_1$, it easily follows that

$$\begin{aligned}
V(v, t, a_2) &= \int_{H(e_2)} U(v(T), a_2) dP(\omega) + \int_{H(e_2)^c} U(v(T), \bar{s}(\omega)) dP(\omega) \\
&= \int_{H(e_1)} U(v(T), a_2) dP(\omega) + \int_{H(e_1)^c} U(v(T), \bar{s}(\omega)) dP(\omega) \\
&+ \int_{e_2 < \bar{e}(t^+, \omega) \leq e_1} (U(v(T), \bar{s}(\omega)) - U(v(T), a_2)) dP(\omega).
\end{aligned} \tag{21}$$

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

$$\begin{aligned}
V(v, t, a_1) - V(v, t, a_2) &= \int_{H(e_1)} U(v(T), a_1) dP(\omega) - \int_{H(e_1)} U(v(T), a_2) dP(\omega) \\
&- \int_{e_2 < \bar{e}(t^+, \omega) \leq e_1} (U(v(T), \bar{s}(\omega)) - U(v(T), a_2)) dP(\omega) \\
&= (U(e_1, a_1) - U(e_1, a_2)) P(H(e_1)) \\
&- \int_{e_2 < \bar{e}(t^+, \omega) \leq e_1} (U(v(T), \bar{s}(\omega)) - U(v(T), a_2)) dP(\omega) \\
&= (U(e_1, a_1) - U(e_1, a_2)) P(H(e_1)) \\
&- \int_{e_2 < \bar{e}(t^+, \omega) \leq e_1} \left(U(\bar{e}(t^+, \omega), \tilde{S}(\bar{e}(t^+, \omega))) - U(\bar{e}(t^+, \omega), a_2) \right) dP(\omega),
\end{aligned} \tag{22}$$

where, as before, the second equality follows from the linearity of U in v and $e_1 = e(v, t) = E_{H(e_1)}(v(T))$, and the third equality follows from the linearity of U and application of Lemma 1.

Consider the last integral. The supermodularity of the U function implies that $\tilde{S}(v)$ must be an increasing function, so for ω such that $e_2 < \bar{e}(t^+, \omega) \leq e_1$ it follows that

$$a_1 \geq \tilde{S}(\bar{e}(t^+, \omega)) = s(\omega) \geq \tilde{S}(e_2) \geq a_2.$$

It also follows from supermodularity and $e_1 \geq \bar{e}(t^+, \omega)$ that

$$U(\bar{e}(t^+, \omega), \tilde{S}(\bar{e}(t^+, \omega))) - U(\bar{e}(t^+, \omega), a_2) \leq U(e_1, \tilde{S}(\bar{e}(t^+, \omega))) - U(e_1, a_2). \tag{23}$$

Finally, since $a_1 = \tilde{S}(e_1)$ it follows that

$$U(e_1, \tilde{S}(\bar{e}(t^+, \omega))) - U(e_1, a_2) \leq U(e_1, a_1) - U(e_1, a_2). \tag{24}$$

Combining (22), (23) and (24), it follows that

$$\begin{aligned} V(v, t, a_1) - V(t, v, a_2) &\geq (U(e_1, a_1) - U(e_1, a_2)) P(H(e_1)) \\ &\quad - \int_{e_2 < \bar{e}(t^+, \omega) \leq e_1} U(e_1, a_1) - U(e_1, a_2) dP(\omega), \end{aligned} \quad (25)$$

and since the set of paths $\{e_2 \leq \bar{e}(t^+, \omega) \leq e_1\}$ is a subset of $H(a_1)$, then $V(v, t, a_1) - V(t, v, a_2) \geq 0$, so the proof is complete.

