Model-free Representation of Pricing Rules as Conditional Expectations*

Sara BIAGINI and Rama CONT

Università degli Studi di Perugia (Italy)
Email: s.biagini@unipg.it
Centre de Mathématiques Appliquées, Ecole Polytechnique (France)
Email: Rama.Cont@polytechnique.org

We formulate an operational definition for absence of *model-free* arbitrage in a financial market, in terms of a set of minimal requirements for the pricing rule prevailing in the market and without making reference to any 'objective' probability measure. We show that any pricing rule verifying these properties can be represented as a conditional expectation operator with respect to a probability measure under which prices of traded assets follow martingales. Our result does not require any notion of "reference" probability measure and is consistent with the formulation of model calibration problems in option pricing.
Key words: model-free arbitrage, martingales, fundamental theorem of asset pricing, pricing of contingent claims.

1. Introduction

1.1 Model-based vs model-free arbitrage

Stochastic models of financial markets represent the evolution of the prices of financial products as stochastic processes defined on some (filtered) probability space $(\Omega, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$, where it is usually assumed [9, 10, 13, 12, 14] that an "objective" probability measure $\mathbb{P}$, describing the random evolution of market prices, is given. Given a set of benchmark assets $(S_t)_{t\geq 0}$, described as semimartingales under $\mathbb{P}$, the gain of a trading strategy $(\phi_t)_{t\geq 0}$ is defined via the stochastic integral $\int \phi dS$ with respect to

*We acknowledge financial support from the European Network on Advanced Mathematical Methods in Finance.

the price processes. Then, one introduces the set of ($\mathbb{P}$-)*admissible trading strategies* as strategies with limited liability i.e. whose value is $\mathbb{P}$-a.s. bounded from below [9, 10]:

$$\phi \text{ is admissible if } \exists c \in \mathbb{R} \text{ such that for all } t, \; \mathbb{P}(\int_0^t \phi \, dS \geq -c) = 1$$

An *arbitrage opportunity* is then defined as an admissible strategy ϕ such that

$$\mathbb{P}(\int_0^T \phi \, dS \geq 0) = 1 \qquad \text{and} \qquad \mathbb{P}(\int_0^T \phi \, dS > 0) > 0, \tag{1}$$

a definition which depends on $\mathbb{P}$ through its null-sets.

The Fundamental Theorem of Asset Pricing [12], which is the theoretical foundation underlying the use of martingale methods in derivative pricing, is then loosely summarized as follows: roughly speaking, in a market where no such arbitrage opportunities exist, there exists a probability measure $\mathbb{Q}$ equivalent to $\mathbb{P}$ such that the (discounted) value $V_t(H)$ of any contingent claim with terminal payoff H is represented by:

$$V_t(H) = E^{\mathbb{Q}}[H|\mathcal{F}_t] \tag{2}$$

Loosely speaking: if the market is arbitrage-free, prices can be represented as conditional expectations with respect to some "equivalent martingale measure" $\mathbb{Q}$.

However, as noted by Kabanov [13], the precise formulation of this fundamental result is quite technical. In the case of market models with an infinite set of market scenarios, absence of arbitrage has to be replaced by a stronger condition known as No Free Lunch with Vanishing Risk [9, 10], which means requiring that, for any sequence of admissible strategies with terminal gains $f_n = \int_0^T \phi_n dS$, such the negative parts f_n^- tend to 0 uniformly and such that $\mathbb{P}(f_n \to f^*) = 1$ then $\mathbb{P}(f^*{=}0) = 1$. Under the NFLVR condition, one obtains [6, 9, 10, 13] the existence of a probability measure $\mathbb{Q}$ equivalent to $\mathbb{P}$ such that the (discounted) value $V_t(H)$ of a contingent claim with terminal payoff H is represented by:

$$V_t(H) = E^{\mathbb{Q}}[H|\mathcal{F}_t] \tag{3}$$

Furthermore, in the case of unbounded price processes the martingale property should be replaced by the weaker local martingale or "σ-martingale" properties [9, 10]. In addition, when asset prices are not locally bounded (as in a model with unbounded price jumps), the only admissible investments are those in the risk-free asset, which makes the above definitions somewhat trivial: the set of strategies needs to to be suitably enlarged [3, 4].

