Compactifying the Payoff Space: Applications in Financial Economics

Stephen F. LeRoy
University of California, Santa Barbara

October 6, 2006

Abstract

Finance models with a finite number of states and dates have properties that do not extend to their infinite counterparts. It is proposed to deal with this situation by appending a date or state 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.

In the simplest analyses of financial economics, portfolios and security prices are analyzed in a setting that specifies that the number of dates and states is finite.[1] The only difference between this setting and that of classical general equilibrium theory is that the objects of choice are securities rather than commodities. Securities are defined by their payoffs, which are taken as given and which generally depend on the date and event. Portfolios—or, in the multidate case, portfolio strategies—are defined as vectors describing security holdings in each non-terminal date and event. Payoffs of portfolio strategies are defined in the obvious way: the payoff at each event equals the value of the portfolio chosen at the preceding node, plus dividends, less the value of the continuation portfolio. The set of portfolio payoffs—the *asset span*, in the terminology of LeRoy and Werner [7]—coincides with the consumption space if markets are complete, and is a proper subspace of the consumption space otherwise.

Assuming that equilibrium security prices satisfy the law of one price, each element of the asset span is associated with a number representing the initial cost of the portfolio strategy (strategies) that generates that payoff. The function describing

[1]This paper draws heavily on Fisher and Gilles [3]. I am indebted to Mark Fisher for many conversations about this material. Thanks to Martin Weitzman for comments, and to Yongli Zhang for correcting an error in an earlier draft.

this association is the *payoff pricing functional.* Each agent's consumption equals the sum of his endowment and the payoff of the chosen portfolio strategy. Assuming that agents have no endowments of securities, each agent's budget constraint implies that the set of feasible portfolio payoffs consists of those payoffs that have nonpositive value under the payoff pricing functional. Each agent, then, choses the zero-value portfolio strategy that maximizes the utility of the resulting consumption. All of this is standard.

1 Sequential Equilibrium

When time is infinite (but maintaining the assumption of discrete time, so that the number of dates is countable), the usual practice is to apply the above framework without substantial modification (see, for example, Santos and Woodford [12]). We will use the term *sequential equilibrium* to describe the above framework when applied to infinite-time settings. What is most relevant in the present context is that under sequential equilibrium it is assumed that the definition given above of payoffs of finite portfolio strategies—portfolio strategies that involve nonzero holdings at only a finite number of events—also applies to infinite portfolio strategies. Thus the payoff is described event-by-event as the current value of the portfolio chosen at the preceding node, plus dividends, less the value of the continuation portfolio.

When applied to infinite portfolio strategies the definition of payoffs leads to anomalies. The easiest way to show this is to look at examples.

1.1 Ponzi Schemes

The initial cost of any finite portfolio strategy (that is, any portfolio strategy that involves nonzero transactions at only a finite number of events) equals the value of its payoff given by the payoff pricing functional, defined above. This may no longer be true for infinite portfolio strategies when portfolio payoffs are generated according to the definition just presented. The classic example is the Ponzi scheme, in which an investor borrows money and rolls over the indebtedness forever. The payoff of the Ponzi scheme is zero at each date, yet its initial cost, equal to the initial borrowing, is negative.

Ponzi schemes as just characterized are arbitrages. The arbitrage arises from the fact that the liability implied by the existence of the loan never appears in the payoff, since at each date the loan is rolled over. The presumption is that agents regard a loan that does not have a specific repayment date as effectively free money. This is a substantive economic assumption about how agents view an infinite future. It may or may not be descriptively accurate.

Optimal portfolios do not exist if (unlimited) arbitrages are available. Therefore the analyst must impose trading restrictions so as to rule out Ponzi schemes. A variety of possible trading restrictions can be imposed. The simplest and most often encountered trading restriction is that the value of the position under any portfolio strategy must be uniformly bounded, which rules out Ponzi schemes if interest rates are positive. Different trading restrictions typically result in different equilibria (see, for example, Huang and Werner [5]). That being so, there does not exist a single canonical version of the sequential equilibrium model under infinite trading that corresponds to the standard Walrasian equilibrium that applies in the finite case.

