

Optimal Control under a Dynamic Fuel Constraint

BY PETER BANK*

First draft: September, 2003

This version: September 1, 2004

Abstract

We present a new approach to solve optimal control problems of the monotone follower type. The key feature of our approach is that it allows to include an arbitrary dynamic fuel constraint. Instead of dynamic programming, we use the convexity of our cost functional to derive a first order characterization of optimal policies based on the Snell envelope of the objective functional's gradient at the optimum. The optimal control policy is constructed explicitly in terms of the solution to a representation theorem for stochastic processes obtained in Bank and El Karoui (2004). As an illustration, we show how our methodology allows to extend the scope of the explicit solutions obtained for the classical monotone follower problem and for an irreversible investment problem arising in economics.

AMS 2000 subject classification. 49J55, 93E20, 60H30, 91B28

Key words and phrases. Monotone follower, Snell envelope, dynamic fuel constraint

* Humboldt University of Berlin and Columbia University in the City of New York. Support of Deutsche Forschungsgemeinschaft through DFG-Research Center "Mathematics for Key Technologies" (FZT 86), and grant BA 2276/1-1 is gratefully acknowledged.

Introduction

Many optimization problems involve so called finite fuel constraints on the allowable control policies, i.e., upper bounds on the resources a control policy can use. The usual methodology to address these optimization problems is to specify a Markovian framework and to compute the problem's value function either by PDE methods based on the problem's Hamilton-Jacobi-Bellman equation or by probabilistic methods and the variational method of switching paths. In some special cases this leads to a more or less explicit solution to the optimization problem.

In any case, the constraint has so far only been specified by a *constant* upper bound for the over all amount of 'fuel' a control policy is allowed to use. *Dynamic* upper bounds, by contrast, are difficult to take into account as their introduction increases the dimensionality of the problem, making it typically impossible to solve the Hamilton-Jacobi-Bellman equation explicitly. On the other hand, it is well-known that in some problems the optimal policy for a (constant) finite fuel constraint can be derived from the optimal policy obtained when disregarding the fuel constraint completely: one just has to follow the unconstraint policy up to the moment when all fuel has been spent; see, e.g., Chow, Menaldi, and Robin (1985), Karatzas (1985) and Fleming and Soner (1993). It is thus natural to conjecture that a suitable variant of this principle should hold true for situations where a dynamic finite fuel constraint is specified by an increasing adapted process. The corroboration of this conjecture and the description of a general framework where it holds true constitute the main results of the present paper.

We consider a convex minimization problem in which a policy θ incurs the costs

$$C(\theta) = \mathbb{E} \int_0^\infty c(t, \theta_t) \mu(dt) + \mathbb{E} \int_0^\infty k_t d\theta_t \quad (\theta \in \mathcal{A}),$$

where $c(t, \cdot)$ describes the (convex) running costs and k_t the control costs at time $t \geq 0$. Our approach is based on a characterization of optimal policies in terms of first order conditions. More specifically, Theorem 1 shows that an optimal control policy will exercise control whenever its impact is maximal as measured by the Snell envelope of the cost functional's subgradient at the optimum; it also shows that actually all available fuel should be spent whenever this Snell envelope tends to decrease. The occurrence of Snell envelopes in this characterization highlights the intimate relationship between singular control and optimal stopping problems which has already been observed in Karatzas and Shreve (1984, 1985) or El Karoui and Karatzas (1988, 1991).

The construction of an optimal policy is achieved in Theorem 2 which relates the dynamic finite fuel problem with a stochastic representation theorem obtained in Bank

and El Karoui (2004). This representation theorem has found a number of other applications ranging from utility maximization to optimal stopping; we refer to Bank and Föllmer (2003) for an overview. Here it provides us with a lower bound which the optimal control policy has to respect if enough fuel is available to do so. This lower bound turns out to be independent from the fuel constraint, thus providing a universal signal process which allows one to construct optimal policies for a whole class of finite fuel problem at the same time.

As an application we provide an explicit solution to the monotone follower problem for Lévy processes with quadratic cost functional in the spirit of Beneš, Shepp, and Witsenhausen (1980/81). We also illustrate how explicit solutions obtained for singular control problems without any fuel constraint, as obtained, e.g., in Kobila (1993) and Scheinkman and Zariphopoulou (2001), can actually be used to describe optimal policies for problems with a dynamic fuel constraint.

Acknowledgment. The author is greatly indebted to Thaleia Zariphopoulou for suggesting this problem.

1 The General Control Problem

A well known problem in stochastic optimization is the problem to control the motion of a particle so as to keep it as close to the origin as possible over some period of time. In the formulation as a monotone follower problem suggested and analyzed by Karatzas and Shreve (1984), one considers a model where the dynamics of the uncontrolled particle is given by standard Brownian motion W and where the control θ is an increasing adapted process θ which specifies the downward displacement of the particle caused by the control. Hence, in this case, the controlled particle would follow the dynamics $W_t + \theta_t$ ($t \geq 0$). The cost incurred by a control policy θ can, for instance, be described as

$$C(\theta) = \mathbb{E} \int_0^\infty \delta e^{-\delta t} \frac{1}{2} (W_t - \theta_t)^2 dt.$$

and one could start studying the optimization problem to minimize C subject to, e.g., a finite fuel constraint on the control θ .

More generally, let $(\Omega, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ be a filtered probability space satisfying the usual conditions of right-continuity and completeness. Controls θ are given by increasing, left-continuous adapted processes starting at $\theta_0 = \underline{\vartheta} \in \mathbb{R}$. We shall impose a dynamic finite fuel constraint, specified by an increasing adapted process $\bar{\vartheta}$ with left continuous

paths and values in $[\underline{\vartheta}, +\infty]$. The class of admissible controls is therefore

$$\mathcal{A} \triangleq \{\theta \text{ incr., left-cont., adapted with } \underline{\vartheta} = \theta_0 \leq \theta_t \leq \bar{\vartheta}_t \text{ for all } t \geq 0 \text{ } \mathbb{P}\text{-a.s.}\}.$$

Remark 1.1 *Note that the lower bound $\underline{\vartheta}$ is assumed to be a real constant, not a process. Assuming a dynamic lower bound would mean that a minimum amount of control must have been exercised up to each point in time, a natural, yet much more demanding extension which is beyond the scope of the present paper.*

The costs incurred by a control policy will be composed by running costs and control costs. The running costs are described by a measurable random field

$$c : \Omega \times [0, \infty) \times \mathbb{R} \rightarrow \mathbb{R}$$

and a positive optional random measure $\mu = \mu(\omega, dt)$ on the time axis satisfying the following convexity and regularity assumption.

