Equivalent Certain Values and Dynamic Irreversibility*

Hugo Hopenhayn[†] UCLA

Maryam Saeedi[‡] Tepper, CMU

November 30, 2025

Abstract

We introduce a tractable methodology for analyzing dynamic decision problems and games with irreversible decisions under uncertainty. Leveraging regularity properties, we propose an intuitive method for solving these problems using equivalent certain values, a concept that extends the idea of certainty equivalence to dynamic environments with irreversibility and delivers properties of optimal strategies and comparative statics. We show that irreversibility is analogous to information loss in the Blackwell order sense, making agents behave as if they had worse information than under reversible actions. As an application, we use our methodology to analyze design features of a class of dynamic auctions.

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, Alessandro Lizzeri, Moritz Meyer-ter-Vehn, Paul Milgrom, Aniko Oery, Ariel Pakes, Aghathe Pernoud, Michael Peters, Robert Porter, Andrea Prat, Debraj Ray, Ali Shourideh, Ennio Stacchetti, Andre Sztutman, Steve Tadelis, Robert Town, Anton Tsoy, Leeat Yariv as well as seminar and conference participants at Northwestern Economics Department Theory seminar, EIEF, Midwest Economic Theory Meetings, IIOC, 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 of research in economics. This class of problems is often difficult to analyze and solve either analytically or numerically. This paper develops a tractable and intuitive methodology for solving a broad class of dynamic problems and games with irreversibilities, exploiting regularity properties that arise naturally in many economic settings to provide valuable insights. We study decision problems with three key features: First, agents make irreversible choices over time—whether sunk R&D investments, capacity allocations, or bids in extended auctions. Second, information arrives throughout the decision period as innovators learn about patent values, firms discover outside opportunities, or bidders observe competing auctions. Third, decision opportunities arise stochastically, capturing frictions in when investments can be made or trades executed.

Our paper makes two main contributions. First, we introduce a novel approach to solving and analyzing this class of problems that is both intuitive and tractable. Second, our methodology characterizes several key properties and comparative statics for both dynamic decision problems and games. In particular, we show that irreversibility and stochastic action timing are analogous to information loss in the Blackwell sense, Blackwell [1951, 1953]. An agent faced with irreversible decisions or random decision opportunities acts as if they had coarser information. Both frictions—irreversibility and stochastic timing—generate mean-preserving contractions in the distribution of decision-relevant states. The implications for expected payoffs thus mirror those of information degradation. In dynamic second-price auctions, for instance, we show that these frictions produce a distribution of final bids that is a mean-preserving contraction relative to the frictionless case.

Our methodology decomposes dynamic decision problems into two components: a dynamic option valuation problem and a static decision problem that maps option values to actions. To illustrate, 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 this value exceeds the entry cost c_e. This value is generally lower than the expected value $v(T)$ at this node, since making the entry decision means giving up the option of waiting. We call this option value the equivalent certain value (ECV). It corresponds to the certainty value at which the decision-maker would make the

same decision absent future opportunities.[1]

The surprising insight is that ECVs depend only on the stochastic structure of information and decision opportunities, and is independent of the specific payoff function in our class.[2] The specific choice problem matters only for the second step, the static decision problem of assigning actions to each ECV. This separation result is particularly useful when considering game design, especially when the design involves repeatedly evaluating changes in the payoff functions.

ECVs are constructed through a self-generating expectation method that depends only on the stochastic processes for information and decision timing, not on payoff functions. To explain this method, we need to explain our modeling framework in more detail. The agent faces a random sequence of decision times which follow a joint Markov process.[3] At each decision time, the agent can adjust their irreversible action only upward (e.g., increasing a bid). Final payoffs depend on the terminal action and realized value. The key insight is that the ECV at any decision node (v_0, t_0) equals the terminal value that would induce the same optimal action in a static problem. Figure 1 illustrates this: the ECV is the terminal value where the agent chooses the same action a_0 as at (v_0, t_0). The indifference curve connecting these points contains all state-time pairs yielding action a_0.

These indifference curves exhibit a self-generating property. From any point (v_0, t_0), consider all future paths where action a_0 remains *final*. This occurs in two cases: (1) *no further* decision opportunities arise, or (2) opportunities arise but *fall below* this indifference curve, where irreversibility prevents desired downward adjustments. The ECV equals the expected terminal value conditional on these paths—hence "self-generating," as the curve defines its own continuation region. This construction depends only on the stochastic structure, not payoffs, explaining why ECVs are the same within our problem class.[4] Additionally, note that the second subset includes *adverse* future paths; therefore, ECVs tend to be lower than the unconditional expected value of agents, hence we get value shading.

These indifference curves are typically downward sloping. Intuitively, the adverse realizations that keep a given action in place are more likely when the agent is far from the deadline

[1] This correspondence satisfies a self-generating property, which is explained below. It has a parallel in the recursive methods in Abreu et al. [1990].

[2] We assume that the payoff function is affine in value.

[3] This 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 in models of optimal election campaigns as in Kamada and Sugaya [2020], among many others.

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

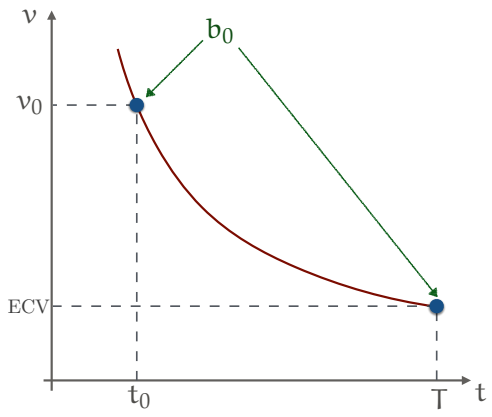

Figure 1: Indifference curves for actions

Note: The x-axis represents the time horizon, and the y-axis represents agent's values. All points on this curve result in the same optimal action value, a_0.

than when they are close to it. This generates stronger shading far from the deadline, weaker shading near it, and thus a gradual increase in actions over time.[5]

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 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 (i.e., its ECV) exceeds $p(t)$.

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 equilibrium strategies can be found by solving a Bayesian Nash equilibrium of an associated static Bayesian game.[6] The equilibrium strategies in the dynamic game are easily derived from

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

[6] Even 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, they can focus solely on dominant strategies, utilizing

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.[7]

Two crucial factors determine the degree of shading: 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 this shading becomes weaker when actions are only partially irreversible—when agents occasionally get opportunities to undo or revise past choices. In that case, outcomes lie between the fully irreversible and fully flexible benchmarks. While in our baseline model the arrival of opportunities is exogenous, we extend the model by considering the case where agents can choose initially how much attention to devote to monitoring and adjusting their actions. Higher arrival rates make behavior more responsive to information, but also increase attention costs, and in strategic environments they shape how aggressively others respond. Taken together, these results show that both irreversibility and infrequent decision opportunities operate much like information frictions in the sense of Blackwell. Tighter irreversibility and rarer opportunities effectively garble the information on which actions are based, while greater flexibility and more frequent opportunities correspond to more informative environments in which behavior better reflects underlying fundamentals.

We apply the method to dynamic auctions in which bidders learn about their private values over time. Long-horizon formats make evolving information first-order, so design must account for both learning and irreversibility. Our setting nests and extends 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 generalized revenue-equivalence result holds in our model even with recurring bid opportunities. The argument relies on reducing the problem to a static auction using the distribution of terminal ECVs.[8] The static reduction also provides insight into optimal reserve prices, in a way analogous to Riley and Samuelson [1981] for static auctions. Their analysis suggests that when values are very diffuse the seller may benefit from higher reserves and our results yield the same comparative statics when bid retraction becomes more likely or bidding opportunities become more frequent.

Our design analysis yields surprising results. Allowing bid retraction can harm bidders while

our method to address the dynamic aspects of these strategies.

[7]Without 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.

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

benefiting auctioneers when participation is high. Without irreversibility, the elimination of shading creates a mean-preserving spread in the bid distribution, raising expected prices with many bidders. Similarly, increasing bidding frequency or adding periodic adjustment opportunities generates comparable spreads with analogous price effects.

Beyond its implications for auctions, our methodology offers insights for broader market design challenges involving irreversible decisions, from procurement contracts to resource allocation systems. By showing that dynamic frictions mirror static information loss, we provide a unified framework for diverse design questions. Our ECV decomposition approach makes these insights implementable, as market designers can apply our insights without repeatedly solving complex dynamic equilibria.

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, as shown in 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.

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 monopolists 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].

A related strand of work studies rational inattention, in which agents optimally choose when

and how precisely to process information subject to attention costs (Sims [2003], Steiner et al. [2017], Zhong [2022], Georgiadis-Harris [2024]). We model attention through endogenous arrival rates showing that higher arrival frequencies generate mean-preserving spreads of terminal ECVs—equivalent to more precise information in the Blackwell sense.

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 4.3 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 times and give properties for the case where values are independent of Poisson arrival opportunities, as well as discussing an application to the identification of anonymous sequential games. All proofs are deferred to the appendix unless specified.

2 A Simple Example

We begin by analyzing a two-period contest that illustrates the main features of our methodology. There are N players and two periods, $t \in \{0, 1\}$. In period 0, each player i observes a private signal v_{i0} drawn from a distribution $H_i(v_{i0})$ and chooses a nonnegative action $a_{i0} \geq 0$ (for example, study effort for an exam or resources allocated to a project). In period 1, the player privately observes a value of winning v_{i1}, drawn from a conditional distribution $F_i(\cdot \mid v_{i0})$. With probability p_i, the player has an opportunity to increase their action to any $a_{i1} \geq a_{i0}$ (e.g., study more or commit additional resources); with probability $1 - p_i$, no revision is possible and a_{i0} remains the final action. The player with the highest final action a_i wins the contest and receives payoff $v_{i1} - a_i$; each other player j receives $-a_j$. Signals and final values are independent across players. For notational simplicity, we suppress the index i when no confusion arises.

A player's strategy specifies choices $a_0(v_0)$ for period 0 and $a_1(v_0, v_1)$ for period 1, subject to the irreversibility constraint $a_1 \geq a_0$. The period-1 choice matters only if the agent receives an opportunity to revise. Given the distribution G of opponents' highest final actions, an agent's expected utility with final value v and action a is:

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

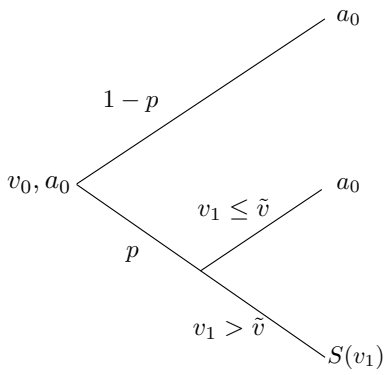

Figure 2: Decision tree

Note: The value of the agent in period 0 is v_0, and their action is a_0. There is a probability $(1-p)$ that the agent does not get a revision opportunity in the next period.

We assume Equation 1 has a unique maximizer that is strictly increasing in v, denoted $S(v)$—the optimal static action.[9] For any first-period action a_0, there exists a unique threshold $\tilde{v}$ satisfying $S(\tilde{v}) = a_0$. This threshold structure generates the decision tree shown in Figure 2.

Figure 2 depicts the agent's decision tree in the two-period game. The top branch represents the case where no revision opportunity arises, making a_0 the final action. The middle branch captures situations where the agent receives a revision opportunity but $v_1 \leq \tilde{v}$. Here, the agent would prefer to lower their action but cannot due to irreversibility, so a_0 remains binding. By construction, $\tilde{v}$ is the threshold where $S(\tilde{v}) = a_0$. The bottom branch represents realizations where $v_1 > \tilde{v}$, prompting the agent to increase their action to $S(v_1)$ in period 1.

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

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

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

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

[9] We make these assumptions and others 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.

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

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

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

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

We call this value $\tilde{v}$ the equivalent certain value (ECV) of v_0. The agent makes the same choice in period 0 when the signal is v_0 as if the final value v_1 were equal to $\tilde{v}$ for sure.

To interpret this, note that the threshold $\tilde{v}$ defines a lottery over final values v_1 under which a_0 remains the agent's final action. This lottery consists of:

1. The agent has no opportunity to revise their first-period choice. This occurs with probability $(1-p)$ and has expected value $\mathbb{E}(v_1|v_0)$.

2. The agent can revise, but the final value is below the threshold $\tilde{v}$, so they keep their initial action. This occurs with probability $pF(\tilde{v})$ and has 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 given by Equation 4, which equals $\tilde{v}$. It also solves a fixed-point problem and can be interpreted as the expected value of v_1 conditional on the future realizations under which if arrival occurs then $v_1 \leq \tilde{v}$. These are the key features that uniquely characterize ECVs in our general model.

As suggested by the example, the partition of the value space into pairs of expected values and their corresponding ECVs, $(v_0, \tilde{v})$, 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 the ECV 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 this ECV. In the complement (i.e., histories where the agent can review their 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 revision probability p defines a distribution for final values $\tilde{F}$ for each player and thus a static Bayesian game. Letting $\tilde{S}_i$ denote an equilibrium strategy for player i in this

static game, we can obtain $a_{i0}(v_0) = \tilde{S}_i(\tilde{v}(v_0))$ and $a_i(v_0, v_1) = \max\left\{\tilde{S}_i(\tilde{v}(v_0)), \tilde{S}_i(v_1)\right\}$ as equilibrium strategies in the dynamic game.

We now highlight a key property of equilibria in this class of games. The possibility of future revision embeds an option in period 0 problem. Equation 4 shows that $\tilde{v} < \mathbb{E}(v_1 \mid v_0)$, so in period 0, the agent acts as if the final value were lower than their conditional expectation; we refer to this as *value shading*. Since actions are monotone in values, this translates into shading of actions as well. Shading becomes more pronounced as the revision probability p increases. In the limit as $p \to 1$, we have $\tilde{v} \to 0$, so the agent acts in period 0 as if the value were arbitrarily close to the lower bound of the support. At the opposite extreme, when $p \to 0$, we obtain $\tilde{v} = \mathbb{E}(v_1 \mid v_0)$ and shading disappears.