With trading restrictions the law of one price is not necessarily satisfied. However, because of the existence of trading restrictions the arbitrage implied by failure of the law of one price cannot be operated on an unlimited scale. Therefore presence of arbitrage is consistent with the existence of optimal portfolios. The role of the trading restrictions is to ensure that associated with any payoff there exists a well-defined minimum among the initial costs of portfolio strategies that attain that payoff. Therefore the payoff pricing functional can be defined as the minimum initial cost at which any payoff can be attained. With the payoff pricing functional well defined, the characterization presented above of optimal consumption—the consumption bundle that maximizes utility subject to it equaling the sum of the endowment plus the payoff of a zero-value portfolio strategy—can be applied.

Analysts view the necessity of imposing trading restrictions to prevent Ponzi schemes as a necessary consequence of assuming an infinite future. In fact it is a consequence of the assumed definition of portfolio payoffs which, as observed above, implies that the indebtedness that the agent rolls over each period in implementing the Ponzi scheme never appears in its payoff.

1.2 Dividend Irrelevance

In finite settings and excluding frictions (specifically, assuming that a firm's reinvested retained earnings generate returns at the same rate 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. Assuming that invested capital generates returns at the same rate as dividend streams are discounted, the two effects cancel (Miller and Modigliani [9]). 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 limiting case the dividend stream under sequential markets is zero at each date. Therefore the firm value justified by dividends (that is, excluding bubbles) is apparently zero. If a zero dividend payout rate is admitted, the payoff pricing functional viewed as a function of the dividend payout rate is seen to be pointwise discontinuous. Under sequential markets the dividend-irrelevance proposition is apparently valid only if zero dividends are ruled out.

It is far from obvious that individuals view the stock of a firm that does not pay dividends as necessarily being worthless. There does not appear to be any economic justification for the assumption that the dividend stream in the limiting case coincides with the pointwise limit (as the dividend payout rate converges to zero) of the dividend streams when firms pay positive dividends. In the infinite case virtually all commonly used utility functions are discontinuous under pointwise convergence, suggesting that pointwise convergence may not give a useful definition of convergence. Yet under the sequential definition of portfolio payoffs there is no way to avoid the conclusion that the limiting payoff is zero, and therefore no way to avoid the discontinuity of valuation.

1.3 The Doubling Strategy

The classic example of the fact that sequential equilibrium implies the existence of arbitrage in infinite portfolio strategies, even when there is no arbitrage under finite portfolio strategies, is the doubling strategy. The doubling strategy is a gambling strategy that initially involves betting $1 on red on a roulette wheel. Assume for simplicity that the roulette wheel is actuarially fair. Upon winning, the gambler stops 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. The doubling strategy is generally treated as an arbitrage. As with Ponzi schemes, analysts take the view that it is necessary to rule out this trading strategy by imposing trading restrictions (for example, see the discussion in Pliska [10]).

The doubling strategy is essentially a stochastic version of the Ponzi scheme. To see this, it is simplest to recast the roulette wheel as a securities market. Suppose that at each date there exists a risky asset with price 1 that generates a payoff of 2 or 0 with equal probability (corresponding to the assumption that the roulette wheel is fair), and a riskless asset with a net interest rate of 0. Then agents can play the doubling strategy only if they can borrow without limit—that is, only if they can operate a Ponzi scheme. Of course, this consideration gives rise to the question why the agent would bother with the doubling strategy if he could operate a Ponzi scheme.

The uncertainty tree that represents the doubling strategy under sequential equilibrium shows a payoff of 1 the first time red appears and zero otherwise. The fact that before the first red appears the agent is in debt by increasing amounts never appears on the uncertainty tree. This occurs because at each intermediate date the agent reborrows the funds necessary to repay the preexisting loan, as with the Ponzi scheme.

Inspection of the uncertainty tree reveals that the event that red never occurs is impossible according to the definition of sequential equilibrium. Since that event does not appear anywhere on the uncertainty tree, a more precise statement would be that there is no event consisting of red never occurring. In fact, red never occurring is a well-defined event, although one that occurs with probability zero. It is not clear that it makes sense to analyze the doubling strategy in a setting in which the event of losing is ruled out by definition.

