

On Gittins' Index Theorem in Continuous Time

BY PETER BANK* AND CHRISTIAN KÜCHLER**

Abstract

We give a new and comparably short proof of Gittins' index theorem for dynamic allocation problems of the multi-armed bandit-type in continuous time under minimal assumptions. This proof gives a complete characterization of optimal allocation strategies as those policies which follow the current leader among the Gittins indices while ensuring that a Gittins index is at an all-time low whenever the associated project is not worked on exclusively. The main tool is a representation property of Gittins index processes which allows us to show that these processes can be chosen to be pathwise lower semi-continuous from the right and quasi lower semi-continuous from the left. Both regularity properties turn out to be crucial for our characterization and the construction of optimal allocation policies.

AMS 2000 subject classification. 60G07, 60G40, 60H25, 91B16, 91B28.

Key words and phrases. Gittins index, multi-armed bandits, representation theorem

* Columbia University in the City of New York. Support of Deutsche Forschungsgemeinschaft through grant BA 2276/1-1, and of NSF through grant DMS-0505021 is gratefully acknowledged.

** Humboldt-Universität zu Berlin. Support of BMBF through project 03SF0312E is gratefully acknowledged.

Introduction

Multi-armed bandit problems arise when a limited amount of resources, usually thought of as working time, has to be allocated to a number of different projects. These offer dynamically varying, random rewards when worked on, and so one has to trade-off those projects which currently offer the highest rewards against those projects which possibly will yield even better payoffs after they have been worked on for some time.

Typically, these problems are rather difficult (if not even impossible) to solve due to their high dimensionality and the complex interdependencies between the projects. For independent projects, however, it was shown first by Gittins and Jones (1974) that there exists a project-specific dynamic performance measure, later called the Gittins index of a project, such that optimal allocations are obtained from an index policy which (essentially) amounts to focussing at each point only on those projects which exhibit a maximal Gittins index. This celebrated result was subsequently extended from Gittins' and Jones' original discrete-time, Markovian framework to a completely general continuous-time setting; see, e.g., Whittle (1980), Varaiya, Walrand, and Buyukkoc (1985), Mandelbaum (1987), Weber (1992), El Karoui and Karatzas (1994, 1997), and Kaspi and Mandelbaum (1995, 1998).

These accounts are based on two equivalent definitions of the Gittins index. First, there is the definition as an indifference threshold with respect to early retirement of a project like, e.g., in Whittle (1980). Also El Karoui and Karatzas (1994, 1997) use this for their martingale approach in continuous time. Second, there is Gittins' forward characterization of his index as the properly discounted maximal expected future rewards per expected time used for obtaining these. In continuous time, this characterization is the starting point for the excursion-theoretic approach pursued by Kaspi and Mandelbaum (1995, 1998).

By contrast, our approach to Gittins' index theorem is based on a characteristic representation property which relates the Gittins index to the accumulated future expected rewards from a given project; see Corollary 2.1. This property, in conjunction with a novel partial-integration argument, not only allows for a comparably short proof of Gittins' index theorem under minimal assumptions, see Theorem 2, it actually also allows us to give necessary and sufficient conditions for optimality, a result already conjectured by Kaspi and Mandelbaum (1998). Another important point in our analysis is the lower semi-continuity of Gittins index processes. This new regularity result is crucial not only

for the continuous-time construction of Gittins indices as proper optional processes, but it actually turns out to be indispensable even for the existence of optimal allocation strategies. Our Theorem 1 shows how this regularity follows from the representation property of Gittins indices. This result also complements the representation theorem of Bank and El Karoui (2004).

Acknowledgment: We are grateful to Mihai Sirbu for pointing out an error in an earlier version of our proof for Lemma 4.1. We also would like to thank the anonymous referee whose careful comments helped to clarify the presentation.

1 Problem Formulation

We shall follow Mandelbaum (1987) in order to formalize our allocation problem as a multi-parameter control problem. To this end, let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space on which we have N real-valued processes $h^p = (h_s^p, s \geq 0)$ satisfying

$$\mathbb{E} \int_0^\infty e^{-\alpha s} |h_s^p| ds < \infty \quad (p = 1, \dots, N).$$

We will interpret h_s^p as the reward rate from project p after we have spent already $s \geq 0$ units of time on this particular project. When working on a project we typically will learn more about its future prospects. This gain of information is modeled by a project-specific filtration $\mathbb{F}^p = (\mathcal{F}_s^p)_{s \geq 0}$ which satisfies the usual conditions of right-continuity and completeness. We assume each h^p is progressively measurable with respect to the filtration \mathbb{F}^p associated with the corresponding project.

An allocation strategy is given by a collection \mathbf{S} of increasing processes $(S_t^p, t \geq 0)$ ($p = 1, \dots, N$) specifying how much time should be spent on each project up to calendar time t . Feasibility is ensured by the constraint

$$\sum_{p=1}^N S_t^p = t \text{ for all } t \geq 0.$$

In addition, we require the strategy to be nonanticipative in the sense that for each time $t \geq 0$ the vector $\mathbf{S}_t = (S_t^p, p = 1, \dots, N)$ can be viewed as a stopping time for the multi-parameter filtration generated by \mathbb{F}^p ($p = 1, \dots, N$):

$$\{\mathbf{S}_t \leq \mathbf{s}\} \triangleq \bigcap_{p=1}^N \{S_t^p \leq s^p\} \in \mathcal{F}_\mathbf{s} \triangleq \bigvee_{p=1}^N \mathcal{F}_{s^p}^p \text{ for any } \mathbf{s} \triangleq (s^1, \dots, s^p) \in \mathbb{R}_+^N.$$

The class of all such allocation strategies \mathbf{S} will be denoted by \mathcal{S} .

With this notation we can now introduce the rewards generated by a dynamic allocation strategy $\mathbf{S} \in \mathcal{S}$ as

$$\mathcal{R}(\mathbf{S}) \triangleq \sum_{p=1}^N \int_0^\infty e^{-\alpha t} h_{S_t^p}^p dS_t^p,$$

where $\alpha > 0$ is a constant discount factor. The optimization problem we want to study amounts to maximizing the expected rewards:

$$\text{Maximize } \mathbb{E}\mathcal{R}(\mathbf{S}) \text{ over } \mathbf{S} \in \mathcal{S}.$$

A naive approach to the above allocation problem would be to always work on the project which currently exhibits the highest reward rate. It is easy to see that, while such a myopic approach may typically produce reasonable results, it need not be optimal in general. Indeed, despite the discounting of future rewards, it may be worthwhile to engage in a project with presently low rewards if progress in this project is likely to lead to significantly higher rewards in the (not too far) future which will then compensate for the foregoing of *initially* higher rewards obtainable from other projects. However, if the rewards are decreasing for all projects, it is well-known that the myopic strategy is indeed optimal as shown by the following lemma whose proof is given in the appendix for the sake of completeness:

Lemma 1.1 *In the deteriorating case where all reward processes h^p ($p = 1, \dots, N$) are right-continuous and decreasing, an allocation strategy \mathbf{S} maximizes $\mathbb{E}\mathcal{R}(\mathbf{S})$ over $\mathbf{S} \in \mathcal{S}$ if and only if it is myopically following the leading reward:*

$$h_{S_t^p}^p = \max_{q=1, \dots, N} h_{S_t^q}^q \text{ whenever } dS_t^p > 0 \quad (p = 1, \dots, N).$$

Here and in the sequel it is convenient to use the following notation.

