UC Berkeley

UC Berkeley Previously Published Works

Title

Optimal stopping with private information

Permalink

<https://escholarship.org/uc/item/6w84z1wf>

Authors

Kruse, T
Strack, P

Publication Date

2015-09-01

DOI

10.1016/j.jet.2015.03.001

Peer reviewed

Optimal Stopping with Private Information

Thomas Kruse and Philipp Strack

March 3, 2015

Abstract

Many economic situations are modeled as stopping problems. Examples include job search, timing of market entry decisions, irreversible investment or the pricing of American options. This paper analyzes optimal stopping as a mechanism design problem with transfers. We show that under a dynamic single crossing condition a stopping rule can be implemented by a transfer that only depends on the realized stopping decision if and only if it is a cut-off rule. We characterize the transfer implementing a given stopping rule using a novel technique based on constrained stochastic processes.

As an application we prove that in any Markovian optimal stopping problem there exists a welfare maximizing mechanism that does not require any communication. We discuss revenue maximization for separable processes.

Keywords: Dynamic Mechanism Design, Optimal Stopping, Dynamic Implementability, Posted-Price Mechanism

JEL codes: D82, C61

Thomas Kruse, Universit d'Evry Val D'Essonne, 23 Boulevard de France, 91025 Evry Cedex, France.
Email: thomas.kruse@univ-evry.fr;

Philipp Strack, University of California, Berkeley, Evans Hall, Berkeley, CA 94709, USA. Web:
<http://www.philippstrack.com>, Email: philipp.strack@gmail.com.

We would like to thank Stefan Ankirchner, Dirk Bergemann, Paul Heidhues, Daniel Krhmer, Benny Moldovanu, Jen Pl, Alessandro Pavan, Sven Rady, Juuso Toikka, anonymous referees and seminar audiences in Austin, Bonn, HU Berlin, Berkeley, Bielefeld, Boston, Caltech, Columbia, Duke, Harvard, Maastricht, Microsoft Research NE, MIT theory camp, Northwestern, NYU, Paris Evry, Toronto, University College London, UPenn, Warwick, Yale for many helpful comments and suggestions. All remaining errors are our own. Financial support by the Bonn Graduate School of Economics, the Bonn International Graduate School of Mathematics, the German Research Foundation (DFG), the Hausdorff Center for Mathematics, the SFB TR 15 and the French Banking Federation through the Chaire "Markets in Transition" is gratefully acknowledged.

1 Introduction

In an optimal stopping problem information arrives over time and a decision maker decides at every point in time whether to stop or continue. The stopping decision is irreversible. Thus, stopping today implies losing the option to stop later with a potentially larger return.

This paper analyzes the question how the behavior of an agent in an optimal stopping problem can be influenced through transfers. Our model can be understood as a dynamic principal-agent model. The agent privately observes a discrete time Markov process and chooses a stopping rule. The principal observes the stopping decision of the agent, but not the realization of the process. In order to influence the agent's stopping decision the principal commits to a transfer.

For example, the agent could be an unemployed worker who receives job offers until she stops the process and accepts an offer. The principal could be the unemployment agency who wants the agent to accept certain offers, but who does not observe the offers that the agent receives. Alternatively, the agent could be a firm that developed a new technology and has to decide when to introduce it to the market place. The firm observes private signals regarding the demand, and this knowledge changes over time. The principal is a social planner who also takes the consumer surplus into account and has, hence, different preferences over stopping decisions than the firm.

1.1 Contribution of the Paper

This paper shows that under a dynamic single-crossing condition all cut-off rules can be implemented without communication. A cut-off rule is a strategy that stops the first time the value of the process exceeds a time-dependent threshold. Implementation without communication means that the transfer depends only on the time the agent stops and does not require her to report any private information.¹ This feature resembles a posted-price mechanism and makes our result especially appealing from an applied perspective. We provide a closed form solution for the transfer as the expectation over a new type of stochastic process, we call *constrained process*.

The aim of our article is different from most articles in the literature as it does not focus on implementing a specific policy (i.e. the welfare or virtual valuation maximizing one),

¹In this sense we understand a contract which offers the agent a menu of time-dependent transfers to choose from as a mechanism with communication since the transfer conditions on the choice made (message sent) by the agent.

but characterize the set of policies that can be implemented in the class of mechanisms without communication. However, this characterization can then be used to implement specific policies. To illustrate this point we show that the welfare maximizing policy in a principal agent problem with interdependent values can always be implemented. Due to the interdependent value structure this case is to the best of our knowledge not covered by any of the existing results in the literature. Furthermore, we show that revenue can be maximized using mechanisms with communication only at time zero in the case of additive random walks.

For example, those results imply in the context of unemployment benefits that for one-dimensional Markovian search processes it is without loss of generality from a welfare perspective to restrict attention to mechanisms in which the worker does not report rejected offers to the mechanism and the benefits paid to her depend only on the time she has been unemployed.

To get an intuition, consider a point in time and a value of the process where it is optimal for the agent to continue. The dynamic single crossing condition implies that continuing is even better for lower values of the process. Thus, to implement a cut-off rule it suffices to provide incentives to the (marginal) cut-off type. Taking future transfers as given, the transfer providing incentives to the marginal type today could be calculated recursively. But as future transfers are endogenous, this recursive approach is not analytically tractable as it requires the calculation of the value function at every point in time. We circumvent these problems attached to the recursive approach by directly constructing transfers using constrained stochastic processes. We define a constrained process as a Markov process that has the same dynamics as the original process but is required to stay below the cut-off whenever the original process exceeds the cut-off.

We show that every cut-off rule can be implemented through a transfer that equals the agent's expected marginal incentive to wait evaluated at the process constrained by the cut-off. The dynamic transfer coincides with the transfers used in static mechanism design if one replaces the constrained process by the true process observed by the agent.

In stopping problems it is in general not optimal to use a myopic stopping rule which only compares the payoff of stopping immediately with the expected payoff of stopping one period later. The gain in expected payoffs from using an optimal policy compared to such a myopic policy is what the literature on optimal stopping calls the *option value*. The option value is the additional payoff the agent gains from the fact that she can update her decision based on new information she learns when observing the process. We modify the

transfers from static mechanism design by replacing the true process with its constrained version and prove that this adjustment compensates the agent exactly for her option value.

As far as we know, this article is the first to use constrained processes in the context of mechanism design or optimal stopping. We also believe that this paper contributes to the literature on optimal stopping as we give a new, completely probabilistic characterization of the option value. In particular, this result has applications in mathematical finance since it leads to a new characterization of the optimal exercise boundary for American options (see Section 4.1).

1.2 Related Literature

Optimal stopping theory has been influential in many areas of economics. In labor economics, the seminal contributions of Stigler [1962] and McCall [1970] established the perspective on job search as an optimal stopping problem. In finance, the pricing of American options and other financial contracts is a classical optimal stopping problem, cf. McKean [1965]. Following McDonald and Siegel [1986] the optimal timing of irreversible investments and market entry decisions are modeled as stopping problems in the literature on industrial organizations, cf. Dixit and Pindyck [2008].

To the best of our knowledge we are the first to analyze general optimal stopping as a principal agent problem with transfers. Several strands of literature consider special cases of optimal stopping problems and take into account incentives. The literature on optimal unemployment insurance analyzes how the search (stopping) behavior of an unemployed can be influenced through transfers (see Shavell and Weiss [1979], Hopenhayn and Nicolini [1997] and Farhi and Werning [2013] for a model with random productivity). A classical optimal stopping problem in the literature on revenue management is the sale of multiple objects to stochastically arriving buyers (cf. Talluri and Van Ryzin, 2005). When the buyers valuations are private information, this leads naturally to a dynamic mechanism design problem, which has been analyzed among others in Gallego and Van Ryzin [1994], Gallien [2006], Gershkov and Moldovanu [2009, 2012] and Board and Skrzypacz [2013]. Board [2007] solves the problem of welfare and revenue maximization for the sale of real options. Oh and Özer [2013] analyze a delegated optimal stopping problem which arises as a consequence of asymmetric demand information in a manufacturer supplier relationship. The decision when to optimally end a manager's contract leads naturally to a stopping problem. Garrett and Pavan [2012] derive optimal managerial compensation schemes and show that the associated separation policy is a cut-off rule with a threshold that depends

only on time. Closely related Tirole [2014] analyzes optimal stopping as mechanism design problem and allows for the situation where the principal has limited commitment.

The approach of this paper differs from the standard mechanism-design approach, which by virtue of the revelation principle restricts attention to truthful (direct) mechanisms. We do not use direct mechanisms, which allows us to bypass technicalities that a formalization of the message (cheap talk) protocol between the principal and the agent would entail: the agent would have to communicate with the principal in every period and, hence, the space of communication strategies would be very rich. As optimal communication strategies are not necessarily Markovian, an optimization over those is a hard problem.

The direct mechanism design approach, has been successfully used in different settings, for example in Bergemann and Välimäki [2010] and Athey and Segal [2013] for welfare maximization, or in Baron and Besanko [1984], Courty and Li [2000], Battaglini [2005], Eső and Szentes [2007], Kakade et al. [2011] and Pavan, Segal, and Toikka [2014] for revenue maximization. Garrett and Pavan [2015] develop a novel approach to revenue maximization based on the local optimality conditions for the principal in the context of managerial compensation.

In contrast to the general dynamic mechanism design approach developed in Pavan, Segal, and Toikka [2014], we restrict here attention to allocations (stopping decisions) that can be implemented by transfers that do not require communication. This restriction is analytically highly convenient: Firstly, it allows for a sharp characterization of the class of allocations that can be implemented using these simple transfers. Secondly, the transfers admit a closed form representation based on constrained processes. On the downside, excluding communication reduces the set of implementable allocations. There are certainly settings where communication is necessary in order to implement optimal allocations. Whether this is the case has to be determined on a case-by-case basis. We present two applications where the restriction is without loss of generality and where either mechanisms without communication (“welfare maximization”, Section 3.6) or with communication only in the first period (“revenue maximization”, Section 3.7) prove to be optimal.

