

# Optimal discrete hedging of American options using an integrated approach to options with complex embedded decisions

J. Gerer<sup>a</sup>, G. Dorfleitner<sup>a</sup>

<sup>a</sup> *University of Regensburg, Department of Finance, 93040 Regensburg, Germany*

---

## Abstract

In order to solve the problem of optimal discrete hedging of American options, this paper utilizes an integrated approach in which the writer's decisions (including hedging decisions) and the holder's decisions are treated on equal footing. From basic principles expressed in the language of acceptance sets we derive a general pricing and hedging formula and apply it to American options. The result combines the important aspects of the problem into one price. It finds the optimal compromise between risk reduction and transaction costs, i.e. optimally placed rebalancing times. Moreover, it accounts for the interplay between the early exercise and hedging decisions.

We then perform a numerical calculation to compare the price of an agent who has exponential preferences and uses our method of optimal hedging against a delta hedger. The results show that the optimal hedging strategy is influenced by the early exercise boundary and that the worst case holder behavior for a sub-optimal hedger significantly deviates from the classical Black-Scholes exercise boundary.

*Keywords:* Transaction costs, early exercise, discrete hedging, American option, embedded decisions, good-deal bounds

JEL classification codes: G13, C61

---

---

*Email addresses:* `ssrn@gerer.name` (J. Gerer), `gregor.dorfleitner@ur.de` (G. Dorfleitner)

*Submitted to SSRN*

*June 2, 2016*

## 1. Introduction

Whenever an option writer hedges an option, their net payoff is given by the option's premium minus the tracking error of the hedging activity. For European options and in a complete market, there is one hedging strategy that will turn the random future tracking error into a constant known at inception, rendering the pricing problem trivial. In reality, markets are neither complete nor friction-less and there are options and other claims whose payoff can be modified by the holder. Thus, in practice, the tracking error is random and can depend on a time-continuum of decisions by both the writer (deciding whether to change the current hedging position) and the holder (e.g. in the case of the American option, deciding whether to exercise or not).

While both aspects enjoy extensive treatment in scientific publications, most contributions only look at one aspect in isolation from the other, i.e. they focus either on realistic (discrete) hedging or on exercise features. The common ad-hoc approach to decisions embedded in option contracts is stretched over its limit, when applied to a complex combination of decisions by both counterparties.

This issue concerns, among other fields, the literature on realistically hedging of American options, which despite its practical relevance comprises only a handful of contributions. These contributions provide important groundwork, and satisfy many requirements that are in our view desirable for a solution of the problem. This paper's aim is to improve on existing work by combining all these requirements. Table 1 contains the requirements numbered from one to eight in the column headers and provides an overview of the requirements satisfied by each contribution.

	<i>American exercise (1)</i>	<i>Worst-case holder strategy (2)</i>	<i>No restriction of hedging times (3)</i>	<i>Transaction costs (4)</i>	<i>Weighing of risk/cost (5)</i>	<i>Risk &amp; costs influence price (6)</i>	<i>Time consistency (7)</i>	
Chalasanani and Jha (2001)	✓	✓	✓	✓			✓	
Bouchard and Temam (2005)	✓	✓	✓	✓			✓	
Coleman et al. (2007)	✓				✓			
Roux and Zastawniak (2008)	✓	✓	✓	✓			✓	
De Vallière et al. (2008)	✓	✓	✓	✓			✓	
Constantinides and Zariphopoulou (2001)	✓		✓	✓	✓	✓		
Constantinides and Perrakis (2007)	✓	✓	✓	✓	✓	✓	✓	n.a.
Ahn and Wilmott (2009)		n.a.				✓		
Gobet and Landon (2014)		n.a.		✓	✓			

**Table 1:** Relevant literature

Due to the limited number of contributions we also include two notable, newer

contributions on optimal hedging of European options, that therefore violate requirement 1.

Contributions failing the second requirement assume some externally given exercise strategy of the holder. As an example take [Constantinides and Zariphopoulou \(2001\)](#), who state that “[the holder’s exercise time is given by some predetermined stopping time  $\tau$ , which] may be the optimal exercise time [...] in the absence of transaction costs”, or [Coleman et al. \(2007\)](#) claiming that “[t]he holder will choose an exercise strategy to maximize the option value to him; hedging decisions of the writer are irrelevant to his exercise decision.” This view neglects the existence of a possible exercise strategy of the holder which is more expensive to hedge and which would thus lead to too low selling prices and insufficient insurance against all possible holder behaviors.

The third requirement is violated if rehedging times are somehow restricted by an external mechanism e.g. by a predetermined number of hedges, or if rebalancing is only allowed after some risk measure exceeds a certain threshold (e.g. as in [Ahn and Wilmott, 2009](#)). Such a restriction is always sub-optimal, because it ignores the fact that at any given point of time, rebalancing is optimal if and only if the reduction in risk does outweigh the cost associated with rebalancing. This kind of sub-optimality implies unrealistic behavior: if a trade leads to a more favorable hedging position *now*, why should any past consideration stop the hedger from executing it?

These considerations also originate the need to include transaction costs (requirement 4)—be it traditional transaction fees, spreads or opportunity costs. Without transaction costs, the optimal hedging strategy always consists of quasi-continuous rebalancing, i.e. rebalancing to the optimal hedging position as fast as practically possible, which is simply unrealistic.