Assumption 1 (i) *The measure μ is atomless and has full support $\text{supp } \mu = [0, +\infty)$ almost surely.*

(ii) *For any $(\omega, t) \in \Omega \times [0, +\infty)$, the mapping $\vartheta \mapsto c(\omega, t, \vartheta)$ is strictly convex with continuous derivative $c'(\omega, t, \vartheta) = \frac{\partial}{\partial \vartheta} c(\omega, t, \vartheta)$ increasing from $c'(\omega, t, -\infty) = -\infty$ to $c'(\omega, t, +\infty) = +\infty$.*

(iii) *For $\vartheta \in \mathbb{R}$ fixed, $(\omega, t) \mapsto c(\omega, t, \vartheta)$ is progressively measurable and $\mathbb{P} \otimes \mu$ -integrable.*

(iv) *The process $(\omega, t) \mapsto \inf_{\vartheta \in [\underline{\vartheta}, \bar{\vartheta}_t]} c(t, \vartheta)$ is $\mathbb{P} \otimes \mu$ -integrable.*

The control costs are described by a stochastic process k with the following properties.

Assumption 2 *The process k is optional and continuous in expectation with $k_\infty = 0$. Moreover, the family of random variables $(\int_0^\infty k_t^- d\theta_t, \theta \in \mathcal{A})$ is bounded in $L^1(\mathbb{P})$.*

These assumptions allow us to consider the cost functional

$$C(\theta) = \mathbb{E} \int_0^\infty c(t, \theta_t) \mu(dt) + \mathbb{E} \int_0^\infty k_t d\theta_t \quad (\theta \in \mathcal{A}).$$

Remark 1.2 *Full support of μ ensures that strict convexity transfers from the random field c to our cost functional C . Our assumptions on the derivative of c will be used when applying a representation theorem obtained in Bank and El Karoui (2004); see Section 3. Integrability of $c(\cdot, \vartheta)$ for $\vartheta \in \mathbb{R}$ means that the decision not to intervene at all will not cause infinite costs. Integrability of $\inf_{\vartheta \in [\underline{\vartheta}, \bar{\vartheta}]} c(\cdot, \vartheta)$ ensures that our minimization problem has a finite value. For the same reason we assume L^1 -boundedness of $(\int_0^\infty k_t^- d\theta_t, \theta \in \mathcal{A})$, which amounts to requiring that the negative (!) ‘costs’ of exercising control must not be ‘too large’.*

The general optimization problem we shall be concerned with in this paper can now be stated as follows:

$$(1) \quad \text{Minimize } C(\theta) \text{ over } \theta \in \mathcal{A}.$$

Remark 1.3 *Observe that our introductory example would be accommodated in this setting by choosing*

$$c(\omega, t, \vartheta) \triangleq \frac{1}{2} (W_t(\omega) - \vartheta)^2, \quad \mu(dt) = \delta e^{-\delta t} dt \quad \text{and} \quad k \equiv 0.$$

Observe furthermore that this setting can also accommodate the monotone follower problems studied in Chow, Menaldi, and Robin (1985), Karatzas (1985), Karatzas and Shreve (1984), as well as the irreversible investment problems solved in Kobila (1993) (see their Criterion (3.2) and Condition (5.1)), Scheinkman and Zariphopoulou (2001) (see Section 4.2), and Karatzas and Baldursson (1997). Settings not covered by our framework include Chiarolla (1997) and Jacka (1999, 2002) since their cost functional is specified in terms of the controlled system instead of the cumulatively exercised control. We also do not cover the ‘cheap monotone follower’ of Chiarolla and Hausmann (1994) as they allow for two-dimensional controls, as well as the finite fuel problem of Beneš, Shepp, and Witsenhausen (1980/81) and Karatzas and Shreve (1985) who consider two-sided controls.

Notation and Conventions. All (in)equalities between random variables are meant to hold true in the \mathbb{P} -a.s.-sense. We shall let \mathcal{T} denote the set of all stopping times and we use $\mathcal{T}(I)$ to denote the class of stopping times almost surely taking values in a given random set I such as, e.g., $I = [S, +\infty]$ with $S \in \mathcal{T}$. A supremum over an empty set is defined to be $\sup \emptyset \triangleq -\infty$. Intervals $[a, b]$ with $b < a$ are interpreted as the empty set. We also put $x^+ \triangleq x \vee 0 = \max\{x, 0\}$ and $x^- \triangleq (-x)^+$.

2 First Order Conditions for Optimality

In this section, we are going to provide a first order characterization of optimal control policies for Problem (1). While the main avenue of approach to achieve this characterization is classical, we need to be a little bit careful to ensure that our Assumptions 1 and 2 suffice to deduce all the integrability requirements we need along the way.

Our first step is to note that the convex functional C is supported by the subgradients

$$(2) \quad \nabla C(\theta)_S \triangleq \mathbb{E} \left[\int_S^\infty c'(t, \theta_t) \mu(dt) \middle| \mathcal{F}_S \right] + k_S \quad (S \in \mathcal{T})$$

in the following sense:

Lemma 2.1 *For any $\theta \in \mathcal{A}$, the optional process $\nabla C(\theta)$ of (2) is well-defined and $\nabla C(\theta)^-$ is $\mathbb{P} \otimes d\theta$ -integrable. If θ satisfies $C(\theta) < +\infty$, also $\nabla C(\theta)^+$ is $\mathbb{P} \otimes d\theta$ integrable and $\nabla C(\theta)$ satisfies the subgradient property*

$$C(\theta') - C(\theta) \geq \mathbb{E} \int_0^\infty \nabla C(\theta)_s d(\theta'_s - \theta_s) \quad \text{for any } \theta' \in \mathcal{A} \text{ with } C(\theta') < +\infty.$$

PROOF :

