Pricing and hedging contingent claims with liquidity costs and market impact

Frédéric Abergel, Grégoire Loeper

Pricing and hedging contingent claims with liquidity costs and market impact

Frederic Abergel and Grégoire Loeper
Chair of Quantitative Finance
Laboratory of Mathematics Applied to Systems,
École Centrale Paris, 92290 Châtenay Malabry, France

September 6, 2013

Abstract

We study the influence of taking liquidity costs and market impact into account when hedging a contingent claim, first in the discrete time setting, then in continuous time. In the latter case and in a complete market, we derive a fully non-linear pricing partial differential equation, and characterizes its parabolic nature according to the value of a numerical parameter naturally interpreted as a *relaxation coefficient* for market impact. We then investigate the more challenging case of stochastic volatility models, and prove the parabolicity of the pricing equation in a particular case.

Introduction

There is a long history of studying the effect of transaction costs and liquidity costs in the context of derivative pricing and hedging. Transaction costs due to the presence of a Bid-Ask spread are well understood in discrete time, see [9]. In continuous time, they lead to quasi-variational inequalities, see e.g. [18], and to imperfect claim replication due to the infinite cost of hedging continuously over time. In this work, the emphasis is put rather on **liquidity costs**, that is, the extra price one has to pay over the theoretical price of a tradable asset, due to the finiteness of available liquidity at the best possible price. A reference work for the modelling and mathematical study of liquidity in the context of a dynamic hedging strategy is [3], see also [15], and our results can be seen as partially building on the same approach.
It is however unfortunate that a major drawback occurs when adding liquidity costs: as can easily be seen in [3] [13] [15], the pricing and hedging equation are not unconditionally parabolic anymore and, therefore, only a local existence and uniqueness of smooth solutions may be available. Note that this drawback can easily be inferred from the very early heuristics in [10]: the formula suggested by Leland makes perfectly good sense for small perturbation of the initial volatility, but is meaningless when the modified volatility becomes negative. An answer to this problem is proposed in [4], where the authors introduce *super-replicating* strategies and show that the minimal cost of a super-replicating strategy solves a well-posed parabolic equation. Still, a partial conclusion is that incorporating liquidity cost leads to ill-posed pricing equation for large option positions, a situation which cannot be considered satisfactory and hints at the fact that some ingredient may be missing in the physical modelling of the market. It turns out that this missing ingredient is precisely the **market impact** of the delta-hedger, as will become clear from our results. This fact is already observed by the second author in [12], where a well posed, fully non-linear parabolic equation is obtained using a simple market impact model.
Motivated by the need for quantitative approaches to algorithmic trading, the study of **market impact** in order-driven markets has become a very active research subject in the past decade. In a very elementary way, there always is an instantaneous market impact - termed *virtual impact* in [17] - whenever a transaction

takes place, in the sense that the best available price immediately following a transaction may be modified if the size of the transaction is larger than the quantity available at the best limit in the order book. As many empirical works show, see e.g. [2] [17], a relaxation phenomenon then takes place: after a trade, the instantaneous impact decreases to a smaller value, the permanent impact. This phenomenon is named **resilience** in [17], it can be interpreted as a rapid, negatively correlated response of the market to large price changes due to liquidity effects. In the context of derivative hedging, it is clear that there are realistic situations - e.g., a large option on an illiquid stock - where the market impact of an option hedging strategy is significant. This situation has already been addressed by several authors, see in particular [16] [7] [6] [14], where various hypothesis on the dynamics, the market impact and the hedging strategy are proposed and studied. One may also refer to [8] [11] [15] for more recent related works. It is however noteworthy that in these references, liquidity costs and market impact are not considered jointly, whereas in fact, the latter is a rather direct consequence of the former. As we shall demonstrate, the level of permanent impact plays a fundamental role in the well-posedness of the pricing and hedging equation, a fact that was overlooked in previous works on liquidity costs and impact. Also, from a practical point of view, it seems relevant to us to relate the well-posedness of the modified Black-Scholes equation to a parameter that can be measured empirically using high frequency data.

This paper aims at contributing to the field by laying the grounds for a reasonable model of liquidity costs and market impact for derivative hedging. We start in a discrete time setting, where notions are best introduced and properly defined, and then move on to the continuous time case. Liquidity costs are modelled by a simple, stationary order book, characterized by its shape around the best price, and the permanent market impact is measured by a numerical parameter γ, $0 \leqslant \gamma \leqslant 1$: $\gamma = 0$ means no permanent impact, so the order book goes back to its previous state after the transaction is performed, whereas $\gamma = 1$ means no relaxation, the liquidity consumed by the transaction is not replaced. This simplified representation of market impact rests on the hypothesis that the characteristic time of the derivative hedger may be different from the relaxation time of the order book, a realistic hypothesis since delta-hedge generally occurs at a lower frequency than does liquidity providing.

What we consider as our main result is Theorem 4.1, which states that, in the complete market case, the range of parameter for which the pricing equation is unconditionally parabolic is $\frac{2}{3} \leqslant \gamma \leqslant 1$. This result, which we find quite nice in that it is explicit in terms of the parameter γ, obviously explains the ill-posedness of the pricing equations in the references [3] [13] that correspond to the case $\gamma = 0$, or [8] [11] that correspond to the case to the case $\gamma = \frac{1}{2}$ within our formulation. In particular, Theorem 4.1 implies that when re-hedging occurs at the same frequency as that at which liquidity is provided to the order book - that is, when $\gamma = 1$ - the pricing equation is well-posed. This result was already obtained by the second author in [12]. Note that, according to recent theoretical as well as empirical work on market impact, see [5], the level of permanent impact should actually be equal to $\frac{2}{3}$, in striking compliance with the constraints Theorem 4.1 imposes!

The paper is organized as follows: after recalling some classical notations and concepts, Section 1 presents the order book model that will be used to describe liquidity costs. Then, in Section 2, we write down the model for the observed price dynamics and study the associated risk-minimizing strategy taking into account liquidity costs and market impact. Section 3 is devoted to the continuous time version of these results. The pricing and hedging equations are then worked out and characterized in the case of a complete market, in the single asset case in Section **??**, and in the multi-asset case in Section **??**. Finally, Section 6 touches upon the case of stochastic volatility models, for which partial results are presented.