The paper proceeds as follows. Section 2 introduces the model. In Section 3 we show that all cut-off rules are implementable using posted-price mechanisms and derive the transfer as an expectation over constrained processes. We derive regularity assumptions under which also the opposite implication holds true: Every implementable stopping rule is a cut-off rule. Moreover, the transfer implementing a given cut-off is unique. We apply the results to principal-agent problems with interdependent values and show that there always

exists a simple welfare maximizing mechanism which does not require any communication. We discuss revenue maximization for the class of additive random walks. In this case a revenue maximizing contract is given by a transfer which only depends on the initial type and the realized stopping decision. In Section 4 we relate the findings to optimal stopping and derive a new characterization of the option value and the optimal cut-off. Section 5 concludes. All proofs not presented in the main text are given in the Appendix.

2 The Model

2.1 Evolution of the Private Information

Time is discrete and indexed by $t \in \{0, 1, \dots, T\} = \mathbb{T}$, for some fixed time horizon $T < \infty$. At every point in time $t \in \mathbb{T}$ the agent privately observes a real-valued Markov process X_t on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}, (\mathcal{F}_t)_{t \in \mathbb{T}})$. The initial value of the process X_0 is distributed according to the distribution function $F : \mathbb{R} \rightarrow [0, 1]$,

$$\mathbb{P}[X_0 \leq z] = F(z).$$

The formulation allows for completely general Markov processes which satisfy the following weak regularity assumptions.

Standing Assumption 1 (Polynomial Growth). *X is of polynomial growth, i.e. there exists a number $p > 0$ and a constant $C > 0$ such that*

$$\mathbb{E}[|X_{t+1}|^p \mid X_t = x] \leq C(1 + |x|^p) \text{ for all } x \in \mathbb{R}.$$

We say that a function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is of polynomial growth if there exists a constant $\tilde{C} > 0$ such that $|\phi(x)| \leq \tilde{C}(1 + |x|^p)$ for all $x \in \mathbb{R}$ with the same p as in Assumption 1. Assumption 1 assures that expected values of polynomial growth functions are finite with respect to the conditional probability measure of the process X .

Standing Assumption 2 (Monotone Transitions). *A higher value of the process $x' \geq x$ at time t leads to a higher value of the process at time $t + 1$ in the sense of first order stochastic dominance*

$$\mathbb{P}[X_{t+1} \leq z \mid X_t = x'] \leq \mathbb{P}[X_{t+1} \leq z \mid X_t = x] \text{ for all } z \in \mathbb{R}. \quad (1)$$

Standing Assumption 3 (Continuous Transitions). *For every $t \in \mathbb{T}$ and for every continuous, polynomial growth function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ the function $x \mapsto \mathbb{E}[\phi(X_{t+1}) \mid X_t = x]$ is continuous.*

All Markov processes commonly used in the economic literature satisfy these assumptions. The following result, proven in the Appendix, shows that for processes with independent, multiplicative or additive random shocks the assumptions are satisfied

Example 1 (Additive Random Walk). Let $X_t = X_0 + \sum_{s \leq t} \epsilon_s$ be the sum of identically distributed, independent shocks $(\epsilon_t)_{t \in \mathbb{T}}$ with finite second moment $\mathbb{E}[\epsilon_t^2] = \sigma^2 < \infty$.

Example 2 (Multiplicative Random Walk). Let $X_t = X_0 \prod_{s \leq t} \epsilon_s$ be the product of identically distributed, non-negative, independent shocks $(\epsilon_t)_{t \in \mathbb{T}}$ with finite second moments $\mathbb{E}[\epsilon_t^2] = \sigma^2 < \infty$.

Proposition 1 (Regularity of Random Walks). *Additive and Multiplicative Random Walks are of polynomial growth of order $p = 2$ and have continuous and monotone transitions.*

Assumptions 1, 2 and 3 are standing assumptions which hold throughout the paper. If explicitly stated we impose additional assumptions. For some of the uniqueness results we assume that the probability measure governing the transitions of X has full support.

Condition 1 (Full Support). For every $x \in \mathbb{R}$, $a < b$ and $t < T$ we have

$$\mathbb{P}[X_{t+1} \in [a, b] \mid X_t = x] > 0.$$

Moreover the distribution function F of the initial value X_0 is absolutely continuous.

2.2 Strategies and Payoffs of the Agent

Based on her past observations of the process the agent decides when to stop. Denote by $\mathcal{T}$ the set of $(\mathcal{F}_t)$ -adapted stopping rules with values in $\mathbb{T}$. A stopping rule is a complete contingent plan which maps every history observed by the agent into a binary stopping decision.

The agents payoff consists of three parts. At any time t before stopping the agent receives a flow payoff $f(t, X_t)$. At the time she stops she receives a final payoff $g(\tau, X_\tau)$. At any time t after stopping she obtains a flow payoff $h(t, X_t)$. The agent is risk-neutral

and her expected payoff $V(\tau)$ when using the stopping rule τ equals

$$V(\tau) = \mathbb{E} \left[\left(\sum_{t=0}^{\tau-1} f(t, X_t) \right) + g(\tau, X_\tau) + \left(\sum_{t=\tau+1}^T h(t, X_t) \right) \right]. \quad (2)$$

The continuous, polynomial growth payoff functions $f, g, h : \mathbb{T} \times \mathbb{R} \rightarrow \mathbb{R}$ depend on time t and the value of the process X_t .

Definition 1 (Marginal Incentive). We define the marginal incentive of the agent $z : \mathbb{T} \times \mathbb{R} \rightarrow \mathbb{R}$ as the expected change in payoffs when instead of stopping in period t she stops in period $t + 1$

$$z(t, x) = f(t, x) + \mathbb{E} [g(t + 1, X_{t+1}) - g(t, x) - h(t + 1, X_{t+1}) \mid X_t = x]. \quad (3)$$

We will show later that z completely determines the agents preferences over stopping rules. In static mechanism design with real-valued private information one often imposes a single crossing condition. This condition ensures that the agent's change in utility from getting a higher allocation is decreasing in her private information. In the dynamic setup, allocations differ in the time dimension. The stopping time τ rather describes *when* the agent is getting the object instead of *how much* of the object the agent is getting. We impose the following analogue of the static single crossing condition, which requires the agent's change in payoffs from stopping one period later to be decreasing in her private information.

Standing Assumption 4 (Dynamic Single Crossing).

The marginal incentive $z(t, x)$ is strictly decreasing in x for every $t \in \mathbb{T}$.

To ensure that Assumption 4 holds one can impose the following monotonicity conditions on f, g and h . If h is increasing and f and $\mathbb{E} [g(t, X_{t+1}) - g(t, X_t) \mid X_t = x]$ are decreasing in x and if there is at least one strict monotonicity, then z is strictly decreasing.

The single crossing condition is satisfied and has a natural interpretation in many economic models. To illustrate this and the general setup, let us briefly describe three examples. In a search model, the dynamic single crossing condition means that it is more costly to reject a good offer than to reject a bad offer. In an irreversible investment problem where the process X represents the demand, it means that the loss in profits from entering the market is higher if the demand is higher. In a two-state one-armed bandit model, the process X is the posterior belief that the state of the world is low. The single

crossing condition is satisfied since the expected payoff is lower the higher the value of X . In the following examples we discuss three models in more detail.²

Here $\delta < 1$ denotes a discount factor.

Example 3 (Search without Recall). In every period $t \in \mathbb{T}$ a worker receives a job-offer. The wages X_t of the offers are independently drawn from a known distribution $H : \mathbb{R}_+ \rightarrow [0, 1]$, with mean $\mu_H = \int_{\mathbb{R}_+} w dH(w)$. The worker aims at maximizing her discounted expected income $V(\tau) = \mathbb{E} [\sum_{t=\tau}^{\infty} \delta^t X_t] = \mathbb{E} [\frac{\delta^\tau X_\tau}{1-\delta}]$. The marginal incentive equals

$$z(t, x) = \mathbb{E} \left[\frac{\delta^{t+1} X_{t+1}}{1-\delta} - \frac{\delta^t X_t}{1-\delta} \mid X_t = x \right] = \frac{\delta^t (\delta \mu_H - x)}{1-\delta}, \quad (4)$$

and thus satisfies the dynamic single crossing condition for every distribution of job-offers H .

Example 4 (Search with Recall). A worker receives a job-offer with wage w_t every period $t \in \mathbb{T}$. The wages are independently drawn from a known distribution $H : \mathbb{R}_+ \rightarrow [0, 1]$. The worker can recall past offers such that at time t she will accept her best previous offer, i.e. $X_t = \max\{w_s : s \leq t\}$. She aims at maximizing her discounted expected income $V(\tau) = \mathbb{E} [\sum_{t=\tau}^{\infty} \delta^t X_t] = \mathbb{E} [\frac{\delta^\tau X_\tau}{1-\delta}]$. The marginal incentive is given by

$$z(t, x) = \mathbb{E} \left[\frac{\delta^{t+1} X_{t+1}}{1-\delta} - \frac{\delta^t X_t}{1-\delta} \mid X_t = x \right] = \frac{\delta^t \left(\delta \int_{\mathbb{R}_+} \max\{w, x\} dH(w) - x \right)}{1-\delta}. \quad (5)$$

As $\partial z(t, x)/\partial x = \delta^t (\delta(1 - H(x)) - 1)/(1 - \delta) < 0$ the dynamic single crossing condition is satisfied for every distribution of job-offers H .

Example 5 (A Bandit Model of Rational Learning). There are two states of the world $\theta \in \{l, h\}$. The returns (R_t) from running a company are distributed according to $H_\theta : \mathbb{R} \rightarrow [0, 1]$ depending on the state θ . A manager observing the returns decides at which time τ to liquidate the firm in order to maximize expected discounted returns $V(\tau) = \mathbb{E} [\sum_{t=0}^{\tau-1} \delta^t R_t]$. As the expected returns depend only on the conditional probability that the world is in low state $X_t = \mathbb{P}[\theta = l \mid \mathcal{F}_t]$, using the law of iterated expectations, we can rewrite the objective function as $V(\tau) = \mathbb{E} [\sum_{t=0}^{\tau-1} \delta^t (\mu_h - X_t(\mu_h - \mu_l))]$. Here, $\mu_\theta = \int_{\mathbb{R}} r dH_\theta(r)$ denotes the expected return conditional on the state of the world θ . Without loss of

²Note that the restriction to one-dimensional processes precludes search models where workers have different types. In a learning model a special signal structure needs to be assumed to ensure that the posterior belief is a one-dimensional Markov process. Examples are the commonly used normal prior with normal signals or any prior which assigns probability only to two states.

generality we assume $\mu_h \geq \mu_l$. The marginal incentive equals $z(t, x) = \delta^t (\mu_h - x(\mu_h - \mu_l))$ and the dynamic single crossing condition is satisfied for every pair of distributions H_l, H_h .