As mentioned above, deriving an optimal hedging strategy means weighing risk against transaction costs (requirement 5).

Requirement 6 might seem superfluous as it could be argued that the derivation of a super-hedging price actually satisfies requirement 5. However, the underlying risk-measure is too risk-averse to be realistic. Option writers *do* accept the risk of hedging losses, hence requirement 6.

Requirement 7 demands the remaining risk and transaction costs to actually contribute to the final option price. This requirement is for example missed by [Gobet and Landon \(2014\)](#). They minimize the product of the number of hedges and the quadratic variance of the hedging error, two quantities undoubtedly influencing the bottom line of a real-world hedging activity. However, their combination into a product is completely arbitrary and does not translate into a consistent option price, for there will always exist hedging strategies that do not optimize the above product and still lead to a lower selling price.

Time consistency (requirement 8), which is neglected in many contributions, is an important property, especially in the context of pricing options with complex decisions. Results obtained by time-inconsistent methods will either assume sub-optimal future choices or do not give the optimal solution from today’s perspective (see [Gerer and Dorfleitner, 2016](#), for more details on the relation of decisions and time consistency in option pricing).

The article of [Constantinides and Perrakis \(2007\)](#) actually satisfies nearly all of our requirements, yet their contribution consists of the derivation of stochastic bounds on option prices for utility maximizing agents and thus has a different focus.

The isolated treatment of the two decisions by the holder and the writer completely ignores one of the core aspects of realistically hedged options, namely their interplay. Questions about the possibility of a holder’s strategy that explicitly exploit

e.g. a fixed number of hedges by the writer or makes the original hedging strategy inhibitingly expensive due to transactions costs are not even considered.

This paper solves the problem of realistically pricing and hedging an American option. It is based on the insight that a realistic hedging theory is always a theory of hedging with transaction costs: it is not the physical impossibility of continuous trading that needs to be addressed.<sup>1</sup> Instead, we allow continuous trading (in the limit), but acknowledge that there is some kind of cost associated with each trade. This cost is weighed against the remaining hedging risk. Expressing both risk and transaction costs monetarily, gives rise to an optimization problem whose solution is a finite number of optimally placed hedging trades and a consistent option price.

Our approach is motivated by the understanding of a “price” as an intrinsically one-dimensional quantity, which does not leave much conceptual freedom. The results from Gerer and Dorfleitner (2016) imply, that under mild assumptions, the decisions can be *formally eliminated* from the problem in a consequent manner without the need to resort to external concepts and without any further motivating argument. Our analysis differs from others in that we seek to calculate the agent’s indifference or reservation price in an uncompromising fashion.

In Section 2, we summarize the formalism and results from Gerer and Dorfleitner (2016) used in this paper. In Section 3, a general duality between pricing functions and acceptance sets for payoffs with decisions is applied to derive a general pricing and hedging principle for options with decision by both parties. Section 4 specializes this principle to a formula for American options, which is then numerically solved in Section 5. Conclusions are given in Section 6.

## 2. Theoretical fundament

The theory is formulated from the perspective of a single market participant, that we will refer to as *agent*, engaging in financial activities and entering contracts with other agents, collectively called her *counterparty*.

All possible evolutions of the world, their physical probabilities and the time-dependence of information are described by a filtered probability space  $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in \mathcal{T}}, \mathbb{P})$ , where all points of time are given by the totally ordered set  $\mathcal{T}$ .

Let  $L_t^G$  represent all  $\mathcal{F}_t$ -measurable random variables into the set  $G \subseteq \overline{\mathbb{R}}$ . We will use the abbreviations  $L_t \equiv L_t^{(-\infty, \infty)}$ ,  $L_t^- \equiv L_t^{[-\infty, \infty)}$ ,  $L_t^+ \equiv L_t^{(-\infty, \infty]}$  and  $L_t^\pm \equiv L_t^{[-\infty, \infty]}$ , and employ the convention  $\infty - \infty = \infty$  on  $\overline{\mathbb{R}}$ . Define also the set of positive *t-premiums*  $V_t \equiv \left\{ x \in L_t^+ \mid 0 \stackrel{\text{a.s.}}{<} x \right\}$ .

We assume decisions happen at predetermined times,  $\mathbf{T}_d \subseteq \mathcal{T}$ . As we will see later, this does not prevent us from describing more complex decision for which the point of time can also be chosen by the agent, like options with *American* exercise. At each point of time  $t \in \mathbf{T}_d$  there can be exactly one decision by either the agent or the counterparty. This poses no limitation, as instantaneous decisions by the same agent can be merged into a tuple of decisions and it ensures that there is always a well-defined order between decisions by different agents, even if the physical time between them can be infinitesimally short. Decisions to be made by the agent happen at times  $\mathbf{T}_a$  and decisions by the counterparty at times  $\mathbf{T}_c \equiv \mathbf{T}_d \setminus \mathbf{T}_a$ .

---

<sup>1</sup>With today’s trend to sub-millisecond order execution, continuous hedging could be approximated sufficiently if there were no transaction costs.

