

Compactifying the Commodity Space: Examples in Financial Economics

Stephen F. LeRoy
University of California, Santa Barbara

May 24, 2006

Abstract

Financial models with a finite number of states and dates have a number of properties that do not extend to their infinite counterparts. It is proposed to deal with this situation by appending a date called ∞ to the payoff index set, and defining a topology such that the payoff index set is compact. This allows a simplified mathematical treatment of a number of topics that are unwieldy when modeled in a setting where the payoff index set is noncompact. Topics discussed include price system bubbles, payoff bubbles, the doubling strategy and non-equivalent martingale measures.

Finance models incorporating a finite number of dates and states have a number of attractive properties not shared by their infinite-dimensional counterparts.¹ Among these are:

- Excluding frictions (specifically, assuming that a firm's reinvested retained earnings earn the same returns as its original capital, and that this return coincides with the factor at which dividends are discounted), dividend policy does not affect firm value. This is so because lower future dividends imply higher retained earnings, and therefore a higher growth rate of invested capital. Because of the assumption that invested capital generates returns at the same rate as dividend streams are discounted, the two effects cancel (Miller and Modigliani [16]). Of course, this argument requires that firms pay a liquidating dividend at the terminal date. With an infinite future this argument fails because there is no terminal date. In the extreme case in which firms never pay a dividend, for example, the firm value justified by dividends is apparently zero. Therefore the dividend-irrelevance proposition is apparently valid only if some dividend strategies are ruled out.

¹This paper draws heavily on Fisher and Gilles [5]. I am indebted to Mark Fisher for many conversations about this material.

- The “Fundamental Theorem of Finance” (Dybvig and Ross [4]) holds that in the absence of arbitrage there always exists a strictly positive continuous valuation operator that assigns value to all contingent claims, whether or not they are payoffs of portfolios. This valuation operator can be expressed as a risk-neutral measure which is equivalent to the original probability measure (in the sense that the two measures have the same null events). The Fundamental Theorem of Finance holds without qualification in finite settings. However, in infinite settings a variety of problems arise: valuation may be discontinuous in any preassigned topology, or the risk-neutral measure may not be equivalent to the original measure.
- In finite settings absence of arbitrage in simple portfolio strategies (portfolio strategies that involve transactions at only one event) implies absence of arbitrage under all portfolio strategies. That proposition appears not to extend to infinite settings. The classic example is the doubling strategy, in which a gambler who is able to double his bets forever can apparently earn a certain profit.

The loss of these attractive properties of finite models as one moves to infinite settings is difficult to interpret. What is the source of these qualitative changes—apparent discontinuities—as one moves to the more general setting? To investigate these questions it is necessary to specify precisely the sense in which (loosely) the infinite is representable as a limit of the finite. In this paper we implement this task by including ∞ in the commodity index set and imposing a topology that defines the sense in which what happens at infinity is a limit of what happens at finite dates (when the commodity index set is interpreted as time).

Specifically, we will restrict attention to a countable commodity index set, which is therefore representable by $N_\infty = N \cup \{\infty\}$, where N denotes the natural numbers. We define a metric $d(x_i, x_j)$ on N_∞ by

$$d(x_i, x_j) = 2^{-i} - 2^{-j}, \tag{1}$$

where $0 < i < j < \infty$, and

$$d(x_i, \infty) = 2^{-i}, \tag{2}$$

$i < \infty$. This metric is the appropriate one if one associates x_1, x_2, x_3, \dots with $1/2, 3/4, 7/8, \dots$, and associates x_∞ with 1. Under the topology associated with the distance measure just defined, ∞ is the limit point of the sequence consisting of the natural numbers, so adding ∞ to N amounts to compactifying N .

A basis for the topology associated with the distance measure defined by (1)-(2) consists of (1) the singleton subsets of N , plus (2) the union of ∞ and the cofinite subsets of N . Thus every open set that includes ∞ also includes the tail of N beyond some $n \in N$. In this sense ∞ is indistinguishable from the very large elements of N .

We will see below that this choice of topology will make infinite economies limits of finite economies.

In the following section we set out the common framework we will use to analyze the three problems just listed in the sections that follow.

1 Two Function Spaces on N_∞

Let C be the space of continuous real-valued functions on N_∞ .² Under the topology on N_∞ just specified, C consists of the functions $f : N_\infty \rightarrow R$ such that (1) $\lim_{n \rightarrow \infty} f(n)$ exists, and (2) $\lim_{n \rightarrow \infty} f(n) = f(\infty)$. We equip C with the topology induced by the norm $\|f\|_\infty = \sup_n |f(n)|$, $n \in N_\infty$.

Consider also $M(N_\infty)$, the space of finite signed measures on the Borel σ -field $B(N_\infty)$, and define on $M(N_\infty)$ the topology of weak-* convergence, so that $\mu_n \rightarrow \mu$ if and only if $\int_{N_\infty} f d\mu_n \rightarrow \int_{N_\infty} f d\mu$ for all $f \in C$. Then C and M are topological duals under

$$\langle f, \mu \rangle = \int_{N_\infty} f d\mu = \sum_{i=1}^{\infty} f_i \mu(\{i\}) + f_\infty \mu(\{\infty\}) \quad (3)$$

for $f \in C$, $\mu \in M$.

In the four applications that follow, N_∞ will be interpreted variously as the set of possible dates and as a probability space. As a consequence, the interpretations of C and M will differ depending on the context. However, the topologies defined on C and M imply a limiting property that will obtain in all four of the applications considered below, so we outline this limiting property before proceeding to the applications.