All these additional technical assumptions are less obvious to justify in economic terms. But perhaps the most important aspect of this characterization of absence of arbitrage in terms of "equivalent martingale measures" is the way an arbitrage opportunity (or free lunch) is defined: the definition explicitly refers to an objective probability measure $\mathbb{P}$. In financial terms, such a strategy is more appropriately termed a *model-based* arbitrage, where the term "model" refers to the choice of $\mathbb{P}$. The absence of arbitrage is then justified by saying that, if such an arbitrage opportunity would appear in the market, market participants ("arbitrageurs") would exploit it and make it disappear. This argument implicitly assumes that market participants are able to detect whether a given trading strategy is an arbitrage. Such a reasoning can be safely applied to model-free arbitrage opportunities: for instance, if discrepancies appear between an index and its components or if triangle arbitrage relations in foreign exchange markets are not respected, market participants will presumably trade on them. In fact this is the basis of many automated "program" trading strategies, which make such arbitrage opportunities short-lived.

But the argument is less obvious when applied to a model-based arbitrage. A model-based arbitrage opportunity is risk-free if the model $\mathbb{P}$ on which it is based is equivalent to the (unknown) one underlying the market dynamics. Once "model risk" – i.e. the possibility that $\mathbb{P}$ is misspecified– is taken into account, a model-based arbitrage is not riskless anymore. However model uncertainty cannot be ignored when dealing with the pricing of derivative instruments [7] and model-based arbitrage strategies can in fact be quite risky. Hence, market participants will attempt to exploit a model-based arbitrage opportunity if they believe that there is some market consensus on the underlying model i.e. that market prices will not move in a way which is precluded in the model.

However, in financial markets, and even more so in the context of derivative pricing, there is no consensus on the "underlying model" $\mathbb{P}$ [7]: the relevance of a definition of arbitrage which relies on the existence of a consensual or "objective" probability measure may thus be questioned.

Market consensus is expressed, not in terms of probabilities, but in terms of prices of various underlying assets and their derivatives traded in the market. It thus seems more natural to formulate the absence of arbitrage in terms of properties of market prices, that is, as constraints linking the relative values of traded instruments. Well-known constraints of this type are cash-and-carry arbitrage relations between spot and forward prices, spot relations between an index and its components, triangle relations between exchange rates, put-call parity relations, arbitrage inequalities linking values of call and put options of different strikes and maturities, in-out parity relations for barrier options.

Characterization of arbitrage-free price systems in terms of equivalent martingale measures also contrasts with the way the martingale pricing approach is commonly used in derivatives markets. Derivative pricing models are usually specified in terms of a (parametric) family $(\mathbb{Q}_\theta, \theta \in E)$ of "martingale measures" and the parameters θ of the pricing model are typically obtained by calibrating them to observed prices of various derivatives. The specification of an objective probability measure typically plays no role in this process. In fact, in most cases (Black-Scholes model, diffusion models, stochastic volatility models,..) the probability measures $(\mathbb{Q}_\theta, \theta \in E)$ are mutually singular so the model selection problem cannot be formulated as a search among martingale measures equivalent to a given measure $\mathbb{P}$ [2]. So, any characterization of absence of arbitrage in terms of *equivalent* martingale measure would appear as inconsistent with the practice of specifying and calibrating pricing rules in this way.

Our goal in the present work is to present a formulation of the martingale approach to derivative pricing which is

- consistent with the way arbitrage constraints are formulated by market participants, namely, in terms of market prices
- consistent with the way derivative pricing models are specified and calibrated in practice, that is, without referring to any "objective" probability measure.

We will start by formulating a set of minimal requirements for a pricing rule which can be interpreted as *absence of model-free arbitrage*. These requirements are formulated in terms of properties of prices (i.e. market observables), which is closer to the way arbitrage constraints are viewed in a financial markets, and without resorting to any reference probability measure.

We will then show that any pricing rule verifying these minimal assumptions can be represented by a conditional expectation operator with respect to a probability measure $\mathbb{Q}$ under which prices of traded assets are martingales ("martingale measure"). Our proof is based on simple probabilistic arguments. Our result can thus be viewed as a model-free version of the fundamental theorem of asset pricing.

1.2 Relation with previous literature

As noted above, classical formulations of the Fundamental Theorem of Asset pricing are based on the absence of *model-based* arbitrage (which includes model-free arbitrage as a special case). It is therefore interesting that one obtains a similar result with weaker assumptions. Since our result does not hinge on the existence of an objective probability measure, it is robust to model misspecification, an important issue in financial modeling. The relation of our framework to classical formulations of the Fundamental Theorem of Asset pricing are further discussed in Section 4.