- (i) As $c'(t, \theta_t) \geq c'(t, \vartheta) \in L^1(\mathbb{P} \otimes \mu)$ by convexity and $\mathbb{P} \otimes \mu$ -integrability of $c(t, \vartheta)$ for $\vartheta \in \mathbb{R}$, the conditional expectation appearing in (2) is well-defined as a random variable taking values in $(-\infty, +\infty]$. As for $\mathbb{P} \otimes d\theta$ -integrability of $\nabla C(\theta)^-$, we note that

$$c'(t, \theta_t)^-(\theta_t - \vartheta) \leq c(t, \vartheta) - \inf_{\vartheta \in [\vartheta, \bar{\vartheta}_t]} c(t, \vartheta) \in L^1(\mathbb{P} \otimes \mu)$$

by Assumption 1 (iii) and (iv), and this yields

$$\mathbb{E} \int_0^\infty \int_s^\infty c'(t, \theta_t)^- \mu(dt) d\theta_s = \mathbb{E} \int_0^\infty c'(t, \theta_t)^-(\theta_t - \vartheta) \mu(dt) < +\infty.$$

by Fubini's theorem. By Assumption 2,

$$\mathbb{E} \int_0^\infty k_s^- d\theta_s < +\infty,$$

and it follows that

$$\nabla C(\theta)_s^- \leq \mathbb{E} \left[\int_s^\infty c'(t, \theta_t)^- \mu(dt) \middle| \mathcal{F}_s \right] + k_s$$

is $\mathbb{P} \otimes d\theta_s$ integrable.

(ii) We next prove $\mathbb{P} \otimes d\theta$ -integrability of $\nabla C(\theta)$ for θ with $C(\theta) < +\infty$. By convexity of $c(t, \cdot)$, we have

$$c'(t, \theta_t)(\theta_t - \underline{\vartheta}) \leq c(t, \theta_t) - c(t, \underline{\vartheta}).$$

The latter difference is $\mathbb{P} \otimes \mu$ -integrable for θ with $C(\theta) < +\infty$ by Assumption 1, and therefore

$$\mathbb{E} \int_0^\infty \int_s^\infty c'(t, \theta_t)^+ \mu(dt) d\theta_s = \mathbb{E} \int_0^\infty c'(t, \theta_t)^+(\theta_t - \underline{\vartheta}) \mu(dt) < +\infty.$$

(iii) Let us now prove the subgradient estimate for θ' with $C(\theta') < +\infty$. It follows that the left side in the convexity estimate

$$c(t, \theta'_t) - c(t, \theta_t) \geq c'(t, \theta_t)(\theta'_t - \theta_t) = c'(t, \theta_t)(\theta'_t - \underline{\vartheta}) - c'(t, \theta_t)(\theta_t - \underline{\vartheta})$$

is $\mathbb{P} \otimes \mu$ -integrable. By (i) the same is true for the last term $c'(t, \theta_t)(\theta_t - \underline{\vartheta})$. It follows as in (ii) that $\int_0^\infty c'(t, \theta'_t)^+ \mu(dt)$ is $\mathbb{P} \otimes d\theta'$ -integrable, and so the expectation

$$\mathbb{E} \left\{ \int_0^\infty \int_s^\infty c'(t, \theta_t) \mu(dt) + k_s \right\} d(\theta'_s - \theta_s) = \mathbb{E} \int_0^\infty \nabla C(\theta)_s d(\theta'_s - \theta_s) \in [-\infty, +\infty)$$

is well-defined and indeed not larger than $C(\theta') - C(\theta)$.

□

Let us denote by $\mathbb{S}(\theta)$ the lower Snell envelope

$$\mathbb{S}(\theta)_S = \operatorname{ess\,inf}_{T \in \mathcal{T}([S, \infty))} \mathbb{E} [\nabla C(\theta)_T | \mathcal{F}_S] \quad (S \in \mathcal{T})$$

of $\nabla C(\theta)$. This is an optional process taking values in $(-\infty, 0]$ since $\nabla C(\theta)_\infty = 0$ by definition and

$$\inf_{T \in \mathcal{T}} \mathbb{E} \nabla C(\theta)_T \geq \mathbb{E} \int_0^\infty c'(t, \underline{\vartheta}) \wedge 0 \mu(dt) > -\infty.$$

We shall use $M(\theta)$ and $A(\theta)$ to denote the martingale and predictable increasing part in the Doob-Meyer decomposition $\mathbb{S}(\theta) = M(\theta) + A(\theta)$ of the submartingale $\mathbb{S}(\theta)$.

After these preliminaries, we now can give the following characterization of optimal policies in terms of first order conditions:

Theorem 1 *Under Assumptions 1 and 2, a control policy $\theta^* \in \mathcal{A}$ is optimal for Problem (1) iff*

(i) θ^* is flat off $\{\nabla C(\theta^*) = \mathbb{S}(\theta^*)\}$ and

(ii) $A(\theta^*)$ is flat off $\{\theta^* = \bar{\vartheta}\}$.

Remark 2.2 Condition (i) requires that control should be exercised only when its marginal impact on future costs is maximal. Condition (ii) reveals that all fuel should be spent at moments where the maximal expected marginal impact tends to decrease.

The proof of this theorem uses the following two lemmata and will be given at the end of this section.

Lemma 2.3 A plan $\theta^* \in \mathcal{A}$ is optimal for Problem (1) iff for any $\theta \in \mathcal{A}$ the process $\nabla C(\theta^*)^-$ is $\mathbb{P} \otimes d\theta$ -integrable and we have

$$\mathbb{E} \int_0^\infty \nabla C(\theta^*)_s d\theta_s^* \leq \mathbb{E} \int_0^\infty \nabla C(\theta^*)_s d\theta_s.$$

In particular, $\nabla C(\theta^*)$ is $\mathbb{P} \otimes d\theta^*$ -integrable in this case.