1 Basic notations and definitions

To ease notations, we will assume throughout the paper that the risk-free interest rate is always 0, and that the assets pay no dividend.

Discrete time setting

The tradable asset price is modelled by a stochastic process S_k, $(k = 0, \cdots, T)$ on a probability space $(\Omega, \mathcal{F}, P)$. $\mathcal{F}_k$ denotes the $\sigma-$field of events observable up to and including time k. S_k is assumed to be adapted and square-integrable.
A contingent claim is a square-integrable random variable $H \in L^2(P)$ of the following form $H = \delta^H S_T + \beta^H$ with δ^H and β^H, $\mathcal{F}_T$-measurable random variables.
A trading strategy Φ is given by two stochastic processes δ_k, $(k = 0, \cdots, T)$ and β_k, $(k = 0, \cdots, T)$. δ_k (resp. β_k) is the amount of stock (resp. cash) held during period k, $(= [t_k, t_{k+1}))$ and is fixed at the beginning of that period, *i.e.* we assume that δ_k (resp. β_k) is $\mathcal{F}_k-$measurable $(k = 0, \cdots, T)$. Moreover, δ and β are in $L^2(P)$.
The theoretical value of the portfolio at time k is given by

$$V_k = \delta_k S_k + \beta_k,\ (k = 1, \cdots, T).$$

A strategy is $H-$*admissible iff* each V_k is square-integrable and $V_T = H$.
In order to avoid dealing with several unnecessary yet involved cases, we assume that no transaction on the stock will take place at maturity: the claim will be settled with whatever position there is in stock, plus a cash adjustment to match its theoretical value (see the discussion in [9], Section 4).

Continuous time setting

In the continuous case, $(\Omega, \mathcal{F}, P)$ is a probability space with a filtration $(\mathcal{F}_t)_{0 \le t \le T}$ satisfying the usual conditions of right-continuity and completeness. $T \in \mathbb{R}^{*+}$ denotes a fixed and finite time horizon. Moreover, $\mathcal{F}_0$ is trivial and $\mathcal{F}_T = \mathcal{F}$.
The risky asset $S = (S_t)_{0 \le t \le T}$ is a strictly positive, continuous $\mathcal{F}_t$-semimartingale, and a trading strategy Φ is a pair of càdlàg and adapted processes $\delta = (\delta_t)_{0 \le t \le T}$, $\beta = (\beta_t)_{0 \le t \le T}$, while a contingent claim is described by a random variable $H \in L^2(P)$, with $H = \delta^H S_T + \beta^H$, δ^H and β^H being $\mathcal{F}_T-$measurable random variables. $H-$*admissible* strategies are defined as follows:

Definition 1.0.1 *A trading strategy will be called H-admissible iff*

$$\begin{cases} \delta_T = \delta^H \ P-a.s. \\ \beta_T = \beta^H \ P-a.s. \\ \delta \text{ has finite and integrable quadratic variation} \\ \beta \text{ has finite and integrable quadratic variation} \\ \delta \text{ and } \beta \text{ have finite and integrable quadratic covariation.} \end{cases}$$

Since market impact is considered, the dynamics of S is not independent from that of the strategy (δ, β), so that this set of assumption can only be verified *a posteriori*. One of the important consequences of Theorem 4.1 is precisely to give sufficient conditions ensuring that admissible trading strategies exist.

Order book, transaction cost and impact

A stationary, symmetric order-book profile is considered around the **logarithm** of the price $\hat{S}_t$ of the asset S at a given time t **before** the option position is delta-hedged - think of $\hat{S}_t$ as a theoretical price in the absence of the option hedger. The relative density $\mu(x) \geqslant 0$ of the order book is the derivative of the function $M(x) \equiv \int_0^x \mu(t)dt \equiv$ number of shares one can buy (resp. sell) between the prices $\hat{S}_t$ and $\hat{S}_t e^x$ for positive (resp. negative) x. This choice of representation using exponential is made to avoid difficulties in the definiteness of costs and impact for large sell transactions.
The instantaneous - virtual in the terminology of [17] - market impact of a transaction of size ϵ is then

$$I_{virtual}(\epsilon) = \hat{S}_t(e^{M^{-1}(\epsilon)} - 1), \tag{1.1}$$

it is precisely the difference between the price before and immediately after the transaction is completed. The level of permanent impact is then measured *via* a parameter γ:

$$I_{permanent}(\epsilon) = \hat{S}_t(e^{\gamma M^{-1}(\epsilon)} - 1). \tag{1.2}$$

The actual cost of the same transaction is

$$C(\epsilon) = \hat{S}_t \int_0^{\epsilon} e^{M^{-1}(y)} dy. \tag{1.3}$$

Denote by κ the function M^{-1}. Since some of our results depend on the simplifying assumption that κ is a linear function

$$\kappa(\epsilon) \equiv \lambda\epsilon \tag{1.4}$$

for some $\lambda \in \mathbf{R}$, the computations in this setting are worked out explicitly in this setting:

$$I_{virtual}(\epsilon) = \hat{S}_t(e^{\lambda\epsilon} - 1), \tag{1.5}$$

$$I_{permanent}(\epsilon) = \hat{S}_t(e^{\gamma\lambda\epsilon} - 1). \tag{1.6}$$

and

$$C(\epsilon) = \hat{S}_t \int_0^{\epsilon} e^{M^{-1}(y)} dy \equiv \hat{S}_t \frac{(e^{\lambda\epsilon} - 1)}{\lambda}. \tag{1.7}$$

This simplifying assumption seems necessary for a rigourous derivation of the local-risk minimizing strategies in the Section 2 and, therefore, for the interpretation of a pseudo-optimal strategy in continuous time. Note however that it plays no role in the case of a complete market studied in sections 4 and 5.

2 Cost process with market impact in discrete time

In this section, we focus on the discrete time case. As said above, the order book is now assumed to be *flat*, so that κ is a linear function as in (1.4).