A similar formulation of properties of pricing rules was proposed by Rogers [14]. In [14], a pricing rule was defined as a map on $L^\infty(\Omega, \mathbb{P})$ for some reference probability measure $\mathbb{P}$. Unlike [14], our formulation avoids any reference to a consensual or "objective" probability measure, and the set of contingent claims i.e. the domain of the pricing rule is determined a posteriori, not imposed a priori. We believe this renders our approach more general and more amenable to financial interpretation. This point is further commented upon in Section 4.2.

1.3 Outline

The article is structured as follows. In Section 2 we discuss some reasonable and financially meaningful requirements for a *pricing rule* and formulate them in mathematical terms. In Section 3 we characterize any pricing rule verifying these requirements as conditional expectation with respect to a martingale measure. Section 4 discusses some implications of our result and its relation to previous literature on arbitrage theorems.

2. Definitions and Notations

Let $(\Omega, (\mathcal{F}_t)_{t\in[0,T]})$ be the set of market scenarios endowed with a filtration $(\mathcal{F}_t)_{t\in[0,T]}$ representing the flow of information with time (in particular, $\mathcal{F}_0$ is trivial). Let $\mathcal{L}^0$ denote the space of $\mathbb{R}$-valued, $\mathcal{F}_T$-measurable random variables, representing payoffs of contingent claims and let $\mathcal{L}^\infty$ denote the subspace of bounded variables.

Let $\mathcal{Y}$ be the set of the non-anticipative processes

$$Y : \Omega \times [0, T] \to \mathbb{R} \cup \{+\infty, -\infty\}$$

i.e. such that for each t, Y_t is $\mathbb{R} \cup \{+\infty, -\infty\}$-valued and $\mathcal{F}_t$-measurable.

A pricing rule can be seen as an operator $\Pi : \mathcal{L}^0 \to \mathcal{Y}$ which assigns a price process $\Pi_t(H)$ to each contingent claim $H \in \mathcal{L}^0$. Note that a pricing rule does not necessarily assign a finite price to all payoffs $H \in \mathcal{L}^0$. Denote by $\mathrm{Dom}(\Pi)$ the *domain* of Π, that is, the set of payoffs with a finite price:

$$\mathrm{Dom}(\Pi) \triangleq \{G \in \mathcal{L}^0 \mid \Pi(G) \text{ is finite valued}\}$$

We can now formulate the minimal requirements for a pricing rule via the following definition:

Definition 2.1. A pricing rule is a mapping

$$\begin{aligned} \Pi : \mathcal{L}^0 &\to \mathcal{Y} \\ H &\mapsto (\Pi_t(H))_{t\in[0,T]} \end{aligned} \tag{4}$$

that satisfies the following properties:

A1 If $G, H \in \mathrm{Dom}(\Pi)$, then $K = \max(G, H) \in \mathrm{Dom}(\Pi)$.

A2 Positivity. For any $H \in \mathcal{L}^0$, if $H \geq 0$, then $\Pi(H) \geq 0$.

A3 $\mathcal{F}_t$-linearity on $\mathrm{Dom}(\Pi)$: For any $H_1, H_2 \in \mathrm{Dom}(\Pi)$ and any bounded $\mathcal{F}_t$-measurable variable λ, $\lambda H_1 + H_2 \in \mathrm{Dom}(\Pi)$ and

$$\Pi_t(\lambda H_1 + H_2) = \lambda \Pi_t(H_1) + \Pi_t(H_2) \tag{5}$$

A4 Time consistency.

$$\forall H \in \mathcal{L}^0, \qquad \Pi_s(\Pi_t(H)) = \Pi_s(H) \;\; 0 \leq s \leq t \leq T$$

A5 Normalization. $\Pi(1) = 1$.

A6 Market consistency. If H is tradable at price $(V_t)_{t \in [0,T]}$ in the market (whence in particular $H = V_T$), then $H \in \mathrm{Dom}(\Pi)$ and

$$\forall t \in [0, T], \forall \omega \in \Omega, \qquad \Pi_t(H)(\omega) = V(t, \omega). \tag{6}$$

A7 Continuity. If $(H_n)_{n \geq 1}$ is an increasing sequence in $\mathcal{L}^0$, uniformly bounded from below, with $H_n \uparrow H$, then $\Pi_0(H_n) \uparrow \Pi_0(H)$.