PROOF : Sufficiency of the first order condition follows as usual from the subgradient estimate $C(\theta) - C(\theta^*) \geq \mathbb{E} \int_0^\infty \nabla C(\theta^*)_s d(\theta - \theta^*)_s$ provided by Lemma 2.1. To ensure that this lemma is indeed applicable we need to prove that $C(\theta^*) < +\infty$. So let us introduce the passage times $T^\vartheta \triangleq \inf\{t \geq 0 \mid \theta_t^* \geq \vartheta\}$ and note that the corresponding stopped control policies $(\theta_{t \wedge T^\vartheta}^*)_{t \geq 0}$ are contained in \mathcal{A} . By convexity of $c(t, \cdot)$, we have

$$\begin{aligned} (3) \quad \mathbb{E} \int_0^{T^\vartheta} c(t, \theta_t^*) \mu(dt) + \mathbb{E} \int_0^{T^\vartheta} k_s d\theta_s^* \\ \leq \mathbb{E} \int_0^{T^\vartheta} c(t, \underline{\vartheta}) \mu(dt) + \mathbb{E} \int_0^{T^\vartheta} c'(t, \theta_t^*)(\theta_t^* - \underline{\vartheta}) \mu(dt) + \mathbb{E} \int_0^{T^\vartheta} k_s d\theta_s^*. \end{aligned}$$

For $\vartheta \uparrow \infty$, the first term on the right side of (3) converges to $C(\underline{\vartheta})$ by dominated convergence. The remaining two terms sum up to

$$\mathbb{E} \int_0^\infty \nabla C(\theta^*)_s d\theta_{s \wedge T^\vartheta}^* = \mathbb{E} \int_0^{T^\vartheta} \nabla C(\theta^*)_s \wedge 0 d\theta_s^* \longrightarrow \mathbb{E} \int_0^\infty \nabla C(\theta^*)_s \wedge 0 d\theta_s^* < +\infty$$

since condition (ii) entails in particular that the measure $d\theta^*$ must be concentrated on the set $\{\nabla C(\theta^*) \leq 0\}$ almost surely. It follows that the right side is uniformly bounded from above in ϑ .

As for the left side of (3), note that for $\vartheta \uparrow \infty$ we have

$$\int_0^{T^\vartheta} k_s^- d\theta_s^* = \int_0^\infty k_s^- d\theta_{s \wedge T^\vartheta}^* \longrightarrow \int_0^\infty k_s^- d\theta_s^* \quad \mathbb{P}\text{-a.s. and in } L^1(\mathbb{P})$$

by L^1 -boundedness of the family $(\int_0^\infty k_s^- d\theta_s, \theta \in \mathcal{A})$. This allows us to use monotone convergence to conclude that

$$\exists \lim_{\vartheta \uparrow \infty} \mathbb{E} \int_0^{T^\vartheta} k_s d\theta_s^* = \mathbb{E} \int_0^\infty k_s d\theta_s^* > -\infty.$$

Similarly, $c(t, \theta_t^*) \geq \inf_{\vartheta \in [\underline{\vartheta}, \bar{\vartheta}_t]} c(t, \vartheta) \in L^1(\mathbb{P} \otimes \mu)$ yields

$$\exists \lim_{\vartheta \uparrow \infty} \mathbb{E} \int_0^{T^\vartheta} c(t, \theta_t^*) \mu(dt) = \mathbb{E} \int_0^\infty c(t, \theta_t^*) \mu(dt) > -\infty.$$

It thus follows that the left side of (3) converges to $C(\theta^*) \in \mathbb{R}$.

Necessity could be derived using arguments from the calculus of variations. We shall pursue an alternative approach here and verify the validity of the first order conditions *ex post*, i.e., directly from our construction of the unique optimal policy. This construction will not use the necessity part of this lemma. \square

The preceding lemma shows that an optimal policy for our convex optimization problem (1) is also a solution to some linear minimization problem. Solutions to this kind of problem are characterized by the following result:

Lemma 2.4 *Let $\phi \leq 0$ be an optional process of class (D) which is continuous in expectation with $\phi_\infty = 0$. Let ψ denote its lower Snell envelope*

$$\psi_S = \operatorname{ess\,inf}_{T \in \mathcal{T}([S, \infty])} \mathbb{E}[\phi_T | \mathcal{F}_S] \quad (S \in \mathcal{T})$$

and consider the corresponding Doob-Meyer decomposition $\psi = M + A$ into a uniformly integrable martingale M and an increasing, predictable process A with $A_0 = 0$.

Then the solutions θ^ to the linear optimization problem*

$$\text{Minimize } \mathbb{E} \int_0^\infty \phi_s d\theta_s \text{ subject to } \theta \in \mathcal{A}$$

are characterized by the properties

- (i) θ^* is flat off $\{\phi = \psi\}$ and
- (ii) A is flat off $\{\theta^* = \bar{\vartheta}\}$.

PROOF : As ϕ is continuous in expectation, so is its Snell envelope ψ . In particular, ψ is an RCLL process and its predictable compensator A has continuous paths almost surely. For any control policy $\theta \in \mathcal{A}$, we thus can derive the following estimate:

$$\begin{aligned}
 \mathbb{E} \int_{[0,\infty)} \phi_t d\theta_t &\geq \mathbb{E} \int_{[0,\infty)} \psi_t d\theta_t = \mathbb{E} \int_{[0,\infty)} M_t d\theta_t + \mathbb{E} \int_{[0,\infty)} A_t d\theta_t \\
 (4) \qquad \qquad \qquad &= \mathbb{E} - A_\infty \theta_\infty + \int_{[0,\infty)} A_t d\theta_t = -\mathbb{E} \int_{(0,\infty]} (\theta_t - \underline{\vartheta}) dA_t \\
 &\geq -\mathbb{E} \int_{(0,\infty]} (\bar{\vartheta}_t - \underline{\vartheta}) dA_t.
 \end{aligned}$$

Indeed, the first estimate follows from $\phi \geq \psi$ and the third equality is due to the martingale property of M and due to our assumption $\phi_\infty = 0$ which implies $0 = \psi_\infty = M_\infty + A_\infty$. The last equality follows by partial integration and the last estimate holds true because $dA \geq 0$ and $\theta \leq \bar{\vartheta}$ by admissibility of θ .

It is easy to see that any θ^* satisfying (i) and (ii) will yield equality everywhere in the above estimate. On the other hand, these two conditions are also necessary for a plan θ^* to minimize $\mathbb{E} \int_0^\infty \phi_s d\theta_s$ over $\theta \in \mathcal{A}$. This follows readily from the above estimate in conjunction with the identity

$$(5) \qquad \qquad \qquad \inf_{\theta \in \mathcal{A}} \mathbb{E} \int_0^\infty \phi_s d\theta_s = -\mathbb{E} \int_0^\infty (\bar{\vartheta}_t - \underline{\vartheta}) dA_t.$$