We argue that there is a difference between a non-event and a zero-probability event. It is true that mathematicians are accustomed to defining payoffs as equivalence sets of random variables, where payoffs that differ on a set of probability zero are identified. Economists, if required to justify the assumption that zero-probability events can be ignored, would point out that (under well-understood conditions) the value of a payoff is generally representable as the product of the payoff and an event price, where the latter is in turn the product of a probability and a marginal utility. If the probability is zero, the product is zero. However, this argument appears to presume that the relevant payoff is finite. In the case of the doubling strategy the payoff is (negative) infinity. Under the usual rules of multiplication under the extended real numbers the product of infinity and zero is undefined.[2] Therefore it is not clear that ruling out the event of red never occurring is the same as treating it as a zero-probability event.

As this analysis suggests, some analysts appear to be ambivalent about labeling the doubling strategy as an arbitrage. Delbaen and Schachermayer [2] concluded from the fact that losses are unbounded that "[e]verybody, especially a casino boss, knows that [the doubling strategy] is a very risky way of winning 1€. This type of strategy has to be ruled out: there should be a lower bound on the player's loss" ([2] p. 130). This passage is interesting on several levels. First, as a general matter there is no need to rule out very risky portfolio strategies. Assuming that agents are risk averse, they will not adopt very risky portfolio strategies even if they are available. Just the opposite: the necessity for ruling out the doubling strategy arises precisely from the fact that the doubling strategy is not risky under the sequential equilibrium definition of payoffs. Second, in labeling the doubling strategy as risky, Delbaen and Schachermayer make clear that they do not believe their own characterization of the payoff of the doubling strategy, which implies that it is riskless. As will become clear, we agree with this judgment.

It turns out that the transactions that constitute a Ponzi scheme under sequential markets do not create an arbitrage under the Arrow-Debreu equilibrium concept. Accordingly, there is no need to rule out such transactions by imposing transactions costs, although one can still impose such restrictions.

---

[2] Wolfram MathWorld and Wikipedia, "Extended Real Number Line" state that $0 \times \infty$ is usually left undefined. The latter adds that in probability and measure theory, in contrast, $0 \times \infty$ is sometimes defined as equaling 0. Royden [11] listed the assumption $0 \times \infty = 0$ separately from the others, identifying it as an "arbitrary *convention*" (emphasis in original). Aliprantis and Border [1], assumed $0 \times \infty = 0$ along with the other arithmetic rules involving ∞.

The rules of arithmetic for expressions involving ∞ are justified in terms of limits. None of the above sources presented a justification along these lines for $0 \times \infty = 0$.

1.4 Assessment

The preceding examples show that under the definition of portfolio payoffs adopted in sequential equilibrium the value of the portfolio at intermediate dates does not appear. This is not a problem with finite portfolio strategies because the value of the portfolio at the penultimate date is reflected in the payoff of the portfolio strategy at the terminal date. The same observation applies to infinite portfolio strategies for which the value of the portfolio (measured in units of consumption at the initial date) converges to zero. These are precisely the portfolio strategies for which the initial cost equals the limiting initial value of finite-date payoffs—that is, the portfolio strategies that do not contain bubbles (here we use the usual definition of a rational bubble as that part of an asset's value that exceeds the present value of its payoff—see LeRoy [**?**]).

For portfolio strategies that do contain bubbles—the examples considered in the preceding subsections all fit this description—the fact that intermediate portfolio values do not converge to zero means that omitting them from the definition of payoffs has implications. These implications are entirely a consequence of the assumed definition of portfolio payoffs under sequential equilibrium. They do not arise as a consequence of economic analysis or, for that matter, analysis of any other kind. There is no reason to suppose that they constitute reliable insights.

In the following section we propose a different definition of the payoffs of infinite portfolio strategies. We implement that definition in a setting that allows reexamination of the Ponzi scheme, the zero-dividend firm and the doubling strategy. It turns out that that the alternative payoff definition allows a more satisfactory treatment of these portfolio strategies than sequential equilibrium.