Notation. *Let $(A_t, t \geq 0)$ be left-continuous and increasing. Following Mandelbaum (1987), we call $t \geq 0$ a point of right increase for S and write*

$$dA_t > 0 \quad \text{if} \quad A_t < A_u \text{ for all } u > t.$$

Similarly, we use expressions like $dA_t < dt$ to indicate that $t \geq 0$ is a point of right increase for $t - A_t$.

Points of right-increase fit neatly together with semi-continuity from the right as pointed out in the following lemma which we will use for our characterization of optimal allocation strategies and whose proof we defer to the appendix.

Lemma 1.2 *Let $(A_t, t \geq 0)$ be left-continuous and increasing. If $(k_t, t \geq 0)$ is non-negative with lower semi-right-continuous paths, i.e., $\liminf_{s \searrow t} k_s \geq k_t$ for all $t \geq 0$, we have*

$$k_t = 0 \text{ whenever } dA_t > 0 \quad \text{if and only if} \quad \int_0^\infty k_t dA_t = 0.$$

2 Gittins index processes and their representation properties

As described by Lemma 1.1, focussing myopically on current reward rates will be optimal in the case of deteriorating rewards. It is well-known, however, that in general one should focus instead on the projects' Gittins indices. These are dynamic performance measures which at each point in time indicate the future prospects of each project, taking into account all the information available gathered about the project sofar.

While the mathematical construction of these Gittins index processes is rather straight forward in discrete time, it becomes surprisingly subtle in continuous time for technical as well as conceptual reasons. Indeed, as is well-known, optimal policies crucially depend on the running infimum of these processes and so these processes sample paths must be sufficiently regular for such infima to become right-continuous and adapted processes. This technical difficulty will be overcome—without any further assumptions on the filtrations—by our Theorem 1 which we shall use to construct Gittins indices as the unique optional solution with lower semi-right-continuous paths of some stochastic representation problem. This pathwise lower semi-continuity from the right will be needed also for our *necessary and sufficient* conditions for optimality of an allocation strategy; cf. the proof of Theorem 2 in the following section.

Even more important on a conceptual level is the fact that one also needs these Gittins index processes to be quasi lower semi-continuity from the left—this is in order to ensure the mere existence of (optimal) index strategies as will become clear in Section 4 below; see in particular Remark 4.2.

2.1 A stochastic representation problem

In order to address the aforementioned regularity issues of Gittins index processes, we shall investigate more deeply their extraordinarily useful connection with certain stochastic representation problems which was already observed but not further inves-

tigated in Bank and El Karoui (2004). The following theorem supplements their main results accordingly; to clarify the structural assumptions, it is formulated in a slightly more general form than necessary for our later applications.

Theorem 1 *Let X be an optional process of class (D) and continuous in expectation with $X_\infty = 0$, defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$ satisfying the usual hypotheses of right-continuity and completeness. Consider an optional random measure μ on $[0, \infty)$ with finite expected total mass $\mathbb{E}\mu([0, \infty)) < \infty$, full support $\text{supp } \mu = [0, \infty)$ and no atoms almost surely.*

Then there exists a unique optional process ξ with lower semi-right-continuous paths, i.e.,

$$\liminf_{u \searrow s} \xi_u \geq \xi_s \text{ for all } s \geq 0 \text{ } \mathbb{P}\text{-a.s.},$$

and the property that for every stopping time S we have $1_{[S, \infty)}(s) \inf_{u \in [S, s]} \xi_u \in L^1(\mathbb{P} \otimes \mu(ds))$ with

$$(1) \quad X_S = \mathbb{E} \left[\int_S^\infty \inf_{u \in [S, s]} \xi_u \mu(ds) \middle| \mathcal{F}_S \right].$$

In addition, ξ is also quasi lower semi-continuous from the left:

$$\liminf_{u \nearrow S} \xi_u \geq \xi_S \text{ for every finite predictable stopping time } S,$$

and we have the alternative equivalent descriptions of ξ as

$$(2) \quad \xi_S = \inf \left\{ m \in \mathbb{R} \middle| \begin{array}{l} m \geq \mathbb{E} [X_S + m\mu((S, T]) | \mathcal{F}_S] \text{ } \mathbb{P}\text{-a.s.} \\ \text{for all stopping times } T \geq S \end{array} \right\}$$

and

$$(3) \quad \xi_S = \text{ess sup}_{T > S \text{ stopping time}} \frac{\mathbb{E} [X_T - X_S | \mathcal{F}_S]}{\mathbb{E} [\mu((S, T]) | \mathcal{F}_S]}$$

at every finite stopping time S .

PROOF : Existence of an optional ξ with (1) follows with $f(t, l) \triangleq -l$ and $\xi \triangleq -L$ from Theorem 3 and Remark 2.1 in Bank and El Karoui (2004). From their Equation (23) we also infer that ξ can actually be chosen via the following threshold principle (2) or, equivalently, via

$$\xi_S = \inf \{ m \in \mathbb{R} \mid Y_S^m = X_S \} \text{ for every finite stopping time } S \geq 0$$

where

$$Y_S^m = \operatorname{ess\,sup}_{T \geq S \text{ stopping time}} \mathbb{E}[X_T + m\mu([S, T)) | \mathcal{F}_S] \quad (m \in \mathbb{R}).$$

Our assumptions on X and μ ensure that these Snell envelopes have right-continuous sample paths. Moreover, they can be chosen in such a way that $m \mapsto Y_s^m(\omega)$ is continuous and increasing for any $\omega \in \Omega, s \geq 0$. Granted the claimed lower semi-continuity from the right (proven below), Theorem 1 in Bank and El Karoui (2004) shows that ξ must satisfy (3) at every finite stopping time S and so ξ is uniquely determined as an optional process. So property (2) implies (1) which in turn implies (3). It is also readily checked that conversely (3) implies (2).