Let $\mu = \{\mu_n\}$ be a continuous bounded function from N_∞ to M . Continuity implies that $\{\mu_n\}$ converges to μ_∞ . This fact implies that μ_∞ can be approximated arbitrarily accurately by some μ_n for $n \in N$. Thus we have a precise sense in which prices or payoffs at ∞ (depending on which M represents; we will consider both specifications below) can be approximated arbitrarily accurately by prices or payoffs at finite dates (when N represents time).

2 Deterministic Applications

The applications discussed in this section take the index set as representing time, implying that there is no uncertainty. Thus adding ∞ to the natural numbers in defining the index set implies that we are formally modeling payoffs in the infinite future and assigning prices to these payoffs.

²For economists, the best source for the mathematics used in this paper is Aliprantis and Border [1].

It should be understood that this exercise involves a purely theoretical construct; in proposing models we are free to assign as large or as small a role to payoffs or prices at infinity as appears appropriate. In particular, there is no presumption that any security necessarily has a nonzero payoff at infinity, nor that individuals' choice sets necessarily include securities with nonzero payoffs at infinity. Indeed, some of the applications below explicitly rule out the possibility that agents can select portfolios with nonzero payoffs at infinity. Even if portfolios with payoffs at infinity are admissible, utility functions may be such that these are never optimal.

Whether it is worthwhile to include payoffs at infinity in one's model depends, as always, on the usefulness of the additional insights made possible. In this paper we make the case that additional useful insights are in fact made available by doing so.

2.1 The Dividend Irrelevance Proposition

In the first application we will represent payoffs by measures on N_∞ .³ For example, the Dirac measure δ_t defined by

$$\delta_t(A) = \begin{cases} 1 & \text{if } t \in A \\ 0 & \text{if } t \notin A \end{cases} \quad (4)$$

for $A \subset N_\infty$ denotes one unit of output delivered at date t . As a special case, δ_∞ denotes one unit of output delivered at date ∞ .

Imposition of the weak-* topology on N_∞ implies that as t approaches ∞ , $\{\delta_t\}$ converges to δ_∞ . To see this, take any arbitrary test function f from C . Then

$$\langle f, \delta_t \rangle = f(t) \text{ and } \langle f, \delta_\infty \rangle = f(\infty). \quad (5)$$

Because f is continuous, $f(t)$ converges to $f(\infty)$, implying that $\langle f, \delta_t \rangle$ converges to $\langle f, \delta_\infty \rangle$, as required for weak-* convergence of δ_t to δ_∞ . Incidentally, it is worth noting that δ_t does not converge to δ_∞ under the total variation norm: we have $\|\delta_t - \delta_\infty\| = 2$ for all $t < \infty$.

In this application the elements of C , the dual space of M , are used to test convergence, not to represent price systems. Instead, we model price systems as measures, just like payoffs. For example, in the application to follow the relevant price system involves discounting at a constant rate ρ , with $\rho > 0$. This price system is represented by a measure λ on N_∞ , defined by

$$\lambda(A) = \sum_{i \in A} (1 + \rho)^{-i}, \quad A \subset N_\infty. \quad (6)$$

It is understood here that $\lambda(\{\infty\}) = 0$, as one would expect from that fact that $(1 + \rho)^{-\infty} = 0$.

³Examples of papers that use measures to represent commodities are MasColell [15] and Jones [11].

Consider a firm that starts with one unit of capital which generates earnings at constant rate ρ at each date. It pays out a proportion γ of earnings as a dividend, using retained earnings to acquire more capital. The newly acquired capital also generates earnings at constant rate ρ . It is easily verified that the firm pays dividend $z_\gamma(t)$ at date t , given by

$$z_\gamma(t) = \gamma\rho(1 + \rho(1 - \gamma))^t. \quad (7)$$

The discounted dividend stream has typical element

$$(1 + \rho)^{-t}z_\gamma(t) = (1 + \rho)^{-t}\gamma\rho(1 + \rho(1 - \gamma))^t. \quad (8)$$

and the value of the sequence $\{z_\gamma(t)\}$ for any $\gamma > 0$ is

$$\sum_{t=1}^{\infty} (1 + \rho)^{-t}z_\gamma(t) = 1, \quad (9)$$

so that for any $\gamma > 0$ dividend policy does not affect the value of the firm (Miller and Modigliani [16]).

What happens when $\gamma = 0$? From (7) it is clear that for each t , $z_\gamma(t)$ equals 0, and the sum of an infinite number of 0s is 0. The apparent conclusion is that the dividend-irrelevance proposition does not apply when γ equals zero. This is a troublesome finding: where did the firm's value go? However, the argument is far from water-tight. The payoff stream when $\gamma = 0$ is taken as the limit as γ goes to 0 of the payoff stream for positive γ , where the limit is taken pointwise. The fact that valuation is discontinuous when limits are taken pointwise raises questions about the appropriateness of pointwise convergence.

The analysis leads in a different direction if we recast the firm model just outlined in the analytical framework of the preceding section, which relies on weak-* limits rather than pointwise limits. The first step is to redefine the discounted dividend stream as a measure. To that end, define the measure μ_γ by

$$\mu_\gamma(A) = \sum_{i \in A} (1 + \rho)^{-i} \gamma \rho (1 + \rho(1 - \gamma))^i, \quad (10)$$

$A \subset N_\infty$. The discounted dividend stream as written in the usual form—that is, as characterized in (8)—is recognized as the derivative of μ_γ with respect to the counting measure. Because μ_γ measures dividends in units of date-0 value, the relevant price system is the sequence which assigns 1 to every element of N_∞ (including ∞ , for reference below). Therefore the value of μ_γ is given by $\langle 1, \mu_\gamma \rangle$. This equals 1 for each $\gamma > 0$.