To prove this identity, we introduce for $n = 1, 2, \dots$ the sequence of stopping times

$$\begin{aligned}
 T_0^n &\triangleq \inf\{t \geq 0 \mid \phi_t = \psi_t\}, \\
 T_j^n &\triangleq \inf\{t \geq T_{j-1}^n \mid \phi_t = \psi_t, \bar{\vartheta}_t > \bar{\vartheta}_{T_{j-1}^n} + 1/n\} \quad (j = 1, 2, \dots)
 \end{aligned}$$

and consider the admissible control policy

$$\theta_t^n \triangleq \sum_{j=0}^{\infty} \bar{\vartheta}_{T_j^n + 1} 1_{(T_j^n, T_{j+1}^n]}(t) \quad (t \geq 0).$$

For $\theta = \theta^n$ we have equality in the first part of estimate (4) and so we obtain

$$\mathbb{E} \int_0^\infty \phi_s d\theta_s^n = -\mathbb{E} \int_{(0,\infty]} (\theta_t^n - \underline{\vartheta}) dA_t.$$

It follows from general results on optimal stopping, that dA is supported by the set $\{t \geq 0 \mid \phi_t = \psi_t\}$ almost surely. By definition of the stopping times T_j^n ($j = 0, 1, \dots$) this entails that almost surely $\theta_t^n \geq \bar{\vartheta}_t - 1/n$ for dA -a.e. t . We thus can conclude that

$$\mathbb{E} \int_0^\infty \phi_s d\theta_s^n = -\mathbb{E} \int_{(0,\infty]} (\theta_t^n - \underline{\vartheta}) dA_t \leq -\mathbb{E} \int_{(0,\infty]} (\bar{\vartheta}_t - \underline{\vartheta}) dA_t + \frac{1}{n} \mathbb{E} A_\infty.$$

For $n \uparrow \infty$ this establishes the desired identity (5), accomplishing our proof. \square

It is now easy to give the

Proof of Theorem 1 Let $\theta^* \in \mathcal{A}$ be a policy such that θ^* is flat off $\{\nabla C(\theta)^* = \psi(\theta^*)\}$ and $A(\theta^*)$ is flat off $\{\theta^* = \bar{\vartheta}\}$. Then $\nabla C(\theta^*)^-$ is $\mathbb{P} \otimes d\theta^*$ -integrable by Lemma 2.1 and θ^* satisfies the conditions of Lemma 2.4 with $\phi \triangleq \nabla C(\theta^*) \wedge 0$. It follows that

$$-\infty < \mathbb{E} \int_0^\infty \nabla C(\theta^*) d\theta^* = \mathbb{E} \int_0^\infty \nabla C(\theta^*) \wedge 0 d\theta^* \leq \mathbb{E} \int_0^\infty \nabla C(\theta^*) \wedge 0 d\theta$$

for all $\theta \in \mathcal{A}$. In particular, $\nabla C(\theta^*)^-$ is $\mathbb{P} \otimes d\theta$ -integrable for all $\theta \in \mathcal{A}$, and we can use Lemma 2.3 to conclude optimality of the policy θ^* in Problem (1). Hence, our first order conditions are indeed sufficient.

Necessity of the first order conditions also follows from combining Lemmas 2.3 and 2.4. We merely need to note that $\nabla C(\theta^*)^-$ is of class (D) by $\mathbb{P} \otimes \mu$ -integrability of $c'(t, \vartheta) \leq c'(t, \theta_t^*)$ and Assumption 2. \square

Remark 2.5 For problems without fuel constraint ($\bar{\vartheta} \equiv +\infty$), the integrability assertion in Lemma 2.3 implies in particular that for an optimal policy $\theta^* \in \mathcal{A}$ the gradient $\nabla C(\theta^*)$ has to be non-negative, i.e.,

$$k_S \geq -\mathbb{E} \left[\int_S^{+\infty} c'(t, \theta_t) \mu(dt) \middle| \mathcal{F}_S \right].$$

Moreover, condition (i) in Theorem 1 implies that equality must hold true in the above relation whenever S is a time of intervention. This is in accordance with the first order characterizations obtained for such problems in Bertola (1998) or Bank and Riedel (2001).

3 Construction of an Optimal Policy

In this section, we shall show how to use the first order characterization of the optimal policy provided by Theorem 1 in order to construct the solution to the finite fuel problem (1). The construction will be given in terms of a progressively measurable random process κ specifying a lower bound which the optimal control should respect granted enough fuel is left to do so. This lower bound is characterized as the optional solution κ to the representation problem

$$(6) \quad k_S = -\mathbb{E} \left[\int_S^\infty c'(t, \sup_{s \in [S, t]} \kappa_s) \mu(dt) \middle| \mathcal{F}_S \right] \quad \text{for any } S \in \mathcal{T}.$$

Assumptions 1 and 2 ensure existence of a solution to this problem; see Theorem 3 in Bank and El Karoui (2004).

Theorem 2 *Under Assumptions 1 and 2, the unique minimizer for Problem (1) is given by*

$$\theta_t^* \triangleq \sup_{s \in [0,t)} \{ \kappa_s \wedge \bar{\vartheta}_s \} \vee \underline{\vartheta} \quad (t \geq 0),$$

where κ is the optional process solving the representation problem (6).

Remark 3.1 *Note that the process κ does not depend on the bounds $\underline{\vartheta}$, $\bar{\vartheta}$ describing the set of admissible policies. As a consequence, the solution κ to the representation problem (6) can be viewed as a universal signal yielding optimal policies for a whole class of finite fuel problems.*

PROOF : Let us verify that the policy $\theta^* \in \mathcal{A}$ satisfies the first order conditions derived in Theorem 1.

We first compute the lower Snell envelope $\mathbb{S}(\theta^*)$ of $\nabla C(\theta^*)$. To this end, consider $S, T \in \mathcal{T}$ with $S \leq T$ and note that by definition of θ^* and (6) we have

$$\begin{aligned} \mathbb{E}[\nabla C(\theta^*)_T \mid \mathcal{F}_S] &= \mathbb{E} \left[\int_T^\infty \{ c'(t, \sup_{s \in [0,t)} \{ \kappa_s \wedge \bar{\vartheta}_s \} \vee \underline{\vartheta}) - c'(t, \sup_{s \in [T,t)} \kappa_s) \} \mu(dt) \mid \mathcal{F}_S \right] \\ &\geq \mathbb{E} \left[\int_T^\infty \{ c'(t, \sup_{s \in [0,t)} \{ \kappa_s \wedge \bar{\vartheta}_s \}) - c'(t, \sup_{s \in [T,t)} \kappa_s) \} \mu(dt) \mid \mathcal{F}_S \right] \\ &\geq \mathbb{E} \left[\int_T^\infty \{ c'(t, \sup_{s \in [0,t)} \{ \kappa_s \wedge \bar{\vartheta}_s \}) - c'(t, \sup_{s \in [S,t)} \kappa_s) \} \mu(dt) \mid \mathcal{F}_S \right] \\ &\geq \mathbb{E} \left[\int_S^\infty \{ c'(t, \sup_{s \in [0,t)} \{ \kappa_s \wedge \bar{\vartheta}_s \}) - c'(t, \sup_{s \in [S,t)} \kappa_s) \} \wedge 0 \mu(dt) \mid \mathcal{F}_S \right]. \end{aligned}$$