2.1 The observed price dynamics

The model for the dynamics of the observed price - that is, the price S_k that the market can see at every time t_k after the re-hedging is complete - is now presented.
A natural modelling assumption is that the price moves according to the following sequence of events:

- First, it changes under the action of the "market" according to some (positive) stochastic dynamics for the theoretical price increment $\Delta\hat{S}_k$

$$\hat{S}_k \equiv S_{k-1} + \Delta\hat{S}_k \equiv S_{k-1}e^{\Delta M_k + \Delta A_k}, \tag{2.1}$$

  where ΔM_k (resp. ΔA_k) is the increment of an $\mathcal{F}$-martingale (resp. an $\mathcal{F}$-predictable process).

- Then, the hedger applies some extra pressure by re-hedging her position, being thereby subject to liquidity costs and market impact as introduced in Section 1. As a consequence, the dynamics of the observed price is

$$S_k = S_{k-1}e^{\Delta M_k + \Delta A_k}e^{\gamma\lambda(\delta_k - \delta_{k-1})}. \tag{2.2}$$

  Since this model is "exponential-linear" - a consequence of the assumption that κ is linear - this expression can be simplified to give

$$S_k = S_0 e^{M_k + A_k} e^{\gamma\lambda\delta_k}. \tag{2.3}$$

  with the convention that M, A, δ are equal to 0 for $k = 0$.

2.2 Incremental cost and optimal hedging strategy

Following the approach developed in [13], the incremental cost ΔC_k of re-hedging at time t_k is now studied. The strategy associated to the pair of processes β, δ consists in buying $\delta_k - \delta_{k-1}$ shares of the asset and rebalancing the cash account from β_{k-1} to β_k at the beginning of each hedging period $[t_k, t_{k+1})$. With the notations just introduced in Section 2.1, there holds

$$\Delta C_k = \hat{S}_k \frac{(e^{\lambda(\delta_k - \delta_{k-1})} - 1)}{\lambda} + (\beta_k - \beta_{k-1}). \tag{2.4}$$

Upon using a quadratic criterion, and under some assumptions ensuring the convexity of the quadratic risk, see e.g. [13], one easily derives the two (pseudo-)optimality conditions for local risk minimization

$$E(\Delta C_k | \mathcal{F}_{k-1}) = 0 \tag{2.5}$$

and

$$E((\Delta C_k)(\hat{S}_k(\gamma + (1-\gamma)e^{\lambda(\delta_k - \delta_{k-1})})|\mathcal{F}_{k-1}) = 0,$$

where one must be careful to differentiate $\hat{S^k}$ with respect to δ_{k-1}, see (2.3).
This expression is now transformed - using the martingale condition (2.5) and the observed price (2.3) - into

$$E((\Delta C_k)(S_k e^{-\lambda\gamma(\delta_k - \delta_{k-1})}(\gamma + (1-\gamma)e^{\lambda(\delta_k - \delta_{k-1})})|\mathcal{F}_{k-1}) = 0 \tag{2.6}$$

Equation (2.6) can be better understood - especially when passing to the continuous time limit - by introducing a modified price process accounting for the cumulated effect of liquidity costs and market impact, as in [13] [3]. To this end, we introduce the

Definition 2.2.1 *The supply price $\bar{S}$ is the process defined by*

$$\bar{S}_0 = S_0 \tag{2.7}$$

and, for $k \geqslant 1$,

$$\bar{S}_k - \bar{S}_{k-1} = S_k e^{-\lambda\gamma(\delta_k - \delta_{k-1})}(\gamma + (1-\gamma)e^{\lambda(\delta_k - \delta_{k-1})}) - S_{k-1}. \tag{2.8}$$

Then, the orthogonality condition (2.6) is equivalent to

$$E((\Delta C_k)(\bar{S}_k - \bar{S}_{k-1})|\mathcal{F}_{k-1}) = 0. \tag{2.9}$$

It is classical - and somewhat more natural - to use the portfolio value process

$$V_k = \beta_k + \delta_k S_k, \tag{2.10}$$

so that one can then rewrite the incremental cost in (2.4) as

$$\Delta C_k = (V_k - V_{k-1}) - (\delta_k S_k - \delta_{k-1} S_{k-1}) + \hat{S}_k \frac{(e^{\lambda(\delta_k - \delta_{k-1})} - 1)}{\lambda}, \tag{2.11}$$

or equivalently

$$\Delta C_k = (V_k - V_{k-1}) - \delta_{k-1}(S_k - S_{k-1}) + S_k\left(\frac{e^{\lambda(\delta_k - \delta_{k-1})} - 1}{\lambda e^{\gamma\lambda(\delta_k - \delta_{k-1})}} - (\delta_k - \delta_{k-1})\right). \tag{2.12}$$

To ease the notations, let us define, for $x \in \mathbf{R}$,

$$\mathbf{g}(x) \equiv \frac{e^{\lambda x} - 1}{\lambda e^{\gamma\lambda x}} - x. \tag{2.13}$$

The function $\mathbf{g}$ is smooth and satisfies

$$\mathbf{g}(0) = \mathbf{g}'(0) = 0, \; \mathbf{g}''(0) = (1 - 2\gamma)\lambda. \tag{2.14}$$

As a consequence, the incremental cost of implementing a hedging strategy at time t_k has the following expression

$$\Delta C_k = (V_k - V_{k-1}) - \delta_{k-1}(S_k - S_{k-1}) + S_k \mathbf{g}(\delta_k - \delta_{k-1}), \tag{2.15}$$

and Equation (2.6) can be rewritten using the value process V and the supply price process $\bar{S}$ as

$$E((V_k - V_{k-1} - \delta_{k-1}(S_k - S_{k-1}) + S_k \mathbf{g}(\delta_k - \delta_{k-1}))(\bar{S}_k - \bar{S}_{k-1})|\mathcal{F}_{k-1}) = 0. \tag{2.16}$$

One can easily notice that Equations (2.5) and (2.6) reduce exactly to Equations (2.1) in [13] when market impact is neglected ($\gamma = 0$) and the risk function is quadratic.