The obvious next step is to define the payoff of the “zero”-dividend firm as the weak-* limit of $\{\mu_\gamma\}$ as γ approaches zero. Because for γ near zero most of the value of the dividend stream occurs in the tail of the dividend sequence—that is, because we have

$$\lim_{\gamma \downarrow 0} \langle 1, \mu_\gamma(\{T, \infty\}) \rangle = 1 \quad (11)$$

for any finite T , it follows that μ_γ converges to δ_∞ . Thus the dividend stream generated by the “zero”-dividend firm is one unit of output (measured in value terms) at ∞ . This payoff has value $\langle 1, \delta_\infty \rangle = 1$. Because valuation is weak-* continuous, it is seen that modeling the payoff of the “zero”-dividend firm as the limit of the payoff of the γ -dividend firm implies that the dividend-irrelevance proposition applies for $\gamma = 0$ as well as for non-limiting values of γ .

The applicability of this analysis requires that one accept the maintained assumption that valuation is weak-* continuous. Standard existence theory suggests that valuation will be weak-* continuous if the same is assumed true of preferences. This might or might not be an acceptable assumption. Whether or not one accepts this assumption, however, the present exercise has the important implication that the standard analysis of the “zero”-dividend firm is surely faulty. Taking the payoff of the zero-dividend firm to be zero reflects a presumption that pointwise convergence gives the appropriate definition of convergence. In view of the fact that virtually none of the standard utility functions are pointwise continuous on infinite-dimensional spaces, there seems little reason to identify the value of the “zero”-dividend firm with the value of the pointwise limit of the payoff of the γ -dividend firm.

In substantive terms, the analysis just presented mirrors exactly that of Gilles and LeRoy [9]. In formal terms, the two analyses differ in that Gilles-LeRoy did not compactify the payoff index set. As a result, the limiting payoff in Gilles-LeRoy’s model took the form of a pure charge (a set function which, unlike a measure, is not countably additive). The fact that we can represent the limiting payoff here as a measure constitutes a major advantage in terms of simplicity: the existence of pure charges requires appeal to the axiom of choice (since proving the existence of pure charges requires the axiom of choice). The fact that pure charges on N are inherently nonconstructive makes them difficult to work with, compared to measures on N_∞ , which are constructive and easy to work with, as we have seen.

Gilles-LeRoy noted that in their setting the Yosida-Hewitt theorem implies that each element of the payoff space—that is, each charge—can be decomposed into the sum of a measure and a pure charge. They interpreted the measure component of the payoff stream as the fundamental payoff and the pure charge component as the bubble. An analogous interpretation is available here: any measure on N_∞ can be decomposed into the sum of a measure that is absolutely continuous with respect to the counting measure and another measure that gives weight to ∞ alone. The former can be identified as the fundamental component of the payoff and the latter as the bubble component. In terms of the example, the dividend streams associated with $\gamma > 0$ are fundamental payoffs, while the dividend associated with $\gamma = 0$ is a pure bubble. It is worth noting that, as the example indicates, a net of payoffs each element of which has a nonzero fundamental and a zero bubble can converge to a

payoff that has zero fundamental and nonzero bubble.

2.2 Price System Bubbles

The application of this subsection, like that of the preceding subsection, is most easily interpreted as taking the index set N_∞ to represent time. The function spaces C and M continue to play a central role. However, in this subsection we represent payoffs as elements of the space C of continuous functions on N_∞ , instead of as elements of the measure space M as in the preceding subsection. In this subsection elements of M will be reserved to represent price systems.

The “Fundamental Theorem of Finance” (Dybvig and Ross [4]) states that the absence of arbitrage implies that there exists a linear functional that values all contingent claims. This functional may be represented in various forms: it may be a set of state prices, a stochastic process (the state price deflator) or a probability measure (the risk-neutral measure). As noted in the introduction, the Fundamental Theorem of Finance holds without qualification in finite settings. However, in infinite settings the requisite linear functional may not exist even in the absence of arbitrage (Kreps [12]). If it does exist, it may not be representable by a measure. For example, Back and Pliska [2], Gilles and LeRoy [8], [10], Werner [17] and others have analyzed examples in which the pricing “measure” is not countably additive, contrary to the definition of a measure.

As in the preceding subsection, the failure of countable additivity in the models of the papers just cited is attributable to the use of a noncompact payoff index set. We saw in the preceding subsection that taking the payoff index set to be compact implied that certain limiting payoffs that would be represented as pure charges in the noncompact case are representable as measures. Precisely the same result occurs here, except that here the measures that replace pure charges represent the equilibrium price system rather than payoffs.

Suppose that an infinite number of securities is available for inclusion in portfolios. Security 0 is a perpetuity with payoff $(1, 1, 1, \dots)$, while security i is the i -th Arrow security (which pays one unit of consumption at date i and zero at other dates), $i = 1, 2, \dots$. If portfolios are assumed to consist of a finite number of securities, then the space of portfolio payoffs consists of a proper subspace of the space of sequences on N that have limits (a proper subspace because any portfolio payoff attains its limit at some finite date). The limiting payoff of any portfolio, of course, equals the holding of security 0. Since a payoff stream is characterized by its payoff at every date and the limiting payoff, it is natural to define the payoff at ∞ as the limiting payoff. Then the space of payoffs is a subspace of the space C of continuous functions on N_∞ defined in Section 1.