By Proposition 2 in Dellacherie and Lenglart (1982), to show that ξ is indeed pathwise lower semi-continuity from the right, it suffices to prove $\lim_n \xi_{S_n} \geq \xi_S$, for any sequence of stopping times $S^n \downarrow S$ such that $\zeta \triangleq \lim_n \xi_{S_n}$ exists almost surely. Now, for such a sequence of stopping times, we indeed have that, for $\varepsilon > 0$,

$$Y_S^{\zeta+\varepsilon} = \lim_n Y_{S_n}^{\zeta+\varepsilon} = \lim_n X_{S_n} = X_S.$$

As $\varepsilon > 0$ was arbitrary, this suffices to conclude $\zeta \geq \xi_S$ from the above threshold description of ξ_S .

In order to prove ξ 's quasi lower semi-continuity from the left, let $S^n \uparrow S$ be an announcing sequence of stopping times for the finite predictable stopping time S . Note that since $S^n < S$ on $\{S > 0\}$ we have

$$\eta^n(s) \triangleq 1_{[S^n, \infty)}(s) \inf_{u \in [S^n, s]} \xi_u \rightarrow 1_{[S, \infty)}(s) \inf_{u \in [S, s]} \xi_u \wedge \liminf_{u \nearrow S} \xi_u \quad \mathbb{P} \otimes \mu\text{-a.e. for } n \uparrow \infty.$$

So, granted $\mathbb{P} \otimes \mu$ -uniform integrability of the family $(\eta^n, n = 1, 2, \dots)$, we can invoke the representation property of ξ and use Lebesgue's theorem to conclude

$$\begin{aligned} \mathbb{E}X_S &= \lim_n \mathbb{E}X_{S^n} = \lim_n \mathbb{E} \int_0^\infty \eta^n(s) \mu(ds) = \mathbb{E} \int_S^\infty \inf_{u \in [S, s]} \xi_u \wedge \liminf_{u \nearrow S} \xi_u \mu(ds) \\ &\leq \mathbb{E} \int_S^\infty \inf_{u \in [S, s]} \xi_u \mu(ds) = \mathbb{E}X_S. \end{aligned}$$

It follows that equality must in fact hold true everywhere in the above calculation and so $\liminf_{u \uparrow S} \xi_u \geq \xi_S$ due to the lower semi-right-continuity of ξ .

So to conclude our claim it remains to prove that the family $(\eta^n, n = 1, 2, \dots)$ is $\mathbb{P} \otimes \mu$ -uniformly integrable. In fact, since $\eta^n(s) \geq 0 \wedge \inf_{u \in [0, s]} \xi_u \in L^1(\mathbb{P} \otimes \mu(ds))$, it

suffices to prove this for $(\eta^n \vee 0, n = 1, 2, \dots)$ or, equivalently,

$$\limsup_{c \uparrow \infty} \sup_n \mathbb{E} \int_0^\infty \eta^n 1_{\{\eta^n > c\}} d\mu = 0.$$

To this end, let us introduce $T_c^n \triangleq \inf\{t \geq S^n \mid \xi_t \leq c\}$, $T_c \triangleq \inf\{t \geq S \mid \xi_t \leq c\}$ and note that $T_c^n \searrow S^n$ as $c \uparrow \infty$ and $T_c^n \leq T_c$ almost surely, with equality holding true on $\{T_c^n \geq S\}$. Moreover, we have

$$\mathbb{E} \int_0^\infty \eta^n 1_{\{\eta^n > c\}} d\mu = \mathbb{E} \int_{S^n}^{T_c^n} \inf_{u \in [S^n, s]} \xi_u \mu(ds) = \mathbb{E} X_{S^n} - \mathbb{E} X_{T_c^n},$$

where for the last identity we used again the representation property of ξ as well as the observation that $\inf_{u \in [S^n, T_c^n]} \xi_u = \xi_{T_c^n}$ by definition of T_c^n and lower semi-continuity from the right. It follows that for $m = 1, 2, \dots$, we obtain the estimate

$$(4) \quad \sup_{n \geq m} \mathbb{E} \int_0^\infty \eta^n 1_{\{\eta^n > c\}} d\mu \leq \sup_{n \geq m} \mathbb{E} X_{S^n} - \inf_{\substack{S^m \leq T \leq S \\ T \text{ a stopping time}}} \mathbb{E} X_T + \mathbb{E} \int_S^{T_c} \inf_{u \in [S, s]} \xi_u \mu(ds).$$

Indeed, due to the previous identity we can write

$$\mathbb{E} \int_0^\infty \eta^n 1_{\{\eta^n > c\}} d\mu = \mathbb{E} X_{S^n} - \mathbb{E} X_{T_c^n \wedge S} + \mathbb{E} [X_S - X_{T_c^n \vee S}]$$

where, for $n \geq m$, $T \triangleq T_c^n \wedge S$ is a stopping time between S^m and S as in the above infimum. Moreover, since $T_c^n = T_c$ on $\{T_c^n > S\}$ are level passage times for ξ after time S , employing the representation property in the same way as above allows us to rewrite the last expectation as

$$\mathbb{E} [X_S - X_{T_c^n \vee S}] = \mathbb{E} \int_S^{T_c^n} \inf_{u \in [S, s]} \xi_u \mu(ds) 1_{\{T_c^n > S\}} \leq \mathbb{E} \int_S^{T_c} \inf_{u \in [S, s]} \xi_u \mu(ds),$$

where the final estimate holds true since the infimum over ξ is nonnegative over $[S, T_c]$ for $c \geq 0$ by definition of T_c .

Now, since X is continuous in expectation and $S^m \uparrow S$, the supremum and the infimum on the right side of (4) both converge to $\mathbb{E} X_S$ as $m \uparrow \infty$ and so their difference will be smaller than any given $\varepsilon > 0$ for m large enough. As $T_c \searrow S$ for $c \uparrow \infty$, we can choose c so large that the last term in (4) becomes less than ε along with the terms $\mathbb{E} \int_0^\infty \eta^n 1_{\{\eta^n > c\}} d\mu$ for n less than our previously chosen m . This accomplishes our proof. \square

2.2 Gittins index processes

We now can use the above representation theorem in order to construct the desired Gittins index processes for every project $p = 1, \dots, N$. To this end, we choose X in Theorem 1 to be the \mathbb{F}^p -optional projection of $(\int_t^\infty e^{-\alpha s} h_s^p ds, t \geq 0)$ and we set $\mu(ds) \triangleq e^{-\alpha s} ds$ and obtain the following immediate corollary:

Corollary 2.1 *For every project $p = 1, \dots, N$, there exists a unique \mathbb{F}^p -optional process g^p which is lower semi-continuous from the right and quasi lower semi-continuous from the left that satisfies the following three equivalent properties:*