For some results we assume in addition that z is unbounded, which ensures that it is optimal to stop for sufficiently high values of the process and optimal to continue for sufficiently low values of the process.

Condition 2 (Unbounded Marginal Incentive).

We have $\lim_{x \rightarrow \infty} z(t, x) = -\infty$ and $\lim_{x \rightarrow -\infty} z(t, x) = \infty$ for every $t \in \mathbb{T}$.

3 Implementable Stopping Rules and Transfers

We want to characterize how the behavior of the agent can be influenced using a transfer $\pi : \mathbb{T} \rightarrow \mathbb{R}$ that only depends on the realized stopping decision, but not on the path of the process X the agent observes. The transfer π is paid to the agent in addition to her payoffs f, g, h and is thereby changing her preferences over stopping times.

Definition 2 (Implementable Stopping Rule). A stopping rule τ^* is implemented by a transfer π if τ^* is the minimal³ stopping rule satisfying

$$\sup_{\tau \in \mathcal{T}} V(\tau) + \mathbb{E}[\pi(\tau)] = V(\tau^*) + \mathbb{E}[\pi(\tau^*)]. \quad (6)$$

We say that τ^* is implementable if there exists a transfer π that implements it.

Such a transfer rule has multiple attractive economic features: As the transfer π is independent of the realization of the process X it can be paid even if the realization of X is unobservable. The only required information is the realized stopping decision. Intuitively, it suffices to know that the agent stopped, instead of for what reasons she stopped. Furthermore, as the transfer depends only on the stopping decision it requires no communication.

The expected payoff V defines a preference relation over stopping times. The agent prefers the stopping time τ over τ' if and only if $V(\tau) \geq V(\tau')$. The next result, proven in the Appendix, shows that the preferences of the agent over stopping times depend only on her marginal incentive to delay the allocation.

³A stopping rule τ is minimal in a set of stopping rules $S \subseteq \mathcal{T}$ if $\tau \leq \tau'$ almost surely for all $\tau' \in S$. Note that minimal stopping times are almost surely unique. If $\tau, \tau' \in S \subseteq \mathcal{T}$ are minimal stopping times in S , then $\tau \leq \tau' \leq \tau$ and hence $\tau = \tau'$ almost surely.

Proposition 2. *The expected payoff of the agent when using the stopping rule τ can be represented as the sum of the payoff of stopping in period zero plus her expected marginal incentives*

$$V(\tau) = \mathbb{E} \left[\sum_{t=0}^{\tau-1} z(t, X_t) \right] + V(0).$$

Proposition 2 shows that the agent's preferences over stopping times are completely determined by the marginal incentive z to delay the allocation to the next period. Note that Proposition 2 is a result about *expected* payoffs. The realized payoffs for a given path of X may differ. As a consequence of Proposition 2 the definition of implementability can be simplified. A stopping time τ^* is implemented by the transfer π if τ is the minimal stopping time satisfying⁴

$$\sup_{\tau \in \mathcal{T}} \mathbb{E} \left[\sum_{t=0}^{\tau-1} z(t, X_t) + \pi(\tau) \right] = \mathbb{E} \left[\sum_{t=0}^{\tau^*-1} z(t, X_t) + \pi(\tau^*) \right]. \quad (7)$$

It is convenient to generalize the agent's payoff V and to allow for arbitrary initial times t and initial values x . Let $\mathcal{T}_{t,T}$ denote the set of all stopping rules with values in the set $\{t, \dots, T\}$. For any stopping rule $\tau \in \mathcal{T}_{t,T}$ we define the agent's expected continuation value by

$$V_{t,x}(\tau) = \mathbb{E} \left[\left(\sum_{s=t}^{\tau-1} f(s, X_s) \right) + g(\tau, X_\tau) + \left(\sum_{s=\tau+1}^T h(s, X_s) \right) \mid X_t = x \right]. \quad (8)$$

For every transfer π we introduce the agent's value function $v_\pi : \mathbb{T} \times \mathbb{R} \rightarrow \mathbb{R}$ as

$$v_\pi(t, x) = \sup_{\tau \in \mathcal{T}_{t,T}} (V_{t,x}(\tau) + \mathbb{E}[\pi(\tau) \mid X_t = x]) - V_{t,x}(t).$$

Hence, $v_\pi(t, x) - \pi(t)$ is the difference between the expected payoff when continuing optimally and stopping immediately. Put differently, the difference $v_\pi(t, x) - \pi(t)$ is the agent's willingness to pay for the option to stop after period t . It follows from the argument in Proposition 2 that v_π is the supremum of the sum of expected future marginal incentives and the transfer, when the agent follows the optimal continuation strategy

$$v_\pi(t, x) = \sup_{\tau \in \mathcal{T}_{t,T}} \mathbb{E} \left[\sum_{s=t}^{\tau-1} z(s, X_s) + \pi(\tau) \mid X_t = x \right]. \quad (9)$$

⁴We use the convention $\sum_{s=t}^{t-1} \cdot = 0$.

Since the agent always has the opportunity to stop immediately, i.e. to choose $\tau = t$, it follows $v_\pi(t, x) - \pi(t) \geq 0$. At each time t the agent faces the binary decision of continuing or stopping. Given a value x of the process at time t she bases her decision on whether the value of the option to continue is positive. If $v_\pi(t, x) - \pi(t) = 0$, then there is no gain from continuing and the agent stops. If $v_\pi(t, x) - \pi(t) > 0$, the agent continues at least one more period. Intuitively, it follows that the minimal optimal stopping rule for the agent is given by⁵

$$\tau^\star = \inf \{t \geq 0 \mid X_t \in D_\pi(t)\} \wedge T, \quad (10)$$

where the so-called stopping region $D_\pi(t)$ is defined by $D_\pi(t) = \{x \in \mathbb{R} \mid v_\pi(t, x) = \pi(t)\}$. If the agent decides to stop at time t , she receives the transfer $\pi(t)$. If it is optimal to continue, she obtains the marginal incentive $z(t, x)$ plus the expected value of continuing optimally in the next period $\mathbb{E}[v_\pi(t+1, X_{t+1}) \mid X_t = x]$. This leads to the dynamic programming principle which represents the value function in recursive form for all $t \in \mathbb{T}$ and $x \in \mathbb{R}$

$$v_\pi(t, x) = \max \{\pi(t), z(t, x) + \mathbb{E}[v_\pi(t+1, X_{t+1}) \mid X_t = x]\}. \quad (11)$$

For a rigorous derivation of Eq.s (10) and (11) we refer to Peskir and Shiryaev [2006, Chapter 1, Theorem 1.9]. The following Lemma establishes the regularity of the value function for every transfer.

Lemma 1. *v_π is non-increasing, continuous and of polynomial growth in x .*

We introduce the notion of cut-off rules. Cut-off rules are stopping rules such that the agent stops the first time the process X exceeds a time-dependent threshold b . A mapping $b : \mathbb{T} \rightarrow \bar{\mathbb{R}}^6$ with $b(T) = -\infty$ is called a cut-off.

Definition 3 (Cut-Off Rule). A stopping rule τ is a cut-off rule if there exists a cut-off b such that almost surely

$$\tau = \inf \{t \geq 0 \mid X_t \geq b(t)\} \wedge T.$$

We denote the cut-off rule corresponding to the cut-off b by τ_b . If $-\infty < b(t) < \infty$ for all $t < T$ we call b a finite cut-off and τ_b a finite cut-off rule.

The next Lemma shows that under the full support condition we can uniquely recover the cut-off b of a cut-off rule τ .

⁵We use the notation $x \wedge y = \min\{x, y\}$.

⁶Here and in the sequel $\bar{\mathbb{R}} = \mathbb{R} \cup \{\infty\} \cup \{-\infty\}$

Lemma 2. *Suppose that the process X has full support, i.e. Condition 1 holds. Then for every cut-off rule τ there exists a unique cut-off b satisfying $\tau = \tau_b$ almost surely.*

In Section 3.1 we show that every implementable stopping rule is a cut-off rule. In Section 3.3 we show that all cut-off rules are implementable. The associated transfer admits an explicit representation in terms of the constrained version of X which we introduce in Section 3.2.

3.1 All Implementable Stopping Rules are Cut-off Rules

It is known that if one imposes the single crossing condition in a static mechanism design problem, an allocation rule can be implemented if and only if it is monotone. In the dynamic setup, cut-off rules play the role of monotone allocations. Indeed, Assumption 2 implies that a cut-off rule τ_b is monotone in the sense of first order stochastic dominance: For any $t < s$ the probability that X exceeds b before time s is increasing in the conditional value of X at time t

$$\mathbb{P}[\tau_b \leq s \mid X_t = x'] \leq \mathbb{P}[\tau_b \leq s \mid X_t = x] \text{ for all } x' < x < b(t).$$

The next Proposition shows that only cut-off rules are implementable.

Proposition 3. *If the stopping rule τ is implementable, then τ is a cut-off rule.*

While the proof of Proposition 3 is given in the Appendix the intuition is straightforward: The gain from continuing v_π equals the sum of future marginal incentives z . The marginal incentive is decreasing in x by Assumption 4. By Assumption 2 higher values of x today lead to higher values in all later periods (in the sense of first order stochastic dominance). Thus, if it is optimal for an agent with value $x = X_t$ to stop the process then stopping is also optimal for all agents with higher values $X_t = x' > x$.

The next result provides a condition on the marginal incentive such that at every point in time at least some agents stop and some continue, i.e. every implementable cut-off $b(t)$ is finite. Under Condition 2 the marginal incentive $z(t, x)$ to delay the stopping decision from time t to $t + 1$ gets arbitrarily large as x decreases. We show that this condition suffices to guarantee that for any transfer π there exists some level $\underline{x} \in \mathbb{R}$ where it is strictly optimal for the agent to continue. Moreover there exists a finite threshold $\bar{x} > \underline{x}$ where it is optimal to stop. In between these two levels there exists a number $b(t) \in [\underline{x}, \bar{x}]$ where the agent is indifferent between stopping and continuing.

Figure 1: The binomial random walk (dashed blue line) and the constrained binomial random walk (solid red line) starting in $X_0 = 8$ constrained by the constant cut-off $b(t) = 10$.