Assume that the price p_i of security i is 2^{-i} , $i = 1, 2, \dots$. Whether or not there exists arbitrage depends on the price of security 0. It is easily checked that there is no arbitrage if the price of security 0 is greater than or equal to 1 (otherwise a

portfolio consisting of a long position in security 0 and an equal short position in securities $1, \dots, n$ will produce an arbitrage for sufficiently high n). Further, there is no approximate arbitrage under the same restriction. Therefore when the price of security 0 is greater than or equal to 1, there exists a strictly positive valuation operator.⁴

Suppose first that $p_0 = 1$. Then the measure μ defined by

$$\mu(A) = \sum_{i \in A} 2^{-i}, \quad (12)$$

$A \subset N_\infty$, is a valuation operator. That is, we have that $\langle f, \mu \rangle$ equals the value of payoff f . Here, for reference in the next paragraph, it is understood that $2^{-\infty} = 0$ as above, so the measure of any set A that includes ∞ equals the measure of A/∞ .

All of this is standard. The example becomes more interesting when we assume that $p_0 > 1$. To be specific, take $p_0 = 2$, and assume that $p_i = 2^{-i}$ for $i = 1, 2, \dots$ as above. The fact that there does not exist arbitrage or approximate arbitrage under this modification again implies that there exists a continuous strictly positive valuation operator. This operator can be represented by the measure ν defined by

$$\nu(A) = \begin{cases} \sum_{i \in A} 2^{-i} & \text{if } \infty \notin A \\ \sum_{i \in A} 2^{-i} + 1 & \text{if } \infty \in A \end{cases}, \quad (13)$$

$A \subset N_\infty$. Thus the value of payoff f under price system ν is given by $\langle f, \nu \rangle = \sum_i 2^{-i} f_i + \lim_i f_i$. With the price system ν , portfolios with a nonzero limiting payoff must include a position in security 0, and p_0 exceeds the limit of the summed prices of the first n securities by 1.⁵

Depending on preferences, either μ or ν could be an equilibrium price system. In a representative agent setting in which preferences are given by $U(c) = \sum_i 2^{-i} c_i$ the equilibrium price system will be μ . Correspondingly, if preferences are given by $U(c) = \sum_i 2^{-i} c_i + \lim_i c_i$, the equilibrium price system will be ν .

As in the preceding subsection, the case in which the equilibrium price measure takes on nonzero weight at ∞ can be connected with bubbles (Gilles [6], Gilles and LeRoy [7], [8]). In the preceding subsection, measures were identified with payoff streams, so the component of any measure associated with ∞ was interpreted as the bubble component of the payoff stream identified with that measure. In this subsection, however, the measure is identified with price systems, not payoffs, so the

⁴For a proof that nonexistence of arbitrage and approximate arbitrage implies the existence of a strictly positive continuous valuation operator, see Clark [3]. Clark's result assumes that the payoff space is a Banach lattice; we delete demonstration that his result applies in the present setting.

⁵Note that there is no violation of continuity here: the summed payoff of the first n securities does not converge to the payoff of security 0 under the sup norm, so continuity does not imply that their summed prices converge to the price of security 0.

interpretation is different: the equilibrium price system has a bubble if the measure associated with the price system takes on nonzero weight at infinity, as is the case with price system ν , but not with μ .

We see that, exactly as in the preceding subsection, adding ∞ to the payoff index set has the benefit that it allows a representation of price systems (payoffs in the preceding subsection) containing bubbles that is just as simple as that of price systems without bubbles. Without the compactification, on the other hand, it is necessary to represent the bubble component of equilibrium price systems using pure charges. As noted above, the nonconstructive nature of charges makes them difficult to interpret and badly suited for applied work.

In the context of the example considered in this subsection, the index set N_∞ can be associated with states instead of dates, so that the model involves uncertainty rather than time. In that case if the probability that state i occurs is assumed to be 2^{-i} , then the utility function $U(c) = \sum_i 2^{-i} c_i$ incorporates risk neutrality. Under the utility function

$$U(c) = \sum_i 2^{-i} c_i + \lim_i c_i, \tag{14}$$

on the other hand, the equilibrium pricing measure, given by (13), is not absolutely continuous with respect to probability measure (because, unlike probability, it has nonzero weight at ∞). However, it is worth noting that the support of ν is the same as that of the probability measure.⁶

It must be observed that the utility function $U(c) = \sum_i 2^{-i} c_i + \lim_i c_i$ is unconventional. In Subsection 3.2 we present another example of an equilibrium price system that is discontinuous with respect to probability. That example will assume standard preferences; the distinctive feature of that example comes from the assumed stochastic process for the aggregate endowment.

3 Stochastic Applications

The two examples to be discussed involve uncertainty. To model uncertainty in a dynamic setting we interpret N_∞ as a probability space and define a filtration that represents agents' information at each date. This, of course, is standard practice in finance.

3.1 The Doubling Strategy

As indicated in the introduction, the doubling strategy is a gambling strategy that involves betting \$1 on red on a roulette wheel. Upon winning, the gambler stops

⁶The support of a measure is the smallest closed subset outside of which the measure takes on zero value. Probability takes on nonzero values on N , a proper subset of N_∞ . However, in the specified topology N is not closed, so the support of probability is N_∞ , the same as that of ν .

betting. Upon losing, the gambler doubles the bet, and continues to do so until he wins. In infinite time, the gambler will win with probability 1, since red will eventually occur except on a set of measure zero. The doubling strategy is generally treated as an arbitrage. Analysts take the view that it is necessary to rule out this trading strategy by invoking trading restrictions.