The decision behavior of the agents will be modeled by *decision procedures*, describing how the choices for a subset of decisions depend on the world state. The set of *decision procedures* for decisions at times  $\mathbf{T} \subseteq \mathbf{T}_d$  is abbreviated by  $\Phi_{\mathbf{T}}$  and defined as the set of stochastic processes whose values at time  $t$  are elements of  $D_t$ :

$$\Phi_{\mathbf{T}} \equiv \left\{ \varphi : \mathbf{T} \times \Omega \rightarrow \bigcup_{t \in \mathbf{T}} D_t \mid \varphi_t : \Omega \rightarrow D_t, \text{ for all } t \in \mathbf{T} \right\} \quad (1)$$

$D_t$  contains all possible choices at time  $t$ . We will use the abbreviation  $\Phi \equiv \Phi_{\mathbf{T}_d}$ .

The payoff of an option with embedded decisions is described by specifying the cumulative discounted cash-flow to be *received by the agent* for any possible combination of choices and world states:

**Definition 2.1** (Payoffs). *Define  $\mathcal{X}_{\mathbf{T}}^t$  as the set of  $\mathcal{F}_t$ -measurable payoffs that only depend on decisions made at times  $\mathbf{T} \subseteq \mathcal{T}$ :*

$$\mathcal{X}_{\mathbf{T}}^t \equiv \left\{ f : \Phi \rightarrow L_t^{\pm} \mid f(\psi) \stackrel{B}{=} f(\varphi), \text{ if } B \in \mathcal{F}_t \text{ and } \psi_t \stackrel{B}{=} \varphi_t \text{ for all } t \in \mathbf{T} \cap \mathbf{T}_d \right\}$$

Putting a set  $B \in \mathcal{F}_{\infty}$  above a comparison operator means conditionally almost surely equal:  $x \stackrel{B}{=} y \Leftrightarrow \mathbb{P}(\{x = y\} | B) = 1$ , with  $\{x = y\} \equiv \{\omega \in \Omega \mid x(\omega) = y(\omega)\}$ .

We will use the abbreviations  $\mathcal{X}_{\mathbf{T}} \equiv \mathcal{X}_{\mathbf{T}}^{\infty}$  and  $\mathcal{X} \equiv \mathcal{X}_{\mathcal{T}}$ .

*Remark 1.* If a random variable  $x \in L_{\infty}^{\pm}$  is used in the context of payoffs, it is understood as the corresponding constant payoff given by  $\psi \mapsto x$ , which is an element of  $\mathcal{X}_{\emptyset}$ , and vice versa.

If not stated differently, all operators, relations and also suprema/infima used on payoffs are the pointwise versions of their  $L_{\infty}^{\pm}$ ,  $\mathbb{P}$ -almost sure, variants:  $fRg \Leftrightarrow \forall \varphi \in \Phi : f(\varphi) \stackrel{\text{a.s.}}{R} g(\varphi)$

Furthermore, we provide an operation to produce the *effective* payoff, that results if an agent or counterparty follows a decision procedure for a certain subset of decisions. These decisions can be considered *fixed* and the effective payoff does not depend on them anymore:

**Definition 2.2** (Effective payoff). *For any payoff  $f \in \mathcal{X}$  and decision procedure  $\varphi \in \Phi_{\mathbf{T}}$  define the effective payoff,  $f[\varphi] \in \mathcal{X}_{\mathcal{T} \setminus \mathbf{T}}$  by*

$$f[\varphi](\psi) \equiv f(\varphi \mathbf{1}_{\mathbf{T}} + \psi \mathbf{1}_{\mathbf{T}_d \setminus \mathbf{T}}), \text{ for all } \psi \in \Phi.$$

The framework aims to provide the tools to build and analyze pricing theories for options with decisions. This is achieved by providing a minimal characterization of acceptance sets and pricing functions, proving their equivalence, and thus making these concepts usable interchangeably.

A  $t$ -acceptance set contains the agent's *acceptable opportunities* at time  $t$ , i.e. payoffs without decisions before time  $t$  that she accepts as zero-cost investments. We require the following property to ensure that an acceptance set can serve as modeling tool for pricing theories:

**Definition 2.3** (Proper acceptance sets). *A  $t$ -acceptance set  $\mathcal{A}$  is called proper if it is  $t$ -compatible (see below) and  $\mathcal{A} = \{f \in \mathcal{X}_{[t, \infty)} \mid \{f + x \mid x \in V_t\} \subseteq \mathcal{A}\}$ .*

**Definition 2.4** ( $t$ -compatibility). *A non empty set  $X$  is  $t$ -compatible, if for all  $\{x_n\} \subseteq X$  and mutually disjoint  $\{B_n\} \subseteq \mathcal{F}_t$  with  $\mathbb{P}(\bigcup_n B_n) = 1$  it holds  $\sum_n x_n \mathbf{1}_{B_n} \in X$ , where  $\mathbf{1}_{B_n}$  is the indicator function of the set  $B_n$ .*

*Remark 2.* In this framework, the result of a pricing function is understood as the highest premium the agent would accept to pay for entering the contract, or her *bid* price. If a contract's *ask* price is wanted, it can be calculated by the negative of the bid price of the reversed contract.

A *t*-pricing function is any function  $\pi : \mathcal{X}_{[t,\infty)} \rightarrow L_t^\pm$ .