Proof of Proposition 2 The value of the repurchase option at price ε in state (v, t) satisfies the following Bellman equation:

$$V(\varepsilon, v, t) = \int_t^T \max(V(\varepsilon, v', \tau'), v' - \varepsilon) dP(v', \tau' | v, t) \quad (26)$$

Guess that the solution to this functional equation is

$$V(\varepsilon, v, t) = v - (W(\varepsilon, v, t) + \varepsilon).$$

Using equation (12)

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

This proves the claim. Using the definition of ECV, It follows that the value of the repurchase option at price $\varepsilon = e(v, t)$, $V(e(v, t), t) = v - e(v, t)$. Moreover, the repurchase option at any future trading window t' is exercised if and only if

$$\begin{aligned} 0 &\leq v' - e(v, t) - V(e(v, t), v' - \varepsilon) \\ &= v' - e(v, t) - (v' - W(e(v, t), v', t') + e(v, t)) \\ &= W(e(v, t), v', t') \end{aligned}$$

or, equivalently, if and only if $e(v', t') \geq e(v, t)$. This is precisely the threshold used to define the

ECV $e(v, t)$ in our definition.

Proof of Theorem 2 Let $U_i(v_{iT}, a_{iT}) = \mathbb{E}_{a_{-iT}} u(v_{iT}, a_{iT}, a_{-iT} | S_{-i})$, that is, the expected final payoff given v_{iT}, a_{iT} after integrating out the strategies of the other players. Assumption 4 implies that U_i is linear in v and supermodular. This payoff function and the stochastic process P_i for values and decision times define a dynamic decision problem that satisfies the assumptions of Theorem 1. Since \tilde{S}_i is a best response for agent i in the Bayesian game, it follows that almost surely for \tilde{v} in the support of Ψ_i

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

for all $a \in A_i$. So, $\tilde{S}_i(\tilde{v})$ is an optimal solution for any ECV \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 .

Proof of Proposition 3 We prove that the property $W(\varepsilon, v, t) \leq E(v_T | v, t) - \varepsilon$ is preserved under the Bellman equation. Suppose that $W(\varepsilon, v', \tau') \leq E(v(T) | v', \tau') - \varepsilon$. Then

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

where the last equality follows from the law of iterated expectation. It immediately follows from $W(e(v, t), v, t) = 0$ that $e(v, t) \leq \mathbb{E}(v_T | v, t)$. By Assumption 5, it follows that $W(\varepsilon, v, t)$ is strictly increasing in v . Together with Assumption 2, it implies that the first inequality is strict. In particular, for $\varepsilon = e(v, t)$ it follows that the first inequality above is strict, so $\mathbb{E}(v_T | v, t) > e(v, t)$.

Proof of Proposition 4 Using Lemma 1 by setting $B = \{\omega | \bar{e}(t^+, \omega) \geq e(v, t)\}$, we have the following:

$$\int_{\bar{e}(t^+, \omega) \geq e(v, t)} (v_T(\omega) - \bar{e}(t^+, \omega)) d\Pi(\omega|v, t) = 0$$

and by the definition of self-generating expectation

$$\int_{\bar{e}(t^+, \omega) < e(v, t)} (v_T(\omega) - e(v, t)) d\Pi(\omega|v, t) = 0$$

Adding the two implies that

$$\begin{aligned} v &= \int v_T(\omega) d\Pi(\omega|v, t) \\ &= e(v, t) \Pi\{\omega | \bar{e}(t^+, \omega) < e(v, t)\} + \int_{\bar{e}(t^+, \omega) \geq e(v, t)} \bar{e}(t^+, \omega) d\Pi(\omega|v, t) \\ &= \mathbb{E}[\max\{\bar{e}(t^+, \omega), e(v, t)\} | \omega] \end{aligned}$$

Proof of Proposition 5 We first show the claim for the first condition. Consider the Bellman equation:

$$W(\varepsilon, v, t) = \int_t^T \left[\int \min(W(\varepsilon, v', \tau'), 0) dP(v'|v) \right] dF(\tau'|t) \quad (27)$$

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

$$= \int_t^T \left[\int \min(W(\varepsilon, v', \tau'), 0) dP(v'|v) \right] dF(\tau'|t) \quad (28)$$

$$+ (1 - F(T|t)) [\mathbb{E}(v_T|v) - \varepsilon]. \quad (29)$$

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, $F(\tau'|t)$ 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 3, $W(\varepsilon, v, t) \leq E(v_T|v, t) - \varepsilon$, implying that the second term on the right hand side of (27) is greater in expectation than the first term. So, the shift in mass also contributes to increasing the overall expectation. This proves the case for the first condition.