2 Arrow-Debreu Equilibrium

We call our alternative to sequential equilibrium "Arrow-Debreu equilibrium" because it is more in the spirit of received static Walrasian equilibrium than is sequential equilibrium. There is no assertion that Arrow and Debreu ever proposed anything along these lines.

Under Arrow-Debreu equilibrium payoffs of infinite portfolio strategies are defined to be the limits of payoffs of finite portfolio strategies. Specifically, the payoff of an infinite portfolio strategy is defined as the limit of the payoffs of the associated truncated—and therefore finite—portfolio strategies, which are defined in the usual way. That is, if h is an infinite portfolio strategy and h_n is the same portfolio strategy with the positions after date n set to zero, then the payoff of h is defined to be the limit of $\{h_n\}$. If the limit does not exist, the portfolio strategy is undefined. Of course, under this alternative the analyst must specify a framework in which the meaning of "limit" is clearly defined. This involves committing to a particular topology so that the limiting payoff is well-defined.

If the payoffs of infinite portfolio strategies are to be modeled as limits of payoffs of finite portfolio strategies, the setting is useful insofar as these limits exist for a wide range of portfolio strategies. This involves specifying a payoff index set that is compact. The natural numbers, however, are not compact in the topology to be specified. We remedy this problem by appending ∞ to the natural numbers, so that we are allowing portfolio payoffs that have nonzero components at ∞. When payoffs and price systems are constructed as function spaces based on the compactified natural numbers, they turn out to have desirable properties. In particular, the anomalies discussed above disappear.

2.1 Two Function Spaces on N_∞

We restrict attention to a countable payoff index set, which is therefore initially representable by N, where N denotes the natural numbers. The compactified version of N is $N_\infty \equiv N \cup \{\infty\}$. We define a metric d on N_∞ by

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

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

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

$i < \infty$. This metric is indicated if one associates $1, 2, 3, ...$ with $1/2, 3/4, 7/8,...$, and associates ∞ with 1. Under the topology associated with the distance measure just defined, ∞ is the only limit point of N_∞, so adding ∞ to N amounts to compactifying N.

A base for the topology associated with the distance measure defined by (1)-(2) consists of (1) the singleton subsets of N, plus (2) the tails of N_∞ (that is, the unions of ∞ and the elements N beyond some $n \in N$), so that, loosely, ∞ is indistinguishable from the very large elements of N. We will see below that this choice of topology will ensure that infinite portfolio strategies have well-defined payoffs under weak boundedness conditions.

Let C be the space of continuous real-valued functions on N_∞.[3] Under the topology on N_∞ just specified, C consists of the functions $f : N_\infty \to \mathcal{R}$ such that (1) $\lim_{n\to\infty} f(n)$ exists, and (2) $\lim_{n\to\infty} f(n) = f(\infty)$. 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 $\mu_n \to \mu$ if and only if $\int_{N_\infty} f d\mu_n \to \int_{N_\infty} f d\mu$ for all $f \in C$. Then C and M are topological duals

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

under

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

for $f \in C$, $\mu \in M$. In the applications that follow, N_∞ will be interpreted variously as the set of possible dates and as a probability space. Accordingly, the interpretations of C and M will differ depending on the context.

In the following subsections Ponzi schemes, the zero-dividend firm problem and the doubling strategy are analyzed using the Arrow-Debreu definition of payoffs.

2.2 Ponzi Schemes

In this section we take the payoff index set N_∞ 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.

It turns out that Ponzi schemes are not arbitrages under the Arrow-Debreu definition of equilibrium. To see this, we refer to the function spaces defined in the preceding section. We specify that payoffs are represented as elements of M (i.e., as measures on N_∞).[4] For example, the Dirac measure δ_t defined by

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

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 ∞.

The Alaoglu theorem states that the unit ball is compact under the weak-* topology. To see the implications of this result for the present analysis, suppose that each node of the uncertainty tree is assigned a portfolio. These portfolios define an infinite portfolio strategy only if the limit of the payoffs of the truncated portfolio strategies exists. The Alaoglu theorem implies that this will be the case as long as the payoffs of the truncated portfolio strategies are uniformly bounded. This is an acceptably weak restriction.