In the preceding subsection it was assumed explicitly that portfolios consist of a finite number of securities. If the analogous restriction—that agents can play the doubling strategy only a finite number of times—is imposed in the present context, nothing of interest remains: the date- n payoff of the doubling strategy (a small gain with high probability combined with a large loss with low probability), far from being an arbitrage, is unattractive under most utility functions. The arbitrage, or apparent arbitrage, appears only if the gambler can double the bet an infinite number of times. Therefore we assume that doing so is possible.

In the real world, of course, it is not possible to conduct an infinite number of transactions, or to implement a portfolio strategy that requires unbounded wealth. However, in theoretical settings nothing inherently rules out an infinite number of trades (and the assumption that an infinite number of trades is possible is made routinely in continuous-time finance). The infinitely-repeated doubling strategy can be implemented in all states of the world only if consumption sets are unbounded below. However, some utility functions (quadratic, linear, negative exponential) are well defined over all negative reals as well as some or all positive reals, in which case the assumption is satisfied. There do not exist optimal portfolios if the infinitely-repeated doubling strategy is feasible and is an arbitrage, so it is necessary to face this issue.

It is far from clear that real-world agents—either gamblers or investors—are so foolish as to believe that they could manufacture certain wealth from each other by engaging in an infinite number of gambles with each other if only they were able to do so. As with the example of Subsection 2.1, the conclusion that they can produce wealth in this way involves presuming without justification that the relevant standard for convergence is pointwise convergence (or, equivalently in this context, convergence in measure). As above, we propose instead that the payoffs of infinite portfolio strategies—in particular, the payoff of the doubling strategy—be modeled as weak-* limits of the payoffs of the corresponding finitely-repeated portfolio strategies (if the limit does not exist, the payoff of the infinite portfolio strategy is undefined, and the portfolio strategy itself is treated as not feasible). In the case of a fair roulette wheel, to which we will restrict our attention, this limit always exists and takes a very simple form. In the general case, so that the house takes a nonzero percentage of the bets (or, in the case of securities, if securities prices reflect risk aversion), the limit may not exist. In that case the interpretation is that the doubling strategy does not exist as a well-defined infinite-time portfolio strategy.⁷

⁷An infinite sequence of transactions is admissible only when the appropriate limit is well defined. For example, consider the transaction purportedly defined by “I give you \$1, then you give me \$1,

The present example, unlike the preceding examples, involves uncertainty as well as infinite time. We take the state of the economy to be the date at which red first occurs. Doing so implies that the state space can be taken to be N_∞ , as in the preceding example, with the element ∞ of N_∞ denoting the event that red never occurs. This event, although occurring with probability zero, cannot be neglected in the analysis. This is so because, as we will see, the payoff of the doubling strategy approaches $-\infty$ in that event. This will turn out to imply that the fact that the probability of red never occurring is zero does not imply that a security's payoff in that event is irrelevant to its price.

We define a filtration \mathcal{F} , which consists of the σ -fields $(\mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_\infty)$. Here \mathcal{F}_0 is the trivial partition, \mathcal{F}_∞ is based on the power set of N_∞ , and \mathcal{F}_n is based on the partition $\{\{1\}, \{2\}, \dots, \{n\}, \{\cup_{\tau=n+1}^\infty \{\tau\} \cup \{\infty\}\}\}$. The interpretation is that at date n the agent knows whether or not red has occurred by date n . By extension, at date ∞ the agent is treated as knowing either the date at which red first occurred, or the fact that it never occurred.⁸

To model the payoff of the doubling strategy, we first specify the measure π representing probabilities as

$$\pi(\{\tau\}) = \begin{cases} 2^{-\tau}, & \tau < \infty \\ 0, & \tau = \infty \end{cases} . \quad (15)$$

Second, we specify the payoff a_n of the doubling strategy: after playing n rounds of the doubling strategy the payoff is

$$a_n\{\tau\} = \begin{cases} 1, & \tau \leq n \\ 1 - 2^n, & n < \tau \end{cases} , \quad (16)$$

a random variable on N . This specification of the payoff of the finitely-repeated doubling strategy amounts to assuming that the roulette wheel is fair (observe that $E(a_n) = 0$, where the expectation is taken with respect to π). Our evaluation of the limiting payoff of the doubling strategy will be seen below to depend on this assumption.

The next step is to find an expression for μ_n . Define a measure ν_n on \mathcal{F}_∞ as the product of a_n and π :

$$\nu_n(\{\tau\}) = a_n(\{\tau\})\pi(\{\tau\}) = \begin{cases} 2^{-\tau}, & \tau \leq n \\ 2^{-\tau}(1 - 2^n), & n < \tau < \infty \\ 0, & \tau = \infty. \end{cases} \quad (17)$$

and we repeat an infinite number of times". Because the appropriate limit is undefined (is the net payoff $(-1, 1)$ or $(0, 0)$?), the above supposed characterization of a transaction is inadmissible.

⁸As noted in the introduction, payoffs that are measurable only with respect to \mathcal{F}_∞ involve a considerable element of abstraction. Rejecting the assumption that there exist payoffs that are measurable with respect to \mathcal{F}_∞ but nonmeasurable with respect to \mathcal{F}_n for finite n , however, is tantamount to assuming that agents cannot play the doubling strategy an infinite number of times. This defines away the problem.