(i) *For every \mathbb{F}^p -stopping time $S \geq 0$,*

$$g_S^p = \inf \left\{ m \in \mathbb{R} \mid m \geq \mathbb{E} \left[\int_S^T e^{-\alpha(s-S)} h_s^p ds + e^{-\alpha(T-S)} m \mid \mathcal{F}_S^p \right] \mathbb{P}\text{-a.s.} \right. \\ \left. \text{for all } \mathbb{F}^p\text{-stopping times } T \geq S \right\}.$$

(ii) *For every \mathbb{F}^p -stopping time $S \geq 0$,*

$$g_S^p = \operatorname{ess\,sup}_{T > S \text{ } \mathbb{F}^p\text{-stopping time}} \frac{\mathbb{E} \left[\int_S^T e^{-\alpha s} h_s^p ds \mid \mathcal{F}_S^p \right]}{\mathbb{E} \left[\int_S^T e^{-\alpha s} ds \mid \mathcal{F}_S^p \right]}.$$

(iii) *For every \mathbb{F}^p -stopping time $S \geq 0$, the process $1_{[S, \infty)}(s) e^{-\alpha s} \inf_{u \in [S, s]} g_u^p$ is $\mathbb{P} \otimes ds$ -integrable and*

$$\mathbb{E} \left[\int_S^\infty e^{-\alpha s} h_s^p ds \mid \mathcal{F}_S^p \right] = \mathbb{E} \left[\int_S^\infty e^{-\alpha s} \inf_{u \in [S, s]} g_u^p ds \mid \mathcal{F}_S^p \right].$$

□

The above process g^p is called the *Gittins index process* of project p . Its quasi lower semi-continuity from the left indicates that Gittins indices weigh the gathering of information about future rewards against the size of current rewards: if new information about a project can be learned at a predictable time, the Gittins index shortly before will be based on the best possible outcome of this information in order to fully reveal the benefit of the additional information which can be gathered with just a little extra effort; once this information has been obtained, the index can only drop from (and will never rise above) its inevitably overoptimistic previous level. This is of course not to

say that Gittins indices cannot jump upwards at all—it only means that upward jumps come as a complete surprise, i.e., at a totally inaccessible stopping time.

The pathwise lower semi-continuity from the right is important for the construction of Gittins indices as proper *optional* processes rather than as a collection of random variables each of which is only determined up to a \mathbb{P} -null set. While the latter construction is suggested by the classical descriptions (i) and (ii) of Gittins indices described in the preceding corollary, it has the drawback that one cannot really consider the running infima of Gittins index processes needed in the discussion of index policies.

As we shall see, only the above property (iii) is needed for our full characterization of the solution to our allocation problem; this representation property was noted before, but actually not used to prove Gittins' index theorem in, e.g., Mandelbaum (1987) and El Karoui and Karatzas (1994). Property (i) in the above corollary is the classical characterization of Gittins indices as the minimal 'retirement reward' which will make one stop a project immediately; see Gittins and Jones (1974), Whittle (1980) and Weber (1992). Property (ii) is Gittins' (1979) forward characterization of his index. It is at the heart of the excursion theoretic proofs in Kaspi and Mandelbaum (1995) and Kaspi and Mandelbaum (1998); see also Varaiya, Walrand, and Buyukkoc (1985) and Presman and Sonin (1990).

3 Gittins' Index Theorem

Let us now focus on this paper's main result, the complete characterization of optimal continuous-time allocation strategies as the class of so-called Gittins index strategies. This result was conjectured by Kaspi and Mandelbaum (1998), whose excursion theoretic approach already showed that index strategies are indeed optimal. As we shall see, our novel approach via the representation property of Gittins indices allows for a rather short proof of this result which reveals in addition that the index strategies are in fact the *only* optimal allocation policies.

A crucial feature of a project's Gittins index process is that it can be calculated without considering any of the other projects. As a well-known consequence, this allows one to reduce the generally multi-dimensional allocation problem to the problem of separately calculating the N Gittins index processes g^p ($p = 1, \dots, N$). This, of course, can only hold true provided that our projects are not inextricably intertwined with each other. We thus make the following standard assumption:

Assumption 3.1 *The σ -algebras \mathcal{F}_∞^p ($p = 1, \dots, N$) are independent.*

We can now give the following ramification of Gittins' celebrated index theorem:

Theorem 2 *Under Assumption 3.1, an allocation strategy $\mathbf{S} \in \mathcal{S}$ is optimal if and only if it is an index strategy, i.e., if and only if almost surely*

(i) *the strategy follows the leading Gittins index in the sense that*

$$g_{S_t^p}^p = \max_{q=1, \dots, N} g_{S_t^q}^q \text{ whenever } dS_t^p > 0 \quad (p = 1, \dots, N),$$

and

(ii) *whenever a project is not engaged full-time, its Gittins index is at an all-time low:*

$$g_{S_t^p}^p = \inf_{u \in [0, S_t^p]} g_u^p \text{ whenever } dS_t^p < dt \quad (p = 1, \dots, N).$$

PROOF : Consider an arbitrary allocation strategy \mathbf{S} . We first use partial integration to rewrite the rewards from project $p = 1, \dots, N$ as follows:

$$\begin{aligned} \int_0^\infty e^{-\alpha t} h_{S_t^p}^p dS_t^p &= - \int_0^\infty e^{-\alpha(t-S_t^p)} d \int_{S_t^p}^\infty e^{-\alpha s} h_s^p ds \\ &= \int_0^\infty e^{-\alpha s} h_s^p ds + \int_0^\infty \left\{ \int_{S_t^p}^\infty e^{-\alpha s} h_s^p ds \right\} de^{-\alpha(t-S_t^p)}. \end{aligned}$$

Hence the expected rewards from project p are

$$\mathbb{E} \int_0^\infty e^{-\alpha t} h_{S_t^p}^p dS_t^p = \mathbb{E} \int_0^\infty e^{-\alpha s} h_s^p ds + \mathbb{E} \int_0^\infty \mathbb{E} \left[\int_{S_t^p}^\infty e^{-\alpha s} h_s^p ds \middle| \mathcal{F}_{\mathbf{S}_t} \right] de^{-\alpha(t-S_t^p)},$$

where the passage from the $\{\dots\}$ -term to its optional projection with respect to $\mathbb{F}^{\mathbf{S}} = (\mathcal{F}_{\mathbf{S}_t})_{t \geq 0}$ is justified since $e^{-\alpha(t-S_t^p)}$ ($t \geq 0$) is a nonincreasing process adapted to this filtration; see, e.g., Théorème (1.33) in Jacod (1979) or Theorem (VI.20.6)(i) in Rogers and Williams (2000).