*Remark 3.* For a general option  $f \in \mathcal{X}$ —possibly including decisions before  $t$ —and a *t*-pricing function  $\pi$ , we use  $\pi(f)$  to denote  $\varphi \mapsto \pi\left(f[\varphi|_{(-\infty,t)}]\right)$ , which is a payoff in the sense of Definition 2.1.

For pricing functions, the property corresponding to properness is *cash invariance*.

**Definition 2.5** (Cash invariance). *A t-pricing function  $\pi$  is called cash invariant if for any payoff  $f \in \mathcal{X}$  and payoff  $g \in \mathcal{X}_{(-\infty,t)}^t$  with no present or future decision and  $g > -\infty$ , it holds  $\pi(f + g) = \pi(f) + g$ .*

We will also use the term normalized pricing function:

**Definition 2.6** (Normalized pricing function). *A pricing function  $\pi$  is called normalized, if  $\pi(0) = 0$ . Every pricing function  $\pi$  with  $|\pi(0)| < \infty$  has a normalized version  $x \mapsto \pi(x) - \pi(0)$ .*

As proved in Gerer and Dorfleitner (2016, Theorem 3.1) there exists a one-to-one correspondence between the set of cash invariant *t*-pricing functions and the set of proper *t*-acceptance sets. The *duality operations* are given in the following definition and correctly replicate the description of pricing functions in Remark 2:

**Definition 2.7** (Duality operations). *For any t-acceptance set  $\mathcal{A}$  define its dual t-pricing function  $P[\mathcal{A}]$  by  $P[\mathcal{A}](f) \equiv \sup\{x \in L_t^- \mid f - x \in \mathcal{A}\}$  for all  $f \in \mathcal{X}_{[t,\infty)}$ .<sup>2</sup> For any t-pricing function  $\pi$  define its dual t-acceptance set  $A[\pi] \equiv \{f \in \mathcal{X}_{[t,\infty)} \mid 0 \leq \pi(f)\}$ .*

This duality enables us to develop our pricing theory for options with decisions in terms of the more directly accessible language of acceptance. Specifically, we will use *conservative acceptance*, which is employed implicitly in most of the option pricing literature. For a given *t*-acceptance set  $\mathcal{A}$  and a set of admissible decision procedures  $S$ ,  $\mathcal{A}^{\vee S}$  and  $\mathcal{A}^{\exists S}$  represent the *conservative acceptance sets for decisions by the counterparty* or *the agent*, respectively.  $\mathcal{A}^{\vee S}$  includes an option if and only if for every possible decision procedure by the counterparty the resulting effective option is acceptable:

$$\mathcal{A}^{\vee S} \equiv \{f \in \mathcal{X}_{[t,\infty)} \mid \forall \varphi \in S : f[\varphi] \in \mathcal{A}\} \quad (2)$$

It is important to understand, that this presumes nothing about the counterparty's actual behavior. For her own decisions,  $\mathcal{A}^{\exists S}$  contains an option if and only if there always is at least one decision procedure the agent could follow to make the effective option almost acceptable, as indicated by the  $+x$ :<sup>3</sup>

$$\mathcal{A}^{\exists S} \equiv \{f \in \mathcal{X}_{[t,\infty)} \mid \forall x \in V_t, \exists \varphi \in S : f[\varphi] + x \in \mathcal{A}\} \quad (3)$$

<sup>2</sup> The supremum is to be understood as the *essential supremum*, i.e. the supremum in the almost sure sense.

<sup>3</sup>See Gerer and Dorfleitner (2016, Example 4.2) for why this definition needs a more complicated form than Eq. (2)

In order to talk about acceptance sets and pricing functions at different times, we introduce *acceptance families* as time indexed families of proper acceptance sets  $\{\mathcal{A}_t\}_t$  written as  $\mathcal{A}$ , and *pricing families*, as time indexed families of cash invariant pricing functions written as  $\pi$ .

An especially important property of such families is time consistency (requirement 8 in Table 1). By Theorem C.6, the following definitions are equivalent:

**Definition 2.8** (Time consistency). *A proper acceptance family  $\mathcal{A}$ , with a normalized dual  $\pi$ , is called time consistent if  $f \in \mathcal{A}_t \iff \pi_s(f) \in \mathcal{A}_t$ , for all  $s \geq t$ .*

*A normalized pricing family  $\pi$ , is called time consistent if for all  $s \geq t$  and  $f \in \mathcal{X}$ :  $\pi_t(\pi_s(f)) = \pi_t(f)$ .*

### 3. Optimal hedging — the general formula

We will treat the hedging activity as decisions within our theory of options with decisions. Between rehedges there can be further decisions for both the agent and its counterparty.

This approach will produce the pricing function for the hedging agent as well as the optimal hedging ratios, i.e. optimally placed rehedges, without the need to formulate an exogenous optimization problem. Instead these results are direct consequences of the construction and imposed properties of the acceptance set and their relation to pricing functions.

We impose no practical limitation on the number of rehedges the agent can perform. For formal reasons, we approximate the set of hedging decisions using a finite, increasing sequence  $(\tau_i)_{i \leq n}$ , where the hedging position is closed on the last date  $\tau_n$ .

*Remark 4.* While this introduces a dependency of the results on the particular choice of this sequence, the time intervals can be made smaller than any physical time scale of our world and thus its practical influence eliminated. Numerical calculations typically will be feasible only for much larger intervals.