3 The continuous-time setting

This section is devoted to the characterization of the limiting equation for the value and the hedge parameter when the time step goes to zero. Since the proofs are identical to those given in [1] [13], we shall only provide formal derivations, limiting ourselves to the case of (continuous) Itō semi-martingales for the driving stochastic equations. However, in the practical situations considered in the last sections of this paper, necessary and sufficient conditions are given that ensure the well-posedness in the classical sense of the strategy-dependent stochastic differential equations determining the price, value and cost processes, so that the limiting arguments can be made perfectly rigourous under these conditions.

3.1 The observed price dynamics

A first result concerns the dynamics of the observed price. Assuming that the underlying processes are continuous and taking limits in *ucp* topology, one shows that the continuous-time equivalent of (2.3) is

$$dS_t = S_t(dX_t + dA_t + \gamma\lambda d\delta_t) \tag{3.1}$$

where X is a continuous martingale and A is a continuous, predictable process of bounded variation.
Equation (4.1) is fundamental in that it contains the information on the strategy-dependent volatility of the observed price that will lead to fully non-linear parabolic pricing equation. In fact, the following result holds true:

Lemma 3.1 *Consider a hedging strategy δ which is a function of time and the observed price S at time t: $\delta_t \equiv \delta(S_t, t)$. Then, the observed price dynamics* (4.1) *can be rewritten as*

$$(1 - \gamma\lambda S_t \frac{\partial \delta}{\partial S})\frac{dS_t}{S_t} = dX_t + dA'_t, \tag{3.2}$$

where A' is another predictable, continuous process of bounded variation.

Proof: use Itō's lemma in Equation (4.1).

3.2 Cost of a strategy and optimality conditions

At this stage, we are not concerned with the actual optimality - with respect to local-risk minimization - of pseudo-optimal solutions, but rather, with pseudo-optimality in continuous time. Hence, we shall use Equations (2.5)(2.16) as a starting point when passing to the continuous time limit.
Thanks to $\mathbf{g}'(0) = 0$, there holds the

Proposition 3.2 *The cost process of an admissible hedging strategy* (δ, V) *is given by*

$$C_t \equiv \int_0^t (dV_u - \delta dS_u + \frac{1}{2} S_u \mathbf{g}''(0) d < \delta, \delta >_u). \tag{3.3}$$

*Moreover, an admissible strategy is (pseudo-)**optimal** iff it satisfies the two conditions*

- C *is a martingale*
- C *is orthogonal to the supply price process* $\bar{S}$*, with*

$$d\bar{S}_t = dS_t + S_t((1-2\gamma)\lambda d\delta_t + \mu d < \delta, \delta >_t) \tag{3.4}$$

and $\mu = \frac{1}{2}(\lambda^2(\gamma^3 + (1-\gamma)^3))$.

In particular, if C is pseudo-optimal, there holds that

$$d < C, \bar{S} >_t \equiv d < V, S >_t - \delta d < S, S >_t + (1-2\gamma)\lambda S_t d < V, \delta >_t - \delta S_t (1-2\gamma)\lambda d < \delta, S >_t = 0. \tag{3.5}$$

4 Complete market: the single asset case

It is of course interesting and useful to fully characterize the hedging and pricing strategy in the case of a complete market. Hence, we assume in this section that the driving factor X is a one-dimensional Wiener process W and that $\mathcal{F}$ is its natural filtration, so that the increment of the observed price is simply

$$dS_t = S_t(\sigma dW_t + \gamma\lambda d\delta_t + dA_t) \tag{4.1}$$

where the "unperturbed" volatility σ is supposed to be constant. We also make the markovian assumption that the strategy is a function of the state variable S and of time.
Under this set of assumptions, perfect replication is considered: the cost process C has to be identically 0, and Equation (3.3) yields the two conditions

$$\frac{\partial V}{\partial S} = \delta, \tag{4.2}$$

and

$$\frac{\partial V}{\partial t} + \frac{1}{2}(\frac{\partial^2 V}{\partial S^2} + S_t \mathbf{g}''(0)(\frac{\partial^2 V}{\partial S^2})^2)\frac{d < S, S >_t}{dt} = 0. \tag{4.3}$$

Applying Lemma 3.1 yields

$$(1 - \gamma\lambda S_t \frac{\partial \delta}{\partial S})\frac{dS_t}{S_t} = \sigma dW_t + dA'_t \tag{4.4}$$

leading to

$$\frac{d < S, S >_t}{dt} = \frac{\sigma^2 S_t^2}{(1 - \gamma\lambda S_t \frac{\partial \delta}{\partial S})^2}. \tag{4.5}$$

Hence, taking (4.2) into account, there holds

$$\frac{\partial V}{\partial t} + \frac{1}{2}(\frac{\partial^2 V}{\partial S^2} + \mathbf{g}''(0) S_t (\frac{\partial^2 V}{\partial S^2})^2)\frac{\sigma^2 S_t^2}{(1 - \gamma\lambda S_t \frac{\partial \delta}{\partial S})^2} = 0 \tag{4.6}$$

or, using (4.2) and the identity $\mathbf{g}''(0) = (1-2\gamma)\lambda$:

$$\frac{\partial V}{\partial t} + \frac{1}{2}(\frac{\partial^2 V}{\partial S^2}(1 + (1-2\gamma)\lambda S_t \frac{\partial^2 V}{\partial S^2}))\frac{\sigma^2 S_t^2}{(1 - \gamma\lambda S_t \frac{\partial^2 V}{\partial S^2})^2} = 0. \tag{4.7}$$

Equation (4.7) can be seen as the pricing equation in our model: any contingent claim can be perfectly replicated at zero cost, as long as one can exhibit a solution to (4.7). Consequently, of the utmost importance

is the parabolicity of the pricing equation (4.7).
For instance, the case $\gamma = 1$ corresponding to a full market impact (no relaxation) yields the following equation

$$\frac{\partial V}{\partial t} + \frac{1}{2}\frac{\partial^2 V}{\partial S^2}\frac{\sigma^2 S_t^2}{(1-\gamma\lambda S_t \frac{\partial^2 V}{\partial S^2})} = 0, \tag{4.8}$$

which can be shown to be parabolic, see [12]. In fact, there holds the sharp result

Theorem 4.1 *Let us assume that $\frac{2}{3} \leqslant \gamma \leqslant 1$. Then, there holds:*

- *The non-linear backward partial differential operator*

$$V \to \frac{\partial V}{\partial t} + \frac{1}{2}((1+(1-2\gamma)\lambda S_t \frac{\partial^2 V}{\partial S^2}))\frac{\sigma^2 S_t^2}{(1-\gamma\lambda S_t \frac{\partial^2 V}{\partial S^2})^2}\frac{\partial^2 V}{\partial S^2} \tag{4.9}$$

  is parabolic.