Due to our independence Assumption 3.1, the optional projection of $(\int_{S_t^p}^\infty e^{-\alpha s} h_s^p ds, t \geq 0)$ w.r.t. $\mathbb{F}^{\mathbf{S}}$ is indistinguishable from the optional projection of $(\int_S^\infty e^{-\alpha s} h_s^p ds, S \geq 0)$ w.r.t. \mathbb{F}^p when the latter is evaluated along the time-change described by $(S_t^p, t \geq 0)$; see Lemma 3.2 below. By Corollary 2.1 (iii), this latter optional projection is given by

$$\mathbb{E} \left[\int_S^\infty e^{-\alpha s} \inf_{u \in [S, s]} g_u^p ds \middle| \mathcal{F}_S^p \right]$$

at every \mathbb{F}^p -stopping time S . Again by Lemma 3.2, the evaluation of this last optional projection along the time-change described by $(S_t^p, t \geq 0)$ is indistinguishable from the $\mathbb{F}^{\mathbf{S}}$ -optional projection of $(\int_{S_t^p}^{\infty} e^{-\alpha s} \inf_{u \in [S_t^p, s]} g_u^p ds, t \geq 0)$. Hence, the expected rewards from project p can be written as

$$\mathbb{E} \int_0^{\infty} e^{-\alpha t} h_{S_t^p}^p dS_t^p = \mathbb{E} \int_0^{\infty} e^{-\alpha s} \inf_{u \in [0, s]} g_u^p ds + \mathbb{E} \int_0^{\infty} \left\{ \int_{S_t^p}^{\infty} e^{-\alpha s} \inf_{u \in [S_t^p, s]} g_u^p ds \right\} de^{-\alpha(t-S_t^p)}.$$

Now consider the auxiliary reward process $\tilde{h}_s^p \triangleq \inf_{u \in [0, s]} g_u^p$ and note that

$$\tilde{h}_s^p \leq \inf_{u \in [S_t^p, s]} g_u^p \text{ for } s \geq S_t^p.$$

Since $e^{-\alpha(t-S_t^p)}$ is nonincreasing, it follows that

$$\begin{aligned} (5) \quad \mathbb{E} \int_0^{\infty} e^{-\alpha t} h_{S_t^p}^p dS_t^p &\leq \mathbb{E} \int_0^{\infty} e^{-\alpha s} \tilde{h}_s^p ds + \mathbb{E} \int_0^{\infty} \left\{ \int_{S_t^p}^{\infty} e^{-\alpha s} \tilde{h}_s^p ds \right\} de^{-\alpha(t-S_t^p)} \\ &= \mathbb{E} \int_0^{\infty} e^{-\alpha t} \tilde{h}_{S_t^p}^p dS_t^p. \end{aligned}$$

Here the last identity follows by repeating our initial calculation for the reward stream \tilde{h}^p instead of h^p which, being right-continuous, decreasing and adapted, coincides with its own Gittins index: $\tilde{g}^p = \tilde{h}^p$.

Finally summing over p in (5) and denoting by $\tilde{\mathcal{R}}(\mathbf{S})$ the rewards generated by \mathbf{S} when viewed as an allocation strategy for the auxiliary projects \tilde{h}^p ($p = 1, \dots, N$) shows that

$$(6) \quad \mathbb{E} \mathcal{R}(\mathbf{S}) \leq \mathbb{E} \tilde{\mathcal{R}}(\mathbf{S}) \leq \sup_{\mathbf{S}' \in \mathcal{S}} \mathbb{E} \tilde{\mathcal{R}}(\mathbf{S}').$$

Let us verify that in fact equality holds true everywhere in (6) if \mathbf{S} is an index strategy, i.e., a strategy with properties (i) and (ii) as stated in our theorem; existence of such a strategy is proven in Theorem 3 below. Indeed, Inequality (5) results from estimating the integrand $I_t^p \triangleq \int_{S_t^p}^{\infty} e^{-\alpha s} \inf_{u \in [S_t^p, s]} g_u^p ds$ against $\tilde{I}_t^p \triangleq \int_{S_t^p}^{\infty} e^{-\alpha s} \inf_{u \in [0, s]} g_u^p ds$. Clearly I^p has lower semi-right-continuous paths and \tilde{I}^p is continuous. Hence, $k^p \triangleq I^p - \tilde{I}^p \geq 0$ is lower semi-right-continuous and, moreover, it vanishes precisely at those times t when $g_{S_t^p}^p = \inf_{u \in [0, S_t^p]} g_u^p$, since g^p is lower semi-right-continuous. By property (ii) of an index strategy this holds true whenever $de^{-\alpha(t-S_t^p)} < 0$. Thus, by Lemma 1.2, equality holds true in (5) for any $p = 1, \dots, N$ and thus also in the first estimate in (6). In addition

such a strategy will also satisfy equality in the second estimate, since following the leader among the g^p ($p = 1, \dots, N$) entails in particular that the strategy also follows the leader among our auxiliary deteriorating reward processes \tilde{h}^p ($p = 1, \dots, N$). The latter condition, however, was already identified as sufficient for equality in the second estimate; see Lemma 1.1.

As a consequence, any index strategy \mathbf{S} yields a maximal expected reward

$$\mathbb{E}\mathcal{R}(\mathbf{S}) = \sup_{\mathbf{S}' \in \mathcal{S}} \mathbb{E}\tilde{\mathcal{R}}(\mathbf{S}') = \sup_{\mathbf{S}' \in \mathcal{S}} \mathbb{E}\mathcal{R}(\mathbf{S}')$$

and is therefore optimal.

Conversely, it follows that an optimal strategy $\mathbf{S} \in \mathcal{S}$ must satisfy equality in both estimates of (6). Equality in the first estimate requires equality in (5) for every $p = 1, \dots, N$, a condition which using Lemma 1.2 as in the previous paragraph is seen to amount to property (ii) of an index strategy. By the same lemma, the second estimate will hold true as an equality only if

$$\inf_{u \in [0, S_t^p]} g_u^p = \max_{q=1, \dots, N} \inf_{u \in [0, S_t^q]} g_u^q \text{ whenever } dS_t^p > 0.$$

Now, when $dS_t^p > 0$, all other projects must have $dS_t^q < dt$ and so, due to the previously established property (ii), $\inf_{u \in [0, S_t^q]} g_u^q = g_{S_t^q}^q$ for $q \neq p$. It follows that

$$g_{S_t^p}^p \geq \inf_{u \in [0, S_t^p]} g_u^p = \max_{q=1, \dots, N} \inf_{u \in [0, S_t^q]} g_u^q \geq \max_{q \neq p} g_{S_t^q}^q \text{ whenever } dS_t^p > 0.$$

So \mathbf{S} also has to satisfy property (i). □

The following lemma is needed in our proof of Gittins' Index Theorem 2. We record it separately as it highlights the importance of the independence assumption 3.1.