Proposition 4. *If Condition 2 holds and the stopping rule τ is implemented by a transfer π , then τ is a cut-off rule with a finite cut-off b . The agent is indifferent between stopping and continuing at the cut-off*

$$v_\pi(t, b(t)) = \pi(t) = \sup_{\tau \in \mathcal{T}_{t+1, T}} \mathbb{E} \left[\sum_{s=t}^{\tau-1} z(s, X_s) + \pi(\tau) \mid X_s = b(t) \right]. \quad (12)$$

3.2 Constrained Processes

In this section we introduce a distortion of the process X that plays the crucial role in designing mechanisms in optimal stopping problems. For a given cut-off b we convert X to a new process $\tilde{X}$ which is called the constrained version of X . It is characterized as the Markov process which evolves as X below b , but is constrained to be on the cut-off b whenever X tries to exceed it. We show that the transfer implementing the given cut-off rule τ_b can be represented as an expectation over the constrained version $\tilde{X}$ of X .

For an illustration we first present the construction of a constrained random walk (see also Figure 1).

Example 6 (Constrained random walk). As in Example 1 let X be a random walk, i.e. $X_t = X_0 + \sum_{s \leq t} \epsilon_s$ with $\mathbb{P}[\epsilon_s = 1] = \mathbb{P}[\epsilon_s = -1] = \frac{1}{2}$. In this case the constrained version can be constructed path-by-path. First, set $\tilde{X}_0 = X_0 \wedge b(0)$ and then define $\tilde{X}$ recursively by $\tilde{X}_{t+1} = (\tilde{X}_t + \epsilon_{t+1}) \wedge b(t+1)$ for all $t < T$.

For general Markov processes we take a different approach to construct the constrained version. We first modify the transition probabilities of the original process X and then we define $\tilde{X}$ as the Markov process having these new transition probabilities. To simplify notation we define the transition kernel $P_{t,s}$ of X for $t < s \in \mathbb{T}$, which acts on polynomial growth, measurable functions by

$$(P_{t,s}\phi)(x) = \mathbb{E}[\phi(s, X_s) \mid X_t = x].$$

Assumption 2 implies that the kernel $P_{t,t+1}$ preserves monotonicity: If $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is non-increasing then $P_{t,t+1}\phi$ is non-increasing as well. Assumptions 1 and 3 ensure that for every continuous, polynomial growth function ϕ the function $P_{t,t+1}\phi$ is continuous and of polynomial growth as well.

For every $t < T$ we define the transition kernel of the constrained process $\tilde{P}_{t,t+1}$ which acts on bounded, measurable functions by

$$(\tilde{P}_{t,t+1}\phi)(x) = \mathbb{E}[\phi(X_{t+1} \wedge b(t+1)) \mid X_t = x]. \quad (13)$$

We extend the family of kernels to a semi-group $(\tilde{P}_{t,s})_{t \leq s}$ via the composition $\tilde{P}_{t,s}\phi = \tilde{P}_{t,t+1}(\tilde{P}_{t+1,t+2} \dots (\tilde{P}_{s-1,s}\phi))$. The operator $\tilde{P}_{t,t}$ is defined to be the identity. Then there exists a Markov process $\tilde{X}$ on some filtered probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}, (\tilde{\mathcal{F}}_t)_{t \in \mathbb{T}})$ with the transition kernel $\tilde{P}$. We call $\tilde{X}$ the constrained version of X .

Above, we defined the constrained process $\tilde{X}$ as the unique Markov process which follows the same transitions as the original process which is constrained by the cut-off

$$\tilde{\mathbb{P}}[\tilde{X}_{t+1} \in A \mid \tilde{X}_t = x] = \mathbb{P}[X_{t+1} \wedge b(t+1) \in A \mid X_t = x].$$

The constrained process has two important properties: First, the transition probabilities of $\tilde{X}$ and X coincide below the cut-off, i.e. for all set below the cut-off $A \subset (-\infty, b(t+1))$

$$\tilde{\mathbb{P}}[\tilde{X}_{t+1} \in A \mid \tilde{X}_t = x] = \mathbb{P}[X_{t+1} \in A \mid X_t = x].$$

Second, the constrained process $\tilde{X}$ never exceeds the cut-off, i.e.

$$\tilde{\mathbb{P}}[\tilde{X}_{t+1} > b(t+1) \mid \tilde{X}_t = x] = 0.$$

3.3 All Cut-Off Rules are Implementable

As shown in Section 3.1 only cut-off rules are implementable. In this section we prove the opposite direction, i.e. every finite cut-off rule is implementable. For a given cut-off rule we define the transfer and explicitly verify that it implements the cut-off rule. In static mechanism design every monotone allocation rule is implemented by a transfer equal to the integral over marginal incentives (cf. Guesnerie and Laffont [1984]). In the dynamic context cut-off rules are the equivalent of monotone allocation rules.⁷ The transfer π implementing a finite cut-off rule τ_b equals the expected future marginal incentive z to delay the allocation evaluated at the process $\tilde{X}$ constrained by the cut-off b

$$\pi(t) = \tilde{\mathbb{E}} \left[\sum_{s=t}^{T-1} z(s, \tilde{X}_s) \mid \tilde{X}_t = b(t) \right] \text{ for } t < T \text{ and } \pi(T) = 0. \quad (14)$$

Theorem 1. *Every finite cut-off rule τ_b is implemented by the transfer defined in Eq. (14).*

While the proof is given in the Appendix the result can be understood intuitively: Suppose the transfer $(\pi(s))_{s \in \{t+1, t+2, \dots\}}$ provides incentives to stop if and only if $X_s \geq b(s)$ for all $s > t$. Then for any agent with a value $X_s \geq b(s)$ it is optimal to stop and thus she behaves exactly as an agent with value $X_s = b(s)$. Therefore, we can identify all agent above the cut-off with the agent on the barrier who finds it also optimal to continue. Consequently, the continuation value if $X_t = b(t)$ is given by the sum of marginal incentives if we replace the original process by the constrained process

$$\sup_{\tau} \mathbb{E} \left[\sum_{s=t}^{\tau-1} z(s, X_s) + \pi(\tau) \mid X_t = b(t) \right] = \mathbb{E} \left[\sum_{s=t}^{\tau-1} z(s, \tilde{X}_s) \mid \tilde{X}_t = b(t) \right].$$

As on the cut-off the agent is indifferent between stopping with the transfer $\pi(t)$ and

⁷Under the single crossing condition and additional regularity assumptions one can use Proposition 2 to establish that the set of cut-off rules coincides with the set of Markovian average monotone allocations in the sense of Corollary 1 from Pavan, Segal, and Toikka [2014]. Theorem 3 of Pavan et al. then implies that cut-off rules can be implemented in a direct mechanism, using transfers that depend on the complete history of reports.

receiving the above continuation value, Eq. (14) follows.

3.3.1 Comparative Statics

The closed form solution for the transfer obtained in Eq. (14) allows to easily derive comparative statics. First, as the transfer is the sum of marginal incentives, increasing the marginal incentive z increases the transfer.

Corollary 1. *The transfer defined in Eq. (14) is nondecreasing in z , i.e. let $z(t, x) \leq \hat{z}(t, x)$ for all t and x and let $\pi, \hat{\pi}$ be the associated solutions of Eq. (14), then $\pi(t) \leq \hat{\pi}(t)$.*

Intuitively, if the incentive to continue increases, the transfer to the agent when she stops has to increase to motivate her to maintain her stopping decision.

If the incentive to delay stopping is nonnegative $z \geq 0$ then the transfer obtained in Eq. (14) is decreasing over time. To see this note that $\pi(t)$ and $\pi(t+1)$ differ in two ways: First, the sum defining $\pi(t)$ includes the marginal incentive $z(t, b(t))$ at time t which is by assumption positive. Second, the expectation over the other elements of the sum is taken conditional on $\tilde{X}_t = b(t)$ instead of $\tilde{X}_{t+1} = b(t+1)$. As $\tilde{X}_{t+1}$ is smaller $b(t+1)$ by definition of the constrained process, the transition operator of the constrained process is monotone and z is monotone by the single crossing assumption the expectation over the difference of the other elements of the sum is positive as well.

Corollary 2. *The transfer defined in Eq. (14) is nonincreasing over time if $z(t, x) \geq 0$ for all (t, x) .*

Intuitively, as the incentive to wait is positive the agent has to be compensated less at later points in time as she has less opportunity to wait.

Due to the monotonicity of z the transfer obtained in Eq. (14) is monotone in the constrained process $\tilde{X}$. As the reflected process $\tilde{X}$ is monotone (in the sense of FOSD) in the cut-off b it thus follows that the transfer is monotone decreasing in the cut-off.

Corollary 3. *The transfer is nonincreasing in the barrier, i.e. let $b(t) \leq \hat{b}(t)$ for all t and $\pi, \hat{\pi}$ the solutions of Eq. (14) for the cut-offs $b, \hat{b}$ then $\pi(t) \geq \hat{\pi}(t)$.*

3.4 Examples

One advantage of the solution method based on constrained processes is the explicit characterization of transfers derived in Eq. (14). The next two examples provide formulas for the transfers in search models with and without recall.