Let us comment on the various elements in this definition.

The requirement that $\Pi(H)$ is non-anticipative simply means that the pricing rule only makes use of information available at t in order to assign the price at time t to a claim.

Also, it is quite natural that $\Pi_t(H)$ is $\mathbb{R} \cup \{+\infty, -\infty\}$ valued. For example, some payoffs H may carry a huge downside risk that no market participant is willing to assume at any price: this formally translates into $\Pi(H) = -\infty$.

A1 This property means that, if H and G are two payoffs priced in the market then the option to exchange them i.e. $\max(H, G)$ is also priced in the market. Together with [A5], it ensures that, if an asset S is priced in the market then the most common derivatives on S, namely calls and puts, also belong to the domain of Π.

A2 Positivity ensures that the pricing rule verifies model-free static arbitrage inequalities. For instance, it guarantees that the price of call options is decreasing and the price of a put option is increasing with respect to its strike.

A3 $\mathcal{F}_t$-linearity on Dom(Π) expresses additivity of prices plus the fact that the value of a position, when computed at time t, scales linearly when we multiply the size of the position by a factor which is known at t (i.e. $\mathcal{F}_t$-measurable). This property obviously implies linearity: Dom(Π) is thus a vector space.
In financial terms, linearity together with (A2) guarantees that the price of call and put options is convex in the strike price.

A4 Time consistency rules out "cash and carry" arbitrage strategies for traded assets. It ensures for instance that forward contracts on traded assets are priced consistently with their underlyings.

A5 Normalization simply means that we are dealing with prices expressed in units of a given numeraire.[1] Since (A2) and (A3) imply that Π is monotone, a consequence of the normalization condition is that $\mathcal{L}^\infty \subset$ Dom(Π).

A6 Market consistency means that the pricing rule is compatible with observed market prices. It reflects the fact that pricing rules used by market operators are "calibrated" to prices of instruments (underlyings, derivatives) whose prices are observed in the market. Together with the linearity condition (A2), it implies put-call parity for calls and puts on traded assets.

A7 By the positivity property, if $(H_n)_{n\geq 1}$ is a monotone (increasing to H) sequence of payoffs then $(\Pi_0(H - H_n))_{n\geq 0}$ is a decreasing and positive sequence so it has a limit. So the continuity condition boils down to requiring continuity from above at zero for $\Pi_0(.)$. This is a rather weak continuity requirement, which excludes unrealistic specifications of pricing rules which would allocate very different prices to very similar payoffs.

Remark 2.1. (Vector lattice property)Properties [A1], [A2], [A3] and [A5] imply that the set Dom(Π) of payoffs with a finite price forms a vector lattice that contains $\mathcal{L}^\infty$ (see [1] for definitions).

3. Pricing Rules as Conditional Expectation Operators

Let us start by showing that, for any market-consistent "martingale" measure $\mathbb{Q}$, the conditional expectation operator with respect to $\mathbb{Q}$ defines a pricing rule in the sense of Definition 2.1:

[1] One could rewrite the whole formalism with the apparently (but not really) more general condition $0 < \Pi(1) \leq 1$.

Proposition 3.1. *Let $\mathbb{Q}$ be a probability measure defined on $(\Omega, (\mathcal{F}_t)_{t\in[0,T]})$ such that the prices $V_t(H)$ of all traded assets are martingales with respect to $\mathbb{Q}$. There exists a pricing rule Λ such that*

1. $\mathrm{Dom}(\Lambda)$ *is the vector space $\mathcal{L}^1(\mathbb{Q})$ of $\mathbb{Q}$-integrable payoffs ;*
2. *For any $H \in \mathrm{Dom}(\Lambda)$,*

$$\Pi_t(H) = E^{\mathbb{Q}}[H|\mathcal{F}_t] \qquad \mathbb{Q} - a.s. \tag{7}$$

Proof. For a $\mathbb{Q}$-integrable payoff H one can define $\Lambda(H)$ as (a version of the) $\mathbb{Q}$-conditional expectation of H, as in (7). To define a pricing rule, we need to extend Λ to the entire space $\mathcal{L}^0$, i.e. also to non-integrable payoffs.