Lemma 3.2 *For any strategy $\mathbf{S} \in \mathcal{S}$, let the σ -algebra*

$$\mathcal{F}_{\mathbf{S}_t} \triangleq \left\{ A \in \bigvee_{p=1}^N \mathcal{F}_{\infty}^p \mid A \cap \{\mathbf{S}_t \leq \mathbf{s}\} \in \mathcal{F}_{\mathbf{s}} \text{ for all } \mathbf{s} \in \mathbb{R}_+^N \right\}$$

describe the information gathered up to time $t \geq 0$ and use $\mathbb{F}^{\mathbf{S}} = (\mathcal{F}_{\mathbf{S}_t})_{t \geq 0}$ to denote the corresponding filtration. Fix $p \in \{1, \dots, N\}$ and let $H^p = (H_s^p, s \geq 0)$ be an $\mathcal{F}_{\infty}^p \otimes \mathcal{B}[0, \infty)$ -measurable bounded stochastic process.

Under Assumption 3.1, the optional projection of $(H_{S_t^p}^p, t \geq 0)$ with respect to $\mathbb{F}^{\mathbf{S}}$ coincides with the \mathbb{F}^p -optional projection G^p of $(H_s^p, s \geq 0)$ evaluated along $(S_t^p, t \geq 0)$:

$$(7) \quad \mathbb{E} \left[H_{S_T^p}^p \mid \mathcal{F}_{\mathbf{S}_T} \right] = G_{S_T^p}^p \text{ for any } \mathbb{F}^{\mathbf{S}}\text{-stopping time } T$$

where G^p satisfies

$$\mathbb{E} [H_S^p \mid \mathcal{F}_S^p] = G_S^p \text{ for any } \mathbb{F}^p\text{-stopping time } S.$$

PROOF : It is easy to see that the class \mathcal{H}^p of those processes H^p for which the assertion of our lemma holds true forms a monotone class. By the monotone class theorem it therefore suffices to show that \mathcal{H}^p contains all $H^p : \Omega \times [0, \infty) \rightarrow \mathbb{R}$ of the form

$$H_s^p(\omega) = B^p(\omega) 1_{[s_0, \infty)}(s)$$

for some bounded, \mathcal{F}_∞^p -measurable B^p and some $s_0 \geq 0$. These H^p are right-continuous and have left limits and so it suffices to verify (7) for deterministic times $T \equiv t \in \mathbb{R}_+$. Moreover, we obtain that G^p is the right-continuous version of $(\mathbb{E} [B^p \mid \mathcal{F}_s^p] 1_{[s_0, \infty)}(s), s \geq 0)$. Now take $A \in \widehat{\mathcal{F}}_{\mathbf{S}_t}$ and note that with $s_i^n \triangleq s_0 + i/n$ ($i = 0, 1, \dots, n = 1, 2, \dots$) we can compute

$$\begin{aligned} \mathbb{E} \left[H_{S_t^p}^p 1_A \right] &= \lim_n \sum_{i=1}^{\infty} \mathbb{E} \left[B^p 1_{A \cap \{s_i^n \leq S_t^p < s_{i+1}^n\}} \right] \\ &= \lim_n \sum_{i=0}^{\infty} \mathbb{E} \left[\mathbb{E} \left[B^p \mid \mathcal{F}_{s_{i+1}^n}^p \vee \bigvee_{q \neq p} \mathcal{F}_\infty^q \right] 1_{A \cap \{s_i^n \leq S_t^p < s_{i+1}^n\}} \right] \\ &= \lim_n \sum_{i=0}^{\infty} \mathbb{E} \left[G_{s_{i+1}^n}^p 1_{A \cap \{s_i^n \leq S_t^p < s_{i+1}^n\}} \right] \\ &= \mathbb{E} \left[G_{S_t^p}^p 1_A \right]. \end{aligned}$$

Indeed, the first identity follows by right-continuity of H^p , the second is due to $A \cap \{s_i^n \leq S_t^p < s_{i+1}^n\} \in \mathcal{F}_{s_{i+1}^n}^p \vee \bigvee_{q \neq p} \mathcal{F}_\infty^q$ by definition of $\widehat{\mathcal{F}}_{\mathbf{S}_t}$, the third identity follows from our assumption of independence 3.1, and the last holds because of right-continuity of G^p . \square

4 Construction of Index Strategies

In this section we shall complete the proof of Theorem 2 by explicitly constructing an index strategy. This will be achieved by combining techniques from El Karoui and Karatzas (1997) and Kaspi and Mandelbaum (1995) with the quasi lower left-continuity of the Gittins index processes. As we shall see, this latter property allows one to avoid any further assumption on the underlying filtrations \mathbb{F}^p ($p = 1, \dots, N$) such as quasi left-continuity (cf., e.g., El Karoui and Karatzas (1997)). As already pointed out by Kaspi and Mandelbaum (1998), the main issue here is that, for index strategies to exist, upward jumps of different Gittins index processes may never occur away from the same all-time low: if such a joint jump was conceivable, the required properties of an index strategy would contradict each other; see Remark 4.2 below. Complementing the rather involved excursion-theoretic Proposition 10 of Kaspi and Mandelbaum (1998), we show how the quasi lower semi-continuity of Gittins index processes from the left allows one to easily rule out the existence of such unpleasant common jumps:

Lemma 4.1 *Assumption 3.1 entails that no two Gittins index processes jump upward from the same new all-time low. More precisely, the random project-specific level sets*

$$J^p \triangleq \left\{ m \in \mathbb{R} \mid m = \inf_{u \in [0, s]} g_u^p \text{ is not attained on } [0, s] \text{ for some } s \in [0, \infty) \right\}$$

are disjoint almost surely:

$$\mathbb{P}[J^p \cap J^q = \emptyset \text{ for all } p \neq q] = 1.$$

PROOF : Note first that the random sets J^p ($p = 1, \dots, N$) are independent by Assumption 3.1 and also at most countable because each g^p has lower semi-right-continuous paths. It thus suffices to prove that $\mathbb{P}[m \in J^p] = 0$ for fixed $m \in \mathbb{R}$ and $p = 1, \dots, N$. To this end, consider the sequence

$$S^{n,p} \triangleq \inf\{s \geq 0 \mid m < g_s^p \leq m + 1/n\}$$

which announces when $\inf_{u \in [0, \cdot]} g_u^p$ creeps down to the level m strictly from above, i.e., $(S^{n,p}, n = 1, 2, \dots)$ is an announcing sequence for

$$S^p \triangleq \inf\{s \geq 0 \mid \inf_{u \in [0, s]} g_u^p = m < g_u^p \text{ for all } u \in [0, s)\}.$$

It follows that S^p is an \mathbb{F}^p -predictable stopping time and, as g^p is quasi lower semi-continuous from the left, we have

$$g_{S^p}^p \leq \liminf_{u \uparrow S^p} g_u^p = \inf_{u \in [0, S^p)} g_u^p = m \quad \text{a.s. on } \{S^p < \infty\}$$

by definition of S^p . It follows that almost surely $m \notin J^p$ whenever this level is at all attained by $\inf_{u \in [0, s]} g_u^p$ for some $s \in [0, \infty)$. \square

Now we can complete the construction of El Karoui and Karatzas (1997) and arrive at the following result.