- *Every european-style contingent claim with payoff Φ satisfying the terminal constraint*

$$\sup_{S\in\mathbf{R}^+}(S\frac{\partial^2 \Phi}{\partial S^2}) < \frac{1}{\gamma\lambda} \tag{4.10}$$

  can be perfectly replicated via *a δ-hedging strategy given by the unique, smooth away from T, solution to Equation* (4.7)*.*

Proof: the parabolic nature of the operator is determined by the monotonicity of the function

$$p \to F(p) = \frac{p(1+(1-2\gamma)p)}{(1-\gamma p)^2}. \tag{4.11}$$

A direct computation shows that $F'(p)$ has the sign of $1 + (2-3\gamma)p$, so that F is globally (in p) monotonic increasing on its domain of definition $]-\infty, \frac{1}{\gamma}[$ whenever $\frac{2}{3} \leqslant \gamma \leqslant 1$. Therefore, the pricing equation is globally well-posed in this parameter range. Now, given that the payoff satisfies the terminal constraint, classical results on the maximum principle for the **second derivative** of the solution ensure that the same constraint is satisfied globally for $t \leqslant T$, and therefore, that the stochastic differential equation determining the price of the asset has classical, strong solutions up to time T. As a consequence, the cost process introduced in Proposition 3.2 is well-defined, and is identically 0. Hence, the perfect replication is possible. Clearly, the constraint on the second derivative is binding, in that it is necessary to ensure the existence of the asset price itself. See however Section 7 for a discussion of other situations.

5 Complete market: the multi-asset case

Consider a complete market described by d state variables $X = X_1, ..., X_d$: one can think for instance of a stochastic volatility model with $X_1 = S$ and $X_2 = \sigma$ when option-based hedging is available. [1]
Using tradable market instruments, one is able to generate d hedge ratio $\delta = \delta_1, ..., \delta_d$ with respect to the independent variables $X_1, ..., X_d$, that is, one can buy a combination of instruments whose price $P(t, X)$ satisfies

$$\partial_{X_i} P = \delta_i. \tag{5.1}$$

We now introduce two market impact matrices, Λ_1 and Λ_2. The first one represents the virtual market impact and the second, the permanent impact. When $d = 1$, they are linked to the previous notations by

$$\Lambda_1 = \lambda S, \Lambda_2 = \gamma\lambda S.$$

[1] In the next section, stochastic volatility is addressed in the context of an incomplete market

Note that here, we proceed directly to the continuous time case, so that the actual shape of the order book plays a role only through its Taylor expansion around 0; hence, the use of the "linearized" impact *via* the matrices Λ_i.
The permanent market impact of a transaction $d\delta$ is given by

$$dX = \Lambda_2 \delta. \tag{5.2}$$

The pricing equation is derived along the same lines as in Section 3: denote by V the option price, and let Δ, Γ be the first and second derivatives of V. The relationship between the "unperturbed " price change $d\hat{X}_t$ and the observed price change dX_t is given by

$$dX_t = d\hat{X}_t + \Lambda_2 d\delta_t + dA_t \tag{5.3}$$

the d-dimensional version of (4.1). As before, a straightforward application of Itō's formula in a markovian setting yields the dynamics of the observed price

$$dX = (I - \Lambda_2\Gamma)^{-1} d\hat{X} + dA'_t. \tag{5.4}$$

The d-dimensional version of (2.15) for the incremental cost of hedging is

$$dC_t = dV_t - \sum_{i=1}^{d} \delta_i dX_t^i + \frac{1}{2}\text{Trace}((\Lambda_1 - 2\Lambda_2) d < \delta, \delta >_t), \tag{5.5}$$

$$\tag{5.6}$$

so that, the market being complete, the perfect hedge condition $dC_t = 0$ yields the usual delta-hedge

$$\frac{\partial V}{\partial X_i} = \delta_i \tag{5.7}$$

together with the pricing equation

$$\partial_t V + \frac{1}{2}\text{Trace}(\Gamma \frac{d < X, X >_t}{dt}) = \text{Trace}(\Gamma(\Lambda_2 - \frac{1}{2}\Lambda_1)\frac{d < X, X >_t}{dt}). \tag{5.8}$$

Using (5.4), the pricing equation becomes

$$\partial_t V + \frac{1}{2}\text{Trace}\left[(\Gamma(I - (2\Lambda_2 - \Lambda_1)\Gamma))(M\Sigma^{tr} M)\right] = 0. \tag{5.9}$$

where we have set $\Sigma = \frac{d<\hat{X},\hat{X}>_t}{dt}$, $M = (I - \Lambda_2\Gamma)^{-1}$.
When $\Lambda_1 = \Lambda_2$ (i.e. no relaxation), the pricing equation becomes

$$\partial_t V + \frac{1}{2}\text{Trace}(\Gamma(I - \Lambda\Gamma)^{-1}\Sigma_t) = 0. \tag{5.10}$$

which, degenerating further to the 1-dimensional case, yields the pricing equation already derived in [12]

$$\partial_t V + \frac{1}{2}\frac{\Gamma}{1 - \lambda\Gamma} S^2\sigma^2 = 0. \tag{5.11}$$

The assessment of well-posedness in a general setting is related to the monotonicity of the linearized operator; in the case of full market impact $\Lambda_1 = \Lambda_2 \equiv \Lambda$, there holds the

Proposition 5.1 *Assume that the matrix Λ is symmetric. Then, Equation (5.10) is parabolic on the connected component of* $\{\det(I - \Lambda\Gamma) > 0\}$ *that contains* $\{\Gamma = 0\}$.