For a positive payoff G, $E_{\mathbb{Q}}[\,G \mid \mathcal{F}_t]$ is always well-defined, with values in $\mathbb{R} \cup \{+\infty\}$. Let us fix a general payoff H and call $(\alpha_t)_t$ a version of $(E_{\mathbb{Q}}[\,|H| \mid \mathcal{F}_t])_t$. For all $t \le T, k \in \mathbb{N}$ consider the $\mathcal{F}_t$-measurable sets

$$A_{k,t} = \{k \le \alpha_t < k+1\}$$

Fix t and for any $A_{k,t}$ select a version $f_{k,t}$ of $E_{\mathbb{Q}}[HI_{A_{k,t}} \mid \mathcal{F}_t]$ and define

$$\Lambda_t(H) = f_{k,t} \text{ on } A_{k,t}$$
$$\Lambda_t(H) = +\infty \text{ on } \Omega - \cup_k A_{k,t},$$

$\Lambda(H)$ thus defines an element of $\mathcal{Y}$. It is very easy to see that $\mathrm{Dom}(\Lambda) = \mathcal{L}^1(\mathbb{Q})$, i.e. it is the space of $\mathbb{Q}$-integrable payoffs. On this space $\Lambda_t(H)$ satisfies (7).

The properties (A1), (A2), (A3), (A4), (A5), (A7) of a pricing rule are easily verified for Λ and to obtain (A6) when H is tradable, simply choose $\Lambda_t(H)$ to be the version of $E^{\mathbb{Q}}[H|\mathcal{F}_t]$ that coincides with $V_t(H)$. □

We now state our main result, which shows that any pricing rule can be represented as a conditional expectation with respect to a "martingale measure" $\mathbb{Q}$:

Theorem 3.1. *Given a pricing rule Π, there exists a probability measure $\mathbb{Q}$ defined on $(\Omega, \mathcal{F}_T)$ such that Π coincides with the conditional expectation with respect to $\mathbb{Q}$. More precisely:*

1. $\mathrm{Dom}(\Pi)$ *is the vector space $\mathcal{L}^1(\mathbb{Q})$ of $\mathbb{Q}$-integrable payoffs ;*
2. *For any $H \in \mathrm{Dom}(\Pi)$,*

$$\Pi_t(H) = E^{\mathbb{Q}}[H|\mathcal{F}_t] \qquad \mathbb{Q} - a.s. \tag{8}$$

3. *Prices of traded assets are $\mathbb{Q}$-martingales.*

Proof. Define $\mathbb{Q}$ on $\mathcal{F}_T$ by

$$\forall A \in \mathcal{F}_T, \qquad \mathbb{Q}(A) = \Pi_0(1_A)$$

It is not difficult to see that $\mathbb{Q}$ is a probability measure. In fact, $\mathbb{Q}$ is positive by positivity of Π, additive by linearity of Π and normalized. Furthermore, the continuity property (A7) of Π_0 implies the monotone convergence property for $\mathbb{Q}$, which is therefore a probability. Define a simple payoff as an element $H \in \mathcal{L}^\infty$ of the form

$$H = \sum_{i=1}^{n} c_i 1_{A_i}, \qquad A_i \in \mathcal{F}_T, \quad c_i \in \mathbb{R}.$$

Since Π is linear, for any simple payoff H we have $\Pi_0(H) = E^{\mathbb{Q}}[H]$. A general $H \in \mathcal{L}^0, H \geq 0$ can be approximated from below by a monotone sequence $(H_n)_{n\geq 1}$ of simple payoffs:

$$H_n \uparrow H$$

Using the monotone convergence theorem for $\mathbb{Q}$-expectation and the continuity property (A7) for Π, we can pass to the limit in $E^{\mathbb{Q}}[H_n] = \Pi_0(H_n)$ and we thus obtain

$$\Pi_0(H) = E^{\mathbb{Q}}[H]$$

If H is in $\mathcal{L}^1(\mathbb{Q})$, both its positive and negative part H^+, H^- have finite $\mathbb{Q}$-expectation and by additivity of $\mathbb{Q}$ and Π we get $\Pi_0(H) = E^{\mathbb{Q}}[H]$. If H is not integrable, then either $\Pi_0(H^+) = E^{\mathbb{Q}}[H^+]$ or $\Pi_0(H^-) = E^{\mathbb{Q}}[H^-]$ is infinite. By property (A1), $\Pi_0(H)$ cannot be finite. In particular, we obtain $\mathrm{Dom}(\Pi) \subseteq \mathcal{L}^1(\mathbb{Q})$ but not equality yet, since we need more properties to control Π_t when $t > 0$.