Theorem 3 *Under Assumption 3.1, there exists an index strategy.*

PROOF : For every project $p = 1, \dots, N$, we consider the \mathbb{F}^p -stopping times

$$\sigma_m^p \triangleq \inf\{s \geq 0 \mid g_s^p \leq m\} \quad (m \in \mathbb{R}),$$

the time for which project p has to be worked on before its Gittins index is less than or equal to m for the first time. Let us in addition denote by

$$\tau_m \triangleq \sum_{p=1}^N \sigma_m^p \quad (m \in \mathbb{R})$$

the time needed by *any* strategy which always follows the leading Gittins index until each of these indices is less than or equal to m . Observe that, in particular, the index strategy to be constructed will and must have these τ_m ($m \in \mathbb{R}$) as level passage times in this sense.

The mappings $m \mapsto \sigma_m^p$ ($p = 1, \dots, N$) and $m \mapsto \tau_m$ are obviously decreasing. In particular, the latter mapping provides a random partition of the time axis $[0, \infty)$ into 'jump intervals' of the form $[\tau_{m+}, \tau_{m-})$ and their complement onto which $m \mapsto \tau_m$ is mapping continuously.

If $t \geq 0$ lies in this complement, there is a unique m such that $\tau_m = t$ and, to ensure that \mathbf{S} follows the leader among the Gittins indices and in accordance with our observation on τ_m , we put $S_t^p \triangleq \sigma_m^p$ so that each project $p = 1, \dots, N$ is either stalled (when $g_{S_t^p}^p < m$) or worked on with Gittins index $g_{S_t^p}^p = m$.

Let us now construct the index strategy on a jump interval like $[\tau_{m+}, \tau_{m-})$. We first focus on the subinterval $[\tau_{m+}, \tau_m)$. This is nonempty if there is at least one project p_m such that $\sigma_{m+}^{p_m} < \sigma_m^{p_m}$, i.e., if there is a project whose Gittins index approaches the level

$m = \inf_{u \in [0, \sigma_{m+}^p]} g_u^{p_m}$ strictly before actually passing it for the first time. Hence, level m is contained in the set J^{p_m} associated to this project by Lemma 4.1. It follows from this lemma that p_m is actually the only project for which m is such a level. Consequently, $\sigma_{m+}^p = \sigma_m^p$ for all $p \neq p_m$ and so our observation on τ_m suggests to put

$$S_t^p \triangleq \begin{cases} \sigma_{m+}^{p_m} + (t - \tau_{m+}) & \text{if } p = p_m, \\ \sigma_{m+}^p = \sigma_m^p & \text{otherwise} \end{cases} \quad \text{for } t \in [\tau_{m+}, \tau_m)$$

in order to ensure the excursion property of an index strategy as well as that we follow the leading Gittins index.

Finally, for the complementary subinterval $[\tau_m, \tau_{m-})$, we follow one of Kaspi and Mandelbaum (1998)'s priority schemes and first activate project $p = 1$ until its Gittins index falls below m , before we abandon this project and engage in project $p = 2$ in the same manner etc. This procedure ensures the excursion property. El Karoui and Karatzas (1997) formalize this procedure by considering the partition $[\tau_m, \tau_{m-}) = \cup_{p=1}^N [\tau_m^{p-1}, \tau_m^p)$ where

$$\tau_m^0 \triangleq \tau_m, \quad \tau_m^p \triangleq \tau_m^{p-1} + (\sigma_{m-}^p - \sigma_m^p) \quad (p = 1, \dots, N)$$

and defining

$$S_t^p \triangleq \begin{cases} \sigma_{m-}^p & p = 1, \dots, q-1 \\ \sigma_m^p + t - \tau^{p-1} & \text{if } p = q \\ \sigma_m^p & p = q+1, \dots, N \end{cases} \quad \text{for } t \in [\tau_m^{q-1}, \tau_m^q) \quad (q = 1, \dots, N).$$

Now S_t^p is defined for every $t \geq 0$ and $p = 1, \dots, N$ and we obtain an allocation strategy $\mathbf{S} \triangleq (S_t^p, t \geq 0) \in \mathcal{S}$. It is readily checked that \mathbf{S} has the desired properties (i) and (ii) and is therefore an index strategy. \square

Remark 4.2 *The above construction shows that the quasi lower semi-continuity from the left of Gittins index processes is essential for the existence of an index strategy. More precisely, as already observed in, e.g., Kaspi and Mandelbaum (1998), the excursion property of an index strategy could otherwise contradict the requirement that an index strategy follows the leading index. Indeed, it would be conceivable that the index processes of two different projects could approach a common new all-time low without actually attaining it. This would forbid us to retire either project because of the excursion*

property. At the same time, however, the 'follow the leader'-property of an index strategy would require us to retire the project which afterwards exhibits the smaller Gittins index—a dilemma which we show cannot actually occur with positive probability.

A Appendix

Proof of Lemma 1.1 (Optimality of myopic strategies in the deteriorating case) The following proof is essentially the time change argument due to El Karoui and Karatzas (1997); we only extend it to allow for possibly negative reward processes. We shall use the same notation as in our proof of Theorem 3 above.

Due to our assumption that the reward processes h^p ($p = 1, \dots, N$) are right-continuous and decreasing the solution of the allocation problem turns out to coincide with the completely anticipating, pathwise solution, and is thus independent from the considered probability measure and filtrations. These are therefore omitted in the sequel for the sake of simplicity.

To see this we use Fubini's theorem and integration by parts to rewrite the (pathwise) reward from a strategy \mathbf{S} as

$$\begin{aligned}
\mathcal{R}(\mathbf{S}) &= \sum_{p=1}^N \int_0^\infty e^{-\alpha t} \left\{ \int_0^\infty 1_{\{m < h_{S_t^p}^p\}} dm - \int_{-\infty}^0 1_{\{h_{S_t^p}^p < m\}} dm \right\} dS_t^p \\
&= \sum_{p=1}^N \int_0^\infty \left\{ \int_0^\infty e^{-\alpha t} d(S_t^p \wedge \sigma_m^p) \right\} dm - \sum_{p=1}^N \int_{-\infty}^0 \left\{ \int_0^\infty e^{-\alpha t} d(S_t^p \vee \sigma_m^p) \right\} dm \\
&= \int_0^\infty \left\{ \int_0^\infty \left(\sum_{p=1}^N S_t^p \wedge \sigma_m^p \right) \alpha e^{-\alpha t} dt \right\} dm \\
&\quad - \int_{-\infty}^0 \left\{ \int_0^\infty \sum_{p=1}^N (S_t^p - \sigma_m^p)^+ \alpha e^{-\alpha t} dt \right\} dm \\
&\leq \int_0^\infty \left\{ \int_0^\infty (t \wedge \tau_m) \alpha e^{-\alpha t} dt \right\} dm - \int_{-\infty}^0 \left\{ \int_0^\infty (t - \tau_m)^+ \alpha e^{-\alpha t} dt \right\} dm \\
&= \frac{1}{\alpha} \int_0^\infty \{1 - e^{-\alpha \tau_m}\} dm - \frac{1}{\alpha} \int_{-\infty}^0 e^{-\alpha \tau_m} dm.
\end{aligned}$$