Proof: let

$$F(\Gamma) = \text{Trace}(\Gamma(I-\Lambda\Gamma)^{-1}\Sigma_t),$$

and

$$H(\Gamma) = \Gamma(I-\Lambda\Gamma)^{-1}.$$

Denoting by $\mathbb{S}_d^+$ the set of d-dimensional symmetric positive matrices, we need to show that for all $d\Gamma \in \mathbb{S}_d^+$, for all covariance matrix $\Sigma \in \mathbb{S}_d^+$, there holds

$$F(\Gamma + d\Gamma) \geq F(\Gamma).$$

Performing a first order expansion yields

$$\begin{aligned} H(\Gamma + d\Gamma) - H(\Gamma) &= \Gamma(I-\Lambda\Gamma)^{-1}\Lambda d\Gamma(I-\Lambda\Gamma)^{-1} + d\Gamma(I-\Lambda\Gamma)^{-1} && (5.12)\\ &= (\Gamma(I-\Lambda\Gamma)^{-1}\Lambda + I)d\Gamma(I-\Lambda\Gamma)^{-1}. && (5.13) \end{aligned}$$

Using the elementary lemma 5.2 - stated below without proof - there immediately follows that

$$\begin{aligned} F(\Gamma + d\Gamma) - F(\Gamma) &= \text{Trace}((I-\Gamma\Lambda)^{-1}d\Gamma(I-\Lambda\Gamma)^{-1}\Sigma) && (5.14)\\ &= \text{Trace}(d\Gamma(I-\Lambda\Gamma)^{-1}\Sigma(I-\Gamma\Lambda)^{-1}). && (5.15) \end{aligned}$$

Then, the symmetry condition on Λ allows to conclude the proof of Proposition 5.1.

Lemma 5.2 *The following identity holds true for all matrices* Γ, Λ*:*

$$\Gamma(I-\Lambda\Gamma)^{-1}\Lambda + I = (I-\Gamma\Lambda)^{-1}.$$

6 The case of an incomplete market

In this section, stochastic volatility is considered. As said earlier, the results in Section 5 directly apply in this context whenever the market is assumed to be completed *via* an option-based hedging strategy. However, it is well known that such an assumption is equivalent to a very demanding hypothesis on the realization of the options dynamics and their associated risk premia, and it may be more realistic to assume that the market remains incomplete, and then, study a hedging strategy based on the tradable asset only. As we shall see below, such a strategy leads to more involved pricing and hedging equations.
Let then the observed price process be a solution to the following set of SDE's

$$\begin{aligned} dS_t &= S_t(\sigma_t dW_t^1 + \gamma\lambda d\delta_t + \mu_t dt) && (6.1)\\ d\sigma_t &= \nu_t dt + \Sigma_t dW_t^2 && (6.2) \end{aligned}$$

where (W^1, W^2) is a two-dimensional Wiener process under $\mathcal{P}$ with correlation ρ:

$$d < W^1, W^2 >_t = \rho dt,$$

and the processes μ_t, ν_t and Σ_t are actually functions of the state variables S, σ.
Consider again a markovian framework, thereby looking for the value process V and the optimal strategy δ as smooth functions of the state variables

$$\begin{aligned} \delta_t &= \delta(S_t, \sigma_t, t)\\ V_t &= V(S_t, \sigma_t, t). \end{aligned}$$

Then, the dynamics of the observed price becomes

$$dS_t = \frac{S_t}{1 - \gamma\lambda S_t \frac{\partial\delta}{\partial S}}(\sigma_t dW_t^1 + \gamma\lambda\frac{\partial\delta}{\partial\sigma}d\sigma_t + dQ_t), \tag{6.3}$$

the orthogonality condition reads

$$(\frac{\partial V}{\partial S} - \delta) d < S, \bar{S} >_t + \frac{\partial V}{\partial \sigma} d < \sigma, \bar{S} >_t = 0 \tag{6.4}$$

and the pricing equation for the value function V is

$$\frac{\partial V}{\partial t} + \frac{1}{2}(\frac{\partial^2 V}{\partial S^2} - \gamma\lambda S_t (\frac{\partial \delta}{\partial S})^2) \frac{d < S, S >_t}{dt} + \frac{1}{2}(\frac{\partial^2 V}{\partial \sigma^2} - \gamma\lambda S_t (\frac{\partial \delta}{\partial \sigma})^2) \frac{d < \sigma, \sigma >_t}{dt} +$$

$$+ (\frac{\partial^2 V}{\partial \sigma \partial S} - \gamma\lambda S_t \frac{\partial \delta}{\partial \sigma} \frac{\partial \delta}{\partial S}) \frac{d < S, \sigma >_t}{dt} + \mathcal{L}_1 V = 0, \tag{6.5}$$

where $\mathcal{L}_1$ is a first-order partial differential operator.
Equations (6.4) and (6.5) are quite complicated. In the next paragraph, we focus on a particular case that allows one to fully assess their well-posedness.

6.1 The case $\gamma = 1$, $\rho = 0$

When $\gamma = 1$, the martingale component of the supply price does not depend on the strategy anymore. As a matter of fact, the supply price dynamics is given by

$$d\bar{S}_t = dS_t + S_t((1 - 2\gamma)\lambda d\delta_t + \frac{1}{2}\mu d < \delta, \delta >_t),$$

see (3.4), and therefore, using (6.1), there holds that

$$d\bar{S}_t = S_t(\sigma_t dW_t^1 + \lambda(1 - \gamma) d\delta_t + dR_t) \equiv S_t(\sigma_t dW_t^1 + dR_t), \tag{6.6}$$

where R is a process of bounded variation. If, in addition, the Wiener processes for the asset and the volatility are supposed to be uncorrelated: $\rho = 0$, the tedious computations leading to the optimal hedge and value function simplify, and one can study in full generality the well-posedness of the pricing and hedging equations (6.4)(6.5).
First and foremost, the orthogonality condition (6.4) simply reads in this case

$$\delta = \frac{\partial V}{\partial S}, \tag{6.7}$$

exactly as in the complete market case. This is a standard result in local-risk minimization with stochastic volatility when there is no correlation.
As for the pricing equation (6.5), one first works out using (6.7) the various brackets in (6.5) and finds that

$$\frac{d < S, S >_t}{dt} = (1 - \lambda S_t \frac{\partial^2 V}{\partial S^2})^{-2} (\sigma_t^2 S_t^2 + \lambda^2 S_t^2 (\frac{\partial^2 V}{\partial S \partial \sigma})^2 \Sigma_t^2), \tag{6.8}$$

$$\frac{d < \sigma, \sigma >_t}{dt} = \Sigma^2 \tag{6.9}$$

and

$$\frac{d < S, \sigma >_t}{dt} = (1 - \lambda S_t \frac{\partial^2 V}{\partial S^2})^{-1} \lambda S_t \Sigma_t^2 \frac{\partial^2 V}{\partial S \partial \sigma}. \tag{6.10}$$