Let then $H \in \mathcal{L}^1(\mathbb{Q})$ and fix $t \in [0, T]$. Applying $\mathcal{F}_t$-linearity and time consistency, for any $A \in \mathcal{F}_t$ we have that $\Pi_0(1_A H)$ is finite and coincides with $\Pi_0(1_A \Pi_t(H))$. Hence, for any $A \in \mathcal{F}_t$

$$E^{\mathbb{Q}}[1_A H] = E^{\mathbb{Q}}[1_A \Pi_t(H)] \tag{9}$$

which characterizes $\Pi_t(H)$ as a version of the $\mathbb{Q}$-conditional expectation of H with respect to $\mathcal{F}_t$. This shows also that $\mathrm{Dom}(\Pi)$ coincides with $\mathcal{L}^1(\mathbb{Q})$. Finally, property (A6) of Π entails that if V is the market price of a traded asset H, then V is a version of the $\mathbb{Q}$-martingale with terminal value H:

$$V_t = \Pi_t(H) = E^{\mathbb{Q}}[H \mid \mathcal{F}_t] \qquad \square$$

Remark 3.1. (Continuity of Π) Inspecting the first part of the above proof shows that we could have recovered $\mathbb{Q}$ also from the restriction of Π to $\mathcal{L}^\infty \subseteq \mathrm{Dom}(\Pi)$. In particular, it would have been enough to consider the (linear, positive) functional $\psi : \mathcal{L}^\infty \to \mathbb{R}$ defined by:

$$\psi(H) = \Pi_0(H)$$

If we endow $\mathcal{L}^\infty$ with the uniform norm, it is a Banach space (in fact, a Banach lattice). Hence, thanks to [1] , ψ is already norm-continuous and so it can be identified with a measure $\mathbb{Q}$ on $(\Omega, \mathcal{F}_T)$. But without any extra condition, $\mathbb{Q}$ is a finitely additive measure but not a probability measure in general. To get countable additivity, we need the continuity condition property (A7), which amounts to requiring order-continuity of ψ.

4. Discussion

We have characterized pricing rules defined on $\mathcal{L}^0$ as conditional expectation operators with respect to a probability measure $\mathbb{Q}$ such that prices of traded assets are $\mathbb{Q}$-martingales. Our characterization does not require any a priori restriction on the domain of the pricing rule or the existence of a reference probability measure. We now examine some of the consequences of this result and its relation with previous characterizations of absence of arbitrage.

4.1 Implications for the specification of derivative pricing models

In contrast with previous formulations of no-arbitrage theorems, our result does not include any reference to an "objective" probability measure $\mathbb{P}$. In particular, we characterize internally consistent pricing models in terms of "martingale measures" without requiring that these martingale measures be *equivalent* to a reference probability measure $\mathbb{Q}$.

This is consistent with the way derivative pricing models are specified and used in the market. In practice, one does not necessarily start by identifying/ specifying an "objective" probability measure $\mathbb{P}$ and then subsequently look for a suitable martingale measure $\mathbb{Q}$ compatible with market prices, among those equivalent to $\mathbb{P}$. Instead, common practice is to specify a derivative pricing model in terms of a (parametric) family $(\mathbb{Q}_\theta, \theta \in E)$ of "martingale measures" and select the parameter θ of the pricing model are typically obtained by calibrating them to observed prices of various derivatives. The specification of an objective probability measure typically plays no role in this process. In fact, in most cases (Black-Scholes model, diffusion models, stochastic volatility models,..) the probability measures $(\mathbb{Q}_\theta, \theta \in E)$ are mutually singular: for example, if $\mathbb{Q}_\sigma$ designates a Black-Scholes model with volatility parameter σ then $\sigma_1 \neq \sigma_2$ entails that $\mathbb{Q}_{\sigma_1}$ and $\mathbb{Q}_{\sigma_2}$ are mutually singular measures. So, the model selection problem cannot be formulated as a search among martingale measures equivalent to a given measure $\mathbb{P}$ [2].

Therefore, while any characterization of absence of arbitrage in terms of equivalent martingale measure would appear as inconsistent with this (commonly used) way of specifying and calibrating pricing models, our result provides a justification for it: it simply reflects the fact that there is no consensus in the market on the "objective" probability and not even on its equivalence class.