The measure μ_n associated with playing n rounds of the doubling strategy is the restriction of ν_n to \mathcal{F}_n . This is

$$\mu_n(\{\tau\}) = \begin{cases} 2^{-\tau}, & \tau \leq n \\ 2^{-n} - 1, & \tau > n \end{cases}, \quad (18)$$

reflecting the fact that

$$\sum_{\tau=n+1}^{\infty} 2^{-\tau}(1 - 2^n) = 2^{-n} - 1. \quad (19)$$

Now we are able to identify the limiting payoff μ_∞ of the infinitely-repeated doubling strategy, expressed as a measure.⁹ As expected from the preceding sections, we define this as just the weak-* limit of μ_n as n goes to ∞ :

$$\mu_\infty = \lim_{n \rightarrow \infty} \mu_n = \begin{cases} 2^{-\tau}, & \tau < \infty \\ -1, & \tau = \infty \end{cases}. \quad (20)$$

This expression for μ_∞ is intuitively plausible inasmuch as it coincides with (18) except that the term 2^{-n} disappears in the limit. The expression may, of course, be justified formally by taking the typical continuous function f on N_∞ as a test function for weak-* convergence.

The payoff of the doubling strategy played n times can be represented equivalently by the measure μ_n , given by (18), or the random variable a_n , given by (16). The point of this exercise is to demonstrate that the former expression is more convenient when it comes to characterizing the limiting payoff. This is so because the payoff of μ_n if $\tau > n$ converges to -1 , reflecting the fact that μ_n weak-* converges to μ_∞ . In contrast, the payoff of a_n converges to $-\infty$ if $\tau > n$, implying that a_n does not converge (in, for example, mean-square) to any random variable as n goes to ∞ . Thus (16) does not apply if $n = \infty$. Nonexistence of a limiting density, of course, reflects the fact that μ_∞ is not absolutely continuous with respect to the counting measure for any finite n , since ∞ is a null set for the counting measure, but not for μ_∞ . Here we see the advantage of working with the underlying measures μ_n : we have convergence as n goes to ∞ instead of nonconvergence.

A complete treatment of the economic theory of the doubling strategy would involve defining the utility of payoffs represented as measures. In an expected utility context this would involve extending expected utility from the space of payoffs that have derivatives with respect to probability to the larger space that includes limit points. We do not pursue this line here. It is worth observing, however, that the

⁹Note that we define payoffs as elements of M . This would be inappropriate if consumption occurs at intermediate dates. In that case payoffs would be modeled as elements of $M \times N_\infty$. We ignore this complication, assuming implicitly that payoffs at intermediate dates can be rolled over one-for-one.

In the following subsection we present an example in which the payoff space is taken to be a stochastic process.

conclusion of earlier sections, that pointwise convergence does not provide a sensible way to define limiting payoffs, applies here also.

If one uses weak-* convergence, as here, rather than pointwise convergence to define limiting payoffs, the demonstration that the doubling strategy generates an arbitrage fails. This is clear from (20): μ_∞ is not a positive measure. This is a sensible conclusion: there is no economic or mathematical justification for the presumption, implicit in pointwise convergence, that the decreasing probability of a loss as n goes to ∞ somehow trumps the size of the loss in the limit. If weak-* convergence is used to define the limiting payoffs, it follows that excluding arbitrage does not imply the necessity of ruling out the doubling strategy.

Again we may compare the present analysis with that of Gilles and LeRoy [9]. These authors, using N as the payoff index set rather than N_∞ , showed that the weak-* limit of the payoff on the finitely repeated doubling strategy is a charge which has a nonzero pure charge component. In contrast, compactifying the payoff index set here allows us to avoid introducing pure charges. Evidently the remarks offered above to the effect that the present treatment is preferable on grounds of simplicity apply here also.

3.2 Non-Equivalent Martingale Measures

The example of this subsection combines features of the models of Subsection 2.2 and Subsection 3.1. As in all the models of this paper, price systems are identified with measures. Unlike the model of Subsection 2.2 but like that of Subsection 3.1, the model of this subsection involves gradual resolution of uncertainty. In the model of Subsection 3.1, prices coincided with probabilities, reflecting the assumption that the roulette wheel was fair. Preferences were not explicitly characterized there. Here, in contrast, agents' preferences are explicitly displayed, and they exhibit risk aversion. Consequently, the measure associated with prices at any date does not coincide with probability: in our example it will turn out that for $n < \infty$ the derivative of the price system at n with respect to probability is a random variable with a nondegenerate distribution (this random variable, of course, is measurable with respect to the information partition at date n).

We will use the term "price system" without a date identification to denote a sequence of measures, one for each date, that value discounted payoffs. Thus the price system is an element of $M \times N_\infty$ rather than M as in Subsection 2.2. Because the price system at each date attaches value to discounted payoffs, it is an extension of the price system from the information partition at the earlier date to the richer information partition at the later date. Equivalently, the price system at any date is a restriction of the price system at any later date. In particular, the price system associated with date ∞ values all discounted payoffs, so it does the work of the entire price system.

The example of this section owes its interest to the fact that the price system

at ∞ which we will be deriving will turn out not to be absolutely continuous with respect to the probability measure, implying nonexistence of an equivalent martingale measure. This is the analogue in a dynamic context of the result of Subsection 2.2 that the equilibrium pricing measure is discontinuous with respect to probability.