Example 7 (Search without Recall continued). In a search problem without recall, as defined in Example 3, the independence of the offers X_t implies that a constrained version of X is given by $\tilde{X}_t = X_t \wedge b(t)$. Using that the marginal incentive equals $(1-\delta)^{-1}\delta^t(\delta\mu_H - x)$ (see Eq. (4)), every finite cut-off rule τ_b can be implemented by the transfer

$$\begin{aligned}\pi(t) &= \tilde{\mathbb{E}} \left[\sum_{s=t}^{T-1} z(s, \tilde{X}_s) \mid \tilde{X}_t = b(t) \right] = \mathbb{E} \left[\sum_{s=t}^{T-1} z(s, X_s \wedge b(s)) \mid X_t = b(t) \right] \\ &= \frac{\delta^t(\delta\mu - b(t))}{1-\delta} + \sum_{s=t+1}^{T-1} \frac{\delta^s \left(\delta\mu_H - \int_{\mathbb{R}_+} w \wedge b(s) dH(w) \right)}{1-\delta}.\end{aligned}$$

Example 8 (Search with Recall continued). In a search problem with recall, defined as in Example 4, the process X of the current best offer is non-decreasing. Therefore a constrained version of X for a decreasing cut-off b is given by $\tilde{X}_t = X_t \wedge b(t)$. Using that the marginal incentive equals $(1-\delta)^{-1}\delta^t \left(\delta \int_{\mathbb{R}_+} \max\{w, x\} dH(w) - x \right)$ (see Eq. (5)), we obtain that every finite cut-off rule τ_b with a decreasing cut-off b can be implemented by the transfer

$$\pi(t) = \mathbb{E} \left[\sum_{s=t}^{T-1} z(s, X_s \wedge b(s)) \mid X_t = b(t) \right] = \sum_{s=t}^{T-1} \frac{\delta^s \left(\delta \int_{\mathbb{R}_+} \max\{w, b(s)\} dH(w) - b(s) \right)}{1-\delta}.$$

3.5 Uniqueness of the Transfer

An important result in static mechanism design and auction theory is the revenue equivalence theorem. It was first observed in Vickrey [1961] that many classical auction mechanisms (first-price auctions, Dutch auctions, English auctions, and second-price auctions) lead to the same expected transfers. This result was later generalized to many other auction setups and mechanism design problems. The revenue equivalence theorem states that the expected transfers implementing a given allocation rule are unique up to a constant. The following proposition shows that revenue equivalence holds in the dynamic model, if one imposes a full support condition similar to the one necessary in static mechanism design problems.⁸

⁸The same result is derived in Theorem 2 in Pavan, Segal, and Toikka [2014] for general “regular” environments. The environment presented in this paper is not regular in the sense of Pavan et al. as g, h are neither assumed to be differentiable (U-D) nor Lipschitz continuous (U-ELC). Furthermore, the processes considered in this paper do not necessarily have bounded impulse responses (F-BIR). The techniques used here to establish the payoff equivalence do not rely on envelope arguments.

Proposition 5 (Payoff Equivalence). *Suppose that X has full support (Condition 1) and let τ be a finite cut-off rule, then the transfer implementing τ is unique up to an additive constant.*

Proof. Let $\pi, \hat{\pi}$ be two payments implementing τ such that $\pi(T) = \hat{\pi}(T)$ and let $v = v_\pi$ and $\hat{v} = v_{\hat{\pi}}$ denote the associated value functions. We show that the two value functions coincide: $v(t, x) = \hat{v}(t, x)$ for all $t \in \mathbb{T}$ and $x \in \mathbb{R}$. This implies uniqueness of the transfer since by Lemma 2, there exists b such that b is the unique cut-off satisfying $\tau = \tau_b$ and hence Eq. (12) holds for v_π as well as for $v_{\hat{\pi}}$ with the same cut-off b . In particular we have $\pi(t) = v(t, b(t)) = \hat{v}(t, b(t)) = \hat{\pi}(t)$. At time T we clearly have $v(T, x) = \hat{v}(T, x)$ for all $x \in \mathbb{R}$. Using this as induction basis we obtain by Eq. (12) for $t < T$

$$\pi(t) = z(t, b(t)) + P_{t,t+1}v(t, b(t)) = z(t, b(t)) + P_{t,t+1}\hat{v}(t, b(t)) = \hat{\pi}(t).$$

Therefore the dynamic programming principle (11) implies $v(t, x) = \hat{v}(t, x)$ for all $x \in \mathbb{R}$. $\square$

The reason why uniqueness fails without the full support condition is the same as in static mechanism design with a discrete type space: As for every realization of the process X that happens with positive probability, the agent might have a strict preference, we can slightly change the transfer without influencing the agent's optimal strategy. Similarly if the cut-off is infinite, i.e. the agent should stop/continue for every realization of X_t , there might be no agent who is indifferent and we can increase/decrease the transfer without changing the agent's behavior.

If Condition 2 is satisfied, i.e. the marginal incentive is unbounded, all implementable stopping rules are finite cut-off rules by Proposition 4 and hence the transfers are always unique up to a constant.

3.6 Welfare in Principal Agent Problems

In this section we use the characterization of implementable stopping rules to construct a simple welfare maximizing mechanism for general optimal stopping problems. Suppose there is a principal who aims at maximizing social welfare

$$W(\tau) = \mathbb{E} \left[\left(\sum_{t=0}^{\tau-1} \hat{f}(t, X_t) \right) + \hat{g}(\tau, X_\tau) + \left(\sum_{t=\tau+1}^T \hat{h}(t, X_t) \right) \right]$$

over stopping rules $\tau \in \mathcal{T}$. As in Section 2 the agent's expected payoff $V(\tau)$ is determined by the triplet of functions f, g, h . As the functions $\hat{f}, \hat{g}, \hat{h}$ do not necessarily coincide with f, g, h , the principal and the agent in general want to use a different stopping rule in order to maximize their payoff. Assume the principal's and the agent's payoffs satisfy the dynamic single crossing condition. Let us denote by $\hat{\tau}$ the minimal stopping time which is optimal for the principal $W(\hat{\tau}) = \sup_{\tau \in \mathcal{T}} W(\tau)$. If the principal could observe the realization of the process X directly, she would use the policy $\hat{\tau}$. A direct implication of our results is that the principal can implement $\hat{\tau}$ even if the agent observes the process X privately.

Theorem 2. *There exists a transfer $\pi : \mathbb{T} \rightarrow \mathbb{R}$ which only conditions on the time the agent stopped that implements the principal optimal stopping time $\hat{\tau}$, i.e.*

$$V(\hat{\tau}) + \mathbb{E}[\pi(\hat{\tau})] = \sup_{\tau \in \mathcal{T}} V(\tau) + \mathbb{E}[\pi(\tau)] .$$

Proof. As the minimal optimal stopping rule is implementable by a transfer of zero the principal's objective function W is maximized by a cut-off rule $\hat{\tau}$ by Proposition 3. By Theorem 1 there exists a transfer π which only conditions on the time the agent stopped that implements $\hat{\tau}$. $\square$

Note that a simple transfer which conditions only on when the agent stopped and does not require any communication, suffices to maximize welfare. This might make the results appealing from an applied perspective. For example, in the context of job search this means that there is no need for the worker to report her job-offers to the unemployment agency to implement a welfare maximizing mechanism.

3.7 Revenue Maximization

While the previous section discussed the case when the principal has no preferences over transfers this section discusses revenue maximization. Transfers without communication cannot be used to implement the revenue maximizing allocation. Still, in the case of an additive random walk we can use the results developed in the previous sections to establish a result similar to Board [2007]: The revenue maximizing contract for selling an investment option is a menu of option contracts and as such only requires communication in the first period. While Board [2007] focuses on the time-homogeneous case the analysis below

shows that this result continues to hold in the time-inhomogeneous case if one allows for time-dependent strike prices.

For simplicity and without affecting the qualitative results we assume that the principal has no intrinsic preferences over stopping rules, i.e. $\hat{f}, \hat{g}, \hat{h} \equiv 0$ and aims at minimizing the expected transfer $\mathbb{E}[\pi]$ subject to the constraint that every time zero type x finds it optimal to participate in the mechanism

$$\mathbb{E} \left[\left(\sum_{t=0}^{\tau-1} f(t, X_t) \right) + g(\tau, X_\tau) + \left(\sum_{t=\tau+1}^T h(t, X_t) \right) + \pi \mid X_0 = x \right] \geq 0. \quad (15)$$

Throughout the section we assume that $X_t = X_0 + \sum_{s=1}^t \epsilon_s$ is an additive random walk (cf. Example 1).⁹ The distribution of the initial value F is absolutely continuous, with strictly positive density F' on $[0, \bar{x})$ with $\bar{x} \leq \infty$, and decreasing hazard rate $d(x_0) = \frac{1-F(x_0)}{F'(x_0)}$. Furthermore, we assume that f, g, h are differentiable in x with bounded derivative. In particular this implies that the marginal incentive z is differentiable and Lipschitz continuous.

While up to now we restricted attention to mechanisms without communication we allow for communication in this section. In each period t after observing the realization of the process X_t the agent reports $\hat{X}_t$. A mechanism maps the agent's reports into a stopping decision τ and a transfer π .

Definition 4 (Incentive Compatible Mechanism). Let $\tau : \mathbb{R}^T \rightarrow \mathbb{T}$ map reported types $\hat{X}$ into stopping decision and $\pi : \mathbb{R}^T \rightarrow \mathbb{R}$ map reported types into payments.¹⁰ A mechanism (τ, π) is incentive compatible if reporting truthfully maximizes the agent's payoffs over the space of all reporting strategies ($\mathcal{F}_t$ -adapted stochastic processes).

The assumptions of this section ensure that the environment is regular in the sense of Pavan, Segal, and Toikka and it follows from Theorem 1 in that paper that in any incentive compatible mechanism the expected payment $\mathbb{E}[\pi]$ satisfies

$$-\mathbb{E}[\pi] = \mathbb{E} \left[\sum_{t=0}^{\tau-1} z(t, X_t) - z_x(t, X_t) \frac{1-F(X_0)}{F'(X_0)} \right] - V_0 \quad (16)$$

where V_0 is the expected value of an agent conditional on $X_0 = 0$. Hence, an upper bound

⁹The case of an multiplicative random walk could be derived analogous using that $\frac{\partial X_t}{\partial X_0} = \frac{X_t}{X_0}$. The second order single crossing condition in this case is given by $x \mapsto z_x(t, x)x$ non-decreasing.

¹⁰Formally τ is a stopping rule if $\tau(X) = \tau(X')$ if $X_t = X'_t$ for all $t \leq \tau$.

on the revenue is given by setting $V_0 = 0$ and maximizing the virtual valuation defined in Eq. (16). To ensure that the stopping rule maximizing Eq. (16) is a cut-off rule for every fixed value of X_0 , we impose the following condition.

Condition 3 (Second Order Single Crossing). The marginal incentive z is convex.

Observe that Condition 3 is for example satisfied in a standard irreversible investment model where the agent is risk-averse (u concave) over the returns X she receives in every period t after investing and the investment cost are independent of X , i.e.

$$V(\tau) = \mathbb{E} \left[\sum_{t=\tau+1}^T u(t, X_t) - c(\tau) \right].$$

Under Condition 3 the virtual valuation defined in Eq. (16) is maximized by a cut-off stopping time, where the cut-off depends not only on time t but also on the initial value X_0 .