The actually number of performed rehedges will usually be smaller than  $n$ , as at each time  $\tau_i$  the agent can decide not to rebalance her hedging portfolio.

The discounted price processes of the assets available for hedging are modeled by an  $N$ -dimensional adapted process  $X = (X_t)_t$  with finite components. A hedging decision consists of choosing the amount of shares to hold from each asset, which we will model using  $N$ -dimensional vectors, i.e.  $D_{\tau_i} \subseteq \mathbb{R}^N$  for all  $i \leq n$ .

Given a decision procedure  $\varphi \in \Phi$ , the future cash-flow of a hedging activity started with an initial position  $\varphi_{\tau_{i-1}}$  at a time  $t \in (\tau_{i-1}, \tau_i]$  is given by

$$H_t(\varphi) \equiv \sum_{j \geq i}^{n+1} \varphi_{\tau_{j-1}} \cdot (X_{\tau_j} - X_{\max\{\tau_{j-1}, t\}}) - C_j(\varphi). \quad (4)$$

The  $\cdot$  marks scalar product between two vectors. The first term calculates the gains from market price movements, and  $C_j(\varphi)$  stands for the finite,  $\mathcal{F}_{\tau_j}$ -measurable transaction costs associated with changing the portfolio from  $\varphi_{\tau_{j-1}}$  to  $\varphi_{\tau_j}$  at time  $\tau_j$ .  $C_{n+1}$  corresponds to the special case of liquidating the last position  $\varphi_{\tau_n}$  at  $\tau_{n+1}$ . It is easy to see that  $H_t$  depends on decisions at time  $\tau_{i-1}$  and later, i.e.  $H_t \in \mathcal{X}_{\{\tau_j\}_{i-1}^n}$ .

In order to derive the general pricing function of the hedging agent for options with decisions, we employ the method of Gerer and Dorfleitner (2016, Section 4.3) to derive the super-replication price for continuous trading. We start with the agent's *internal acceptance family*,  $\mathcal{A}$ , containing payoffs or zero-cost investments she accepts

“as is”, i.e. payoffs that cannot be modified by her beyond the decisions contained in the payoff, especially not be hedged against.  $\mathcal{A}_t$  needs to capture the agent’s risk aversion, business model and regulatory requirements. In this section we treat it as given, for it is completely independent from the aspect of hedging; a separation of concerns made possible by the development of the proposed framework.

Next, this acceptance family is transformed into the agent’s *external acceptance family*,  $\mathcal{B}_t$ , by subtraction of any modifications, which are not part of the original contract specifications. In the current setting this means subtracting the proceeds of her hedging activity:

$$\mathcal{B}_t(\varphi) \equiv \left\{ f - H_t[\varphi|_{\langle -\infty, t \rangle}] \mid f \in \mathcal{A}_t \right\} \quad (5)$$

$\mathcal{B}_t$  depends on the decision procedure  $\varphi$ , because  $H_t$  depends on past decisions, more specifically on the most recent hedging decision. Making this dependency explicit ensures, that  $\mathcal{B}_t(\varphi)$  itself contains only options with no past decisions.

Eq. (5) can also be read in the following way: The agent accepts a contract with another party, if and only if, she *internally* accepts the contract’s payoff plus the result of her hedging activity.

Through the duality we know that these acceptance sets uniquely define the hedging agent’s prices – denoted by  $\eta_t$ , which can be calculated using the duality operation from Definition 2.7:

$$\eta_t(f)(\varphi) \equiv P[\mathcal{B}_t(\varphi)](f)(\varphi) \quad (6)$$

for any option with decisions  $f \in \mathcal{X}$  and decision procedure  $\varphi \in \Phi$ .

Of course,  $\eta_t$  can also be expressed using  $\mathcal{A}_t$ ’s dual pricing family  $\pi_t = P[\mathcal{A}_t]$ :

**Lemma 3.1.**  $\eta_t(f) = \pi_t(f + H_t)$  for all  $t$  and  $f \in \mathcal{X}$ .

*Proof.* See Appendix B.1. □

This is in agreement with the expected result that the price of an option is given by the internal price of the hedged option.

*Remark 5 (Normalization).* The above construction of  $\eta_t$  will in general not yield normalized pricing functions, i.e. the price of the zero payoff is different from zero:  $\eta_t(0) = \pi_t(H_t) \neq 0$ . Depending on the specific nature of  $\mathcal{A}_t$  and  $\pi_t$ , it is possible that the agent assigns a positive net present value to the proceeds of the trading activity  $H_t$ , i.e.  $\pi_{\tau_i}(H_t) \geq 0$ . This can happen for example, if  $\pi_t(X_s) > X_t$  (for  $s > t$ ), i.e. buying or selling the market assets represents an acceptable or even arbitrage opportunity.

If the market is arbitrage free, or the agent’s transaction costs destroy any acceptable or arbitrage opportunity, then  $\pi_{\tau_i}(H_{\tau_i})(\varphi)$  is zero, if  $\varphi_{\tau_{i-1}} = 0$  and even negative for  $\varphi_{\tau_{i-1}} \neq 0$ , due to the unavoidable costs for closing the current position eventually.