The intuition again comes from the two sets of events in which the period 0 action is final. In the first event, where the agent never has a chance to revise, the relevant expected value is $\mathbb{E}(v_1 \mid v_0)$. In the second, the agent could revise but chooses not to do so, which necessarily corresponds to low realizations of v_1 and pulls the effective value below $\mathbb{E}(v_1 \mid v_0)$ which results in shading. Irreversibility and the option to delay thus generate a negative option value in period 0. Equivalently, this can be interpreted as adverse selection against the agent's future self, which leads to value shading.

3 The Decision Problem

We first consider the general structure of a dynamic decision problem. Then we show that our result can be embedded in a wide class of dynamic games as well.

First, we provide an informal description of the model and then provide the formal details. Time is continuous in the interval $[0, T]$. Decision times $0 \leq \tau_0 < \tau_1 < ... < \tau_m \leq T$ are random and exogenous to the agent's choices. At each decision time τ_n, an agent chooses an action a_n from a set $A \subset \mathcal{R}^l$ with the restriction that $a_{n+1} \geq a_n$, which is the source of irreversibility in the model. Payoffs $U(v_T, a_T)$ are received at time T, as a function of the realization of a random variable v_T, where a_T corresponds to the action taken in the last decision time. At each decision time τ_n, the agent observes a signal v_n with the property that $E_{\tau_n}(v_T) = v_n$. A strategy specifies a feasible action a_n at each decision time to maximize the expected final payoff.

We now provide a formal description. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space, endowed with a filtration $\{\mathcal{F}_t\}_{t\in[0,T]}$, the σ-algebra of events observable up to time t, $\mathcal{F}_t \subset \mathcal{F}_{t+s} \subset \mathcal{F}$

for all $t, s \geq 0$ with $t + s \leq T$. We assume that the usual hypothesis holds: $\{\mathcal{F}_t\}$ is right-continuous, and $\mathcal{F}_0$ contains all $\mathbb{P}$-null sets.[10]

The agent observes a continuous-time information process $\{\tilde{v}(t)\}_{t\in[0,T]}$, where $\tilde{v} : [0,T] \times \Omega \to \mathbb{R}$ is an $\{\mathcal{F}_t\}$-adapted process satisfying the martingale property: $\mathbb{E}[\tilde{v}(T) \mid \mathcal{F}_t] = \tilde{v}(t), \quad \forall t \in [0,T], \quad \mathbb{P}$-a.s. The decision times $\{\tau_n\}_{n\in\mathbb{N}_0}$ are defined as the jump times of a right-continuous, non-decreasing counting process $\eta : [0,T] \times \Omega \to \mathbb{N}_0$, as is the case for a Poisson process. Specifically, we can express $\tau_n(\omega) = \inf\{t \in [0,T] : \eta(t,\omega) \geq n\}$. Without loss of generality, we assume that $\{\mathcal{F}_t\}_{t\in[0,T]}$ is the filtration generated by the pair of stochastic processes $\{\eta(t,\omega), \tilde{v}(t,\omega)\}$, ensuring that $\mathcal{F}_t = \sigma(\{\eta(s,\cdot), \tilde{v}(s,\cdot) : 0 \leq s \leq t\})$. This formulation guarantees that the realization of these processes constitutes the complete information set available at time t and the martingale property holds $\mathbb{P}$-almost surely.

Since information arrivals are relevant only at decision nodes, we restrict attention to the joint process $\{(v_n, \tau_n)\}_{n\in\mathbb{N}_0}$, where $v_n(\omega) = \tilde{v}(\tau_n(\omega), \omega)$, which represents the observed signal value at each decision time. Formally, this is the process $\tilde{v}(t,\omega)$ subordinated to the arrival process $\eta(t,\omega)$. We denote the law of these combined processes by a probability measure P and impose the following assumptions.

Assumption 1. *The transition process for values and decision times satisfies the Markov property: for all $n \in \mathbb{N}_0$, $v' \in \mathbb{R}$, and $\tau' \in [0,T]$, where P is the probability measure,*

$$P(v_{n+1} = v', \tau_{n+1} = \tau' \mid \mathcal{F}_{\tau_n}) = P(v_{n+1} = v', \tau_{n+1} = \tau' \mid v_n, \tau_n) \quad \mathbb{P}\text{-almost surely.}$$

Assumption 1 will be maintained for the rest of the paper.

While decision times are exogenously determined, our framework accommodates correlation between decision times and information arrival. This flexibility allows us to model scenarios where agents exhibit state-dependent attention, such as increased responsiveness following significant information updates or heightened vigilance as the terminal time T approaches.

By the Markov property established in Assumption 1, we can identify the state of the decision problem at any point by the pair (v_n, τ_n), representing the most recently observed signal value and its corresponding observation time. This state-space representation allows us to define a decision strategy $s : \mathbb{R} \times [0,T] \to A$ as a mapping that specifies, for each possible state (v_n, τ_n), a desired action $s(v_n, \tau_n) \in A$ that the agent would select if unconstrained by previous choices. Due to the irreversibility constraint, the prevailing action at time $t \in [0,T]$ is given by $a(s,t) = \sup\{s(v_n, \tau_n) : \tau_n \leq t, n \in \mathbb{N}_0\}$, which represents the maximum of all

[10] Protter [2004] discusses this hypothesis in more detail.

actions chosen at decision times prior to or at t. Therefore, the terminal action $a(s,T)$ is the final implemented choice that determines the agent's payoff. We denote by $\mathcal{S}$ the functional space of all admissible strategies, formally defined as follows: $\mathcal{S} = \{s : \mathbb{R} \times [0,T] \to A \mid s \text{ is } \mathcal{F}_t\text{-measurable for all } t \in [0,T]\}$.

For any strategy $s \in \mathcal{S}$ and sample path $\omega \in \Omega$, the realized terminal payoff is given by $U(v(T,\omega), a(s,T,\omega))$, where $a(s,T,\omega) = \sup\{s(v_n(\omega), \tau_n(\omega)) : n \in \mathbb{N}_0, \tau_n(\omega) \leq T\}$. The agent's objective is to identify an optimal strategy $s^* \in \mathcal{S}$ that solves the following maximization problem:

$$\sup_{s \in \mathcal{S}} \mathbb{E}_0[U(v(T), a(s,T))]. \tag{5}$$

Throughout the paper, we make the following assumption on the payoff function.

Assumption 2. *The payoff function $U : \mathbb{R} \times A \to \mathbb{R}$ satisfies the following properties: $U(v,a)$ is affine in v, supermodular in v and a, and admits a maximum with respect to a for all v.*

Equivalently, the payoff function has the representation $U(v,a) = vf(a) + g(a)$ for $f, g : A \to \mathbb{R}$. Supermodularity implies that the maximizer $a(v)$ of this function is increasing in v; it further specializes to the condition $f'(.) \geq 0$ when $A \subseteq \mathbb{R}$ and the function f is differentiable. Linearity in v can be relaxed, as explained in Section 4.1.

In Section 4, we establish the existence conditions for an optimal solution to (5) and develop a methodological approach that transforms this dynamic optimization problem into an equivalent static formulation. This transformation substantially enhances analytical tractability, facilitating both estimation procedures and enable us to extend our framework to dynamic games. Before proceeding to the formal analysis, we illustrate the versatility of our framework through several applications.

3.1 Examples

This general setting embeds various notable applications of dynamic decision problems. As explained in Section 4.2, under certain conditions, this general setting can also be extended to a class of dynamic games. In this section, we provide some examples, before presenting the main results. A reader who prefers to first see the main results can skip this subsection.

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 agent's 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 we have previously described. 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 other players.

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 $a_1, ..., a_N$ denote the final cumulative effort of all players, final payoffs have the form $U_i(v_{iT}, a_i, a_{-i})$ satisfying Assumption 2. 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 the arrival time and price are random, following a joint Markov process. Final payoffs are $\sum_\tau p_\tau q_\tau - C(Q)$, where $Q = \sum_\tau q_\tau$, and C is a strictly increasing and convex function. As an example, a utility company might face opportunities to sell 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 aforementioned 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.[11]

[11] These results extend to the case where p_t is not a martingale. This is done by backloading payoffs,

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 $u\left(\sum a_\tau\right) - \sum c_\tau a_\tau$, where u is increasing and concave. Procrastination occurs because an agent might put a lower level of effort anticipating 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 example. 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 both 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 between no entry and entry. 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 cannot retract past bids. Examples include eBay and GovDeals auctions.[12] In these auctions, bidders frequently place multiple bids over time and increase them as the auction progresses. To model these auctions in the aforementioned class of decision problems, suppose that at random times τ and with an expected final value v, the agent can place (or increase) a bid b_τ. The final expected payoff in the auction will depend on the final value v_T, the final bid b_T of this agent, and the bids of the other agents. Integrating over the bids of others, we find that the expected final value has the form $\left[v_T - \mathbb{E}\left(b_2 | b_2 \leq b_T\right)\right] Prob\left(b_T \text{ is the highest bid}\right)$. 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

as mentioned, and redefining payoffs in those histories where there are no further arrivals to pick up the difference $E(v_\tau | v) - v$. Details of this procedure are available upon request.

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

to maximize

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

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 $U(v,a) = Bu(v,a)$. Let $P(t+1|t) = \beta$ and $P(t'|t) = 0$ for all $t' > t+1$. As of time zero, this implies that the probability of no arrivals is $(1-\beta)$, and that of only n arrivals, $(1-\beta)\beta^n$. Expected final value at time zero is

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

The corresponding final actions are $a_0, a_1, ...$ and

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

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 aforementioned 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.5.

4 Main Results

In this section, we present our main theoretical contribution. We start with dynamic decision problems and then show how our methodology extends to a class of dynamic games with incomplete information. Finally, we analyze the informational and behavioral properties of optimal actions.

4.1 Solving Dynamic Decision Problems Using ECV

In this section, we generalize our findings from Section 2 to the broader mathematical framework established in Section 3. While we build upon the ECV concept introduced earlier, the development follows a different logical sequence. First, we formally define the self-generating property for a function $e : \mathbb{R} \times [0, T] \to \mathbb{R}$, which serves as the mathematical analogue to Equation 4. Subsequently, under the regularity conditions specified in Assumption 3, we prove both the existence and uniqueness of such a function (Proposition 1). Finally, we demonstrate that $e(v, t)$ indeed satisfies the equivalent certain value property: the optimal dynamic action at state (v, t) coincides with the optimal static action when there are no future decision opportunities and the terminal value equals $e(v, t)$ (Theorem 1).

Before introducing the self-generating property, we require additional notation for the decision process. For any sample path $\omega \in \Omega$ in our probability space, let $D(\omega) \subset \mathbb{R} \times [0, T]$ denote the set of all decision nodes (v, t) that occur along path ω.

For every arrival state (v, t), any arbitrary real value $\varepsilon \in \mathbb{R}$, and any arbitrary function $e : \mathbb{R} \times [0, T] \to \mathbb{R}$, let $H_e(\varepsilon, v, t)$ denote the set of sample paths $\omega \in \Omega$ that satisfy the following two conditions: first, there exists an arrival at time t with signal value v; and second, at all subsequent arrivals (v', t') along these paths, the function $e(v', t')$ does not exceed ε. Formally,

$$H_e(\varepsilon, v, t) = \{\omega \in \Omega \mid (v, t) \in D(\omega) \text{ and } \forall t' > t, (v', t') \in D(\omega) \Rightarrow e(v', t') \leq \varepsilon\}.$$

Definition. A function $e : \mathbb{R} \times [0, T] \to \mathbb{R}$ possesses the *self-generating expectation property* with respect to the transition process P if and only if it satisfies the following condition for all $(v, t) \in \mathbb{R} \times [0, T]$:

$$e\left(v, t\right) = \mathbb{E}_{H_e(e(v,t),v,t)} v_T. \tag{6}$$

This property establishes a fixed-point characterization analogous to Equation 4.[13] For ease of notation, henceforth we drop the subscript e from the notation of set H. We next show that a function with a self-generating property exists and is unique. To establish these results, we require additional notations and regularity conditions on the transition process. Let $N(v, t) \subset \Omega$ denote the set of sample paths for which an arrival occurs at state (v, t) and no subsequent arrivals occur before the terminal time T. Formally,

[13] The set $H_e\left(e\left(v, t\right), v, t\right)$ includes all paths after time t, in which no future arrival has a higher value for $e(v', t')$ than that of $e(v, t)$. Borrowing the future result from Theorem 1, $H\left(e\left(v, t\right), v, t\right)$ represents all paths with no future higher optimal action than that of time t; therefore, the current optimal action will prevail.

$$N(v,t) = \{\omega \in \Omega | (v,t) \in D(\omega) \text{ and } \nexists (v',t') \in D(\omega) \text{ such that } t' > t\}. \tag{7}$$

Let $\Pi(\omega|v,t)$ denote the conditional probability of ω given $(v,t) \in D(\omega)$. The following assumptions on the transition process P will be maintained throughout our analysis.

Assumption 3. *The transition process P satisfies the following properties:*

1. There exists a constant $\delta > 0$ such that $\Pi(N(v,t)|v,t) > \delta$ for all $(v,t) \in \mathbb{R} \times [0,T]$.
2. The integral $\int_{N(v,t)} v_T(\omega)\, d\Pi(\omega|v,t)$ is continuous in v and t.
3. The transition process, $P(v',t'|v,t)$, is continuous in (v,t) with respect to the topology of weak convergence.

The first condition of this assumption states that there is a positive probability bounded away from zero that the current decision node is the last one. This ensures that the agent faces a nontrivial probability of having to commit to their current action at each decision point. The second and third conditions impose standard continuity requirements on the stochastic process governing signal evolution. In the special case where arrival times and signal values are stochastically independent, the integral in the second condition simplifies to $v\Pi(N(v,t)|v,t)$, implying that the condition is satisfied if the probability of no further arrivals before time T is a continuous function of t. This condition is satisfied, for example, in the Poisson arrival process where $\Pi(N(v,t)|v,t) = \exp(-\lambda(T-t))$ for arrival rate parameter $\lambda > 0$.