Note that the last expression does no longer depend on $T \geq S$, thus providing a lower bound for the Snell envelope $\mathbb{S}(\theta^*)$. In fact, it coincides with this envelope since we have equality in any of the above estimates for $T = T_S \triangleq \inf\{t \geq S \mid \kappa_t > \bar{\vartheta}_{t+}\}$. This is easy to see for the first of these estimates as, by definition of T_S , $\sup_{s \in [0, T_S]} \{ \kappa_s \wedge \bar{\vartheta}_s \} \geq \bar{\vartheta}_{T_S+} \geq 0$, whence $\sup_{s \in [0,t)} \{ \kappa_s \wedge \bar{\vartheta}_s \} \geq 0$ for $t > T_S$. Equality for the second estimate can be deduced from the observation that T_S is a point of increase for $\sup_{s \in [S,t)} \kappa_s$ which yields $\sup_{s \in [S,t)} \kappa_s = \sup_{s \in [T_S,t)} \kappa_s$ for $t \in (T_S, \infty]$. Finally, equality for the third estimate holds true as

$$c'(t, \sup_{s \in [0,t)} \{ \kappa_s \wedge \bar{\vartheta}_s \}) - c'(t, \sup_{s \in [S,t)} \kappa_s) \quad \text{is} \quad \begin{cases} \geq 0 & \text{for } t \in (S, T_S], \\ \leq 0 & \text{for } t \in (T_S, +\infty], \end{cases}$$

again by definition of T_S . For later use, let us also note here that the stopping time T_S is actually the largest stopping time which attains $\mathbb{S}(\theta^*)_S =$

$\text{ess inf}_{T \in \mathcal{T}([S, \infty])} \mathbb{E} [\nabla C(\theta^*)_T \mid \mathcal{F}_S]$ since the above difference is always non-positive and actually strictly negative on some non-trivial time interval starting at T_S whenever $T_S < +\infty$.

Let us now verify the flat off conditions characterizing optimal plans as described in Theorem 1. If S is point of increase for θ^* , we have $\kappa_S \geq \theta_{S+}^* = \sup_{s \in [0, T_S]} \{\kappa_s \wedge \bar{\vartheta}_s\} \geq 0$ and so $\sup_{s \in [S, t]} \kappa_s \geq \sup_{s \in [0, t]} \{\kappa_s \wedge \bar{\vartheta}_s\} \geq 0$ for $t \in (S, \infty]$. This allows us to conclude

$$\begin{aligned} \mathbb{S}(\theta^*)_S &= \mathbb{E} \left[\int_S^\infty \{c'(t, \sup_{s \in [0, t]} \{\kappa_s \wedge \bar{\vartheta}_s\}) - c'(t, \sup_{s \in [S, t]} \kappa_s)\} \wedge 0 \mu(dt) \middle| \mathcal{F}_S \right] \\ &= \mathbb{E} \left[\int_S^\infty \{c'(t, \sup_{s \in [0, t]} \{\kappa_s \wedge \bar{\vartheta}_s\}) - c'(t, \sup_{s \in [S, t]} \kappa_s)\} \mu(dt) \middle| \mathcal{F}_S \right] \\ &= \mathbb{E} \left[\int_S^\infty c'(t, \sup_{s \in [0, t]} \{\kappa_s \wedge \bar{\vartheta}_s\}) \vee \underline{\vartheta} \mu(dt) \middle| \mathcal{F}_S \right] = \nabla C(\theta^*)_S. \end{aligned}$$

So, θ^* is indeed flat off $\{\mathbb{S}(\theta^*) = \nabla C(\theta^*)\}$.

If, on the other hand, S is a point of increase for the predictable compensator $A(\theta^*)$ of $\mathbb{S}(\theta^*)$ then, by classical results on optimal stopping, the only stopping time in $\mathcal{T}([S, +\infty])$ attaining $\mathbb{S}(\theta^*)_S$ is S itself. In particular, the maximal such stopping time T_S determined above coincides with S : $T_S = S$, i.e., $S = \inf\{t \geq S \mid \kappa_t > \bar{\vartheta}_{t+}\}$ almost surely. This implies $\kappa_S \geq \bar{\vartheta}_{S+}$ and so $\theta_{S+}^* \geq \bar{\vartheta}_{S+}$ \mathbb{P} -a.s. We deduce that almost surely $\theta_t^* = \bar{\vartheta}_t$ for any joint point of continuity t for both θ^* and $\bar{\vartheta}$, i.e., for all but at most countably many points t .

Hence, in order to deduce the desired flat off condition $\theta_t^* = \bar{\vartheta}_t$ for $dA(\theta^*)$ -a.e. t , it now suffices to note that $A(\theta^*)$ has continuous sample paths. This, however, holds true because $\nabla C(\theta^*)^-$ and thus also $\mathbb{S}(\theta^*)$ are continuous in expectation. \square

4 Applications

As an immediate consequence of Theorem 2 we obtain an extension of a result in Karatzas (1985) from the Brownian case to our larger class of finite fuel problems:

Corollary 4.1 *The optimal control policy in Problem (1) with finite fuel ($\bar{\vartheta} \equiv \text{const.}$) is just the optimal control policy with infinite fuel ($\bar{\vartheta} \equiv +\infty$) abandoned after all fuel has been exhausted.* \square

To obtain a more general result, let us note the following version of the dynamic programming principle:

Corollary 4.2 *For each stopping time $S \in \mathcal{T}$, the process*

$$\theta_t^S \triangleq \sup_{s \in [S, t)} \{ \kappa_s \wedge \bar{\vartheta}_s \} \vee \underline{\vartheta}$$

attains

$$(7) \quad \operatorname{ess\,inf}_{\theta \in \mathcal{A}, \theta_S = \underline{\vartheta}} \mathbb{E} \left[\int_S^\infty c(t, \theta_t) \mu(dt) + \int_S^\infty k_t d\theta_t \middle| \mathcal{F}_S \right].$$

In particular, $\theta_{S+}^S = \underline{\vartheta} \vee \limsup_{t \searrow S} \kappa_t \wedge \bar{\vartheta}_{S+}$ describes the initial policy decision one has to take when starting to minimize costs as of time S .