To show the claim for the second condition, we show that monotonicity is preserved under the Bellman equation. Assume by way of induction that $W(\varepsilon, v', \tau')$ is increasing in its last

argument. Using the first assumption in the proposition, equation (12) can be rewritten as

$$\begin{aligned} W(\varepsilon, v, t) &= \int_0^{T-t} \int \min(W(\varepsilon, v', t+x), 0) dP_v(v'|v, x) dF(x) \\ &\quad + \int_{x>T-t} \int (v' - \varepsilon) dP_v(v'|v, x) dF(x). \end{aligned}$$

Take $\tau_2 > \tau_1$.

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

where the first inequality follows from the induction hypothesis. As in the proof of Proposition 3, $W(\varepsilon, v', \tau_1+x) \leq \mathbb{E}(v_T|v', \tau_1+x) - \varepsilon = v' - \varepsilon$ by the martingale assumption. As a consequence, the term subtracted in the last line above is negative. It follows that $W(\varepsilon, v, \tau_2) \geq W(\varepsilon, v, \tau_1)$.

Proof of Proposition 6 As a first step, it should be clear that starting at any node (v, t) the conditional distribution of final values is unbiased, so its expectation equals $E(v_T|v, t)$. This follows from iterated expectation: For $\tau \geq t$ let $\bar{v}(\tau|v, t)$ denote the expected final value conditional on (v, t) and τ being the last arrival time after this node. Let $P(\tau|v, t)$ be the corresponding conditional distribution for this event. By the law of iterated expectation, it follows that $v = \int \bar{v}(\tau, v, t) dP(\tau|v, t)$. By Proposition 4, the distribution of expected final values starting from (v, t) is also unbiased. We now show that the distribution of $\bar{v}(\omega)$ conditional on (v, t) is more dispersed than the corresponding conditional distribution for

$\bar{e}(\omega)$. Let $A = \{\omega \succeq (v, t) \mid \bar{e}(\omega) = e(v, t)\}$ and $A^c = \{\omega \succeq (v, t) \mid \bar{e}(\omega) > e(v, t)\}$. This is a partition of the paths following (v, t) where A corresponds to all paths such that the ECV $e(v, t)$ remains the maximum, while the set A^c is the complement. It follows that

$$\begin{aligned} v &= \Pi(A) \mathbb{E}(v_T|A) + \Pi(A^c) \mathbb{E}(v_T|A^c). \\ &= \Pi(E) e(v, t) + \Pi(A^c) \mathbb{E}(v_T|A^c) \end{aligned} \quad (30)$$

where the equality follows from the definition of ECVs. Starting from any $(v', t') \in A^c$, Proposition 4 implies that the conditional distribution of terminal ECVs $\bar{e}(\omega \mid (v', t'))$ is unbiased, integrating over all such decision nodes, it follows that $E(\bar{e}(\omega) \mid A^c) = E(v_T \mid A^c)$ and by a similar argument, the same holds for $E(\bar{v}(\omega) \mid A^c)$. It then follows from equation (30) that $e(v, t) = E(\bar{v}(\omega) \mid A)$ so the conditional distribution for $\bar{v}(\omega)$ in A is a mean preserving spread of the corresponding distribution of $\bar{e}(\omega)$ which is a point mass on $e(v, t)$. Repeating the same argument for all $(v', t') \in A^c$, it follows that the conditional distribution of $\bar{v}(\omega)$ in A^c is also a mean preserving of the corresponding distribution of $\bar{e}(\omega)$, which completes the proof.