Proposition 1. *For any transition process P satisfying Assumption 3, there exists a unique function $e : \mathbb{R} \times [0,T] \to \mathbb{R}$ that satisfies the self-generating expectation property defined in Equation 6.*

The proof of this proposition is relegated to the appendix, but we outline the major steps here. We establish existence and uniqueness through a constructive approach using an auxiliary functional equation. Specifically, the function $W : \mathbb{R} \times \mathbb{R} \times [0,T] \to \mathbb{R}$ solves

$$\begin{aligned} W(\varepsilon, v, t) &= \int_t^T \min\left(W\left(\varepsilon, v', \tau'\right), 0\right) dP\left(v', \tau' | v, t\right) \\ &\quad + \int_{N(v,t)} \left(v_T(\omega) - \varepsilon\right) d\Pi(\omega|v,t) \end{aligned} \tag{8}$$

with the boundary condition $W(\varepsilon, v, T) = v - \varepsilon$.

We prove that this functional equation is a contraction mapping, and, thus, by the Banach fixed point theorem, it admits a unique fixed point. The solution has several properties: it is strictly decreasing in ε, continuous in all arguments, nonnegative at $\varepsilon = 0$, and negative for sufficiently large ε.

By the intermediate value theorem, for each $(v, t) \in \mathbb{R} \times [0, T]$, there exists a unique value $e(v, t) \in \mathbb{R}$ such that $W(e(v, t), v, t) = 0$. We then show that this implicitly defined function $e(v, t)$ satisfies the self-generating expectation property given in Equation 6. The last step of the proof is to show that we can construct a function W that solves 8 from any function $e(v, t)$ that satisfies the self-generating property.

This functional equation provides a computationally tractable approach to solving for $e(v, t)$. One can iteratively apply the contraction mapping starting from any initial function, with convergence guaranteed by the Banach fixed point theorem. The terminal condition $W(\varepsilon, v, T) = v - \varepsilon$ provides a natural starting point for a backward induction procedure. Furthermore, under additional regularity conditions on the arrival process, one can transform this functional equation into a partial differential equation amenable to standard numerical methods.[14]

Next, we establish the connection between the self-generating expectation function $e(v, t)$ and the optimal decision strategy for the dynamic decision problem. Consider the static optimization problem

$$\max_{a \in A} U(v, a). \tag{9}$$

Let $\tilde{S}(v) : \mathbb{R} \to A$ be a (weakly) increasing selection from the argmax correspondence, which is guaranteed to exist by the supermodularity condition in Assumption 2. This selection function represents the optimal action for an agent who must make a final decision at time T with terminal value $v_T = v$.

Theorem 1. *For any payoff function $U(v, a)$ satisfying Assumption 2 and transition process P satisfying Assumption 3, the strategy $S : \mathbb{R} \times [0, T] \to A$ defined by $S(v, t) = \tilde{S}(e(v, t))$ is an optimal strategy for the dynamic decision problem 5, where the function $e(v, t)$ is the unique self-generating expectation function satisfying Equation 6 corresponding to the process P.*

[14] In Hopenhayn and Saeedi [2020], a companion empirical paper where we apply this method to the case of online auctions, we assume that arrival times follow a Poisson distribution and values change using a Brownian motion. Given these assumptions, Equation 8 becomes a simple partial differential equation, similar to the heat transfer PDF that is easy to compute in any programming language.

This theorem establishes a powerful dimensional reduction result for the dynamic decision problem. Furthermore, we can now formally define the ECV. The ECV at point (v, t) is equal to value of function $e(v, t)$ that uniquely solves the self-generating expectation fixed point of Equation 6; computationally, it is the unique root $W\big(e(v,t), v, t\big) = 0$ of the contraction W in Equation 8, which depends only on the transition P. Additionally, the optimal action at at any state (v, t) is precisely equivalent to the optimal action for a hypothetical terminal-period decision with value $e(v, t)$: $S\left(v, t\right) = \tilde{S}\left(e\left(v, t\right)\right)$.

The significance of this characterization lies in its transformation of a complex dynamic optimization problem into a straightforward two-step procedure: (i) compute the self-generating expectation function $e(v, t)$ as the unique solution to the fixed-point Equation 6, and (ii) determine the optimal static strategy $\tilde{S}$ for problem 9 and apply it to the computed ECVs, yielding $S(v, t) = \tilde{S}(e(v, t))$.

In essence, Theorem 1 reveals that the ECV function $e(v, t)$ fully encapsulates all relevant forward-looking information needed for optimal dynamic decision-making, distilling the multi-dimensional stochastic process into a single sufficient statistic at each decision point.

A crucial insight from this result is that the computation of $e(v, t)$ depends solely on the stochastic properties of the probability measure P of the transition process, and is entirely independent of the specific payoff function $U(v, a)$, as long as the payoff function satisfy Assumption 2. This separation principle provides significant computational advantages, particularly in contexts where the payoff function may be unknown or may vary across applications. Moreover, in strategic settings, as shown in Section 4.2, where the payoff function depends on the strategies of other agents, this decomposition allows us to isolate the stochastic analysis from the strategic considerations so the potentially intensive computation of $e(v, t)$ is performed only once.

This separation property relies on linearity in the value, as required by Assumption 2, ensuring that $\mathbb{E}[U(v_T, a)] = U(\mathbb{E}[v_T], a)$. Our methodology extends beyond linear utilities to cases where the value of agent enters through monotone transformations. For utility functions of the form $U(v, a) = f(v)h(a) + g(a)$ with strictly monotone $f : \mathbb{R} \to \mathbb{R}$, we can reformulate using $\hat{v} = f(v)$ to obtain $\hat{U}(\hat{v}, a) = \hat{v}h(a) + g(a)$ and apply our methodology to this transformed value. While in this case the specific ECV depends on transformation f and requires recalculation for different utility specifications, the fundamental principle—solving the stochastic component once and applying it to various decision problems—remains intact within each transformation class.

These results are further discussed in Online Appendix C.6.

4.2 Embedding in Games

As our leading example suggests, our analysis of dynamic decision problems extends to a class of games with 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 earlier. This best response can be characterized by maximizing expected payoffs at the corresponding ECVs, as established in Theorem 1. The strategic setting introduces an additional equilibrium requirement: each player's strategy must be a best response to the strategies of all other players. To address this complexity, we formulate a static Bayesian game in which each player's type distribution corresponds to the distribution of that player's ECVs, and the strategies map these values into their corresponding action sets. Finally, we establish that any equilibrium of this static Bayesian game induces an equilibrium in the original dynamic game.

Let $\Gamma=\left(I,\{A_i\}_{i\in I},\{Z_i\}_{i\in I},\{P_i\}_{i\in I},\{u_{iT}\}_{i\in I}\right)$ denote a dynamic game with the following structure. There is a fixed set of players $I=\{1,...,N\}$. Each player i faces a stochastic process for values $v\in Z_i$ and decision times in $[0,T]$ characterized by a Markov transition function $P_i(v,t)$, with these processes being independent across players. Final payoffs are given by utility functions $u_{iT}(v_{iT},a_{iT},a_{-iT})$, where v_{iT} represents the vector of final values for player i, and (a_{iT},a_{-iT}) denotes the vector of terminal actions selected from totally ordered sets $A_1,...,A_N$. We assume that each player's information set contains only that player's own history, which restricts strategy profiles to the form $S_i: Z_i\times[0,T]\to A_i$, specifying action choices as functions of individual histories. Without loss of generality, we can focus on Markov strategies $S_i(v,t)$. Let $\mathbf{S_i}$ denote the set of admissible strategies for player i, and define the expected utility function $u_i(S_i,S_{-i})=\mathbb{E}_0\left[u_{iT}(v_{iT},a_{iT},a_{-iT}|S_i,S_{-i})\right]$.

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

Finding the Nash equilibria of this dynamic game seems a formidable challenge, given the high-dimensionality of the strategy space. However, we now demonstrate how this problem can be reduced to finding equilibrium in a simple static Bayesian game where strategies are one-dimensional and an appropriate mapping to the dynamic game.

Consider a player i. For every sample path ω, we can identify a unique value corresponding to the maximum of the ECV reached along that path as $\bar{e}_i(\omega)=\max\{e_i(v_n(\omega),\tau_n(\omega))\}$. This procedure uniquely determines a distribution Ψ_i of the ECVs for this player that depends only on the corresponding Markov process P_i. Using these distributions, we define a static

Bayesian game as follows: set of players $I = \{1, ..., N\}$, distribution of values for each player $\{\Psi_1, ..., \Psi_N\}$, strategy sets $\{A_1, ..., A_N\}$, and payoff function $\{u_{iT}(v_i, a_i, a_{-i})\}_{i \in I}$.

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

We need to make the following analogous assumption to Assumption 2.

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

Similar to the dynamic decision problem, we can relax Assumption 4 to allow for utilities to have a monotone function of v_i rather than being linear in the values. We now present our main result, which establishes a precise equivalence between the equilibria in the dynamic game and those in the associated static Bayesian game.

Theorem 2. *Consider a game Γ that satisfies Assumption 4 and its associated static Bayesian game Γ_B. Furthermore, assume that the Markov processes $P_i(v, t)$ satisfy Assumption 3, for all i. For any profile of equilibrium strategies $\left\{\tilde{S}_i\right\}_{i \in N}$ of Γ_B, the strategy profile defined by $S_i(v, t) = \tilde{S}_i(e_i(v, t))$ constitutes an equilibrium for Γ, where the function $e_i(v, t)$ is the self-generating expectation for player i, satisfying Equation 6.*

This theorem provides a powerful decomposition of the equilibrium analysis into two distinct components: (1) a dynamic decision problem—determining the ECVs—and (2) a static equilibrium problem in the static Bayesian game. This decomposition offers several key advantages beyond computational efficiency. First, it separates the stochastic and strategic components of the problem. The ECV computation—a purely stochastic calculation—need only be performed once per player, independent of strategic interactions. This property facilitates comparative statics analyses, allowing researchers to examine how payoff structure changes affect equilibrium behavior without recalculating dynamic elements.

Second, this result clarifies how dynamic considerations influence strategic behavior. The players in the dynamic game respond to their ECVs rather than directly to current signals and decision times. This provides a formal characterization of forward-looking behavior: the players act as if in a static environment but with transformed valuations incorporating the option value of future decisions.

A limitation of our framework is its reliance on information independence across players. When they observe each other's actions or receive correlated signals, it becomes hard to separate the strategic and stochastic components of the problem. Nevertheless, the ECV

concept remains instructive in understanding how dynamic elements shape strategic behavior.

Additionally, for some special cases, this informational restriction 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 earlier remains an equilibrium of the dynamic game with any added information about the 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, the sellers have incomplete information about their rivals: the sellers observe the aggregated revenue process but not the actions of their individual rivals.

4.3 Value Shading and Informational Properties of ECVs

Building upon our characterization of equivalent certain values and their application to strategic environments, we now examine a fundamental property that emerges in dynamic problems with irreversibility constraints. Agents systematically shade their valuations—acting as if terminal values were lower than conditional expectations—due to the option value associated with future decision opportunities.[15] This section establishes three key results regarding this phenomenon: (i) the existence and general properties of value shading under regularity conditions; (ii) the monotonicity of shading with respect to the temporal proximity to the deadline; and (iii) the relationship between the frequency of decision opportunities and the magnitude of shading. We also establish an analogue between decision-making under irreversibility and decision-making with reduced information: agents facing irreversibility constraints behave as if they were in an environment with reversible actions but less precise information. We formalize this relationship through mean-preserving contraction comparisons of the relevant value distributions. Throughout, we show how the self-generating expectation property of the ECVs provides a tractable framework for analyzing these effects while also yielding insights into the informational content of the agents' observed choices. These results have significant implications for market design and empirical identification, which we will explore in subsequent applications.

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

[15] Value 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.

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 condition would immediately hold under the martingale property if expected values were independent of arrival times. The second condition would hold if the process were independent of the arrival time t', for example, if the probability of a future arrival were strictly positive and the conditional distribution for value were nondegenerate, i.e., $P(v' = v|v,t) < 1$. We first establish the fundamental relationship between ECVs and expected values, demonstrating that optimal forward-looking behavior systematically involves value shading.

Proposition 2. *The equivalent value is less than or equal to the expected value:* $e(v,t) \leq \mathbb{E}[v_T|v,t]$. *If Assumption 5 holds, then the inequality is strict.*

This result has immediate implications for optimal actions: value shading translates directly into action shading through the monotonicity properties of the optimal policy function. Specifically, $S(v,t) = \tilde{S}(e(v,t)) \leq \tilde{S}(\mathbb{E}[v_T|v,t])$, where the inequality follows from the supermodularity assumption of the utility function.

This proposition implies that at any time before the terminal time, values and the corresponding choices are shaded. One might wrongly conclude that the distribution of ECVs and the corresponding choices is downward biased. In fact, ECVs are unbiased, as established in the following proposition. There is an intuitive reason: while the values are shaded according to Proposition 2, they remain unbiased for the set of paths where these values are the terminal ECVs. This follows directly from the definition of self-generating expectation given by Equation 6. The following proposition formally proves this intuition.

Proposition 3. *The terminal ECV is unbiased, namely, at any decision node* $(v,t) \in D(\omega)$*:*

$$\mathbb{E}\left[max\left\{e(v',t'), t' \geq t, (v',t') \in D(\omega)\right\} | (v,t) \in D(\omega)\right] = \mathbb{E}\left[v_T|v,t\right].$$

This proposition establishes that starting from any decision node (v,t), the expected value of the terminal ECV—defined as the maximum equivalent certain value reached along the sample path—equals the conditional expectation of the terminal value. In essence, while agents shade their valuations at each individual decision node, the terminal ECV distribution has the same mean as the distribution of terminal values.

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

Implications of Irreversibility Proposition 3 establishes that the distribution of ECVs is unbiased. To further analyze the impact of irreversibility constraints, we now compare our setting to an alternative environment where actions are fully reversible. For any sample path ω, let $\bar{e}(\omega)$ denote the maximum ECV achieved on that path under irreversibility. For the alternative environment with reversible actions, define the *final value* for the path ω as

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

If arrival times and valuation are stochastically independent values, then $\bar{v}(\omega)$ will simply be $v(T)$. This quantity represents the only relevant value for determining the final action on path ω in the reversible case, with the optimal action given by $a(\omega) = \max_a U(a, \bar{v}(\omega))$. The following proposition formalizes the relationship between the terminal value distributions with and without irreversibility constraints.