Now consider the portfolio strategy consisting of borrowing \$1 at date 0, investing it at the risk-free rate for n periods and paying off the loan at date n. This is, of course, the finite counterpart of the Ponzi scheme. The initial cost of this portfolio strategy is $-\delta_0$, as under sequential markets, and its payoff measured in units of date-0 consumption is $-\delta_n$. The infinitely-repeated version of this portfolio strategy has payoff $-\lim_{n\to\infty} \delta_n$. Imposition of the weak-* topology on N_∞ implies that as n approaches ∞, $\{-\delta_n\}$ converges to $-\delta_\infty$. To see this, take any arbitrary test function f from C. Then

[4] Examples of papers that use measures to represent commodities are MasColell [8] and Jones [6].

$$< f, \delta_n > = f(n) \text{ and } < f, \delta_\infty > = f(\infty). \qquad (5)$$

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

We see that the portfolio strategy that produces an arbitrage in the form of a Ponzi scheme under sequential markets does not produce an arbitrage under Arrow-Debreu markets. This is so because a loan that is paid back at date n in the finite case is paid back at ∞ in the limit. That being so, there is no need to include trading restrictions to make possible existence of optimal portfolios.

Under what we have called the Arrow-Debreu definition of payoffs, agents are treated as paying off (or receiving) at ∞ any value that remains in a portfolio strategy after the finite dates. In the case of the example the negative payoff at ∞ is quantitatively equal (in units of date-0 value) to the date-0 initial cost of the portfolio strategy. In contrast to the case for sequential markets, under Arrow-Debreu markets the value of the payoff of any portfolio strategy, finite or infinite, equals the initial cost of the portfolio strategy.

Because the limit is taken under units of date-0 value, the Arrow-Debreu definition of payoffs as we have defined them automatically imposes continuity of valuation. This is so because the value of any payoff is the initial cost of the portfolio that generates that payoff, and this cost is the same for all portfolios h_n that are truncated versions of some portfolio strategy h. Because the payoff of h is the limit of the payoffs of the h_n, and h is defined to have the same initial cost as the h_n, we have continuous valuation.

For portfolio strategies to which the present value formula applies, the sequential and Arrow-Debreu definitions of payoffs are substantively the same. In those cases the portfolios h_n have payoffs consisting of dividends up to date n and a terminal value. As n rises the terminal value (measured in units of date-0 consumption) converges to zero, by the definition of present value. Therefore adding ∞ to the payoff index set changes nothing, since the indicated portfolio payoff under the Arrow-Debreu definition has zero in this place. Thus the payoff implied by the Arrow-Debreu definition does not differ materially from the payoff as defined under sequential markets. It is only in portfolios like the Ponzi scheme, in which the limit does not converge to zero, that the two definitions differ.

Choosing between these two payoff definitions is, of course, a substantive modeling judgment. The Arrow-Debreu definition implies that agents treat money that is borrowed with no specific term as having to be paid back at some time in the future, possibly at ∞. It enforces this by requiring that any value not paid out in finite dates is paid out at ∞. In contrast, sequential equilibrium implies that, in the absence of trading restrictions, this sum never has to be paid (or is never received). Certainly the Arrow-Debreu formalization is more convenient analytically, since it dispenses with the requirement for trading restrictions.

2.3 Dividend Irrelevance

In the preceding example the elements of C, the dual space of M, were used to test convergence, not to represent price systems. Here we make explicit that price systems are modeled 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 the measure λ on N_∞ defined by

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

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

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

$$z_\gamma(t) = \gamma\rho(1+\rho(1-\gamma))^{t-1}. \tag{7}$$

The discounted dividend stream has typical element

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

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

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

so that for any $\gamma > 0$ dividend policy does not affect the value of the firm (Miller and Modigliani [9]).

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 (Miller and Modigliani [9], note 12). This is a troublesome finding: where did the firm's value go? However, the argument is far from water-tight. The payoff stream when $\gamma = 0$ coincides with 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, and therefore with the definition of portfolio payoffs as specified in sequential equilibrium.