PROOF : Define $\bar{\vartheta}_t^S \triangleq \underline{\vartheta}$ for $t \leq S$ and $\bar{\vartheta}_t^S \triangleq \bar{\vartheta}_t$ for $t > S$ and note that by Theorem 2 θ^S is the optimal policy in the set of admissible policies \mathcal{A}^S corresponding to $\bar{\vartheta}^S$ instead of $\bar{\vartheta}$. It thus attains $\inf_{\theta \in \mathcal{A}^S} C(\theta)$ and also

$$\inf_{\theta \in \mathcal{A}, \theta_S = \underline{\vartheta}} \mathbb{E} \left[\int_S^\infty c(t, \theta_t) \mu(dt) + \int_S^\infty k_t d\theta_t \right] = \mathbb{E} \left[\int_S^\infty c(t, \theta_t^S) \mu(dt) + \int_S^\infty k_t d\theta_t^S \right].$$

It is easy to see that this infimum is actually the expectation of $\operatorname{ess\,inf}$ in (7). This, however, allows us to conclude our assertion since this essential infimum is always less than $\mathbb{E} \left[\int_S^\infty c(t, \theta_t^S) \mu(dt) + \int_S^\infty k_t d\theta_t^S \middle| \mathcal{F}_S \right]$ almost surely. \square

The preceding corollary can be used in two ways. On the hand, it shows that the process κ of (6) can be used to describe optimal solutions not only as of time 0, but actually from any arbitrary initial time $S \in \mathcal{T}$. On the other hand, it allows us to deduce κ (at least partially) from the policies attaining (7). This observation yields the following corollary.

Corollary 4.3 *Let K be an optional process such that, for $S \in \mathcal{T}$, $K_S = \theta_{S+}^S$ is the initial value of the optimal policy for problem (7) when working under the fuel constraint $\theta_t^S \in [\underline{\vartheta}, \bar{\vartheta}_t]$ ($t \geq S$). Then the optimal policy for Problem (1) with fuel constraint $\bar{\vartheta}' \leq \bar{\vartheta}$ is given by*

$$\theta_t' \triangleq \sup_{s \in [0, t)} \{ K_s \wedge \bar{\vartheta}'_s \} \vee \underline{\vartheta}.$$

In particular, the solutions to the problem without fuel constraint ($\bar{\vartheta} \equiv +\infty$) suffice to determine the optimal policies for the problem with an arbitrary dynamic fuel constraint.

4.1 Monotone follower problems

Let us now come back to the special case of a monotone follower problem studied by Karatzas and Shreve (1984) which we used to motivate the formulation of our general finite fuel problem (1) in Section 1: We wish to determine a control policy $\theta \in \mathcal{A}$ which minimizes

$$C(\theta) \triangleq \mathbb{E} \int_0^\infty \delta e^{-\delta t} \frac{1}{2} (W_t - \theta_t)^2 dt,$$

where W is a standard Brownian motion. It follows from Theorem 2 that we can solve this problem explicitly for an arbitrary $\bar{\vartheta}$ by providing a solution to the representation problem (6). In our present setting, this amounts to finding a progressively measurable κ such that

$$(8) \quad \mathbb{E} \left[\int_S^\infty \delta e^{-\delta t} \sup_{s \in [S, t)} \kappa_s dt \middle| \mathcal{F}_S \right] = e^{-\delta S} W_S \quad \text{for all } S \in \mathcal{T}.$$

It is intuitively clear (and has been established formally in Karatzas (1985)) that the optimal policy consists in reflecting the controlled Brownian motion at a certain threshold c . This suggests to consider the Ansatz $\kappa_s \triangleq W_s - c$ for some constant $c \in \mathbb{R}$. Indeed, this Ansatz can easily be vindicated: We can use the independence and time-homogeneity of the increments of W to conclude that (8) will be satisfied if we choose

$$c \triangleq \mathbb{E} \int_0^\infty \delta e^{-\delta t} \sup_{s \in [0, t)} W_s dt,$$

In fact, looking back, we see that the above reasoning will apply not only for Brownian motion but actually for any Lévy process satisfying suitable integrability properties:

Corollary 4.4 *Let X be a Lévy process such that $\mathbb{E} \int_0^\infty \delta e^{-\delta t} X_t^2 dt < +\infty$. Then the optimal policy for the monotone follower problem*

$$\text{Minimize } C(\theta) \triangleq \mathbb{E} \int_0^\infty \delta e^{-\delta t} \frac{1}{2} (X_t - \theta_t)^2 dt \text{ over } \theta \in \mathcal{A}$$

is given by

$$\theta_t^* = \sup_{s \in [0, t)} \{(X_s - c) \wedge \bar{\vartheta}_s\} \vee \vartheta$$

where $c \triangleq \mathbb{E} \int_0^\infty \delta e^{-\delta t} \sup_{s \in [0, t)} X_s dt < +\infty$.

PROOF : We merely have to note that $c < +\infty$ follows from the square-integrability condition on X and Doob's maximal inequality for the martingale $(X_t - t\mathbb{E}X_1)_{t \geq 0}$. \square

Remark 4.5 *While Beneš, Shepp, and Witsenhausen (1980/81) study a practically identical cost functional they allow for downward and upward displacement of the particle: Controls merely have to be of bounded variation. As they show, in this situation the amount of fuel left becomes crucial for the optimal control decision so that there is no longer a universal process like κ describing the optimal policy. Indeed, one will accept larger distances of the controlled processes from the origin with little fuel left than with a lot.*

More generally the above approach will allow us to explicitly describe optimal control policies whenever the representation problem (8) can be solved explicitly for a given process W , not necessarily Brownian motion. This is indeed possible for a large class of diffusions as shown in Bank and Föllmer (2003).

If, on the other hand, one wishes to consider a non-quadratic cost functional under a dynamic finite fuel constraint, we can use Corollary 4.2 to reduce the construction of an optimal policy to the unconstrained case with infinite fuel and make use of the results of Chow, Menaldi, and Robin (1985) or Karatzas (1985).