Proposition 4. *The distribution of final equivalent values $\{\bar{e}(\omega)\}$ is a mean-preserving contraction of the distribution of final values $\{\bar{v}(\omega)\}$.*

The intuition behind this result is that irreversibility constrains the actions of the agent, limiting their ability to condition their choices on information as it arrives. When actions must be increasing, the agent cannot fully exploit unfavorable information realizations about lower terminal values, creating an informational distortion relative to the unconstrained case. This distortion manifests as a reduction in the dispersion of the distribution of decision-relevant values.

The decrease in spread has an additional intuitive interpretation related to how informed the agent's decisions are. Consider the benchmark of complete information, where agents directly observe the terminal value $v_T(\omega)$ before selecting their action. In the reversible case, the agent's decision is based on $\bar{v}(\omega)$, the expected value at the last decision node, which is a mean-preserving contraction of $v_T(\omega)$. However, in our dynamic setting with irreversibility, the distribution $\bar{e}(\omega)$ represents the implicit information embedded in the agent's decision process. 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, which implies that $\bar{e}(\omega)$ has a lower covariance with the true values $v_T(\omega)$ than $\bar{v}(\omega)$. This reduced covariance formalizes our key insight: agents with irreversible actions behave as if they possessed less precise information than agents with full action flexibility.

The specific implications for optimal actions depend on specific characteristics of the utility function $U(a, v)$. For example, if the optimal policy $a(v) = \arg\max_a U(a, v)$ is concave (resp., convex) in v, then the expected action will be higher (resp., lower) in the presence of irreversibility constraints compared to the fully reversible benchmark. The comparison becomes more complex in strategic settings, as the utility function itself depends on the equilibrium strategies of the other players. We explore a concrete application of this insight in Section 5, where we analyze a dynamic second-price auction.

Partial Irreversibility Our baseline model assumes complete irreversibility of actions: once an action is chosen, it cannot be decreased at future decision nodes. We now extend this framework to incorporate partial reversibility through random opportunities for agents to freely adjust their actions without constraints from past choices.

Consider a modified environment where, in addition to the standard decision opportunities governed by the transition process P, agents face random arrivals of reversibility opportunities following a Poisson process with intensity $\rho \geq 0$. When such an opportunity arrives, the agent can freely reset their action to any level, effectively removing the irreversibility constraint for that decision.

Formally, let $\{R_k\}_{k=1}^{\infty}$ denote the arrival times of reversibility opportunities, following a Poisson process independent of the decision process. At any decision node (v, t), the agent's action choice depends on whether a reversibility opportunity has occurred since that agent's last decision:

- With reversibility (probability $1 - \exp(-\rho\Delta t)$ over interval Δt): The agent can choose any action $a \in A$ without constraint.
- Without reversibility (probability $\exp(-\rho\Delta t)$): The agent faces the standard monotonicity constraint $a \geq a_{\text{previous}}$.

Under partial reversibility, the ECV function $e_\rho(v, t)$ represents the equivalent certain value when the reversibility rate is ρ. This function satisfies a modified self-generating expectation

property where adverse future paths (i.e., those requiring downward action adjustments) become less constraining as ρ increases.

The corresponding value function $W_\rho(\varepsilon, v, t)$ solves

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

The key modification is the exponential discount factor $\exp(-\rho(\tau' - t))$ applied to future decision nodes. This reflects the probability that no reversibility opportunity occurs between the current time t and future decision time τ'. When a reversibility opportunity does occur, the irreversibility constraint is lifted, eliminating the negative contribution of adverse future realizations.

Proposition 5. *The ECV function $e_\rho(v, t)$ exists, is unique, and satisfies the following conditions:*

1. $e_\rho(v, t)$ is increasing in ρ for all (v, t).
2. $\lim_{\rho \to 0} e_\rho(v, t) = e(v, t)$ (baseline irreversible case).
3. $\lim_{\rho \to \infty} e_\rho(v, t) = E[v_T | v, t]$ (fully reversible case).
4. The distribution of terminal ECVs under ρ is a mean-preserving contraction of terminal values.

As ρ increases, $e_\rho(v, t)$ rises and the shading $s_\rho(v, t) \equiv \mathbb{E}[v_T \mid v, t] - e_\rho(v, t)$ falls. Intuitively, shading hedges against adverse future states in which the agent can update only one way because past actions are irreversible. Allowing partial reversibility introduces bi-directional adjustments, reducing the incidence of such states and, therefore, the incentive to shade.

Moreover, partial reversibility preserves the static reduction (policy equivalence) and yields clean comparative statics: the irreversible case ($\rho = 0$) is a lower-bound benchmark, the fully reversible case ($\rho \to \infty$) an upper bound, and design objects move monotonically in ρ.

Frequency of Arrival Times The frequency of decision opportunities represents a critical design parameter in many dynamic mechanisms. Platform designers often influence this frequency by sending reminders about making choices or by providing updates about the

time remaining before a deadline. Such interventions can either facilitate or hinder action adjustments, depending on their implementation. In this section, we analyze how variations in arrival frequency affect optimal actions through their impact on ECVs.

Intuitive reasoning suggests that increasing the frequency of decision opportunities should increase the option value associated with future choices, thereby decreasing ECVs and amplifying value shading. To derive formal results on this relationship, we introduce a structural assumption that isolates the effect of arrival frequency from other aspects of the stochastic environment.

Assumption 6. *The transition process is* $P\left(v', t'|v, t\right) = P_v\left(v'|v, t'\right) F\left(t'|t\right)$.

Additionally, we assume that v follows a martingale, so $E(v'|v) = v$. We say that the distribution $\tilde{F}\left(t\right)$ exhibits higher arrival frequency than $F\left(t\right)$ if $F\left(t'|t\right) \leq \tilde{F}\left(t'|t\right)$ for all $t' > t$, meaning that $F\left(t'|t\right)$ first-order stochastically dominates $\tilde{F}\left(t'|t\right)$.

Proposition 6. *Suppose* $F\left(t'|t\right) \leq \tilde{F}\left(t'|t\right)$ *and* $F\left(t', t\right)$ *is non-increasing in* t*, for all* t, t'*. Then the corresponding ECVs satisfy* $e\left(v, t\right) \geq \tilde{e}\left(v, t\right)$*. Furthermore, this implies that the distribution of terminal ECVs* $\bar{e}\left(\omega\right)$ *under* $\tilde{F}$ *is a mean-preserving spread of the distribution corresponding to* F*.*

This proposition formalizes the intuition that higher arrival frequency leads to lower ECVs and thus more aggressive value shading. The mechanism operates through the option value channel: more frequent opportunities to adjust actions enhance the value of waiting because there is a higher chance of next arrival, incentivizing more conservative initial choices. The mean-preserving spread relationship between terminal ECV distributions indicates that, from an informational perspective, higher arrival frequency enables agents to make decisions that more closely reflect their true terminal values. This relationship parallels our earlier finding about reversible actions: increased arrival frequency, similar to reduced irreversibility, allows agents to behave as if they had more precise information about terminal values, resulting in a higher covariance between terminal ECVs and true terminal values.

To build intuition for Proposition 6, consider the simple two-period example from Section 2. An increase in arrival frequency corresponds to an increase in p, the probability of a second-period decision opportunity. Let $\tilde{p} > p$ denote the higher arrival rate, with corresponding first-period ECVs $\tilde{e}_0$ and e_0 respectively.

The higher arrival rate $\tilde{p}$ creates two effects. First, it reduces the initial ECV: $\tilde{e}_0 < e_0$. This occurs because the agent anticipates more frequent opportunities to adjust their action,

increasing the option value of waiting and thus inducing more aggressive shading. Second, it changes how final ECVs are distributed across different paths.

To see the mean-preserving spread, partition the paths based on the final ECV under the original arrival rate p. Consider paths where e_0 remains the final ECV—these occur either when no second-period arrival happens (probability $1-p$) or when arrival occurs but $v_1 < e_0$. Under the higher arrival rate $\tilde{p}$, three changes occur: (i) fewer paths end without arrival (probability $1 - \tilde{p} < 1 - p$), (ii) the update threshold is lower (paths update when $v_1 > \tilde{e}_0$ rather than $v_1 > e_0$), and (iii) both update regions lead to the same terminal value v_1 when updates occur. Therefore, this set which under arrival rate p, has only one final ECV, e_0, under $\tilde{p}$, has multiple values. Additionally, the expected value of this final values in both set are equal to the expected final values which is the same.

The key insight is that while both distributions have the same mean (equal to $\mathbb{E}[v_1]$), the distribution under $\tilde{p}$ concentrates probability at more extreme values. Some probability mass shifts from the intermediate value e_0 to the lower value $\tilde{e}_0$ (paths without arrival), while other mass shifts to higher values (additional paths that now update because $\tilde{e}_0 < v_1 < e_0$) and more arrivals which will result in further dispersion of final ECV values. This redistribution from intermediate to extreme values, while preserving the mean, generates the mean-preserving spread. The higher arrival frequency essentially allows agents to make more informed final decisions.

From a mechanism design perspective, these results suggest that platform operators can influence the agents' behavior by modifying the frequency of decision opportunities. Increasing this frequency generally leads to more informed decisions and potentially higher efficiency, though the welfare implications depend on the specific objective function of the designer and the distribution of values among participants. In auction settings, for instance, higher arrival frequency leads to decreased shading, which can result in higher prices; therefore, counterintuitively, it can benefit the auctioneer and harm the bidders.

Shading Over Time Shading arises from the option value associated with future decision opportunities. This option value is affected by the probability of having additional opportunities to adjust actions and the potential for new information arrival regarding the terminal value. When future decision opportunities are highly likely or when significant information is expected to arrive, the option value increases, causing agents to shade their valuations more aggressively. As a natural consequence, one might expect value shading to diminish as the terminal period approaches, since both the likelihood of future opportunities and the potential for new information decrease over time.

However, in our general framework, this monotonic reduction in shading is not guaranteed without additional structural assumptions. Two potential complications arise: First, the distribution of future decision opportunities may exhibit complex dependencies, where the occurrence of one decision point alters the likelihood of subsequent arrivals. Second, the information structure might be nonstationary, with the precision or relevance of incoming information varying systematically over time. To establish the conditions under which shading decreases over time, which we formalize as ECVs increasing with time, we consider two specialized environments that accommodate many applications of interest.

Proposition 7. *The function $e(v,t)$ is increasing in t if either of the following conditions holds.*

1. Values and arrival times are independent of each other, and $F(\tau'|t)$ is (weakly) decreasing in t for all τ'.
2. $P(v',\tau'|v,t) = P_v(v'|v,\tau'-t)\,F(\tau'-t)$.

The first condition covers environments where value evolution and decision opportunities follow independent processes, with the additional requirement that the probability of future arrivals (weakly) decreases as time progresses. This condition is satisfied when arrival hazard rates are independent of previous arrivals, as in Poisson processes. Importantly, this specification accommodates time-varying arrival rates, including cases where agents become more attentive as the deadline approaches (i.e., increasing $\lambda(t)$).

While the first condition 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 condition 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 followed a standard Poisson process with a constant hazard rate.

An immediate corollary of this proposition is that optimal actions, holding fixed the observed value v, increase over time. This monotonicity result aligns with empirical observations in auction and trading settings, where participants tend to increase their bids or offers as deadlines approach. From a mechanism design perspective, this insight suggests that temporal patterns in agent behavior can be systematically influenced through the appropriate specification of arrival processes and information structures.

4.4 Rational Inattention via Arrival Frequency

In this section, we show how our main characterization result can be used to study endogenous attention. To do so, we allow the agent to choose the distribution of the timing of choice at a cost in the spirit of the rational inattention literature pioneered by Sims [2003] and recently extended by Steiner et al. [2017] and Zhong [2022]. As we will show, the use of our ECV framework allows us to provide sharp predictions about agent's optimal attention strategies.

For the sake of convenience, we focus on Poisson arrival processes with constant arrival rates where the rate is chosen by agents at a cost.[16] Specifically, suppose arrivals follow a Poisson process with rate $\lambda \in \Lambda \subset \mathbb{R}_+$ and that the process for values satisfies $P(v', t'|v, t) = P_v(v'|v, t, t') F(t'|t)$, i.e., Assumption 6 holds. In other words, at any point in time t, the next arrival time has distribution $F_\lambda(\tau'|t) = 1 - \exp(-\lambda(\tau' - t))$ where λ is chosen at time–0 by the agent. Higher rates correspond to earlier arrivals: if $\lambda'' > \lambda'$, then $F_{\lambda''}(\cdot|t)$ first-order stochastically dominates $F_{\lambda'}(\cdot|t)$, satisfying Assumption 6 and stochastic dominance in Proposition 6. As a result we can apply the result of this proposition for different arrival rates: an increase in arrival rates results in mean-preserving spread of distribution of terminal ECV values.

Thus the agent chooses λ at time 0 at cost $C(\lambda)$. This approach follows Maćkowiak et al. [2018] in treating information acquisition as forward-looking optimization. We assume the following:

Assumption 7. *The cost function $C(\lambda)$ is increasing and convex in $\lambda \in \Lambda$.*

This cost function can have different forms, such as Shannon entropy or Kullback-Leibler divergence between the chosen and baseline information processes, which yield the required monotonicity and convexity in λ.

Let us start with single agent decision problem. For any fixed λ, let $W_\lambda(\varepsilon, v, t)$ denote the W-equation with arrival distribution indexed by λ, and let $e_\lambda(v, t)$ be the unique root $W_\lambda\big(e_\lambda(v, t), v, t\big) = 0$. Analogously define *terminal ECV* along path ω by

$$\bar{e}_\lambda(\omega) := \max\left\{e_\lambda(v', t') : (v', t') \in D(\omega)\right\}.$$

Let $G(v) := \max_{a \in A} U(v, a)$ be the indirect static value and $\tilde{S}(v) \in \arg\max_a U(v, a)$ the static best reply. For any fixed arrival rate λ, the optimal dynamic policy is $S_\lambda(v, t) =$

[16] The results presented below can be extended to more general arrival processes for which Assumption 6 and the assumption in Proposition 6 are satisfied; for example, Poisson processes with variable arrival rates or renewal processes can also be considered..