As pointed out in Subsection 2.2, a drawback of the example there is that the assumed equilibrium is supported by unconventional preferences (see (14)). In this section we present an example in which the same characterization of the equilibrium price system holds, but in which preferences are completely standard. The example here, being a general equilibrium representative-agent model in which securities are in zero net supply, has the convenient property that markets are effectively complete under very weak assumptions (see LeRoy and Werner [13] for discussion of effectively complete markets). In such settings the zero portfolio strategy is always optimal. Therefore one can analyze valuation without even specifying the traded securities. Accordingly, in this section the price system rather than the asset span will be the principal focus of study.

To establish notation we will begin with the finite case, in which there is no problem in deriving a unique equivalent martingale measure (this is a standard result in financial economics; see LeRoy and Werner [13] or any of a number of sources for discussion), or of characterizing the martingale measure. The representative agent maximizes

$$\sum_{t=1}^T 2^{-t} E[\ln(c_t)]. \quad (21)$$

The agent's endowment at date t equals 2^t , $t = 1, 2, \dots, T$, if the roulette wheel produces red at t or at any date prior to t , and 1 otherwise. This specification of the endowment produces the distinctive features of the equilibrium. As above, let τ denote the first date that red occurs. Define probability as a measure π on N_∞ :

$$\pi(\{\tau\}) = 2^{-\tau}, \quad (22)$$

with $\pi(\{\infty\}) = 0$, so that successive draws of red and black are independent and equally likely.

Assume, both in the finite case discussed here and the infinite case discussed below, that portfolios consist of a finite number of securities and that each security has a payoff that is nonzero at only a finite number of states and dates. Therefore the payoffs of all portfolio strategies are at most finitely nonzero. This assumption will rule out Ponzi schemes in the infinite version of the model. Portfolio payoffs are adapted stochastic processes on the filtration \mathcal{F} that was defined above. From the representative agent assumption, the zero portfolio strategy is always optimal in this setting. Security prices can therefore be calculated using the event prices that support the representative agent's endowment. We now calculate these event prices.

The representative agent's marginal utility—equivalently, the state price deflator ρ_t —is given by

$$\rho_t = \begin{cases} 2^{-2t} & \text{if } \tau \leq t, \\ 2^{-t} & \text{if } \tau > t \end{cases}. \quad (23)$$

Note here that, because of the presence of the discount factor in the utility function (21), the state price deflator declines with t . However, because of the growth in the endowment after the first occurrence of red, the deflator declines an order of magnitude faster when $\tau < t$ than otherwise.

The state price deflator at date t is the density of the measure representing prices at date t with respect to the probability measure restricted to \mathcal{F}_t . The relevant event probabilities are

$$\pi_t = \begin{cases} 2^{-\tau} & \text{if } \tau \leq t, \\ 2^{-t} & \text{if } \tau > t \end{cases}. \quad (24)$$

The event prices p_t at date t are the prices of one unit of consumption contingent on each of the events in \mathcal{F}_t , equal to the product of ρ_t and π_t :

$$p_t = \rho_t \pi_t = \begin{cases} 2^{-(2t+\tau)} & \text{if } \tau \leq t, \\ 2^{-2t} & \text{if } \tau > t \end{cases}. \quad (25)$$

Let r_t be the gross one-period interest rate from $t - 1$ to t . It is given by

$$r_t = \begin{cases} 4 & \text{if } \tau < t, \\ 4(1 + 2^{-t})^{-1} & \text{otherwise.} \end{cases} \quad (26)$$

This expression is the ratio of the event price at any node of the event tree to the sum of the event prices at its successor nodes. This is the usual expression for the one-period gross interest rate.

The money-market account is the cumulated value of one unit of consumption invested at the one-period interest rate at date 0 and rolled over at each date. Its date- t value b_t is

$$b_t = \prod_{i=1}^t r_i = \begin{cases} 2^{2t} \prod_{i=1}^{\tau} (1 + 2^{-i})^{-1} & \text{if } \tau \leq t, \\ 2^{2t} \prod_{i=1}^t (1 + 2^{-i})^{-1} & \text{if } \tau > t. \end{cases} \quad (27)$$

Define $M = \{m_t\}$ as the product of ρ and B . From (23) and (27), M_T is given by

$$m_T \equiv \rho_T b_T = \begin{cases} \prod_{i=1}^{\tau} (1 + 2^{-i})^{-1} & \text{if } \tau \leq T, \\ 2^T \prod_{i=1}^T (1 + 2^{-i})^{-1} & \text{if } \tau > T, \end{cases} \quad (28)$$

with preceding terms given by $m_t = E_t(m_T)$, where the expectation is taken under the natural probabilities.

The process M is a martingale under the natural probabilities.¹⁰ Its value at T is the Radon-Nikodym derivative of the change of measure from natural probabilities to risk-neutral probabilities: The risk-neutral probabilities π_T^* associated with a claim on one unit of consumption at T equal m_T multiplied by the relevant natural probabilities. From the definitions of ρ_T and m_T they can also be written as the event prices multiplied by the money-market account:

$$\pi_T^* = m_T \pi_T = p_T b_T = \begin{cases} 2^{-\tau} \prod_{i=1}^{\tau} (1 + 2^{-i})^{-1} & \text{if } \tau \leq T, \\ \prod_{i=1}^T (1 + 2^{-i})^{-1} & \text{if } \tau > T \end{cases}. \quad (29)$$

For finite T the risk-neutral probability measure π_T^* is obviously equivalent to the natural measure π , since both are nonzero at all events. When $T = \infty$ the risk-neutral measure is given by

$$\pi^* = \begin{cases} 2^{-\tau} \prod_{i=1}^{\tau} (1 + 2^{-i})^{-1} & \text{if } \tau < \infty, \\ \prod_{i=1}^{\infty} (1 + 2^{-i})^{-1} & \text{if } \tau = \infty \end{cases}. \quad (30)$$

the weak-* limit of π_T^* as T approaches ∞ .