4.2 Irreversible investments

Let us finally illustrate how Corollary 4.4 can be used to extend the closed form solutions obtained for certain irreversible investment problems in Kobila (1993) and Scheinkman and Zariphopoulou (2001) to incorporate a dynamic finite fuel constraint.

The paper Kobila (1993) studies the problem to maximize a reward functional of the type

$$R(\theta) \triangleq \mathbb{E} \int_0^\infty e^{-\delta t} \Pi(X_t, \theta_t) dt + \int_0^\infty k_t d\theta_t$$

where X is a geometric Brownian motion and $\Pi = \Pi(x, \vartheta) : (0, +\infty) \times \mathbb{R} \rightarrow \mathbb{R}$ describes the reward function. Apart from a number of technical conditions, Π is assumed to be strictly concave in ϑ ; see Condition (5.1) in Kobila (1993) and note that our control cost process k does not depend on the control in contrast to the setting chosen in Kobila (1993). All increasing, left-continuous processes θ with $\theta \geq \underline{\vartheta} = \theta_0$ are considered admissible controls, i.e., $\bar{\vartheta} \equiv +\infty$.

Taking a dynamic programming approach the authors set up and explicitly solve the Hamilton-Jacoby-Bellman equation for this problem. It turns out that the optimal policy consists in keeping the problem's state process (X, θ) away from a 'forbidden' region \mathcal{R} of the form

$$\mathcal{R} = \{(x, \vartheta) \in (0, +\infty) \times \mathbb{R} \mid \phi(x) > \vartheta\}$$

for an explicitly given continuous function ϕ . In particular, when starting at time $S \in \mathcal{T}$ in $\theta_S^S = \underline{\vartheta} \in \mathbb{R}$ the optimal policy requires an initial jump to $K_S = \underline{\vartheta} \vee \phi(X_S)$, i.e., to the minimal $\vartheta \geq \underline{\vartheta}$ such that $(X_S, \vartheta) \notin \mathcal{R}$.

It now follows from Corollary 4.4 that the region \mathcal{R} computed in Kobila (1993) can actually be used to solve the same problem with an arbitrary dynamic fuel constraint $\bar{\vartheta} \neq +\infty$: The optimal policy consists still in keeping the state process away from the region \mathcal{R} , at least as long as enough fuel is left to do so. If this is not the case, one has to wait until further supply of fuel becomes available (i.e. until $\bar{\vartheta}$ increases) and then use this fuel to move the state process as close as possible to the complement of \mathcal{R} .

The problem studied in Scheinkman and Zariphopoulou (2001) can be viewed as the problem in Kobila (1993) with an additional finite fuel constraint: $\underline{\vartheta} = 0$ and $\bar{\vartheta} = 1$. Given our previous observation it is now very easy to reduce the explicit solution computed in Scheinkman and Zariphopoulou (2001) to the results of Kobila (1993).

References

- BANK, P., AND N. EL KAROUÏ (2004): “A stochastic representation theorem with applications to optimization and obstacle problems,” *Ann. Probab.*, 32(1B), 1030–1067.
- BANK, P., AND H. FÖLLMER (2003): “American options, multi-armed bandits, and optimal consumption plans: a unifying view,” in *Paris-Princeton Lectures on Mathematical Finance, 2002*, vol. 1814 of *Lecture Notes in Math.*, pp. 1–42. Springer, Berlin.
- BANK, P., AND F. RIEDEL (2001): “Optimal Consumption Choice with Intertemporal Substitution,” *Annals of Applied Probability*, 3, 750–788.
- BENEŠ, V. E., L. A. SHEPP, AND H. S. WITSENHAUSEN (1980/81): “Some solvable stochastic control problems,” *Stochastics*, 4(1), 39–83.
- BERTOLA, G. (1998): “Irreversible Investment,” *Research in Economics*, 52, 3–37.
- CHIAROLLA, M. B. (1997): “Singular stochastic control of a singular diffusion process,” *Stochastics Stochastics Rep.*, 62(1-2), 31–63.
- CHIAROLLA, M. B., AND U. G. HAUSSMANN (1994): “The optimal control of the cheap monotone follower,” *Stochastics Stochastics Rep.*, 49(1-2), 99–128.

- CHOW, P. L., J.-L. MENALDI, AND M. ROBIN (1985): “Additive control of stochastic linear systems with finite horizon,” *SIAM J. Control Optim.*, 23(6), 858–899.
- EL KAROUI, N., AND I. KARATZAS (1988): “Probabilistic aspects of finite-fuel, reflected follower problems,” *Acta Appl. Math.*, 11(3), 223–258.
- (1991): “A new approach to the Skorohod problem, and its applications,” *Stochastics Stochastics Rep.*, 34(1-2), 57–82.
- FLEMING, W. H., AND H. M. SONER (1993): *Controlled Markov processes and viscosity solutions*, vol. 25 of *Applications of Mathematics (New York)*. Springer-Verlag, New York.
- JACKA, S. (2002): “Avoiding the origin: a finite-fuel stochastic control problem,” *Ann. Appl. Probab.*, 12(4), 1378–1389.
- JACKA, S. D. (1999): “Keeping a satellite aloft: two finite fuel stochastic control models,” *J. Appl. Probab.*, 36(1), 1–20.
- KARATZAS, I. (1985): “Probabilistic aspects of finite-fuel stochastic control,” *Proc. Nat. Acad. Sci. U.S.A.*, 82(17), 5579–5581.
- KARATZAS, I., AND F. M. BALDURSSON (1997): “Irreversible Investment and Industry Equilibrium,” *Finance and Stochastics*, 1, 69–89.
- KARATZAS, I., AND S. E. SHREVE (1984): “Connections between optimal stopping and singular stochastic control. I. Monotone follower problems,” *SIAM J. Control Optim.*, 22(6), 856–877.
- (1985): “Connections between optimal stopping and singular stochastic control. II. Reflected follower problems,” *SIAM J. Control Optim.*, 23(3), 433–451.
- KOBILA, T. Ø. (1993): “A class of solvable stochastic investment problems involving singular controls,” *Stochastics Stochastics Rep.*, 43(1-2), 29–63.
- SCHEINKMAN, J. A., AND T. ZARIPHPOULOU (2001): “Optimal Environmental Management in the Presence of Irreversibilities,” *Journal of Economic Theory*, 96(1–2), 180–207.