Proposition 6 (Virtual Valuation Maximizing Allocation). *Under Condition 3 there exists a function $b : \mathbb{T} \times \mathbb{R}_+ \rightarrow \mathbb{R}$ such that the virtual valuation*

$$\mathbb{E} \left[\sum_{t=0}^{\tau-1} z(t, X_t) - z_x(t, X_t) d(X_0) \right]$$

is maximized by the stopping rule $\tau^ = \min\{t \in \mathbb{T} : X_t \geq b(t, X_0)\} \wedge T$. Furthermore, $b(t, x_0)$ is non-increasing in x_0 .*

Proof. For every x_0 the function $x \mapsto z(t, x) - z_x(t, x) d(x_0)$ is decreasing and thus Proposition 3 implies that for a fixed x_0 the optimal stopping rule is a cut-off stopping rule.

Note, that $z(t, x) - z_x(t, x) d(x_0)$ is decreasing in x_0 as the generalized hazard rate d is decreasing and $z_x < 0$. Consequently, for every (t, x) such that it is optimal to stop if the initial value was x_0 , i.e.

$$\sup_{\tau} \mathbb{E} \left[\sum_{s=t}^{\tau-1} z(s, X_s) - z_x(s, X_s) d(x_0) \right] = 0,$$

it is also optimal to stop for all higher $x'_0 \geq x_0$. □

The next theorem establishes that the virtual valuation maximizing allocation of Proposition 6 can be implemented using simple transfers that only depend on the initial value of the process and the realized stopping time.

Theorem 3 (Revenue Maximizing Mechanism). *There exists a mechanism maximizing revenue (in the class of all incentive compatible mechanisms) where the agent reports only her initial value X_0 and is free to choose any stopping time. The transfer depends only on X_0 and the realized stopping time τ .*

The Proof of Theorem 3 given in the Appendix shows that an agent with a higher initial value X_0 has a stronger preference to use an early stopping time and thus one can construct payments that make truthful reporting at time zero incentive compatible.

Similar to Theorem 3 Garrett and Pavan [2015] show in a model of managerial compensation where private information evolves according to an autoregressive process that communication in the first period suffices to implement the revenue maximizing allocation.

Example 9 (Selling an Irreversible Investment Option). Consider a monopolistic seller of an investment opportunity, which can only be exploited by the buyer. We model the investment problem of the buyer as a standard irreversible investment problem (c.f. Dixit and Pindyck, 2008). If the buyer invests at time τ she pays investment costs $c(\tau)$ and receives profits X_t depending on a privately observed state X_t at all times $t \geq \tau$ such that the overall utility of the buyer is given by

$$\sum_{t=\tau+1}^T \delta^t u(X_t) - \delta^\tau c(\tau).$$

In this problem the marginal incentive is given by

$$z(t, x) = \delta^t (\delta c(t+1) - c(t) - \delta \mathbb{E}[u(x + \epsilon_t)])$$

and thus satisfies the first and second order single crossing condition if u is increasing and concave. By Theorem 3 a revenue maximizing contract is given by a menu of time dependent price schedules such that the buyer selects one at time zero and makes a payment only depending the time of investment τ .

4 Relation to Optimal Stopping

In this section we provide a new closed form characterization of the option value in general optimal stopping problems as an expectation over constrained processes. We consider the standard optimal stopping problem, where the agent optimizes her expected payoff $V(\tau)$ over the set of stopping rules $\tau \in \mathcal{T}$.¹¹ We say that τ^* is the minimal optimal stopping rule, if it is the minimal stopping rule satisfying $V(\tau^*) = \sup_{\tau \in \mathcal{T}} V(\tau)$.¹²

The notion of implementability introduced in Definition 2 generalizes the notion of optimality in the sense that every minimal optimal stopping rule is implementable by a transfer of zero. Hence, we get the following immediate corollary of Proposition 3, which reproduces the well known result from optimal stopping theory (see e.g. Jacka and Lynn, 1992).

Corollary 4. *The minimal optimal stopping rule is a cut-off rule.*

Proof. As every minimal optimal stopping rule is implementable, the optimal stopping rule is a cut-off rule by Proposition 3. $\square$

We define the agent's option value $w : \mathbb{T} \times \mathbb{R} \rightarrow \mathbb{R}$ as the difference from stopping at time t with the value x and continuing with the optimal continuation strategy¹³

$$w(t, x) = \sup_{\tau \in \mathcal{T}_{t,T}} V_{t,x}(\tau) - V_{t,x}(t).$$

Similar arguments as in the proof of Proposition 2 show that w equals the sum of marginal incentives z if the agent uses the optimal stopping rule, i.e.

$$w(t, x) = \sup_{\tau \in \mathcal{T}_{t,T}} \mathbb{E} \left[\sum_{s=t}^{\tau-1} z(s, X_s) \mid X_t = x \right].$$

Note that by definition the option value $w(t, x)$ is non-negative as the agent can always choose $\tau = t$ and stop immediately. The option value plays an important role in many economic applications of optimal stopping, especially in irreversible investment and search theory. In particular it characterizes the minimal optimal stopping rule as the first time the option value equals zero $\tau^* = \min\{t : w(t, X_t) = 0\}$.

¹¹Recall that the value V and its generalization $V_{t,x}$ are defined in Eq.s (2) and (8), respectively.

¹²Throughout the section we maintain the Standing Assumptions 1, 2, 3 and 4.

¹³In the mathematical finance literature w is called the time value of V .

The next result is a new closed form representation of the option value in terms of constrained processes. Moreover, it provides a new characterization of the optimal stopping cut-off b .

Proposition 7 (Characterization of the Option Value).

Let b be a finite cut-off. If b satisfies

$$0 = \tilde{\mathbb{E}} \left[\sum_{s=t}^{T-1} z(s, \tilde{X}_s) \mid \tilde{X}_t = b(t) \right], \quad (17)$$

for all $t \in \mathbb{T}$, then τ_b is the optimal stopping rule and the option value of waiting is the expected sum over future marginal incentives evaluated at the constrained process, i.e. for all $x \in \mathbb{R}$ and all $t \in \mathbb{T}$

$$w(t, x) = \tilde{\mathbb{E}} \left[\sum_{s=t}^{T-1} z(s, \tilde{X}_s) \mid \tilde{X}_t = x \wedge b(t) \right]. \quad (18)$$

Under Condition 1 the converse holds true as well. If τ_b is optimal, then b satisfies Eq. (17) for all $t \in \mathbb{T}$.

Proof. Assume that Eq. (17) holds. Then Theorem 1 implies that τ_b is implemented by the zero transfer which means that τ_b is the minimal optimal stopping time. Eq. (18) follows from Eq. (21).

Now assume that τ_b is a minimal optimal stopping rule. Hence, it is implemented by the zero transfer. By Theorem 1 it is also implemented by the transfer $t \mapsto \sum_{s=t}^{T-1} \tilde{\mathbb{E}}[z(s, \tilde{X}_s) \mid \tilde{X}_t = b(t)]$. Then Proposition 5 implies Eq. (17). $\square$

4.1 A New Pricing Formula for the American Option

A classical problem in finance is the pricing of American options and the description of the optimal exercise strategy. No analytic solution exists for the optimal exercise strategy and the valuation of American options. In seminal articles Kim [1990], Jacka [1991] and Carr et al. [1992] derived an integral equation that characterizes the optimal exercise boundary in a Black Scholes framework. As we show in this section, our results on optimal stopping lead to a new characterization of the exercise boundary and the price of American options as expectations over constrained price processes.

The American put option gives its holder the right - but not the obligation - to sell a share of a given stock at a prespecified strike price $K > 0$ at any time before maturity

$T - 1$. If the market is free of arbitrage the discounted price process is a martingale under some measure $\mathbb{P}$, i.e. if the price process $(X_t)_{t \in \mathbb{T}}$ is Markov then $\mathbb{E}[\frac{X_{t+1}}{1+r} | X_t = x] = x$. Given the price process of the stock and a constant interest rate r the fair price of the American option is given by

$$\sup_{\tau \in \mathcal{T}} \mathbb{E} [\mathbf{1}_{\{\tau < T\}} (1+r)^{-\tau} (K - X_\tau)] .$$

Note that if the agent does not exercise the option before time T she receives a payoff of zero. The marginal incentive is weakly increasing and given by

$$z(t, x) = \begin{cases} -\frac{rK}{(1+r)^{t+1}} & \text{if } t < T - 1 \\ -\frac{K-x}{(1+r)^{T-1}} & \text{if } t = T - 1 \end{cases} .$$

In the following proposition we apply our characterization of the optimal stopping boundary to the American option. This leads to a new simple representation of the price of the American put option as the expected discounted value of the constrained price at the time of maturity. Furthermore, it is worth noting that this characterization holds for any one-dimensional Markov price process.

Proposition 8 (Exercise Boundary of the American Option). *Any barrier b solving*

$$\frac{K}{(1+r)^t} = \mathbb{E} \left[\frac{\tilde{X}_{T-1}}{(1+r)^{T-1}} | \tilde{X}_t = b(t) \right] . \quad (19)$$

is an optimal exercise barrier and the arbitrage free price of the American option given the stock price equals x at time t equals

$$P(t, x) = \mathbb{E} \left[\frac{\tilde{X}_{T-1}}{(1+r)^{T-1-t}} | \tilde{X}_t = x \right] - x .$$

Proof. By Corollary 4, there exists a barrier b such that it is optimal to exercise the option the first time the price of the stock falls below b .¹⁴ By Proposition 7 the barrier b is

¹⁴All arguments presented in the paper before hold for a lower barrier if the marginal incentive is strictly increasing. If the marginal incentive is only weakly monotone a closer inspection of the proofs yields that Proposition 3, Proposition 5 and Corollary 4 still hold. Theorem 1 and Proposition 7 only hold in a weaker version where we do not impose minimality.

characterized by the equation

$$0 = \mathbb{E} \left[\sum_{s=t}^{T-1} z(s, \tilde{X}_s) \mid \tilde{X}_t = b(t) \right] = - \sum_{s=t}^{T-2} \frac{r K}{(1+r)^{s+1}} - \mathbb{E} \left[\frac{K - \tilde{X}_{T-1}}{(1+r)^{T-1}} \mid \tilde{X}_t = b(t) \right]$$

This implies Eq. (19). The arbitrage-free price of the American option is given by

$$\begin{aligned} P(t, x) &= (1+r)^t \sup_{\tau \geq t} \mathbb{E} \left[\mathbf{1}_{\{\tau < T\}} (1+r)^{-\tau} (K - X_\tau) \mid X_t = x \right] \\ &= (1+r)^t \left(\mathbb{E} \left[\frac{\tilde{X}_{T-1}}{(1+r)^{T-1-t}} \mid \tilde{X}_t = x \right] - \frac{K}{(1+r)^t} + \frac{K-x}{(1+r)^t} \right) \\ &= \mathbb{E} \left[\frac{\tilde{X}_{T-1}}{(1+r)^{T-1-t}} \mid \tilde{X}_t = x \right] - x. \quad \square \end{aligned}$$

5 Discussion and Conclusion

We characterized transfers as expectations over constrained processes. In setups where one can construct the constrained process explicitly, this leads to a closed form representation of the transfer. We illustrate this point by constructing constrained processes for search with and without recall (Examples 3 and 4) which leads to closed form solutions of the transfer implementing a given cut-off rule. For continuous time diffusion processes one can show that our definition of constrained processes coincides with reflected processes defined in the probability theory literature. As for the Brownian Motion and the Geometric Brownian Motion a closed form solution for the reflected process is known, a further characterization of the transfer can be obtained in those cases.