Plugging these expressions in (6.5) yields the pricing equation for V

$$\frac{\partial V}{\partial t} + \frac{1}{2} \frac{\partial^2 V}{\partial S^2} (1 - \lambda S_t (\frac{\partial^2 V}{\partial S^2}))^{-1} (\sigma_t^2 S_t^2 + \lambda^2 S_t^2 (\frac{\partial^2 V}{\partial S \partial \sigma})^2 \Sigma_t^2) + \frac{1}{2}(\frac{\partial^2 V}{\partial \sigma^2} - \lambda S_t (\frac{\partial^2 V}{\partial S \partial \sigma})^2) \Sigma^2 +$$

$$\lambda S_t \Sigma_t^2 (\frac{\partial^2 V}{\partial S \partial \sigma})^2 + \mathcal{L}_1 V = 0, \tag{6.11}$$

or, after a few final rearrangments,

$$\frac{\partial V}{\partial t}+\frac{\sigma_t^2 S_t^2}{2(1-\lambda S_t(\frac{\partial^2 V}{\partial S^2}))}\frac{\partial^2 V}{\partial S^2}+\frac{1}{2}\frac{\partial^2 V}{\partial \sigma^2}\Sigma^2+\frac{1}{2}\frac{\lambda S_t\Sigma^2}{(1-\lambda S_t(\frac{\partial^2 V}{\partial S^2}))}(\frac{\partial^2 V}{\partial\sigma\partial S})^2+\mathcal{L}_1 V=0. \tag{6.12}$$

The main result of this section is the

Proposition 6.1 *Equation* (6.12) *is of parabolic type.*

Proof: one has to study the monotocity of the operator

$$\mathcal{L}:\, V\to\mathcal{L}(V)\equiv\frac{\sigma_t^2 S_t^2}{2(1-\lambda S_t(\frac{\partial^2 V}{\partial S^2}))}\frac{\partial^2 V}{\partial S^2}+\frac{1}{2}\frac{\partial^2 V}{\partial \sigma^2}\Sigma^2+\frac{1}{2}\frac{\lambda S_t\Sigma^2}{(1-\lambda S_t(\frac{\partial^2 V}{\partial S^2}))}(\frac{\partial^2 V}{\partial\sigma\partial S})^2. \tag{6.13}$$

Introducing the classical notations

$$p\equiv\begin{pmatrix}p_{11} & p_{12}\\ p_{21} & p_{22}\end{pmatrix} \tag{6.14}$$

with $p_{11}=\frac{\partial^2 V}{\partial S^2}$, $p_{12}=p_{21}=\frac{\partial^2 V}{\partial S\partial\sigma}$ and $p_{22}=\frac{\partial^2 V}{\partial\sigma^2}$ and defining

$$\mathbf{L}(S,\mathbf{p})\equiv\frac{\sigma_t^2 S_t^2 p_{11}}{(1-\lambda S_t p_{11})}+\Sigma^2 p_{22}+\frac{\lambda S_t\Sigma^2}{(1-\lambda S_t p_{11})}p_{12}^2, \tag{6.15}$$

one is led to study the positivity of the 2×2 matrix

$$\begin{pmatrix}\frac{\partial\mathbf{L}}{\partial p_{11}} & \frac{1}{2}\frac{\partial\mathbf{L}}{\partial p_{12}}\\ \frac{1}{2}\frac{\partial\mathbf{L}}{\partial p_{12}} & \frac{\partial\mathbf{L}}{\partial p_{22}}\end{pmatrix}. \tag{6.16}$$

Setting $F(p_{11})=\frac{\sigma^2 S^2 p_{11}}{1-\lambda S p_{11}}$ and $D(p_{11})=1-\lambda S p_{11}$, one needs to show that the matrix $\mathbf{H}(\mathbf{p})$

$$\begin{pmatrix}F'(p_{11})+(\lambda S\Sigma)^2\frac{p_{12}^2}{D^2} & \lambda S\Sigma^2\frac{p_{12}}{D}\\ \lambda S\Sigma^2\frac{p_{12}}{D} & \Sigma^2\end{pmatrix} \tag{6.17}$$

is positive. This result is trivially shown to be true by computing the trace and determinant of $\mathbf{H}(\mathbf{p})$:

$$Tr(\mathbf{H}(\mathbf{p}))=F'(p_{11})+\Sigma^2+(\lambda S\Sigma)^2\frac{p_{12}^2}{D^2} \tag{6.18}$$

and

$$Det(\mathbf{H}(\mathbf{p}))=\Sigma^2 F'(p_{11}) \tag{6.19}$$

and using the fact that F is a monotonically increasing function.
This ends the proof of Proposition 6.1.
As a final remark, we point out that the condition on the payoff for (6.12) to have a global, smooth solution, is exactly the same as in the one-dimensional case: stochastic volatility does not impose further constraints, except the now imperfect replication strategy.

7 Concluding remarks

In this work, we model the effect of liquidity costs and market impact on the pricing and hedging of derivatives, using a static order book description and introducing a numerical parameter measuring the level of asymptotic market impact. In the complete market case, a structural result characterizing the well-posedness of the strategy-dependent diffusion is proven. Extensions to incomplete markets and nonlinear hedging strategies are also considered.
We conclude with a discussion of the two conditions that play a fundamental role in our results.