4.2 The domain of the pricing rule

Another common feature of previous formulations of the absence of arbitrage is that the set of contingent claims is chosen in advance, either as $L^\infty(\Omega, \mathbb{P})$ or $L^p(\Omega, \mathbb{P}),\ p \geq 1$ for some reference measure $\mathbb{P}$. In practice the set of payoffs is defined independently from any probability measure: it typically contains unbounded payoffs whose integrability with respect to a given probability measure is not determined a priori, so this approach does not seem very natural.

In our approach, a pricing rule is defined on $\mathcal{L}^0$ -the set of *all* possible payoffs- and the domain of the pricing rule is determined a posteriori.

We find this approach financially meaningful. In fact, the simplest derivatives –call options– have unbounded payoffs and are priced on the market, so taking the set of payoffs to be $L^\infty(\mathbb{P})$ —as in [14]— seems restrictive. Of course, the pricing operator defined in this way can be then extended but this may lead to further mathematical issues (which should be the right extension to use? Is the resulting extension market-consistent?). In our setting, market consistency is guaranteed a priori and as a consequence of our result $\mathrm{Dom}(\Pi)$ turns out a posteriori to be the space $\mathcal{L}^1(\mathbb{Q})$.

4.3 Introduction of a set of benchmark assets

Suppose that a pricing rule Π is given on the market. Consider now a set of d benchmark price processes $S_1, \cdots, S_d$ (the so-called *underlyings*). We will now show how to recover the local-martingale or $\sigma-$martingale properties for $S = (S^1, \cdots, S^d)$ (see e.g. [4, 13, 10]) within our framework.

In the usual approach, to build gain processes of trading strategies as stochastic integrals one requires that S is an $\mathbb{R}^d$-valued semimartingale with respect to the reference probability $\mathbb{P}$. In our model-free context the natural counterpart is the assumption that S is a $\mathbb{Q}$-semimartingale. One can then introduce stochastic integrals with respect to S and define a notion of replicating strategy:

Definition 4.1. Given a pricing rule Π on the market, represented by a martingale measure $\mathbb{Q}$ and an $\mathbb{R}^d$-valued $\mathbb{Q}$-semimartingale S, a payoff $H \in \mathcal{L}^0$ is said to be S-replicable if there exist a $x \in \mathbb{R}$ and a predictable process (*strategy*) φ such that:

1. φ is S-integrable under $\mathbb{Q}$.
2. $\mathbb{Q}(\ \Pi_t(H) = x + \int_0^t \varphi dS\ ) = 1$.

Remark 4.1. In the above definition and in what follows probabilistic notions are induced by the pricing rule through its representing $\mathbb{Q}$.

Delbaen and Schachermayer [10] linked the No Free Lunch with Vanishing Risk property under $\mathbb{P}$ with the existence of a probability measure equivalent to $\mathbb{P}$ under which S is a σ-martingale, a notion introduced in [5]. We will now show how the σ-martingale property appears in our context.

Let us recall the following result from Emery [11]:

Proposition 4.1. *[11] Let S be a d-dimensional semimartingale on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{Q})$ and denote by $L(S)(\mathbb{Q})$ the set of predictable and S-integrable processes under the probability measure $\mathbb{Q}$. The following assertions are equivalent:*

1. *there exist a d-dimensional $\mathbb{Q}$- martingale N and a positive (scalar) process $\psi \in \cap_{1\leq i\leq d} L(N^i)(\mathbb{Q})$ such that $S^i = \int \psi\, dN^i$;*
2. *there exists a countable predictable partition $(B_n)_n$ of $\Omega \times \mathbb{R}_+$ such that $\int I_{B_n} dS^i$ is a $\mathbb{Q}$-martingale for every i, n;*
3. *there exist (scalar) processes η^i with paths that $\mathbb{Q} - a.s.$ never touch zero, such that $\eta^i \in L(S^i)(\mathbb{Q})$ and $\int \eta^i dS^i$ is a $\mathbb{Q}$-local martingale.*

Definition 4.2. We say that S is a σ-martingale under $\mathbb{Q}$ if it satisfies any of the equivalent conditions of the above Proposition.

The above equivalences illustrate that the σ-martingale property is a generalization of the local martingale property.

Remark 4.2. Whenever the the $(B_n)_n$ can be written as stochastic intervals $]T_n, T_{n+1}]$ where T_n is a sequence of stopping times increasing to $+\infty$, then the previous definition coincides with that of local martingale.