This yields an upper bound for the reward from an arbitrary allocation strategy \mathbf{S} . In fact, this estimate is sharp. Indeed, equality will hold true everywhere in the above

calculation iff for any $t \geq 0$ the conditions

$$\sum_{p=1}^N S_t^p \wedge \sigma_m^p = t \wedge \tau_m \text{ for all } m \geq 0$$

and

$$\sum_{p=1}^N (S_t^p - \sigma_m^p)^+ = (t - \tau_m)^+ \text{ for all } m < 0$$

are satisfied, i.e., iff

$$S_{\tau_m}^p = \sigma_m^p \quad \text{for all } m \in \mathbb{R}, p = 1, \dots, N.$$

The last condition is known as *synchronization identity* and was used for the first time in Dynamic Allocation Theory by El Karoui and Karatzas (1997). For the deteriorating reward processes considered here, it is clearly satisfied if and only if \mathbf{S} follows the leading Gittins index. Such strategies which follow the leader among the Gittins indices do in fact exist as shown by Theorem 3. As a consequence, the above upper bound is indeed attained and so we can conclude the asserted characterization of optimal allocation policies. \square

Proof of Lemma 1.2 (Semi-continuity and points of increase) Let us first prove the ‘only if’-part of our assertion and take a lower semi-right-continuous $k \geq 0$ for which there is a $t \geq 0$ with $k_t > 0$ but $dA_t > 0$. Then we actually have $k_u \geq \varepsilon$ for $u \in [t, t + \varepsilon)$ for some sufficiently small $\varepsilon > 0$ and we also have $A_t < A_{t+\varepsilon}$. So

$$\int_0^\infty k_u dA_u \geq \int_{[t, t+\varepsilon)} k_u dA_u \geq \varepsilon(A_{t+\varepsilon} - A_t) > 0.$$

For the ‘if’-part we note that any lower semi-right-continuous $k \geq 0$ is the limit of an increasing sequence of right-continuous, piecewise constant $k^n \geq 0$ ($n = 1, 2, \dots$). If, in addition, k vanishes at times t when $dA_t > 0$, these $0 \leq k^n \leq k$ will inherit this property. But then k^n vanishes on each of its intervals of constancy which contains a point of increase of A . So

$$0 = \int_0^\infty k_t^n dA_t \nearrow \int_0^\infty k_t dA_t \quad \text{as } n \uparrow \infty$$

by monotone convergence. \square

References

- BANK, P., AND N. EL KAROUI (2004): "A stochastic representation theorem with applications to optimization and obstacle problems," *Ann. Probab.*, 32(1B), 1030–1067.
- DELLACHERIE, C., AND E. LENGLART (1982): "Sur des problèmes de régularisation, de recollement et d'interpolation en théorie des processus," in *Séminaire de Probabilités XVI*, ed. by J. Azema, and M. Yor, no. 920 in Lecture Notes in Mathematics, pp. 298–313. Springer.
- EL KAROUI, N., AND I. KARATZAS (1994): "Dynamic Allocation Problems in Continuous Time," *Annals of Applied Probability*, 4, 255–286.
- (1997): "Synchronization and optimality for multi-armed bandit problems in continuous time," *Mat. Apl. Comput.*, 16(2), 117–151.
- GITTINS, J. C. (1979): "Bandit processes and dynamic allocation indices (with discussion)," *J. Roy. Statist. Soc. Ser. B*, 41, 148–164.
- GITTINS, J. C., AND D. M. JONES (1974): "A dynamic allocation index for the sequential design of experiments," in *Progress in statistics (European Meeting Statisticians, Budapest, 1972)*, pp. 241–266. Colloq. Math. Soc. János Bolyai, Vol. 9. North-Holland, Amsterdam.
- JACOD, J. (1979): *Calcul Stochastique et Problèmes de Martingales*, no. 714 in Lecture Notes in Mathematics. Springer.
- KASPI, H., AND A. MANDELBAUM (1995): "Lévy bandits: multi-armed bandits driven by Lévy processes," *Ann. Appl. Probab.*, 5(2), 541–565.
- KASPI, H., AND A. MANDELBAUM (1998): "Multi-armed bandits in discrete and continuous time," *Ann. Appl. Probab.*, 8(4), 1270–1290.
- MANDELBAUM, A. (1987): "Continuous multi-armed bandits and multiparameter processes," *Ann. Probab.*, 15(4), 1527–1556.
- PRESMAN, È. L., AND I. N. SONIN (1990): *Sequential control with incomplete information*, Economic Theory, Econometrics, and Mathematical Economics. Academic

Press Inc., San Diego, CA, The Bayesian approach to multi-armed bandit problems, Translated from the 1982 Russian original, edited and with a preface by E. A. Medova-Dempster and M. A. H. Dempster.

ROGERS, L. C. G., AND D. WILLIAMS (2000): *Diffusions, Markov processes, and martingales. Vol. 1 & 2*, Cambridge Mathematical Library. Cambridge University Press, Cambridge, Foundations, Reprint of the second (1994) edition.

VARAIYA, P. P., J. C. WALRAND, AND C. BUYUKKOC (1985): "Extensions of the multiarmed bandit problem: the discounted case," *IEEE Trans. Automat. Control*, 30(5), 426–439.

WEBER, R. (1992): "On the Gittins index for multiarmed bandits," *Ann. Appl. Probab.*, 2(4), 1024–1033.

WHITTLE, P. (1980): "Multi-armed bandits and the Gittins index," *J. Roy. Statist. Soc. Ser. B*, 42(2), 143–149.

PETER BANK

Department of Mathematics

Columbia University

2990 Broadway, Mail Code 4433

New York, NY 10027

U.S.A.

e-mail: pbank@math.columbia.edu

URL: www.math.columbia.edu/~pbank

CHRISTIAN KÜCHLER

Institut für Mathematik

Humboldt–Universität zu Berlin

Unter den Linden 6

D–10099 Berlin

Germany

e-mail: ckuechler@math.hu-berlin.de

URL: www.math.hu-berlin.de/~ckuechler