The condition $\gamma \in [\frac{2}{3}, 1]$

Of interest is the interpretation of the condition on the resilience parameter: $\frac{2}{3} \leqslant \gamma \leqslant 1$.
The case $\gamma > 1$ is rather trivial to understand, as one can easily see that it leads to arbitrage by a simple round-trip trade. The case $\gamma < \frac{2}{3}$ is not so simple. The loss of monotonicity of the function $F(p) = \frac{p(1+(1-2\gamma)p)}{(1-\gamma p)^2}$ for $\gamma < \frac{2}{3}$ yields the existence of p_1, p_2 such that $p_1 < p_2$ but $F(p_1) > F(p_2)$, which will lead to an arbitrage opportunity as we now show.
Recall that the price of the replicating strategy solves the equation

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 SF\left(S\frac{\partial^2 V}{\partial S^2}\right) = 0, \tag{7.1}$$

and assume that there exists $p \in \mathbb{R}$ with $F'(p) < 0$. One can then find two values $p_1 < p_2$ such that $F(p_1) > F(p_2)$. Consider now two contingent claims Φ_1, Φ_2 satisfying $S\frac{\partial^2 \Phi_i}{\partial S^2} \equiv p_i, i = 1, 2$, together with $\frac{\partial \Phi}{\partial S}(S_0) = 0, \Phi_i(S_0) = 0$ for some given $S_0 > 0$. Under these assumptions, $\Phi_2(S) \geq \Phi_1(S)$ for all S. Then, there exist explicit solutions $V_i(t, S)$ to (7.1) with terminal conditions Φ_i, $i = 1, 2$, given simply by translations in time of the terminal payoff:

$$V_i(t, S) = \Phi_i(S) + (T-t)\frac{\sigma^2}{2}SF(p_i). \tag{7.2}$$

Consider the following strategy: sell the terminal payoff Φ_1 at price $V_1(0, S_0)$, without hedging, and hedge Φ_2 following the replicating strategy given by (7.1). The final wealth of such a strategy is given by

$$\text{Wealth}(T) = \underbrace{(\Phi_2(S_T) - V_2(0, S_0))}_{\text{hedge strategy}} + \underbrace{(V_1(0, S_0) - \Phi_1(S_T))}_{\text{option sold}}. \tag{7.3}$$

Using (7.2), one obtains

$$\text{Wealth}(T) = T\frac{\sigma^2}{2}S_0(F(p_1) - F(p_2)) + (\Phi_2(S_T) - \Phi_2(S_0)) - (\Phi_1(S_T) - \Phi_1(S_0)), \tag{7.4}$$

which is always positive, given the conditions on Φ_1, Φ_2, and thereby generates an arbitrage opportunity.
Note that this arbitrage opportunity exists both for $\gamma > 1$ and $\gamma < 2/3$, since it just requires that F be locally decreasing. However, in the case $\gamma > 1$, since round-trip trades generate money, it is the price dynamics itself that create arbitrage opportunities, whereas in the case $\gamma < 2/3$, it is the option prices generated by exact replication strategies that lead to an arbitrage.
In this case it makes sense to look for super-replicating strategies, in the spirit of [4], this will be the object of a forthcoming work.

The condition $S\frac{\partial^2 V}{\partial S^2} < \frac{1}{\gamma\lambda}$

Another important question has been left aside so far: the behaviour of the solution to the pricing equation when the constraint is violated at maturity - after all, this is bound to be the case for a real-life contingent claim such as a call option ! From a mathematical point of view, see the discussion in [12], there is a solution which amounts to replace the pricing equation $\mathcal{P}(D)(V) = 0$ by $\text{Max}(\mathcal{P}(D)(V), S\frac{\partial^2 V}{\partial S^2} - \frac{1}{\gamma\lambda}) = 0$, but of course, in this case, the perfect replication does not exist any longer - again, one should use a super-replicating strategy.

References

[1] F. Abergel and N. Millot. Non quadratic local risk-minimization for hedging contingent claims in incomplete markets. *SIAM Journal on Financial Mathematics*, 2(1):342–356, 2011.

[2] R. Almgren, C. Thum, E. Hauptmann, and H. Li. Direct estimation of equity market impact. *working paper.*

[3] U. Cetin, R. Jarrow, and P. Protter. Liquidity risk and arbitrage pricing theory. *Finance and Stochastics*, 8:311–341, 2004.

[4] U. Çetin, H. M. Soner, and N. Touzi. Option hedging for small investors under liquidity costs. *Finance Stoch.*, 14(3):317–341, 2010.

[5] J. D. Farmer, A. Gerig, F. Lillo, and H. Waelbroeck. How efficiency shapes market impact. 2013.

[6] R. Frey. Perfect option hedging for a large trader. 1996.

[7] R. Frey and A. Stremme. Market volatility and feedback effects from dynamic hedging. *Mathematical Finance*, 7(4):351–374, 1997.

[8] P. Henry-Labordère. *working paper.*

[9] D. Lamberton, H. Pham, and M. Schweizer. Local risk-minimization under transaction costs. *Mathematics of Operations Research*, 23:585–612, 1997.

[10] H. E. Leland. Option pricing and replication with transactions costs. *Journal of Finance*, 40(5):1283–1301, 1985.

[11] H. Liu and J. M. Yong. Option pricing with an illiquid underlying asset market. *Journal of Economic Dynamics and Control*, 29:2125–2156, 2005.

[12] G. Loeper. Option pricing with market impact and non-linear black and scholes pde's. *http://arxiv.org/abs/1301.6252*, 2013.

[13] N. Millot and F. Abergel. Non quadratic local risk-minimization for hedging contingent claims in the presence of transaction costs. *Available at SSRN 1881175*, 2011.

[14] E. Platen and M. Schweizer. On feedback effects from hedging derivatives. *Mathematical Finance*, 8(1):67–84, 1998.

[15] A. Roch. Liquidity risk, volatility and financial bubbles. *PhD Thesis.*

[16] P. J. Schönbucher and P. Wilmott. The feedback effect of hedging in illiquid markets. *SIAM J. Appl. Maths*, 61(1):232–272, 2000.

[17] P. Weber and B. Rosenow. Order book approach to price impact. *Quantitative Finance*, 5(4):357–364, 2005.

[18] V. Zakamouline. European option pricing and hedging with both fixed and proportional transaction costs. *working paper.*