Economically meaningful prices are obtained from the normalized version:

$$f \mapsto \eta_t(f) - \eta_t(0)$$

This calibration ensures that the price of any sure payoff equals the payoff itself ( $\eta_t(g) = g$  if  $g \in L_t^+$ , which follows from cash invariance), and it is plausible with the two cases described above: If the agent would pay a positive amount  $\eta_t(0) > 0$  for the situation he already is in, normalization decreases all bid prices by that amount, which could be interpreted as compensation for giving up the current favorable position upon entering the new contract. In other words, normalization erases the additional value



$\eta_t(f)$  assigns to the possibility of trading in the market, which the agent can also do without  $f$  and whose value is thus given by  $\eta_t(0)$ .

For an initial unfavorable position  $\varphi_{\tau_{i-1}} \neq 0$ ,  $\eta_t(0)[\varphi]$  would be a negative number representing the negative of the cost associated with optimally closing that position. In this case normalization would increase the agent's naked bid price  $\eta_t(f)$ , because entering and optimally hedging  $f$  would spare her the cost of closing her current position.  $\square$

In addition to the hedging decisions, we explicitly add times for general counterparty decisions, which will then be used for the early exercise decision in the next section. Between any two hedging times  $\tau_i$  and  $\tau_{i+1}$  there is a decision of the counterparty located at  $s_i$ :

$$\tau_i < s_i < \tau_{i+1}$$

Furthermore, for each decision we need a set of admissible decision procedures, denoted by  $R_i \subseteq \Phi_{\{\tau_i\}}$  and  $S_i \subseteq \Phi_{\{s_i\}}$  for all  $i \leq n$ .

The general result needs the following axioms. The first assures that hedging decision cannot see in the future. The additional requirement of  $t$ -compatibility derives from the technical differences between conservative acceptance for the agent and the counterparty (cf. Theorems C.4 and C.5):

**Axiom 1.** *For all  $i \leq n$ ,  $R_i$  contains only  $\mathcal{F}_t$ -adapted decision procedures and is  $\tau_i$ -compatible (cf. Definition 2.4).*

The second axiom specifies how decisions are treated. We use conservative acceptance (as defined in Eqs. (2) and (3)) at the time of each particular decision. This, together with time consistency, will be enough to eliminate all future decisions from the pricing problem.

**Axiom 2** (Conservative acceptance).  $\mathcal{A}_{\tau_i} = \mathcal{A}_{\tau_i}^{\exists R_i}$  and  $\mathcal{A}_{s_i} = \mathcal{A}_{s_i}^{\forall S_i}$  for all  $i \leq n$ .

Feeding all this into our formalism yields the pricing formula for an optimally hedged option:

**Theorem 3.2** (Optimal hedging). *Let a  $\mathcal{A}$  be a time consistent internal acceptance family with a normalized dual pricing family  $\pi$ , and assume Axioms 1 and 2. At the end of the hedging activity, the price of an option  $f \in \mathcal{X}$  can be calculated directly by*

$$\eta_{\tau_{n+1}}(f) = \pi_{\tau_{n+1}}(f) - C_{n+1}, \quad (7)$$

and earlier prices for  $i \leq n$  can be calculated recursively:

$$\eta_{\tau_i}(f) = \sup_{\varphi \in R_i} \pi_{\tau_i} \left( \inf_{\psi \in S_i} \pi_{s_i}(\eta_{\tau_{i+1}}(f)[\varphi][\psi] + \varphi_{\tau_i} \cdot (X_{\tau_{i+1}} - X_{\tau_i})) \right) - C_i[\varphi] \quad (8)$$

*Proof.* See Appendix B.2.  $\square$

For a hedging strategy  $a$  and decision procedure of the counterparty  $b$ , the net payoff of the whole hedging activity is given by the difference of the realized tracking error and the option's upfront premium:

$$\delta \equiv (f + H_t)[a][b] - \eta_t(f)$$

Conservative acceptance ensures, that for any decision procedure  $b$  and  $\epsilon \in V_t$  there exists a hedging strategy  $a$  such that  $\delta + \epsilon$  is acceptable.

#### 4. Optimal hedging of American options

In this section we specialize the results from the previous section in order to price and hedge an American option from the perspective of the writer. The holder of an *American option with discounted exercise value process*  $g$  has the right to exercise it at any time  $t$  prior or equal to the expiration date  $T$  and receive the amount  $g_t$ .

This exercise right gives rise to infinitely many decisions: At every instant the holder can decide to exercise or not to exercise. To model these decision we use the finite set of decision times  $\{s_i\}$  introduced in the last section, with  $s_n \equiv T$ . Following the reasoning from Remark 4 this implies no loss of generality. We define

$$D_{s_i} \equiv \{1, 0\}, \text{ where } 1 \text{ stands for "exercise" and } 0 \text{ for "do not exercise"}, \quad (9)$$

for all  $i$ . The payoff of an American option that has not been exercised before time  $s_i$  is then for any  $\varphi \in \Phi$  and  $\omega \in \Omega$  given by:

$$f_i(\varphi) \equiv g_{t^*} \text{ with the stopping time } t^*(\omega) = \min\{t \geq s_i \mid \varphi_t(\omega) = 1\} \quad (10)$$

This definition shows the expected behavior, if the current decision is fixed using a constant decision procedure:

$$f_i[s_i \mapsto 1] = g_{s_i} \text{ for exercise, and } f_i[s_i \mapsto 0] = f_{i+1} \text{ for continuation.} \quad (11)$$

Basic theory of stochastic processes ascertains that  $f$  is measurable if  $g$  is progressively measurable and  $\varphi$  is adapted. And thus due to Corollary A.1 and  $f_i$ 's pointwise definition it is a payoff according to our Definition 2.1, or more precisely  $f_i \in \mathcal{X}_{\{s_j\}_{j=i}^n}$ , as it only depends on decision at times  $\{s_i, \dots, s_n\}$ .