Proof of Proposition 7 Consider the recursive formulation

$$\begin{aligned} W(\varepsilon, v, t) &= \int_t^T \int \min(W(\varepsilon, v', \tau'), 0) dP(v' \mid v, \tau') dF(\tau' \mid t) \\ &\quad + (1 - F(T \mid t))(v - \varepsilon) \end{aligned} \quad (31)$$

First note that this functional equation is monotone: increasing pointwise the integrated function W increases the integral. Also note that from the proof of Proposition 5, the second assumption implies that $W(\varepsilon, v, t)$ is nondecreasing in t . Let $\tilde{W}(\varepsilon, v, t)$ denote the value function under distribution $\tilde{F}(t' \mid t)$. We will prove, recursively, that $\tilde{W}(\varepsilon, v, t) \leq W(\varepsilon, v, t)$. So assume this is true for $\tau' > t$. Then

$$\begin{aligned} \tilde{W}(\varepsilon, v, t) &= \int_t^T \int \min(\tilde{W}(\varepsilon, v', \tau'), 0) dP(v' \mid v, \tau') d\tilde{F}(\tau' \mid t) + (1 - \tilde{F}(T \mid t))(v_T - \varepsilon) \\ &\leq \int_t^T \int \min(W(\varepsilon, v', \tau'), 0) dP(v' \mid v, \tau') d\tilde{F}(\tau' \mid t) + (1 - \tilde{F}(T \mid t))(v_T - \varepsilon) \\ &\leq \int_t^T \int \min(W(\varepsilon, v', \tau'), 0) dP(v' \mid v, \tau') dF(\tau' \mid t) + (1 - F(T \mid t))(v_T - \varepsilon) \end{aligned}$$

The first inequality follows from the inductive assumption, while the second inequality follows from the fact that $W(\varepsilon, v', t')$ is nondecreasing in t and F stochastically dominates \tilde{F} and that $W(\varepsilon, v, \tau) \leq v - \varepsilon$.

C Extension and Special Cases (For online publication)

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

C.1 Random Termination

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

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

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

It is useful to examine the conditions of Proposition 5 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 5), the result follows. This is now an assumption regarding the difference between the two random variables, τ' and T . The following corollary gives sufficient conditions for this assumption to hold.

Corollary 2. *Let $H(y|t) = P(T - t \leq y|t)$ denote the cdf for the remaining time of the decision problem conditional on $T \geq t$. Let $G(x|t) = P(\tau' - t \leq x|t)$ denote the conditional cdf of the time to the next arrival. Assume that τ' and T are conditionally independent given t and*

1. $H(y, t)$ is weakly increasing in t and

2. $G(x|t)$ is weakly decreasing in t .

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

Proof. Note that $\tau' - T = \tau' - t - (T - t)$. Let $(T - t) = x$ so $P(\tau' - T \leq z|t, x) = P(\tau' - t \leq z + x) = G(z + x|t)$. Integrating over x results in $F(z|t) = \int G(z + x|t) dH(x|t)$. By the second assumption, the integrand is point-wise decreasing in t . By the first assumption, the distribution H is stochastically decreasing in t and, since G is an increasing function, it also implies that the integral is decreasing in t . This proves that $F(z|t)$ is decreasing in t . The second conclusion follows directly from Proposition 5. \square

The assumptions of this corollary have an intuitive interpretation. The second one is the analogue of the assumption made in Proposition 5. 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 $e(v, t)$ will also be independent of t , so the level curves will be flat. Shading will still occur but will not change over time.

C.2 Independent Increments

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

Proposition 8. *Assume $P(v' = v + \delta, t'|v, t)$ is independent of v for all δ and all t , and decision times are independent of v . Then*

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

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

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

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

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

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

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

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

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

Proposition 9. *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

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

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

Shading over Time with Independent Increments While the propositions derived in Section C.2 apply to this special case, an additional intuitive and useful result can be proved. A natural question is how the variance of new values affects shading, as it affects the option value of future actions. In the extreme, if variance were zero so $v(T) = v$ with probability one, there would be no shading. We prove a monotonicity result for the case of independent increments considered in Proposition 8.