Appendix

Proof of Proposition 1. We first consider the multiplicative case. By definition

$$\mathbb{E} [|X_{t+1}^2| | X_t = x] = x^2 \mathbb{E} [\epsilon_{t+1}^2] = x^2 \sigma^2 \leq \sigma^2(1 + |x|^2)$$

and hence X is of polynomial growth of order $p = 2$. For every $x < x'$ and any increasing function ϕ we have that

$$\mathbb{E} [\phi(X_{t+1}) | X_t = x'] - \mathbb{E} [\phi(X_{t+1}) | X_t = x] = \mathbb{E} [\phi(\epsilon_{t+1}x') - \phi(\epsilon_{t+1}x)] \geq 0.$$

Setting $\phi(x) = \mathbf{1}_{\{x \geq z\}}$ yields that the process X has monotone transitions. Next, let ϕ be a continuous, polynomial growth function. Appealing to the dominated convergence theorem yields

$$\begin{aligned} & \lim_{h \rightarrow 0} |\mathbb{E} [\phi(X_{t+1}) | X_t = x + h] - \mathbb{E} [\phi(X_{t+1}) | X_t = x]| \\ & \leq \lim_{h \rightarrow 0} \mathbb{E} [|\phi(\epsilon_{t+1}(x + h)) - \phi(\epsilon_{t+1}x)|] = \mathbb{E} \left[\lim_{h \rightarrow 0} |\phi(\epsilon_{t+1}(x + h)) - \phi(\epsilon_{t+1}x)| \right] = 0. \end{aligned}$$

Hence X has continuous transitions.

Let us now turn to the additive random walk. We have

$$\mathbb{E} [|X_{t+1}^2| | X_t = x] = \mathbb{E} [(x + \epsilon_{t+1})^2] \leq 2(x^2 + \sigma^2)$$

and hence X is of polynomial growth of order $p = 2$. For every $x < x'$ and any increasing function ϕ we have that

$$\mathbb{E} [\phi(X_{t+1}) | X_t = x'] - \mathbb{E} [\phi(X_{t+1}) | X_t = x] = \mathbb{E} [\phi(x' + \epsilon_{t+1}) - \phi(x + \epsilon_{t+1})] \geq 0.$$

As in the case of a multiplicative random walk setting $\phi(x) = \mathbf{1}_{\{x \geq z\}}$ yields that the process X has monotone transitions. Finally, let ϕ be a continuous, polynomial growth function. Again we employ the dominated convergence theorem to obtain

$$\begin{aligned} & \lim_{h \rightarrow 0} |\mathbb{E} [\phi(X_{t+1}) | X_t = x + h] - \mathbb{E} [\phi(X_{t+1}) | X_t = x]| \\ & \leq \mathbb{E} \left[\lim_{h \rightarrow 0} |\phi(x + h + \epsilon_{t+1}) - \phi(x + \epsilon_{t+1})| \right] = 0. \square \end{aligned}$$

Proof of Proposition 2. Let $\tau \in \mathcal{T}$ be an arbitrary stopping rule. We will show that

$$V(\tau) = \mathbb{E} \left[\sum_{t=0}^{\tau-1} z(t, X_t) \right] + \mathbb{E} \left[g(0, X_0) + \sum_{t=1}^T h(t, X_t) \right]$$

which yields the claim. To this end we rewrite $V(\tau)$ as follows

$$V(\tau) = \mathbb{E} \left[\left(\sum_{t=0}^{\tau-1} f(t, X_t) - h(t, X_t) \right) + g(\tau, X_\tau) - h(\tau, X_\tau) \right] + \mathbb{E} \left[\sum_{t=0}^T h(t, X_t) \right]. \quad (20)$$

Using the tower property of conditional expectations we can represent the expected payoff $\mathbb{E}[g(\tau, X_\tau) - h(\tau, X_\tau)]$ as a sum of flow payoffs as follows

$$\begin{aligned} & \mathbb{E}[g(\tau, X_\tau) - h(\tau, X_\tau)] \\ &= \mathbb{E}[g(0, X_0) - h(0, X_0)] + \mathbb{E} \left[\sum_{t=0}^{\tau-1} g(t+1, X_{t+1}) - h(t+1, X_{t+1}) - g(t, X_t) + h(t, X_t) \right] \\ &= \mathbb{E}[g(0, X_0) - h(0, X_0)] \\ & \quad + \mathbb{E} \left[\sum_{t=0}^{\tau-1} \mathbb{E}[g(t+1, X_{t+1}) - h(t+1, X_{t+1}) - g(t, X_t) + h(t, X_t) | \mathcal{F}_t] \right] \\ &= \mathbb{E}[g(0, X_0) - h(0, X_0)] + \mathbb{E} \left[\sum_{t=0}^{\tau-1} \tilde{z}(t, X_t) \right], \end{aligned}$$

with

$$\begin{aligned} \tilde{z}(t, x) &= \mathbb{E}[g(t+1, X_{t+1}) - h(t+1, X_{t+1}) | X_t = x] - g(t, x) + h(t, x) \\ &= z(t, x) - f(t, x) + h(t, x). \end{aligned}$$

Then Eq. (20) implies

$$V(\tau) = \mathbb{E} \left[\left(\sum_{t=0}^{\tau-1} f(t, X_t) - h(t, X_t) + \tilde{z}(t, X_t) \right) \right] + \mathbb{E} \left[g(0, X_0) - h(0, X_0) + \sum_{t=0}^T h(t, X_t) \right],$$

which yields the claim. $\square$

Proof of Lemma 1. We proceed by backward induction. At time T the value function $v_\pi(T, x) = \pi(T)$ does not depend on x and hence satisfies the assertions of the Lemma. At

time $t < T$ the induction hypothesis yields that the function $x \mapsto P_{t,t+1}v(x)$ is continuous, non-increasing and of polynomial growth of order p . Then the dynamic programming principle (11) and the single crossing condition yield the claim for v_π . $\square$

Proof of Lemma 2. Suppose that Condition 1 holds. We show that b is unique. To this end assume that there exist two cut-offs b and $\hat{b}$ such that $\tau = \tau_b = \tau_{\hat{b}}$ and $b(t) < \hat{b}(t)$ for some $t \in \{0, \dots, T-1\}$. By conditioning on $\mathcal{F}_{t-1}$ and using the Markov property of X we obtain

$$\mathbb{P}[\tau_b = t] = \mathbb{P}[X_t \geq b(t), \tau_b > t-1] = \mathbb{E}[\mathbb{E}[\mathbf{1}_{\{X_t \geq b(t)\}} | X_{t-1}] \mathbf{1}_{\{\tau_b > t-1\}}].$$

Similar considerations for $\tau_{\hat{b}}$ yield

$$\begin{aligned} 0 &= \mathbb{P}[\tau_b = t] - \mathbb{P}[\tau_{\hat{b}} = t] = \mathbb{E}[\mathbb{E}[\mathbf{1}_{\{X_t \geq b(t)\}} - \mathbf{1}_{\{X_t \geq \hat{b}(t)\}} | X_{t-1}] \mathbf{1}_{\{\tau_b > t-1\}}] \\ &= \mathbb{E}[\mathbb{E}[\mathbf{1}_{\{b(t) \leq X_t < \hat{b}(t)\}} | X_{t-1}] \mathbf{1}_{\{\tau_b > t-1\}}]. \end{aligned}$$

By the full support condition the random variable $\mathbb{E}[\mathbf{1}_{\{b(t) \leq X_t < \hat{b}(t)\}} | X_{t-1}]$ is strictly positive. Moreover, the full support assumption implies that the event $\{\tau_b > t-1\}$ happens with positive probability. This leads to the contradiction

$$\mathbb{E}[\mathbb{E}[\mathbf{1}_{\{b(t) \leq X_t < \hat{b}(t)\}} | X_{t-1}] \mathbf{1}_{\{\tau_b > t-1\}}] > 0. \quad \square$$

Proof of Proposition 3. Denote by π the transfer implementing τ . Let τ_D be the first hitting time of X of the stopping region

$$\tau_D = \inf\{t \geq 0 \mid X_t \in D_\pi(t)\} = \inf\{t \geq 0 \mid v_\pi(t, X_t) = \pi(t)\}.$$

Then Peskir and Shiryaev [2006, Theorem 1.9] yields that τ_D is a minimal optimal stopping rule for the agent's stopping problem given the transfer π . Fix a point in time t and a value of the process $x \in D_\pi(t)$ such that it is optimal to stop. By Lemma 1 the value function v_π is non-increasing and hence for every point $x' \geq x$ we have $\pi(t) = v_\pi(t, x) \geq v_\pi(t, x')$.

By definition the value function v_π is bounded from below by $\pi(t)$ and hence we have $v_\pi(t, x') = \pi(t)$. Thus every value $x' \geq x$ is in the region $x' \in D_\pi(t)$ where it is optimal for the agent to stop. This implies that the stopping region $D_\pi(t)$ is an interval which is unbounded on the right. Again by Lemma 1 the function $x \mapsto v_\pi(t, x)$ is continuous and hence $D_\pi(t)$ is closed. Therefore there exists some $b(t) \in \bar{\mathbb{R}}$ such that $D_\pi(t) = [b(t), \infty)$.