While we restrict the admissible hedging procedures through Axiom 1 from the last section, the only restriction placed on exercise procedures is their being adapted:

$$S_i \equiv \Phi_{\{s_i\}} \cap \{\varphi \mid \varphi \text{ is adapted to } \mathcal{F}\}, \text{ for all } i. \quad (12)$$

So far the optimizations in Theorem 3.2 have to be performed over random variables making a direct numerical implementation infeasible. However, as they are limited to "current time" decisions, they can be simplified. We give sufficient (but not necessary) conditions under which the essential supremum over the set of procedures can be simplified to a pointwise supremum directly over the set of choices:

**Lemma 4.1** (Countable present time decisions). *If  $D_t$  or  $\Omega$  is countable and  $S \equiv \Phi_{\{t\}} \cap \{\varphi \mid \varphi \text{ is adapted to } \mathcal{F}\}$ , then for every cash invariant  $t$ -pricing function  $\pi$  and  $f \in \mathcal{X}$  it holds:*

$$\sup_{\varphi \in S} \pi(f[\varphi]) = \sup_{a \in D_t} \pi(f[t \mapsto a])$$

*Proof.* See Appendix B.3. □

As the payoffs involved do not have any decisions besides the hedging and exercise decisions, we can give a result that reduces the pricing problem to classical option pricing theory for options without decisions. To make this explicit we will use  $\mathcal{A}^0$  and  $\pi^0$  to denote the restrictions of  $\mathcal{A}$ , and  $\pi$ , to options without decisions. Formally, we define  $\mathcal{A}^0 \equiv \mathcal{A} \cap \mathcal{X}_\emptyset$  and  $\pi^0 = P[\mathcal{A}^0]$ . Lemma C.1 proves the expected connection between  $\pi^0$  and  $\pi$ .

Furthermore, we assume for all  $i$  that the agent's hedging decision at  $\tau_{i+1}$  happens instantly after the exercise decision  $s_i$ . Formally, these two times collapse for quantities that do not depend on the decision at  $s_i$ , i.e.

$$g_{s_i} = g_{\tau_{i+1}}, \text{ and } \pi_{s_i}(f) = \pi_{\tau_{i+1}}(f), \text{ if } f \in \mathcal{X}_{[\tau_{i+1}, \infty)}. \quad (13)$$

Now we can derive the main result:

**Theorem 4.2.** *Given  $X$ ,  $C$ ,  $\mathcal{A}$ ,  $\{R_i\}$  and  $\eta$ , as in Theorem 3.2,  $\{S_i\}$ ,  $f$ ,  $\mathcal{A}^0$  and  $\pi^0$  as defined above, we define  $p_i$  as the ask price (cf. Remark 2) at time  $\tau_i$  of an optimally hedged American option:*

$$p_i \equiv -\eta_{\tau_i}(-f_i), \text{ for all } i \leq n.$$

The price after expiration is given by

$$p_{n+1} = C_{n+1}, \quad (14)$$

and earlier prices for  $i \leq n$  can be calculated recursively:

$$p_i = \inf_{\varphi \in R_i} -\pi_{\tau_i}^0(-\max\{g_{\tau_{i+1}} - \eta_{\tau_{i+1}}(0), p_{i+1}\}[\varphi] + \varphi_{\tau_i} \cdot (X_{\tau_{i+1}} - X_{\tau_i})) + C_i[\varphi] \quad (15)$$

*Proof.* See Appendix B.4. □

As expected, the writer chooses the most favorable hedge and it is most expensive for her if the holder exercises as soon as the payoff exceeds the price of the continued option.

The terms  $C_{n+1}$  and  $-\eta_{\tau_i}(0) = -\pi_{\tau_i}^0(H_i)$  occurring above could be identified with the cost of optimally closing the current hedging position. As already discussed in Remark 5, they are non-negative in a market without acceptable opportunity and thus add to the payment of  $g_{\tau_i}$  faced by the hedging option writer upon exercise by the holder.

Their appearance is a consequence of the fact that  $\eta$  is not normalized (in general) and it can be trivially checked, that normalizing the result — i.e. calculating  $p_i + \eta_{\tau_i}(0)$  — would remove both terms, whilst introducing a similar term in the continuation value. We did not state the normalized result, as it would complicate the recursive calculations, which are more naturally expressed in unnormalized values.

## 5. Numerical demonstration

In this section we produce numerical results from Theorem 4.2. We will use a simple market model consisting of a riskless money market account with interest rate  $r$  and a single stock whose discounted price process  $X$  is a geometric Brownian motion with drift  $\mu > r$  and volatility  $\sigma$ . Besides analytical tractability and intuitiveness, it reveals interesting features of the pricing and hedging problem. It should be noted that our hedging formula can be applied to any model for the price process  $X$ .