If for some i the process S_i is not the market price of a traded asset (but, for instance, a non-traded risk factor such as an instantaneous forward rate or instantaneous volatility process) then S_i is *not necessarily a martingale*. However, the result by Emery allows us to recover the σ-martingale features of S under $\mathbb{Q}$ from a straightforward analysis of *the market spanned by S.* Roughly speaking, there must be a traded derivative H with underlying S, which is S-replicable via a hedging strategy that is always non zero:

Proposition 4.2. *Suppose that for all i there exists an S^i-replicable derivative H^i traded in the market with a strategy $(\varphi_t^i)_{t\in[0,T]}$ that $\mathbb{Q}$-a.s. never touches zero. Then S is a σ-martingale under $\mathbb{Q}$.*

Proof. Since H^i is traded with market price $V^i = \Pi(H^i)$, our Theorem 3.1 implies that the gain $\int \varphi^i dS^i$ is a $\mathbb{Q}$-martingale. Then, given the assumption on the φ^is, S is a σ-martingale under $\mathbb{Q}$ from a direct application of item 3, Proposition 4.1. □

Remark 4.3. (The 'No Free Lunch with Vanishing Risk' property) If S is indeed a $\sigma-$martingale under $\mathbb{Q}$, then the market spanned by S satisfies the NFLVR condition with respect to $\mathbb{Q}$ (and henceforth with respect to any $\mathbb{P} \sim \mathbb{Q}$). In fact, consider the $\mathbb{Q}$-admissible strategies φ i.e. whose gain processes are $\mathbb{Q}$ -almost surely bounded from below:

$$\exists c > 0, \mathbb{Q}(\int \varphi dS \geq -c) = 1$$

If S is a σ-martingale under $\mathbb{Q}$, such strategies give rise to gain processes which are $\mathbb{Q}$-supermartingales (see e.g. [10]). Hence absence of arbitrage obviously holds, since $E_{\mathbb{Q}}[\int_0^T \varphi dS] \leq 0$. An application of Fatou's Lemma then shows that NFLVR also holds.

References

1. Aliprantis, C. D. and Border, K. C. (1999) *Infinite-dimensional analysis: a hitchhiker's guide*. Springer, Berlin.
2. S. Ben Hamida and R. Cont, Recovering volatility from option prices by evolutionary optimization, *Journal of Computational Finance*, 8 (2005), pp. 43–76.
3. Biagini, S. (2005). Convex Duality in Financial Theory with General Semimartingales. PhD Thesis, Scuola Normale Superiore.
4. Biagini, S. and Frittelli, M. (2005) *Utility Maximization in Incomplete Markets for Unbounded Processes*. Finance and Stochastics 9, 493–517.
5. Chou, C. S. (1979) Caractérisation d'une classe de semimartingales. *Sminaire de probabilités XIII*, p. 250–252, Berlin: Springer.
6. Cherny, A. and Shiryaev, A. (2002) Vector stochastic integrals and the fundamental theorem of asset pricing, *Proceedings of the Steklov Institute of mathematics*, 237, pp. 6–49.
7. Cont, R. (2006) Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance*, 16, No. 3, 519–547.
8. Dalang, R. C., Morton, A., and Willinger, W. (1990) Equivalent martingale measures and no-arbitrage in stochastic securities market models, *Stochastics Stochastics Rep.* 51, no. 1–2, 41–49.

9. Delbaen, F. and Schachermayer, W. (1994) A General Version of the Fundamental Theorem of Asset Pricing. *Math. Ann.*, Vol. 300, pp. 463–520.
10. Delbaen, F. and Schachermayer, W. (1998) The fundamental theorem of asset pricing for unbounded stochastic processes, *Math. Ann.*, 312, pp. 215–250.
11. Emery M. (1980) Compensation de processus à variation finie non localement intégrables, *Sém. Prob. XIV, Lecture Notes in Mathematics* Vol. 784, 152–160, Berlin: Springer.
12. Harrison, J. M. and Pliska, S. R. (1981), *Martingales and stochastic integrals in the theory of continuous trading*, Stochastic Process. Appl., 11, pp. 215–260.
13. Kabanov, Yuri (2001) Arbitrage theory, in: J Cvitanic, E Jouini and M Musiela (eds.): *Option pricing, interest rates and risk management*, Cambridge University Press.
14. Rogers, L. C. G. (1998) The origins of risk-neutral pricing and the Black-Scholes formula. *Risk Management and Analysis*, 2, 81–94.
15. Ross, Stephen (1978) A Simple Approach to the Valuation of Risky Streams. *Journal of Business*, vol. 51, No. 3, 453–475.