This implies that τ_D is a cut-off rule with cut-off b . For every minimal optimal stopping rule we have $\tau = \tau_D$ almost surely and hence τ is a cut-off rule. $\square$

Proof of Proposition 4. Let b be the cut-off from Proposition 3 such that $D_\pi(t) = [b(t), \infty)$ and $\tau = \tau_b$ almost surely. Under Condition 2, since v_π is non-increasing in x , there exist $\underline{x}, \bar{x} \in \mathbb{R}$ such that

$$z(t, \bar{x}) + \mathbb{E}[v_\pi(t+1, X_{t+1}) | X_t = \bar{x}] \leq \pi(t) < z(t, \underline{x}) + \mathbb{E}[v_\pi(t+1, X_{t+1}) | X_t = \underline{x}].$$

Hence, at $\bar{x}$ it is optimal for the agent to stop and we have $\bar{x} \in D_\pi(t)$. In particular it follows that $b(t) \leq \bar{x} < \infty$. The agent strictly prefers continuing to stopping at $\underline{x}$ which means that $\underline{x} \notin D_\pi(t)$ and consequently $b(t) > \underline{x} > -\infty$. We deduce that τ is a finite cut-off rule.

Since $b(t)$ is the minimal element of $D_\pi(t)$ we have that $v_\pi(t, b(t)) = \pi(t)$. On the other hand we have $v_\pi(t, x) = z(t, x) + P_{t,t+1}v_\pi(x)$ for all $x < b(t)$. Now taking the limit $x \nearrow b(t)$ yields $v_\pi(t, b(t)) = z(t, b(t)) + P_{t,t+1}v_\pi(b(t))$, where we used Lemma 1 and Assumption 3. The definition of v_π establishes Eq. 12. $\square$

Proof of Theorem 1. First, we claim that for the transfer defined in Eq. (14) the value function defined in Eq. (9) satisfies

$$v_\pi(t, x) = \sum_{s=t}^{T-1} \tilde{P}_{t,s} z(x \wedge b(t)). \quad (21)$$

To prove Eq. (21) we proceed via backward induction. At time T we have by definition $\pi(T) = 0$ and $v_\pi(T, x) = 0$ for all $x \in \mathbb{R}$. At time $t < T$ first observe that we have by induction hypothesis

$$\mathbb{E}[v_\pi(t+1, X_{t+1}) | X_t = x] = \sum_{s=t+1}^{T-1} \mathbb{E}[\tilde{P}_{t+1,s} z(X_{t+1} \wedge b(t+1)) | X_t = x] = \sum_{s=t+1}^{T-1} \tilde{P}_{t,s} z(x).$$

The dynamic programming principle implies

$$v_\pi(t, x) = \max\left\{\pi(t), z(t, x) + \mathbb{E}[v_\pi(t+1, X_{t+1}) | X_t = x]\right\} = \max\left\{\pi(t), \sum_{s=t}^{T-1} \tilde{P}_{t,s} z(x)\right\}.$$

By the single crossing condition and Assumption 2 we obtain that the mapping $x \mapsto \tilde{P}_{t,s} z(x)$

is non-increasing for every $s > t$. As we have by definition $\pi(t) = \sum_{s=t}^{T-1} \tilde{P}_{t,s} z(b(t))$ this yields Eq. (21).

The fact that $x \mapsto z(t, x)$ is strictly decreasing implies that $x \mapsto \sum_{s=t}^{T-1} \tilde{P}_{t,s} z(x)$ is strictly decreasing as well. From Eq. (21) we conclude that the stopping region is equal to the interval $D_\pi(t) = [b(t), \infty)$. Then Peskir and Shiryaev [2006, Theorem 1.9] yields that τ_b is a minimal stopping rule implemented by π , i.e. solves Eq. (7). $\square$

Proof of Theorem 3. By Theorem 1 we can find a transfer $\phi(\tau, \hat{x}_0)$ depending only on the realized stopping time and the initial report that implements the cut-off stopping time

$$\tau(\hat{x}_0) = \min\{t \in \mathbb{T} : X_t \geq b(t, \hat{x}_0)\} \wedge T,$$

for every not necessarily truthful time zero report $\hat{x}_0$. Thus, it only remains to verify that we can make reporting truthfully at time zero incentive compatible. We do so by adding a constant $\psi(\hat{x}_0)$ depending on $\hat{x}_0$ but independent of the realized stopping time τ to the transfer ϕ .

Denote by $W(x_0, \hat{x}_0)$ the value of the agent if her initial type is x_0 and she deviates by reporting $\hat{x}_0$ and after time zero uses the optimal strategy $\tau(\hat{x}_0)$. The derivative of W with respect to her initial type x_0 if she reported to be of type $\hat{x}_0$ is given by

$$\begin{aligned} \frac{\partial}{\partial x_0} W(x_0, \hat{x}_0) &= \frac{\partial}{\partial x_0} \mathbb{E} \left[\left(\sum_{t=0}^{\tau(\hat{x}_0)-1} f(t, X_t) \right) + g(\tau(\hat{x}_0), X_{\tau(\hat{x}_0)}) \right. \\ &\quad \left. + \left(\sum_{t=\tau(\hat{x}_0)+1}^T h(t, X_t) \right) + \phi(\tau(\hat{x}_0), \hat{x}_0) \mid X_0 = x_0 \right] \\ &= \mathbb{E} \left[\sum_{t=0}^{\tau(\hat{x}_0)-1} z_x(t, X_t) \mid X_0 = x_0 \right] + \frac{\partial}{\partial x_0} V(0, x_0). \end{aligned}$$

Note that, $\frac{\partial}{\partial x_0} V(0, x_0)$ is independent of the reported type $\hat{x}_0$. As the stopping time $\tau(\hat{x}_0)$ is pathwise non-increasing in $\hat{x}_0$ by Proposition 6 and $z_x < 0$ by the single crossing assumption it follows that $\frac{\partial}{\partial x_0} W(x_0, \hat{x}_0)$ increases in $\hat{x}_0$. Define the payment

$$\psi(\hat{x}_0) = W(\hat{x}_0, \hat{x}_0) + \int_0^{\hat{x}_0} \frac{\partial}{\partial x_0} W(y, y) dy.$$

The function $\Phi(x_0, \hat{x}_0) = W(x_0, \hat{x}_0) - \psi(\hat{x}_0)$ satisfies the assumptions of Lemma 1 from Pavan, Segal, and Toikka and hence the transfer $\pi(\tau, \hat{x}_0) = \phi(\tau, \hat{x}_0) - \psi(\hat{x}_0)$ implements the allocation from Proposition 6 and (τ, π) is a revenue maximizing mechanism. $\square$

References

- S. Athey and I. Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.
- D. Baron and D. Besanko. Regulation and information in a continuing relationship. *Information Economics and Policy*, 1(3):267–302, 1984.
- M. Battaglini. Long-term contracting with markovian consumers. *American Economic Review*, pages 637–658, 2005.
- D. Bergemann and J. Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.
- S. Board. Selling options. *Journal of Economic Theory*, 136(1):324–340, 2007.
- S. Board and A. Skrzypacz. Revenue management with forward-looking buyers. *Forthcoming in Journal of Political Economy*, 2013.
- P. Carr, R. Jarrow, and R. Myneni. Alternative characterizations of american put options. *Mathematical Finance*, 2(2):87–106, 1992.
- P. Courty and H. Li. Sequential screening. *The Review of Economic Studies*, 67(4):697–717, 2000.
- R.K. Dixit and R.S. Pindyck. *Investment under Uncertainty*. Princeton university press, 2008.
- P. Eső and B. Szentes. Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies*, 74(3):705–731, 2007.
- E. Farhi and I. Werning. Insurance and taxation over the life cycle. *The Review of Economic Studies*, 80(2):596–635, 2013.
- G. Gallego and G. Van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8):999–1020, 1994.
- J. Gallien. Dynamic mechanism design for online commerce. *Operations Research*, 54(2):291–310, 2006.

- D. Garrett and A. Pavan. Managerial turnover in a changing world. *Journal of Political Economy*, 120(5):879–925, 2012.
- D. Garrett and A. Pavan. Dynamic managerial compensation: On the optimality of seniority-based schemes. *Forthcoming in Journal of Economic Theory*, 2015.
- A. Gershkov and B. Moldovanu. Learning about the future and dynamic efficiency. *American Economic Review*, 99(4):1576–1587, 2009.
- A. Gershkov and B. Moldovanu. Optimal search, learning and implementation. *Journal of Economic Theory*, 147(3):881–909, 2012.
- R. Guesnerie and J.J. Laffont. A complete solution to a class of principal-agent problems with an application to the control of a self-managed firm. *Journal of Public Economics*, 25(3):329–369, 1984.
- H. Hopenhayn and J. P. Nicolini. Optimal unemployment insurance. *Journal of Political Economy*, 105(2):412–438, 1997.
- S. Jacka. Optimal stopping and the american put. *Mathematical Finance*, 1(2):1–14, 1991.
- S. Jacka and J. Lynn. Finite-horizon optimal stopping, obstacle problems and the shape of the continuation region. *Stochastics: An International Journal of Probability and Stochastic Processes*, 39(1):25–42, 1992.
- S. M. Kakade, I. Lobel, and H. Nazerzadeh. Optimal dynamic mechanism design and the virtual pivot mechanism. *Operations Research*, 61:255–284, 2011.
- I. Kim. The analytic valuation of american options. *Review of Financial Studies*, 3(4):547–572, 1990.
- J. J. McCall. Economics of information and job search. *The Quarterly Journal of Economics*, 84(1):113–126, 1970.
- R. McDonald and D. Siegel. The value of waiting to invest. *The Quarterly Journal of Economics*, 101(4):707–727, 1986.
- H. P. McKean. Appendix: A free boundary problem for the heat equation arising from a problem in mathematical economics. *Industrial Management Review*, 6(2):32–39, 1965.

- S. Oh and Ö. Özer. Mechanism design for capacity planning under dynamic evolutions of asymmetric demand forecasts. *Management Science*, 59(4):987–1007, 2013.
- A. Pavan, I. Segal, and J. Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653, 2014.
- G. Peskir and A. Shiryaev. *Optimal Stopping and Free-Boundary Problems*, volume 10. Birkhauser, 2006.
- S. Shavell and L. Weiss. The optimal payment of unemployment insurance benefits over time. *Journal of Political Economy*, 87(6):1347–1362, 1979.
- G. J. Stigler. Information in the labor market. *Journal of Political Economy*, 70(5):94–105, 1962.
- K. T. Talluri and G. Van Ryzin. *The theory and practice of revenue management*. Springer, 2005.
- J. Tirole. From bottom of the barrel to cream of the crop: Sequential screening with positive selection. *Unpublished Manuscript*, 2014.
- W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.