The analysis leads in a different direction if we recast the firm model just outlined in the analytical framework of Section 2.1, 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

$$\mu_\gamma(A) = \sum_{t \varepsilon A} (1+\rho)^{-t} \gamma \rho (1 + \rho(1-\gamma))^{t-1}, \tag{10}$$

$A \subset N_\infty$. 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 ∞. Therefore the value of μ_γ is given by $< 1, \mu_\gamma >$. 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

$$\lim_{\gamma \downarrow 0} < 1, \mu_\gamma(\{T, \infty\}) > = 1 \tag{11}$$

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 $< 1, \delta_\infty > = 1$, which equals $< 1, \mu_\gamma >$ for any $\gamma > 0$. Because valuation is weak-* continuous, it is seen that modeling the payoff of the "zero"-dividend firm as the weak-* 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, since it is this assumption that allows us to extend valuation from the payoffs of finite portfolio strategies to payoffs with nonzero components at ∞. 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. Under the sequential equilibrium concept the payoff of the zero-dividend firm is the pointwise limit of the payoffs of the γ-dividend firm. Thus the presumption implicit in sequential equilibrium is 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 support for this presumption.

In substantive terms, the analysis just presented mirrors exactly that of Gilles and LeRoy [4]. 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 a 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.

2.4 The Doubling Strategy

The third example involves uncertainty in an essential way. 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 the theory of stochastic processes.

It is clear that, given the analysis above of the Ponzi scheme, the doubling strategy introduces nothing new. We discuss it primarily to demonstrate the application of the Arrow-Debreu equilibrium concept in a stochastic setting. Also, the doubling strategy provides a convenient context in which to make the point that sequential equilibrium involves some implicit presumptions that are, at best, questionable. Analyzing the doubling strategy as a sequential equilibrium presumes that real-world agents—either gamblers or investors—believe that they could manufacture certain wealth from each other by engaging in an infinite number of gambles, each of which is fair for each agent, if only they were able to do so. There is no reason to suppose that anyone is so foolish as to believe this, so why are we creating models that imply it?

Under the Arrow-Debreu equilibrium concept, on the other hand, the possibility of losing is a determinant of the final payoff even though this is a zero-probability event. Assuming that utility is weak-* continuous, so that it assigns utility to the payoffs of infinite portfolio strategies equal to the limit of the utilities of the corresponding finite portfolio strategies, the doubling strategy, far from being an arbitrage, is an extremely unattractive portfolio strategy under most utility functions that incorporate strict risk aversion. Therefore, as with Ponzi schemes, there is no need to impose trading restrictions so as to render the doubling strategy infeasible. Surely this is a more reasonable analysis than that implied by sequential equilibrium.

We now set out the specifics of the analysis of the doubling strategy under the Arrow-Debreu equilibrium concept. First note that, significantly, under the sequential equilibrium concept there does not exist a state corresponding to the event of red never occurring. Under the Arrow-Debreu equilibrium concept this event cannot be neglected even though it occurs with probability zero. Therefore we augment the probability space to allow for an infinite series of black spins. As above, we model the payoffs of infinite portfolio strategies—in particular, the payoff of the doubling strategy—as weak-* limits of the payoffs of the corresponding finitely-repeated portfolio strategies. In the case of a fair roulette wheel, to which we restrict our attention, this limit always exists and takes a very simple form, as we will see.[5]

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. We define a filtration $\mathcal{F}$, which consists of the σ-fields $(\mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2, ..., \mathcal{F}_\infty)$. Here $\mathcal{F}_0$ is the trivial partition, $\mathcal{F}_\infty$ is based on the power set of N_∞, and $\mathcal{F}_n$ is based on the partition $\{\{1\}, \{2\}, ..., \{n\}, \{\cup_{\tau=n+1}^{\infty}\{\tau\} \cup \{\infty\}\}\}$. The interpretation is that at date n the agent knows whether or not red has occurred by date n, and, if so, at what date this happened. 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.[6]

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

$$\pi(\{\tau\}) = \begin{cases} 2^{-\tau}, & \tau < \infty \\ 0, & \tau = \infty \end{cases} . \tag{12}$$

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

$$a_n\{\tau\} = \begin{cases} 1, & \tau \leq n \\ 1 - 2^n, & n < \tau \end{cases} , \tag{13}$$

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

[5] 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.