$\tilde{S}\big(e_\lambda(v,t)\big)$. Consequently, the optimal choice of arrival frequency solves

$$\max_{\lambda\in\Lambda} E\left[G\left(\bar{e}_\lambda\right)\right] - \kappa C(\lambda).$$

Blackwell/convex-order comparative statics. If $\lambda'' > \lambda'$, Proposition 6 implies that the distribution of $\{\bar{e}_{\lambda''}\}$ is a mean-preserving spread of $\{\bar{e}_{\lambda'}\}$. Since agents' final actions depend only on these terminal ECVs, this distributional relationship establishes a Blackwell ordering: higher arrival frequency generates more informative decision-relevant values. Additionally, given Assumption 2, G is weakly convex and increasing in v. Therefore $E[G(\bar{e}_\lambda)]$ is weakly increasing in λ; abstracting from monitoring costs, higher arrival frequency is weakly preferred.

Proposition 8. *The optimal arrival frequency λ^* is weakly decreasing in attention cost κ.*

This result mirrors static rational inattention models, where lower information costs induce higher information precision. Here, higher λ plays the role of greater precision: it generates a mean-preserving spread of effective values $\bar{e}_\lambda$, so comparative statics in λ correspond directly to comparative statics in rational inattention precision. Another obvious implication of lowering κ is that it leads to higher expected utility for the agent.

Dynamic Games In strategic environment, each player's choice of arrival rate will potentially impact the strategies of other players, creating strategic interdependence absent in the single-agent case. Consider a dynamic Bayesian game with endogenous arrival rates: $\Gamma_{\boldsymbol{\lambda}} = (I, \{A_i\}_{i\in I}, \{Z_i\}_{i\in I}, \{P_{i,\lambda_i}\}_{i\in I}, \{u_i\}_{i\in I})$, where $P_{i,\lambda_i}(v,t)$ denotes player i's Markov process conditional on chosen arrival rate $\lambda_i \in \Lambda_i \subset \mathbb{R}_+$. The cost function is as discussed above, $\kappa_i C_i(\lambda_i)$. The game proceeds in two stages. First, players simultaneously choose arrival rates $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_N)$. Second, given these choices, they play the dynamic game characterized in Section 4.2. The associated static Bayesian game is: $\Gamma_{B,\boldsymbol{\lambda}} = (I, \{\Psi_{i,\lambda_i}\}_{i\in I}, \{A_i\}_{i\in I}, \{u_{iT}\}_{i\in I})$, where $\{\Psi_{i,\lambda_i}\}$ is the distribution of terminal ECVs $\{\bar{e}_{i,\lambda_i}\}$ under arrival rate λ_i. By Theorem 2, equilibrium strategies satisfy $S_{i,\lambda_i}(v,t) = \tilde{S}_i(e_{i,\lambda_i}(v,t))$.

The equilibrium arrival rate profile $\boldsymbol{\lambda}^*$ must satisfy:

$$\lambda_i^* \in \arg\max_{\lambda_i\in\Lambda_i} \mathbb{E}\left[u_{iT}\left(\bar{e}_{i,\lambda_i}, \tilde{S}_i(\bar{e}_{i,\lambda_i}), \tilde{S}_{-i}(\bar{e}_{-i,\lambda^*_{-i}})\right)\right] - \kappa_i C_i(\lambda_i) \tag{10}$$

An increase in arrival rate for one of users, λ_i, leads to a mean-preserving spread of $\{\Psi_{i,\lambda_i}\}$. This increase in arrival rate is akin to an increase in information on the value of this player

in the static game Γ_B. Therefore, the strategic response as a result of change in arrival rate in dynamic Bayesian game, is analogous to the strategic response in the corresponding static game when the information of this player increases. Furthermore, an increase in arrival rate of one player will alter the equilibrium strategies of all players $\{\tilde{S}_j\}_{j\neq i}$ including their choice of arrival rates, $\{\lambda_j\}_{j\neq i}$.

First let us consider an auction format in which the person with the highest terminal ECV wins the object, such as first-price auction or second-price auction. In these auctions, bids of players are based on the distribution of their terminal values $\{\Psi_{i,\lambda_i}\}_{i\in I}$. Let us denote the CDF of the terminal values by $\{H_{i,\lambda_i}\}$. Using Myerson [1981] and assuming no reserve price, we can write the interim value of a bidder with terminal value e_i as:

$$U_i\left(e_i\right) = U_i(0) + \int_0^{e_i} \prod_{j\neq i} H_{j,\lambda_j}\left(t\right) dt$$

Let us assume that $U_i(0) = 0$. First consider the two-bidder case. An increase in λ_2, results in a mean-preserving spread of H_{2,λ_2}, and hence an increase in interim payoff of bidder 1 at all possible ECV values. However, this relationship may not hold for more than two bidders.

Given the complicated interaction of these strategies, there is little knowledge of endogenous information acquisition in game settings beyond some basic characteristics. In a recent working paper, Quint and Kojima [2024], study a symmetric equilibria of a pre-auction game of investment. They show existence of symmetric equilibrium and uniqueness under some additional conditions. They further show that under a symmetric equilibrium a decrease in cost of information acquisition will result in higher information acquisition. We can extend these results to the case of dynamic auction where bidders choose λ_i.

Proposition 9. *In auctions, in a symmetric equilibrium, the optimal arrival frequency λ^* is weakly decreasing in attention cost κ.*

This example illustrates, how one can extend the results related to information acquisition in static games to the ones in dynamic games with endogenous arrival rates. In general, when studying a dynamic game $\Gamma_{\boldsymbol{\lambda}}$, a first-order increase in arrival rates will transform the static game, $\Gamma_{B,\boldsymbol{\lambda}}$, to a one with a mean-preserving spread of distribution of players values, $\{\Psi_{i,\lambda_i}\}_{i\in I}$.

Our analysis reveals a fundamental principle: dynamic frictions—whether from irreversibility constraints or limited decision opportunities—operate through information degradation à la Blackwell. This unified perspective provides analytical tractability through our ECV

decomposition while generating portable insights across economic environments. Policymakers evaluating decisions about reversibility or timing can apply information-theoretic logic: both frictions make agents act as if they had coarser information about payoff-relevant states. This equivalence suggests that classic results about information precision in static settings translate directly to predict effects of these dynamic frictions.

4.5 Advantages of the ECV Method

The ECV method provides a conceptually and computationally streamlined alternative to classical Bellman dynamic programming. It replaces iterative Bellman recursions with a two-step decomposition: a dynamic "option value" component and a static decision component. This separation reduces the overall complexity of the problem. Crucially, the dynamic component can be solved once and reused across different environments, allowing researchers to avoid repeated re-optimization for each new parameter configuration. This feature not only simplifies the computation of policy functions but also facilitates comparative statics and structural estimation, since payoff parameters influence only the static stage. In strategic settings, such as dynamic games, each player's problem reduces to a static Bayesian game once ECVs are computed. This enables the use of static tools to analyze dynamic interactions and extends familiar results from static environments to dynamic frameworks.

Beyond its tractability, the ECV approach yields intuitive economic insights that are often difficult to extract from traditional methods. In particular, the method highlights how irreversibility effectively degrades information: agents behave as if they had less precise knowledge about their future payoffs. The framework also allows for transparent analysis of how changes in irreversibility or the frequency of decision opportunities affects outcomes. Together, these features make ECV a broadly applicable and analytically transparent tool for studying dynamic problems involving irreversible actions.

5 Optimal Dynamic Auction Design

We now apply our theoretical framework to analyze dynamic second-price auctions. As discussed in Section 3.1, many auctions take place over extended periods of time. This application demonstrates how our ECV methodology transforms complex strategic bidding problems into tractable equilibrium characterizations, with clear implications for mechanism design. In the online 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, the 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.[17] 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 (as in the case of spectrum auctions, see Börgers and Dustmann [2005] and Bulow et al. [2009]), and alternative demands for use of resources in the face of capacity constraints (as in procurement auctions, see Jofre-Bonet and Pesendorfer [2003]).

5.1 Model Specification and Equilibrium Bidding

Consider a second-price auction conducted over time interval $[0, T]$. Each bidder i observes a sequence of private signals about their valuation according to a Markov process satisfying Assumption 1, with opportunities to adjust their bid arriving stochastically. Bids are restricted to be non-decreasing over time, reflecting the standard irreversibility constraints in many auction platforms.

Using Theorem 2, the optimal bid for player i at state (v, t) in the dynamic auction, equals the optimal bid in a static auction for $e_i(v, t)$, where $e_i(v, t)$ is the bidder's ECV at state (v, t). In the case of static second-price auctions, the optimal bid is trivial. Therefore, the optimal bid is simply the corresponding ECV $b_i(v, t) = e_i(v, t)$.

When Assumption 5 holds and the changes in values and arrival times are independent of each other, we can use the results in Proposition 2 and Proposition 7. The propositions imply that bids are strictly shaded below expected values and increase monotonically over time for a fixed value. The magnitude of shading is determined by solving the functional

[17] There 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 bidders' reservation price 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.

equation (8) given the stochastic processes governing valuations and bidding opportunities.[18]

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

5.2 Revenue Equivalence

The revenue equivalence principle from static auction theory extends to our dynamic setting under appropriate conditions. When the bidders have independent private values with symmetric distributions and are risk-neutral, the expected revenue depends only on the distribution of final allocation and the expected utility of the bidders with the lowest possible valuation. In the dynamic setting, the same conclusion holds by replacing values for ECVs.

If the bidders' terminal ECV distributions $\{\bar{e}_i(\omega)\}_{i\in I}$ are independent and identical, then all auction formats that allocate to the bidder with the highest terminal ECV and give zero expected surplus to the bidders with the lowest possible ECV yield the same expected revenue. This result allows us to analyze revenue implications across various dynamic auction formats by focusing on the distributions of terminal ECVs rather than the underlying stochastic processes for values and arrival times.

5.3 Design Implications

Our framework enables a systematic analysis of several key auction design features:

Optimal Reserve Price: We now consider how to establish optimal reserve prices in dynamic auctions, using the ECV insights. In our framework, the bidders are indifferent between winning the object at their submitted bid and losing the auction at this bid. This property, which carries over from the static second-price format, implies that introducing a dynamic reserve price does not alter the bidders' incentives in changing their bids. Rather, the reserve price affects the probability of sale and the expected payment conditional on sale.

[18] In Hopenhayn and Saeedi [2020], this model is estimated with 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 the bidders' welfare and sellers' revenue.

When the seller places zero value on retaining the auctioned item, determining the optimal reserve involves a fundamental trade-off: setting a higher reserve price increases revenue conditional on sale but reduces the probability that the item will be sold.[19] Specifically, if the reserve price exceeds the highest bidder's terminal ECV, no sale occurs; if it falls between the highest and second-highest terminal ECVs, the reserve becomes binding and increases the payment made by the winning bidder. These forces correspond exactly to those present when examining the optimal reserve price for static auctions. Consequently, the solution to this optimization problem follows the same mathematical structure as in the static case but with a critical distinction: the relevant distribution for calculating the optimal reserve is the distribution of terminal ECVs $\{\bar{e}(\omega)\}$ rather than the distribution of final values of the bidders, by setting the reserve price equal to the value with zero virtual value with respect to the ECV distribution.

The distribution of terminal ECVs reflects the dynamic aspects of the bidding environment, including value shading induced by irreversibility, the stochastic processes of information arrival, and the frequency of bidding opportunities. Thus, optimal reserve prices in dynamic auctions differ from their static counterparts even when the bidders' underlying value distributions are identical.

When the seller places a positive value v_0 on retaining the auctioned item and this value remains constant throughout the auction, the optimal reserve price becomes the maximum of v_0 and the value with zero virtual value. This extends the familiar static result to our dynamic setting. However, a more sophisticated scenario arises when the seller's value for the item evolves stochastically during the auction period, perhaps due to the same market forces that affect the bidders' valuations.

If the seller can adjust the reserve price both upward and downward throughout the auction, the optimal policy simply tracks the seller's current value (or the value that yields zero virtual value, whichever is higher). More intriguing is the case when the reserve price can only be decreased over time, which is a common restriction in many auction platforms designed to prevent seller manipulation. Under this constraint, the seller faces a problem analogous to the bidders' but with reversed irreversibility: the seller must set an initially higher reserve to account for the option value of future reserve price reductions.

Specifically, if the seller's value follows a stochastic process with the same structural properties as the bidders' values but the reserve price can only decrease, the seller will set an

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

initial reserve price above both the seller's current value and the standard optimal reserve. This upward adjustment—contrary to the downward adjustment (shading) that characterizes bidder behavior—reflects the opposite direction of the irreversibility constraint. The magnitude of this adjustment can be determined by applying our framework with appropriate modifications to account for the reversed irreversibility direction.

Allowing Bid Retraction A fundamental design feature in dynamic auctions is whether to permit bidders to subsequently retract or decrease their bids. 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 bidders are allowed to retract their bids, they would submit their expected value of the auctioned item at any given opportunity to update their bid. This expected value might fluctuate due to the arrival of new information, but it does not result in the shading discussed earlier.

This modification affects both the distribution of bids and market outcomes. We can characterize these effects using our theoretical results from Section 4. The terminal bid distribution under retraction corresponds to $\{\bar{v}(\omega)\}$, which is the distribution of values conditioned on the last arrival time, while the distribution under irreversibility corresponds to $\{\bar{e}(\omega)\}$, which is the distribution of maximal ECVs. Proposition 4 indicates that $\{\bar{e}(\omega)\}$ is a mean-preserving contraction of $\{\bar{v}(\omega)\}$. This mean-preserving contraction has significant implications for expected revenues and bidder welfare.

The winner's payment in a second-price auction corresponds to the second-highest order statistic of the bid distribution. When the number of bidders is small, the increased dispersion under bid retraction typically reduces the expected second-highest bid, lowering seller revenue. Additionally, bid retraction improves allocative efficiency by allowing bidders to adjust their bids downward upon receiving unfavorable information. In such scenarios, allowing bid retraction results in an increase in bidders' expected revenue while decreasing the expected revenue for the auctioneer.

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 generally worse off on average.