By evaluating the $\tau = \infty$ term in (30), it is seen that the event that R never occurs, which has probability zero under the natural measure, has probability $\prod_{i=1}^{\infty} (1 + 2^{-i})^{-1} = 0.4194$ under the risk-neutral measure. The interpretation is that even though the event that R never occurs has probability zero under the natural measure, the payoff of a security or portfolio in that event is a major contributor to its value. This is so because individuals have an increasingly strong desire to insure against the event that R never occurs, since the aggregate endowment at date t is 1 in that case instead of 2^t .

The fact that π^* is not equivalent to π reflects the fact that M is not uniformly integrable, implying that it is not a closed martingale (defined as a martingale which has the property that its pointwise limit has the same expectation as the other terms).¹¹ In this case the expectation of the pointwise limit of m_t is $1 - 0.4194 = 0.5806$.

An observation suggested by the example is that the conditions under which a martingale measure exists are broader than those under which an equivalent martingale measure exists. Specifically, Alaoglu's theorem implies that if the sequence of price systems at each date is uniformly bounded, a weak assumption, then it has a

¹⁰For example, the date-0 value of M is 1, its date-1 values are $m_R = 2/3$, $m_B = 4/3$, its date-2 values are $m_{RR} = m_{RB} = 2/3$, $m_{BR} = 8/15$ and $m_{BB} = 32/15$.

¹¹In continuous time analyses, such as Loewenstein and Willard [14], the analogue of the martingale is the local martingale, and the analogue of the closed martingale is the martingale.

subnet that weak-* converges. The limiting price system may or may not be absolutely continuous with respect to probability; Alaoglu's theorem does not imply this, and in the example the limiting measure turned out not to be continuous. The conclusion is that even in infinite economies pricing measures exist under very weak assumptions, amounting to absence of arbitrage and weak boundedness conditions. The conditions under which these measures are equivalent to probabilities are more restrictive.

4 Conclusion

The examples discussed in this paper have a common structure: in each case a sequence of measures representing price systems or payoffs at finite dates converged to a non-equivalent measure representing a price system or payoff at infinity. This reflects a phenomenon that can occur only in models incorporating an infinite number of dimensions: measures with the same support can fail to be equivalent. In our first two examples we were interested in a sequence of measures μ_n each of which is absolutely continuous with respect to π . The weak-* limit of this sequence of measures is nonzero at infinity. As such, its support is N_∞ , but it is evidently not equivalent to π . In the finite-dimensional case, in contrast, this cannot occur: two measures are equivalent if and only if they have the same support.

References

- [1] Charalambos Aliprantis and Kim Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide, Second Edition*. Springer-Verlag, Berlin, 1999.
- [2] Kerry Back and Stanley R. Pliska. On the fundamental theorem of asset pricing with an infinite state space. *Journal of Mathematical Economics*, 20:1–18, 1991.
- [3] Stephen A. Clark. Arbitrage on a lattice of derivative securities. reproduced, University of Kentucky, 2006.
- [4] Philip H. Dybvig and Stephen A. Ross. Arbitrage. In J. Eatwell, M. Milgate, and P. Newman, editors, *The New Palgrave: A Dictionary of Economics*. McMillan, 1987.
- [5] Mark Fisher and Christian Gilles. Weak-* convergence and the doubling strategy. reproduced, Federal Reserve Bank of Atlanta, 2004.
- [6] Christian Gilles. Charges as equilibrium prices, and asset bubbles. *Journal of Mathematical Economics*, 18:155–167, 1989.

- [7] Christian Gilles and Stephen F. LeRoy. Asset price bubbles. In John Eatwell, Murray Milgate, and Peter Newman, editors, *The New Palgrave Dictionary of Money and Finance*. Macmillan Press, 1992.
- [8] Christian Gilles and Stephen F. LeRoy. Bubbles and charges. *International Economic Review*, 33:323–339, 1992.
- [9] Christian Gilles and Stephen F. LeRoy. Rational bubbles as payoffs at infinity. *Economic Theory*, 9:261–281, 1997.
- [10] Christian Gilles and Stephen F. LeRoy. Arbitrage, martingales and bubbles. *Economics Letters*, 60:357–362, 1998.
- [11] Larry E. Jones. A competitive model of commodity differentiation. *Econometrica*, 52:507–530, 1984.
- [12] David M. Kreps. Arbitrage and equilibrium in economies with infinitely many commodities. *Journal of Mathematical Economics*, pages 15–35, 1981.
- [13] Stephen F. LeRoy and Jan Werner. *Principles of Financial Economics*. Cambridge University Press, Cambridge, 2001.
- [14] Mark Loewenstein and Gregory A. Willard. Rational equilibrium asset-pricing bubbles in continuous-trading models. *Journal of Economic Theory*, 91:17–58, 2000.
- [15] Andreu MasColell. A model of equilibrium with differentiated commodities. *Journal of Mathematical Economics*, 2:263–295, 1975.
- [16] Merton H. Miller and Franco Modigliani. Dividend policy, growth, and the valuation of shares. *Journal of Business*, 34:411–433, 1961.
- [17] Jan Werner. Arbitrage, bubbles and valuation. *International Economic Review*, 38:453–464, 1997.