[6] Payoffs that are measurable only with respect to $\mathcal{F}_\infty$ involve a considerable element of abstraction. Rejecting the assumption that there exist payoffs that are measurable with respect to $\mathcal{F}_\infty$ 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 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, the payoff of the finitely-repeated doubling strategy expressed as a measure. Define a measure ν_n on $\mathcal{F}_\infty$ as the product of a_n and π:

$$\nu_n(\{\tau\}) = a_n(\{\tau\})\pi(\{\tau\}) = \begin{cases} 2^{-\tau}, & \tau \leq n \\ 2^{-\tau}(1-2^n), & n < \tau < \infty \\ 0, & \tau = \infty. \end{cases} \tag{14}$$

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

$$\mu_n(\{\tau\}) = \begin{cases} 2^{-\tau}, & \tau \leq n \\ 2^{-n} - 1, & \tau > n \end{cases}, \tag{15}$$

reflecting the fact that

$$\sum_{\tau=n+1}^{\infty} 2^{-\tau}(1-2^n) = 2^{-n} - 1. \tag{16}$$

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

$$\mu_\infty = \lim_{n\to\infty} \mu_n = \begin{cases} 2^{-\tau}, & \tau < \infty \\ -1, & \tau = \infty \end{cases}. \tag{17}$$

This expression for μ_∞ is intuitively plausible inasmuch as it coincides with (15) 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 (15), or the random variable a_n, given by (13). 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 (13) 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 probability, since $\{\infty\}$ is a null set for probability, 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 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 (17): μ_∞ is not a positive measure. This is a sensible conclusion: there is no economic or mathematical justification for the presumption, implicit in the sequential definition of equilibrium, 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. This is, of course, exactly the same conclusion as applied to Ponzi schemes.

We may compare the present analysis with that of Gilles and LeRoy [4]. 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.

3 Conclusion

In many papers the author presumes that the definitions of payoffs and budget constraints specified in sequential equilibrium is just as appropriate when the setting incorporates infinite portfolio strategies as in the finite case. Santos and Woodford [12] is one example. We question this presumption: sequential equilibrium has implications that may or may not be acceptable. We hae proposed an alternative characterization of equilibrium, one closer to the Walrasian paradigm, and have argued that this alternative framework produces a very different analysis—and, in our view, an analysis that is much more sensible—in cases where the two paradigms differ.

References

[1] Charalambos Aliprantis and Kim Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide, Second Edition.* Springer-Verlag, Berlin, 1999.

[2] Freddy Delbaen and Walter Schachermayer. *The Mathematics of Arbitrage.* Springer, 2006.

[3] Mark Fisher and Christian Gilles. Weak-* convergence and the doubling strategy. reproduced, Federal Reserve Bank of Atlanta, 2004.

[4] Christian Gilles and Stephen F. LeRoy. Rational bubbles as payoffs at infinity. *Economic Theory*, 9:261–281, 1997.

[5] Kevin X. D. Huang and Jan Werner. Asset price bubbles in Arrow-Debreu and sequential equilibrium. *Economic Theory*, 15:253–278, 2000.

[6] Larry E. Jones. A competitive model of commodity differentiation. *Econometrica*, 52:507–530, 1984.

[7] Stephen F. LeRoy and Jan Werner. *Principles of Financial Economics*. Cambridge University Press, Cambridge, 2001.

[8] Andreu Mas-Colell. A model of equilibrium with differentiated commodities. *Journal of Mathematical Economics*, 2:263–295, 1975.

[9] Merton H. Miller and Franco Modigliani. Dividend policy, growth, and the valuation of shares. *Journal of Business*, 34:411–433, 1961.

[10] Stanley R. Pliska. *Introduction to Mathematical Finance: Discrete Time Models*. Oxford University Press, Oxford, 1997.

[11] H. L. Royden. *Real Analysis: Second Edition*. Macmillan, New York, 1968.

[12] Miguel Santos and Michael Woodford. Rational asset pricing bubbles. *Econometrica*, 65:19–57, 1997.