Counterintuitively, despite the apparent freedom that bid retraction offers, bidders as a group may be worse off under this design. This seemingly paradoxical result resembles a prisoner's dilemma: each individual bidder benefits from the ability to retract bids, but collectively

they suffer from the strategic responses that having this ability induces in equilibrium. When many bidders participate, the resulting higher expected prices outweigh the benefits from improved allocation efficiency from each bidder's perspective.

Another way to interpret these results is similar to the interpretation of common-value auctions and the winner's curse. In the presence of common values, bidders will shade their bids; therefore, in equilibrium, instances of regret are relatively infrequent. This is also the case for dynamic auctions. When retracting bids is not possible, bidders are more cautious and shade their earlier bids to prevent future regrets, and therefore, the instances of regret are infrequent, especially as the number of bidders increases. Consequently, the harm of higher prices outweighs the benefits of improved allocation efficiencies.

Increasing Arrival Rates The frequency with which bidders can adjust their bids represents another design parameter in dynamic auctions. Platform designers can influence this frequency through various mechanisms, such as sending reminders about auction deadlines or directly controlling the technological infrastructure that processes bid submissions. We now analyze how variations in arrival rates affect equilibrium bidding behavior and auction outcomes.

Assuming that the process for arrival rate and the process for change in values are independent, by Assumption 6, we can use the result of Proposition 6. This proposition establishes that higher arrival frequency leads to more aggressive value shading. Intuitively, when bidders anticipate more frequent opportunities to adjust their bids before the auction concludes, the option value of waiting increases, causing them to bid more conservatively at earlier decision nodes. Additionally, this proposition states that the distribution of terminal ECVs under higher arrival frequency is a mean-preserving spread of the distribution under lower arrival frequency.

This distributional relationship has significant implications for auction outcomes that parallel our findings regarding bid retraction. The mean-preserving spread in terminal ECVs translates into a mean-preserving spread in the distribution of final bids. When few bidders participate, this typically reduces the expected second-highest bid (i.e., the price paid by the winner in a second-price auction), thereby decreasing seller revenue. Conversely, in the presence of many bidders, the increased dispersion tends to raise the expected second-highest bid, enhancing seller revenue.

From a mechanism design perspective, these results suggest that the optimal arrival frequency depends on the anticipated number of auction participants. When the number of bidders is known with certainty, the designer should select either minimal frequency (effec-

tively a static auction) or maximal frequency, depending on whether the number of bidders is small or large, respectively. However, in environments where the number of active bidders is uncertain, an interior solution may be optimal. In such cases, the designer must balance the revenue-reducing effect in low-participation states against the revenue-enhancing effect in high-participation states.

The parallels between our analyses of bid retraction and arrival frequency are not coincidental. Both design features (i.e., allowing bid decreases and increasing bid adjustment opportunities) fundamentally alter the option value associated with future decisions. As Proposition 4 established for bid retraction and Proposition 6 for arrival frequency, both modifications generate mean-preserving spreads in the distribution of relevant values at decision nodes.

Our analysis provides testable implications for empirical research, with the predicted relationship between arrival frequency and bid shading potentially explaining observed patterns in online auctions, such as increasing bids throughout the auction period and particularly near its end. This demonstrates how our theoretical approach addresses practical auction design questions. Specifically, by decomposing dynamic strategic interactions into equivalent certain values and analyzing their distributional properties, we offer guidance for mechanism designers seeking to optimize dynamic auctions across different market environments.

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 in 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. These advantages are particularly useful in the class of dynamic games considered, where solving for best responses reduces to a static Bayesian game. Similarly, when estimating the parameters of the payoff functions, the dynamic option value problem needs to be solved only once, as opposed to every time parameters are changed as in the standard approach.

This separation result is also particularly useful when considering game design, 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.

A second advantage of our approach is that it 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 it had 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-generating 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-generating expectations would require to be solved jointly across all players.

The analysis of dynamic games is a difficult problem, and 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 provides 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. 3, 6

Kenneth J Arrow and Anthony C Fisher. Environmental preservation, uncertainty, and irreversibility. *The Quarterly Journal of Economics*, 88(2):312–319, 1974. 6

Robert B Ash. *Real Analysis and Probability: Probability and Mathematical Statistics: a Series of Monographs and Textbooks*. Academic press, 1972. 46

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. 34

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. 6

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. 2

David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953. 2

Simon Board and Andrzej Skrzypacz. Revenue management with forward-looking buyers. *Journal of Political Economy*, 124(4):1046–1087, 2016. 34

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. 34

Alessandro Bonatti, Gonzalo Cisternas, and Juuso Toikka. Dynamic oligopoly with incomplete information. *The Review of Economic Studies*, 84(2):503–546, 2017. 22

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. 34

Jeremy Bulow, Jonathan Levin, and Paul Milgrom. Winning play in spectrum auctions. Technical report, National Bureau of Economic Research, 2009. 34

Guillermo A Calvo. Staggered prices in a utility-maximizing framework. *Journal of monetary Economics*, 12(3):383–398, 1983. 3

Dominic Coey, Bradley J Larsen, and Brennan C Platt. Discounts and deadlines in consumer search. *American Economic Review*, 110(12):3748–3785, 2020. 34

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. 6

Avinash K Dixit, Robert K Dixit, Robert S Pindyck, and Robert Pindyck. *Investment under uncertainty*. Princeton university press, 1994. 6

Wedad Elmaghraby and Pınar Keskinocak. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management science*, 49(10):1287–1309, 2003. 6

Drew Fudenberg and David K Levine. Open-loop and closed-loop equilibria in dynamic games with many players. *Journal of Economic Theory*, 44(1):1 – 18, 1988. ISSN 0022-0531. doi: https://doi.org/10.1016/0022-0531(88)90093-2. URL http://www.sciencedirect.com/science/article/pii/0022053188900932. 5

Alkis Georgiadis-Harris. Preparing to act. Working paper, University of Warwick, 2024. 7

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. 4

John Gittins, Kevin Glazebrook, and Richard Weber. Multi-armed bandit allocation indices. *John Wiley & Sons*, 2011. 6

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. 34

Claude Henry. Investment decisions under uncertainty: the" irreversibility effect". *The American Economic Review*, 64(6):1006–1012, 1974. 6

Hugo Hopenhayn and Maryam Saeedi. Bidding dynamics in auctions. 2020. 14, 18, 35

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. 22

Koichiro Ito and Mar Reguant. Sequential markets, market power and arbitrage. *American Economic Review*, 106(7):1921–1957, 2016. 6, 13, 34

Mireia Jofre-Bonet and Martin Pesendorfer. Estimation of a dynamic auction game. *Econometrica*, 71(5):1443–1489, 2003. 34

Yuichiro Kamada and Michihiro Kandori. Cooperation in revision games and some applications. *Global Economic Review*, 49(4):329–348, 2020a. 5

Yuichiro Kamada and Michihiro Kandori. Revision games. *Econometrica*, 88(4):1599–1630, 2020b. 3, 5

Yuichiro Kamada and Takuo Sugaya. Optimal timing of policy announcements in dynamic election campaigns. *The Quarterly Journal of Economics*, 135(3):1725–1797, 2020. 3

Adam Kapor and Sofia Moroni. Sniping in proxy auctions with deadlines. 2016. 5

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. 4

Steven A Lippman and John J McCall. The economics of job search: A survey. *Economic inquiry*, 14(2):155–189, 1976. 6

Qingmin Liu, Konrad Mierendorff, Xianwen Shi, and Weijie Zhong. Auctions with limited commitment. *American Economic Review*, 109(3):876–910, 2019. 36

Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Dynamic rational inattention: Analytical results. *Journal of Economic Theory*, 176:650–692, 2018. 30

Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica*, 62(1): 157–180, 1994. 58

Benny Moldovanu and Aner Sela. The optimal allocation of prizes in contests. *American Economic Review*, 91(3):542–558, 2001. 13

Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981. 32

Phillip E Protter. Stochastic integration and differential equations, 2004. 11

Daniel Quint and Fuhito Kojima. Symmetric equilibrium in pre-auction investment. *Available at SSRN 4969462*, 2024. 32, 58

John G Riley and William F Samuelson. Optimal auctions. *The American Economic Review*, 71(3):381–392, 1981. 5

Maher Said. Sequential auctions with randomly arriving buyers. *Games and Economic Behavior*, 73(1):236–243, 2011. 34

Christopher A Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003. 7, 30

Jakub Steiner, Colin Stewart, and Filip Matějka. Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica*, 85(2):521–553, 2017. 7, 30

Robert Zeithammer. Forward-looking bidding in online auctions. *Journal of Marketing Research*, 43(3):462–476, 2006. 34

Wen Zhao and Yu-Sheng Zheng. Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management science*, 46(3):375–388, 2000. 6

Weijie Zhong. Optimal dynamic information acquisition. *Econometrica*, 90(4):1537–1582, 2022. doi: 10.3982/ECTA17787. 7, 30

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.

$H(\varepsilon,v,t) = \{\omega|\,(v,t) \in D(\omega) \text{ and } \bar{e}(t^+,\omega) \leq \varepsilon\}$, 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 $e(v',t') > \epsilon$ exist for $\forall t' > t$.

$A(\varepsilon,v,t) = \{\omega|\,(v,t) \in D(\omega) \text{ and } \bar{e}(t^+,\omega) \leq \varepsilon, \exists(v',t') \in D(\omega)$,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

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

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

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

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

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.

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

Ψi : 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 $(\Omega, \mathcal{F}, P)$, and $X : (\Omega, \mathcal{F}) \to (\Omega', \mathcal{F}')$ a random object. If $E(Y)$ exists, there is a function $g : (\Omega', \mathcal{F}') \to \left(\bar{R}, \mathscr{B}\right)$ such that for each $A \epsilon \mathcal{F}'$,*

$$\int_{\{X \epsilon A\}} Y dP = \int_A g(x) \, dP_x(x), \tag{11}$$

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

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

Lemma 1. *At any arrival (v, t) and any Borel Set B, we have:*

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

Where $\bar{e}(t^+, \omega) = \max\{e(v', t') \,|\, (v', t') \in D(\omega) \text{ and } t' > t\}$ and if $\nexists((v', t') \in D(\omega) \text{ and } t' > t$, then $\bar{e}(t^+, \omega)$ is an arbitrary low value and not in B.

Proof. Consider the following two random variables: $\bar{e}(t^+, \omega) : [0, T] \times \Omega \to \mathbb{R}$ as defined above, and $\sigma(t, \omega) : [0, T] \times \Omega \to [0, T]$ to be the time at which it was reached, so on ω, there is an arrival at $\sigma(t, \omega)$ and the ECV at that point in time is $\bar{e}(t^+, \omega)$. In accordance with Theorem B we define the random object $X : [0, T] \times \Omega \to [0, T] \times \mathbb{R}$; $X(t, \omega) = (\sigma(t, \omega), \bar{e}(t^+, \omega))$. Now define the extended random variable $Y(t, \omega) = v_T(\omega) - \bar{e}(t^+, \omega)$. First note that for all $(\tau, e) \in X$,

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

the last equality is coming from the definition of self-generating expectation. For any measurable subset $A \subset \{t < \tau \leq T\}$ and Borel subset B of $\mathbb{R}$, let $P_x(A \times B) = \Pi(\sigma(t, \omega) \in A, \bar{e}(t^+, \omega) \in B | v, t)$. Substituting $x = (\tau, \varepsilon)$ and using (11),

$$\int_{\{\sigma(t,\omega) > t, \bar{e}(t^+,\omega) \epsilon B\}} \left(v_T(\omega) - \bar{e}(t^+, \omega)\right) dP(\omega) = \int_{\{(\tau,e) \in X\}} E(Y | \tau, e) \, dP_x(\tau, e) = 0.$$

□

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-generating expectation. To show that it is strictly decreasing in ε, consider $\delta > 0$. Then

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

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

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

$$H(\varepsilon,v,t) = \{\omega|\,(v,t) \in D(\omega) \text{ and } \forall t' > t, (v',t') \in D(\omega),\ e(v',t') \leq \varepsilon\ \} \tag{12}$$

It is convenient to partition the set $H(\varepsilon,v,t)$ into two subsets as follows:

1. The set that 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. The set that contains all ω such that for all next arrivals $(v',t') \in D(\omega)$, $e(v',t') \leq \varepsilon$. Call this set $A(\varepsilon,v,t)$.

We will prove that there is a unique threshold function satisfying property (12) and that it solves $W(e(v,t),v,t) = 0$, where $W(\varepsilon,v,t)$ is the unique solution to the following Bellman equation:

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

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 (12) 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 12.

Step 1. Necessity Take the candidate value function:

$$W(\varepsilon, v, t) = \int_{H(\varepsilon,v,t)} (v_T(\omega) - \varepsilon)\, d\Pi(\omega|v,t). \tag{14}$$

We show that $W(\varepsilon, v, t)$ is a solution to (13). Substituting into (13) 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 (12) 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 (6) 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 (13) 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 (13) 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 (13):

$$\begin{aligned}
&\int_t^T \min\left(W(\varepsilon, v', \tau') + a, 0\right) dP(v', \tau'|v, t) \quad + \int_{N(v,t)} (v_T(\omega) - \varepsilon)\, d\Pi(\omega|v, t) \\
&\leq \int_t^T \min\left(W(\varepsilon, v', \tau'), 0\right) dP(v', \tau'|v, t) \quad + a\left(1 - \Pi\left(N(v, t)|v, t\right)\right) + \int_{N(v,t)} (v_T(\omega) - \varepsilon)\, d\Pi(\omega|v, t) \\
&= W(\varepsilon, v, t) + a\left(1 - \Pi\left(N(v, t)|v, t\right)\right).
\end{aligned} \tag{15}$$

By Assumption 3 $\Pi\left(N(v, t)|v, t\right) > \delta$ 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 form the definition of $e(v, t)$ given in equation (6) 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 (13) maps weakly decreasing functions into strictly decreasing ones. So, assume that the W function on the right hand side of the Bellman equation (13) is weakly decreasing. Letting $\varepsilon' > \varepsilon$,

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

which is strictly less than $TW(\varepsilon, v, t)$ since by Assumption 3 $\Pi\left(N(v, t)|v, t\right) > 0$. 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 $\Pi\left(N(v, t)|v, t\right) > \delta$, so for large enough ε this term will dominate. The same argument can be used by making ε small enough (negative if needed) to make the second term become positive enough to dominate.