Proposition 10. *Under the assumptions of Proposition 8, W is concave in v . A mean preserving increase in spread of the distribution of increments decreases $e(v, t)$.*

Proof. Consider the dynamic programming equation for W ,

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

Given the assumption of scale invariance and that arrival times and values are independent, it follows that

$$\begin{aligned} \int_{N(v, t)} (v_T(\omega) - \varepsilon) d\Pi(\omega | v, t) &= \mathbb{E}[v_T | v, t, \tau' > T] \Pi(N(v, t) | v, t), \\ &= (v + \mathbb{E}[v_T | 0, t, \tau' > T]) \Pi(N(v, t) | v, t) \end{aligned}$$

where $\Pi(N(v, t) | v, t)$ is the probability of no arrivals, which is independent of v given the assumptions. 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 \square

$$\begin{aligned}
& \int_t^T \min(W(\varepsilon, v', \tau'), 0) dP(v', \tau' | (\alpha v_2 + (1 - \alpha)v_1), t) \\
= & \int_t^T \min(W(\varepsilon, \alpha(v' + v_2) + (1 - \alpha)(v' + v_1), \tau'), 0) dP(v', \tau' | 0, t) \\
\geq & \int_t^T (\alpha \min\{W(\varepsilon, (v' + v_2), \tau'), 0\} + (1 - \alpha)W(\varepsilon, (v' + v_1), \tau')) dP(v', \tau' | 0, t) \\
= & \alpha \int_t^T \min\{W(\varepsilon, v', \tau'), 0\} dP(v', \tau' | v_2, t) + (1 - \alpha) \int_t^T \min\{W(\varepsilon, v', \tau'), 0\} dP(v', \tau' | v_1, t).
\end{aligned}$$

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

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

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

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

and the variance is $1/(1/\sigma_0 + t/\sigma)$. Together with an independent arrival process for decision nodes, this formula can be used recursively to define the Markov process $P(v', t' | v, t)$ that satisfies the assumption of independent increments in Proposition 8. As the variance of the increments decreases over time, Proposition 10 implies that shading decreases over time.

C.3 Anonymous Sequential Game

Consider a stationary population of agents. Time is discrete. Each period an agent continues in the game with probability δ and a value v that follows a Markov process with conditional distribution $F(v' | v)$. Exiting agents are replaced by new ones with values v drawn from

some initial distribution $G(v)$. At each of these decision nodes, the agent chooses whether to increase its capital k by Δk at unit cost c and is then faced with a random match to a subset of other players in the population. Profits in the period (gross of investment costs) are given by $\pi(v, k, \tilde{k})$, where \tilde{k} is the vector of capital of other competing agents. Assume this is linear in v and supermodular in v, k . For simplicity, suppose the Markov process has independent increments so that $v' = v + z$, where z has cdf $\Phi(z)$. Define the payoff function as follows: $u(v, k, \tilde{k}) = \frac{1}{1-\delta} \left(\pi(v, k, \tilde{k}) - ck \right)$.

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

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

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

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

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

1. Derive a stationary distribution of values $F(v)$.
2. Outer fixed point:
 - (a) Choose the estimated parameters θ of the payoff function.
 - (b) Inner fixed point:
 - i. Guess a strategy $k(v)$ for all players.
 - ii. Solve the dynamic problem for an agent to get the best response strategy. Define this as a new guess for strategy.

- iii. Get new strategies for all players.
- (c) Adjust parameters θ until a good match to the data is obtained.

Using our method, the steps would be:

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

$$k(v) = \arg \max_k \int u(k, \tilde{k}(\tilde{v}), v - s) dF(\tilde{v}).$$

- iii. Get new strategies for all players.
- (c) Adjust parameters θ until a good match to the data is obtained.

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

A simple solution to our first step can be found for the following stochastic process. Assume that with probability $(1 - p)$ the value continues to be the same, while with probability p it is drawn again from distribution $F(v)$.²² It easily follows that

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

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

²²Note that this stochastic process does not have independent increments.