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 $a_1 = \tilde{S}(e_1)$. First consider the case where $a_2 > a_1$ and let $e_2 = \sup\left\{v|\tilde{S}(v) \leq a_2\right\}$.

Let $V(v,t,a)$ denote the expected utility of choosing a at this state and following the (candidate) optimal policy in the future. We need to prove that $V(v,t,a_1) \geq V(v,t,a_2)$. Let $\bar{s}(\omega,v,t) = \max\{S(v',t') \,|\, (v',t')\,\epsilon D(\omega) \text{ and } (v',t') \neq (v,t)\}$. 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 4.1 and $H(e,v,t)^c$ its complement in the set of paths following $(v,t) : \{\omega|\,(v,t) \in \mathrm{D}(\omega)\}$. To simplify notation, from now on we drop the argument (v,t) from these functions. We can decompose the histories following $a_i, i\epsilon\{1,2\}$ 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 $H(e_2) = H(e_1) \cup \{\omega|e_1 < \bar{e}(t^+,\omega) \leq e_2\}$ so

$$\begin{aligned} V(v,t,a_2) &= \int_{H(e_1)} U(v(T),a_2)\,dP(\omega) + \int_{e_1<\bar{e}(t^+,\omega)\leq e_2} U(v(T),a_2)\,dP(\omega) \\ &+ \int_{H(e_2)^c} U(v(T),\bar{s}(\omega))\,dP(\omega). \end{aligned} \tag{16}$$

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

$$\begin{aligned} V(v,t,a_2) - V(v,t,a_1) &\leq \int_{H(e_1)} [U(v(T),a_2) - U(v(T),a_1)]\,dP(\omega) \\ &= U\left(E_{H(e_1)}v(T),a_2\right) - U\left(E_{H(e_1)}v(T),a_1\right) \leq 0, \end{aligned} \tag{17}$$

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

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

$$\begin{aligned} V(v,t,a_2) &= \int_{H(e_2)} U(v(T),a_2)\,dP(\omega) + \int_{H(e_2)^c} U(v(T),\bar{s}(\omega))\,dP(\omega) \qquad (18)\\ &= \int_{H(e_1)} U(v(T),a_2)\,dP(\omega) + \int_{H(e_1)^c} U(v(T),\bar{s}(\omega))\,dP(\omega)\\ &+ \int_{e_2<\bar{e}(t^+,\omega)\leq e_1} (U(v(T),\bar{s}(\omega)) - U(v(T),a_2))\,dP(\omega). \end{aligned}$$

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

$$\begin{aligned} V(v,t,a_1) - V(v,t,a_2) &= \int_{H(e_1)} U(v(T),a_1)\,dP(\omega) - \int_{H(e_1)} U(v(T),a_2)\,dP(\omega)\\ &- \int_{e_2<\bar{e}(t^+,\omega)\leq e_1} (U(v(T),\bar{s}(\omega)) - U(v(T),a_2))\,dP(\omega) \qquad (19)\\ &= \int_{H(e_1)} (U(e_1,a_1) - U(e_1,a_2))\,P(H(e_1))\\ &- \int_{e_2<\bar{e}(t^+,\omega)\leq e_1} (U(v(T),\bar{s}(\omega)) - U(v(T),a_2))\,dP(\omega)\\ &= \int_{H(e_1)} (U(e_1,a_1) - U(e_1,a_2))\,P(H(e_1))\\ &- \int_{e_2<\bar{e}(t^+,\omega)\leq e_1} \left(U\left(\bar{e}(t^+,\omega),\tilde{S}\left(\bar{e}(t^+,\omega)\right)\right) - U\left(\bar{e}(t^+,\omega),a_2\right)\right)dP(\omega), \end{aligned}$$

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

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

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

$$U\left(\bar{e}(t^+,\omega),\tilde{S}\left(\bar{e}(t^+,\omega)\right)\right) - U\left(\bar{e}(t^+,\omega),a_2\right) \leq U\left(e_1,\tilde{S}\left(\bar{e}(t^+,\omega)\right)\right) - U(e_1,a_2). \qquad (20)$$

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

$$U\left(e_1, \tilde{S}\left(\bar{e}(t^+, \omega)\right)\right) - U(e_1, a_2) \quad \leq \quad U(e_1, a_1) - U(e_1, a_2). \tag{21}$$

Combining (19), (20) and (21), it follows that

$$\begin{aligned} V(v, t, a_1) - V(t, v, a_2) \quad &\geq \quad (U(e_1, a_1) - U(e_1, a_2)) P(H(e_1)) \\ &- \quad \int_{e_2 < \bar{e}(t^+, \omega) \leq e_1} U(e_1, a_1) - U(e_1, a_2)\, dP(\omega), \end{aligned} \tag{22}$$

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 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

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

for all $a \epsilon A_i$. So, $\tilde{S}_i(\tilde{v})$ is an optimal solution for any 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 2 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

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

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

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

and by the definition of self-generating expectation

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

Adding the two implies that

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

Proof of Proposition 4 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 3, 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 $A = \{\omega | (v,t) \in D(\omega), \bar{e}(\omega) = e(v,t)\}$ and $A^c = \{\omega | (v,t) \in D(\omega), \bar{e}(\omega) > e(v,t)\}$. This is a partition of the paths following (v,t) where A corresponds to all paths such that the ECV at time t,$e(v,t)$, remains the maximum, while the set A^c is the complement. It follows that

$$\begin{aligned} v &= \Pi(A)\,\mathbb{E}(v_T|A) + \Pi(A^c)\,\mathbb{E}(v_T|A^c). \\ &= \Pi(A)\, e(v,t) + \Pi(A^c)\,\mathbb{E}(v_T|A^c) \end{aligned} \tag{23}$$

where the equality follows from the definition of ECVs. Starting from any $(v',t') \in A^c$, Proposition 3 implies that the conditional distribution of terminal ECVs $\bar{e}(\omega|(v',t'))$ is unbiased, integrating over all such decision nodes, it follows that $E(\bar{e}(\omega)|A^c) = E(v_T|A^c)$ and by a similar argument, the same holds for $E(\bar{v}(\omega)|A^c)$. It then follows from equation (23) that $e(v,t) = E(\bar{v}(\omega)|A)$ 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 5

Proof. Existence and uniqueness follow from the same contraction mapping arguments as Proposition 1, one just need to set $d\tilde{P}(v',\tau'|v,t) = \exp(-\rho(\tau'-t))\, dP(v',\tau'|v,t)$.

For monotonicity in ρ, consider $\rho' > \rho$. The functional equation for $W_{\rho'}$ applies a smaller weight $\exp(-\rho'(\tau'-t)) < \exp(-\rho(\tau'-t))$ to future adverse outcomes. Since $\min(W_\rho(\varepsilon, v', \tau'), 0) \leq 0$, reducing these weights increases the overall value function. By the strict monotonicity of W_ρ in its first argument (established in the baseline model), this implies $e_{\rho'}(v,t) > e_\rho(v,t)$.

The limiting cases follow directly: as $\rho \to 0$, reversibility opportunities become arbitrarily rare, recovering the baseline model. As $\rho \to \infty$, reversibility opportunities arrive continuously, eliminating all irreversibility constraints and making the ECV equal to the conditional expectation.

The mean-preserving contraction property follows from the same logic as Proposition 4. □

Proof of Proposition 6 Consider the recursive formulation

$$\begin{aligned}\tilde{W}(\varepsilon, v, t) &= \int_t^T \int \min\left(\tilde{W}(\varepsilon, v', \tau'), 0\right) dP(v'|v, \tau')\, d\tilde{F}(\tau'|t) \\ &\quad + \left(1 - \tilde{F}(T|t)\right)(v - \varepsilon)\end{aligned} \tag{24}$$

First note that this functional equation is monotone: increasing pointwise the integrated function $\tilde{W}$, increases the integral. Also note that from the proof of Proposition 7, the second assumption implies that $\tilde{W}(\varepsilon, v, t)$ is nondecreasing in t. Additionally, as shown before it is a contraction. Let's start with $W_0 = W$, where W is the solution corresponding to stochastic arrival rate F which first-order stochastically dominates $\tilde{F}$, or for any $t' > t$, $F(t'|t) \leq \tilde{F}(t'|t)$. Using the functional equation, we will iterate over W.

$$W_{i+1}(\varepsilon, v, t) = \int_t^T \int \min\left(W_i(\varepsilon, v', \tau'), 0\right) dP(v'|v, \tau')\, d\tilde{F}(\tau'|t) + \left(1 - \tilde{F}(T|t)\right)(v - \varepsilon)$$

Given that $W \leq max\{0, v - \varepsilon\}$, in each iteration, $W_{i+1}(\varepsilon, v, t) \leq W_i(\varepsilon, v, t)$, for positive $(v - \varepsilon)$. This condition is preserved under the iteration. Therefore, given that the functional equation is a contraction, the relationship holds for W and $\tilde{W}$. As a result at any (v, t) and $e(v, t)$ for F distribution (given that $e(v, t) \leq v$ by Proposition 2), we will have: $\tilde{W}(e(v, t), v, t) \leq W(e(v, t), v, t) = 0$ which results in $\tilde{e}(v, t) \leq e(v, t)$.

Given Proposition 4, the mean of the two distribution are the same, to show that $\tilde{e}$ is more dispersed, we will apply Theorem B once again. Consider a decision node at (v, t), we had established that $\tilde{e}(v, t) \leq e(v, t)$. Consider all histories under F that $e(v, t)$ remains the final ECV on those path. First note that in the complement set, in which this $e(v, t)$ will not stay the final ECV; $\tilde{e}(v, t)$ will not stay the final ECV as well. The reason for it is as follows. First for ECV to increase we need a new arrival, if there is an arrival under F at any future (v', t') there is a higher probability of arrival under $\tilde{F}$. For ECV to increase under F, v' should be higher than the value that makes ECV at time t' to be equal to $e(v, t)$; namely v_1 in $e(v, t) = e(v_1, t')$ as depicted in Figure 3. Under $\tilde{F}$, in order for ECV to increase v' should be higher that the value that makes ECV at t' to be equal to $\tilde{e}(v, t)$, namely v_2 in $\tilde{e}(v, t) = \tilde{e}(v_2, t')$ as depicted in Figure 3.

Additionally, given that $\tilde{F}$ puts more weights on times closer to t, the indifference curve becomes steeper at v under $\tilde{F}$, which results in $v_2 \leq v_1$. Therefore if v' is high enough to increase ECV under F, it is high enough under $\tilde{F}$. Additionally if $v_2 \leq v' \leq v_1$, we will change ECV under $\tilde{F}$ but not under F.

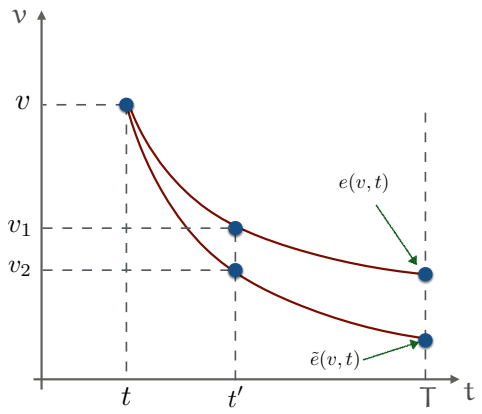

Figure 3: Indifference curves for F and $\tilde{F}$.

Note: The x-axis represents the time horizon, and the y-axis represents agent's values.

Therefore, the set of histories in which $e(v,t)$ remains the final values under F, includes potentially multiple values of ECV under $\tilde{F}$. Using Theorem B, expectation of $e(v,t)$ and $\tilde{e}(v,t)$ on this set is equal to expected final values on the set, so they are equal, with more dispersion under $\tilde{F}$.

Proof of Proposition 7 For $e(v,t)$ to be increasing in t, we should show that W is increasing in t, note that from proof of Proposition 1, we know that $e(v,t)$ is the zero of the W function and Wfunction is decreasing in v. We first show the claim for the first condition. Using the first condition, the Bellman equation can be written:

$$\begin{aligned} W(\varepsilon,v,t) &= \int_t^T \left[\int \min\left(W(\varepsilon,v',\tau'),0\right) dP(v'|v)\right] dF(\tau'|t) \qquad (25)\\ &\quad + \int_{N(v,t)} (v_T(\omega)-\varepsilon)\, d\Pi(\omega|v,t) \\ &= \int_t^T \left[\int \min\left(W(\varepsilon,v',\tau'),0\right) dP(v'|v)\right] dF(\tau'|t) \qquad (26)\\ &\quad + (1-F(T|t))\left[\mathbb{E}(v_T|v)-\varepsilon\right]. \qquad (27) \end{aligned}$$

We show that if we start with a W function that is increasing in t, applying the Bellman equation preserved this condition. 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 2, $W(\varepsilon,v,t) \leq E(v_T|v,t)-\varepsilon$, implying that the second term on the right hand side of (25) 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 second condition in the proposition, equation (8) can be rewritten as

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

Take $\tau_2 > \tau_1$.

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

where the first inequality follows from the induction hypothesis. As in the proof of Proposition 2, $W(\varepsilon, v', \tau_1 + x) \leq \mathbb{E}(v_T|v', \tau_1 + x) - \varepsilon = v' - \varepsilon$ 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 8 The proof is by applying Milgrom-Shannon single crossing theorem. We have:

$$\frac{\partial \left\{E\left[G\left(\bar{e}_\lambda\right)\right] - \kappa C(\lambda)\right\}}{\partial \lambda \partial \kappa} = -C'(\lambda) \leq 0$$

Since $C'(\lambda) > 0$, the objective exhibits decreasing differences in (λ, κ). Therefore, by Theorem 4 in Milgrom and Shannon 1994, the argmax correspondence is monotone nonincreasing in κ. With strict convexity, the argmax is single-valued and the inequality is strict.

Proof of Proposition 9 Following Quint and Kojima [2024], we can define a potential game to the auction as:

$$Q(\boldsymbol{\lambda}) = -\int \prod_{i \in I} H_{i,\lambda_i}(t)\, dt - \sum_{i \in I} \kappa_i C(\lambda_i)$$

The first-order conditions of Q coincide with best-response conditions in symmetric environments, so any maximizer of Q is a symmetric equilibrium of the λ-stage.

In a symmetric equilibrium the first-order condition becomes:

$$\begin{aligned}\frac{\partial^2 Q(\boldsymbol{\lambda})}{\partial \lambda_i \partial \kappa} &= -\int \frac{\partial \prod_{i \in I} H_{i,\lambda_i}(t)}{\partial \lambda_i \partial \kappa} dt - C'(\lambda_i) \\ &= -C'(\lambda_i) \leq 0\end{aligned}$$

Which results in an identical expression as the previous proposition which completes the proof.

C Additional Examples, 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 Additional Examples

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

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

$$V(\varepsilon,v,t) = \int_t^T \max\left(V\left(\varepsilon,v',\tau'\right),v'-\varepsilon\right) dP\left(v',\tau'|v,t\right) \tag{28}$$

Guess that the solution to this functional equation is

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

Using equation (8)

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

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

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

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

□

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.

C.2 Random Termination

In some applications, the decision process may end unexpectedly, forcing players to act before obtaining complete information, this can be the case in patent races or spectrum auction environments. Our baseline model can be extended to accommodate such scenarios with random termination.

Model Setup

Consider a discrete-time setting with $t = 0, 1, \ldots, T$ where $T \in \mathbb{N} \cup \{\infty\}$. At each period t, the decision-maker observes a signal v_t of their final value and can choose to increase their action $a_t \succeq a_{t-1}$ (maintaining the irreversibility constraint). The key modification is that at each period t, the game terminates with probability $(1 - p_t)$, where p_t represents the continuation probability.

The value process follows a Markov structure with conditional distribution $F(v_{t+1}|v_t)$, and utility functions satisfy the same assumptions as in our baseline model.

Two-Period Illustration

Consider the case $T = 1$ to illustrate the main ideas. At $t = 0$, the decision-maker observes initial value v_0 and selects action a_0. With probability p_0, the game continues to period 1, where the decision-maker observes updated value v_1 and may increase their action. With probability $(1 - p_0)$, the game terminates immediately.

The equivalent certain value (ECV) logic applies directly: if $e(v_0, 0) > v_1$, the initial action a_0 is maintained; otherwise, the action is increased. The estimated value for the ECV is going to be the conditional expected value of the histories in which the game has ended in the first period and also the set the values in the second period is lower than the ECV. This mirrors our simplified example, with the key difference that p_0 represents continuation probability rather than arrival probability.

General Multi-Period Extension

In the general setting, if the game reaches period τ and the initial action has not been revised, it will be updated if and only if $e(v_0, 0) < e(v_\tau, \tau)$.

Let $\tilde{v}_0 = e(v_0, 0)$ and define the threshold $\tilde{v}_\tau$ implicitly by the condition $e(\tilde{v}_\tau, \tau) = \tilde{v}_0$. The set of realizations under which the initial action a_0 remains unchanged can be characterized by the self-generating expectation:

$$\tilde{v}_0 = e(v_0, 0) = \frac{\sum_{k=0}^{\infty} \text{Prob}(\text{terminate at } k) \cdot E[v_k|\text{no update through } k]}{\sum_{k=0}^{\infty} \text{Prob}(\text{terminate at } k \text{ and no update through } k)} \tag{29}$$

More explicitly:

$$\tilde{v}_0 = \frac{(1-p_0)v_0 + \sum_{k=1}^{\infty} \left(\prod_{j=0}^{k-1} p_j\right)(1-p_k) \int^{\tilde{v}_k} \cdots \int^{\tilde{v}_1} v_k \, dF_k(v_k|v_{k-1}) ... dF_1(v_1|v_0)}{(1-p_0) + \sum_{k=1}^{\infty} \left(\prod_{j=0}^{k-1} p_j\right)(1-p_k) \int^{\tilde{v}_k} \cdots \int^{\tilde{v}_1} dF_k(v_k|v_{k-1}) ... dF_1(v_1|v_0)} \tag{30}$$

where $F_k(\cdot|v_0)$ denotes the k-step ahead conditional distribution.

Functional Equation and Main Results

The auxiliary functional equation becomes:

$$W(\varepsilon, v, t) = p_t \int \min(W(\varepsilon, v', t+1), 0)\, dF(v'|v) + (1 - p_t)(v - \varepsilon) \tag{31}$$

with boundary condition $W(\varepsilon, v, T) = v - \varepsilon$.

Assumption 8. *We maintain these assumptions:*

1. The termination probabilities satisfy $0 < p_t < 1$ for all $t < T$, and
2. the value process $\{v_t\}$ follows a Markov chain with transition probabilities $F(v_{t+1}|v_t)$ that are continuous in v_t.
3. There exists $\delta > 0$ such that $(1 - p_t) > \delta$ for all t, ensuring a minimum probability of termination at each period.

Theorem 3. *Under Assumptions 8, and Assumption 2 from the main text, there exists a unique function $e : \mathbb{R} \times \{0, 1, \ldots, T\} \to \mathbb{R}$ satisfying the self-generating expectation property. Moreover, the strategy $S(v, t) = \tilde{S}(e(v, t))$ is optimal for the random termination decision problem, where $\tilde{S}(v)$ is the optimal static strategy.*

Proof. [Proof Sketch] The proof follows the same contraction mapping argument as Proposition 1. Assumption RT2 ensures that the mapping $TW(\varepsilon, v, t) = p_t \int \min(W(\varepsilon, v', t + 1), 0) dF(v'|v) + (1 - p_t)(v - \varepsilon)$ is a contraction with modulus $(1 - \delta) < 1$. The solution concept follows the same steps as that of Theorem 1. □

Theorem 4. *Consider a dynamic game with N players where each player i faces the random termination process described above with independent termination probabilities $\{p_{it}\}$. Under Assumptions RT1, RT2, and Assumption 4 from the main text, for any profile of equilibrium strategies $\{\tilde{S}_i\}_{i \in I}$ in the associated static Bayesian game with ECV distributions, the strategy profile $\{S_i(v, t) = \tilde{S}_i(e_i(v, t))\}_{i \in I}$ constitutes an equilibrium for the dynamic game with random termination.*

Proof. [Proof Sketch] The proof parallels Theorem 2. For each player i, fixing the strategies of other players, the best response problem becomes a single-agent random termination problem. By Theorem 3, the optimal strategy is $S_i(v, t) = \tilde{S}_i(e_i(v, t))$. Since this holds for all players simultaneously, the strategy profile constitutes an equilibrium. □

Connection to Baseline Model

The key insight is that random termination can be interpreted as a limiting case of our baseline model where "arrival opportunities" represent chances to make final decisions before termination. When $p_t \to 0$ (high termination probability), we approach a static decision problem. When $p_t \to 1$ (low termination probability), the problem resembles our baseline model with very frequent arrival opportunities.

This extension demonstrates the robustness of the ECV methodology to different sources of uncertainty in the timing of decision-making opportunities.

C.3 Random Termination with Random Arrival

We have assumed that the decision problem lasts for a fixed time $[0, T]$. Our formulation allows for random termination without modification. Consider equation (8), repeated below, which is the key equation used to find the self-generating expectation:

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

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 7 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 7), 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 7. □

The assumptions of this corollary have an intuitive interpretation. The second one is the analogue of the assumption made in Proposition 7. 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-generating expectations are decreasing, as depicted in Figure 1. In the special—time stationary—case where both conditional distributions G and H are independent of t, $W(\varepsilon, v, t)$ and $\varepsilon(v,t)$ will also be independent of t, so the level curves will be flat. Shading will still occur but will not change over time.

C.4 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 11. *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 (8) is a contraction mapping. Assume that W has the property stated above. It follows that

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

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

Letting $\delta = -\varepsilon$, the above proposition implies that $W(\varepsilon, v, t) = W(0, v-\varepsilon, t)$. Letting $s = v - \varepsilon$, functional equation (8) can be written as

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

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 12. *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

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

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-generating expectation. □

Shading over Time with Independent Increments While the propositions derived in Section C.4 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 11.

Proposition 13. *Under the assumptions of Proposition 11, 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,

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

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

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

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 $\square$

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

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 $1/\left(1/\sigma_0 + t/\sigma\right)$. Together with an independent arrival process for decision nodes, this formula can be used recursively to define the Markov process $P\left(v', t' | v, t\right)$ that satisfies the assumption of independent increments in Proposition 11. As the variance of the increments decreases over time, Proposition 13 implies that shading decreases over time.

C.5 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 $\pi\left(v, k, \tilde{k}\right)$, 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: $u\left(v, k, \tilde{k}\right) = \frac{1}{1-\delta}\left(\pi\left(v, k, \tilde{k}\right) - ck\right)$.

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

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

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 $t = 1, ...$ 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

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

   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 $F(v)$.[20] It easily follows that

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

[20] Note that this stochastic process does not have independent increments.

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 where lower.

C.6 Intuitive Analysis

This section provides an intuitive analysis of our main results in Section 4 especially those of ECV and how $e(v,t)$ is determined.

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 is obtained. We can partition the set of value and time pairs (v,t) into indifference classes that can be identified by a final value $\tilde{v}$. This final value is going to be found using a self-generating function, $e(v,t)$, which is defined in Section 4. These have the property that optimal actions are identical for all pairs on 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 4. In the paths shown, the first decision node is (v_0,t_0) where the agent chooses an action a_0. After an arrival at (v_0,t_0), there are many possible realization of path of values and arrival in some, the action at this node will remain the final action, in some it will be increased (given the irreversibility it can only increase). This action is 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) which lies on a lower indifference curve. The last two panels of Figure 4 represent cases where the agent has an opportunity to update its action at a time period in which the agent's value lies on a higher indifference curve. In these two cases, the agent chooses a higher action, so the choice made at decision node (v_0,t_0) no longer binds. More generally, for any path with decision nodes $\{(v_0,t_0),(v_1,t_1),...,(v_n,t_n)\}$ the final action is based on $\max\{e(v_0,t_0),e(v_1,t_1),...,e(v_n,t_n)\}$, i.e., the value associated to the highest indifference curve reached during decision times $t_0,...,t_n$.

To find the mapping between (v_0,t_0) and $e(v_0,t_0)$ note that the action in this node is only important if this action stays the final action, which is the set in which there is no next opportunity to update action or the opportunity arises only when the $e(v',t') < e(v_0,t_0)$.

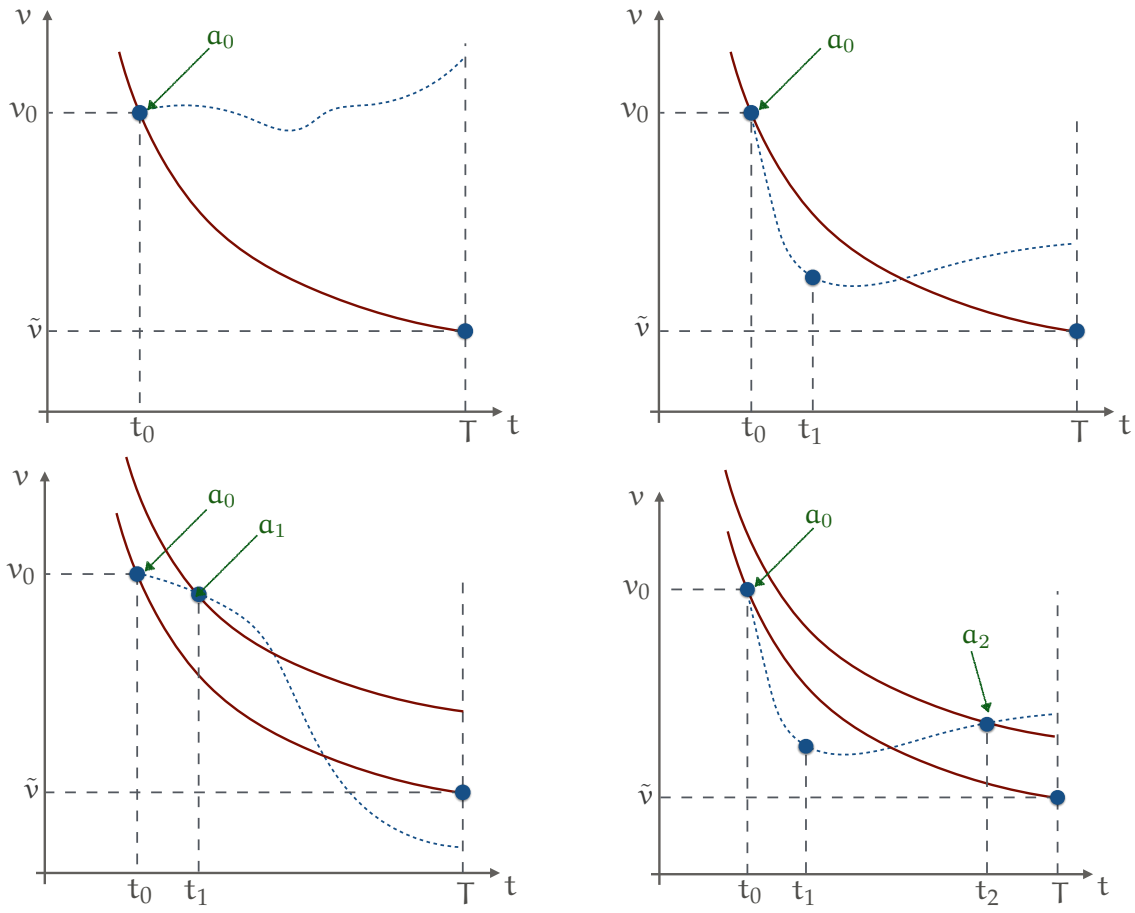

Figure 4: Indifference curves and increasing actions

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.

Therefore, $e(v_0, t_0)$ is the expected final value v_T on this set. This is the case in the top two panels in Figure 4 but not in the last two panels. This will give us the self-generating property of $e(v, t)$. We have also shown, in Theorem 1, the ECV property of this function, namely that the optimal dynamic action at (v_0, t_0) is the same as the static optimal action at $e(v_0, t_0)$, so it is the equivalent certain value for the node.