We are pricing an American put with strike  $K$ , i.e. a payoff  $f$  as defined in Eq. (10) with discounted exercise value written as  $g_t(X_t) = X_t - Ke^{-rt}$ .

The transaction costs consist of a fixed component  $k_0$  and a component proportional to the transaction value (with factor  $k_1$ ). Its discounted value is calculated as follows:

$$C_i(\varphi) \equiv c_i(\varphi_{\tau_i} - \varphi_{\tau_{i-1}}, X_{\tau_i}) \text{ with } c_i(q, x) = e^{-r\tau_i} k_0 \mathbf{1}_{q \neq 0} + k_1 x |q|$$

### 5.1. Selecting a pricing function

Before we can actually implement a numerical program, we need to devise the agent's internal pricing family for options without decisions,  $\pi^0$ .

Let us first state the requirements to be met by  $\pi^0$ . There is, of course *cash invariance* (requirement I), the basic property imposed by our formalism, and *time consistency* (requirement II) upon which the results of the previous section rely.

In addition to these two requirements concerning  $\pi^0$  directly, we place three further restrictions on the resulting external pricing family  $\eta$ . (cf. Lemma 3.1). To ensure consistency with existing results we require that without transaction costs and in the limit of infinitely many hedging times the well-known arbitrage-free prices, i.e. *risk-neutral expectation values*, are obtained for continuously replicable payoffs (requirement III).

As noted in Remark 5,  $\eta$  is not normalized. Thus,  $\eta_t(0) = \pi_t(H_t)$  can be negative due to transaction costs for closing the current position or positive, if trading in the market constitutes an acceptable opportunity for the agent. While these effects can be handled satisfactorily by normalizing the result, normalization is only meaningful if  $|\eta_t(0)| < \infty$ , or informally stated, if the agent *cannot extract infinite wealth from his trading activity* (requirement IV).

Besides these theoretical requirements, for the purpose of this demonstration we need *readily implementable, numerical algorithms* (requirement V).

To satisfy requirement V we exclude all pricing functions or risk-measures whose calculation relies on Monte Carlo methods. We are aware of the existence of Monte Carlo methods suitable for American options, but extending and implementing them for our problem — while deemed possible — would go beyond the scope of this paper. This excludes all candidates containing the value-at-risk and its variants or derivatives like the expected shortfall or conditional value-at-risk, most of which also violate requirement II, time consistency (cf. Cheridito and Stadje, 2009).

We use

$$\pi_t^0(f) \equiv \frac{-1}{\gamma} \ln(\mathbb{E}[e^{-\gamma f} | \mathcal{F}_t]), \quad (16)$$

for some positive degree of risk aversion  $\gamma$ . This function is the indifference price of the exponential utility function, also known as the negative of the conditional entropic risk measure and has gained much attention in the field of utility indifference pricing, among others.

For the remainder of this paper we use the pricing family  $\eta$  as defined in Eq. (6) for an acceptance family  $\mathcal{A}$ , satisfying Axiom 2 and  $\mathcal{A} \cap \mathcal{X}_\emptyset = \mathcal{A}[\pi^0]$  with  $\pi^0$  as defined above in Eq. (16). We also define the normalized pricing family  $\bar{\eta}$ :

$$\bar{\eta}_t(f) \equiv \eta_t(f) - \eta_t(0), \text{ for all } f \in \mathcal{X}$$

It is well-known (see e.g. Cheridito and Kupper, 2009, Eq. 3.3) that  $\pi^0$  is cash invariant (requirement I) and time consistent (requirement II). It also satisfies requirement V, because PDE discretization methods for the calculation of conditional expectations are widely-used and can be applied directly to solve Eq. (15) from Theorem 4.2.

We do not present a formal proof of requirement III for  $\bar{\eta}$ , but instead point out two supporting facts. First, for continuous trading strategies without transaction costs it has been shown that  $\bar{\eta}$  yields the risk-neutral expectation value (cf. Davis et al., 1993, Theorem 1, or for a more recent presentation Becherer, 2003, Eq. 3.8, who calls this *elementary no-arbitrage consistency*). Secondly, we confirmed by numerical calculations that making the time between two rehedges short enough will result in

prices sufficiently close to the Black-Scholes price and optimal strategies coinciding with the Black-Scholes delta.

Again without a formal proof, requirement IV follows from another well-known result (see e.g. [Henderson and Hobson, 2002](#), Eq. 2): In the case of continuous trading without transaction costs, it holds that  $\eta_t(0) < \infty$  and the optimal strategy is given by:

$$Z_t(X_t) \equiv \frac{\mu - r}{\gamma \sigma^2 X_t} \quad (17)$$

Due to the monotonicity of the supremum in [Theorem 3.2](#) and the monotonicity of  $\pi^0$ , restricting to discrete strategies and introducing transaction costs results in even smaller prices.

### 5.2. Translating [Theorem 4.2](#) into a computational procedure

In order to perform the hedging optimization numerically we only consider a finite number of different hedging positions, i.e.  $D_{\tau_i}$  finite for all  $i \leq n$ . Then [Lemma 4.1](#) and the Markov property of  $X$ , enable us to write [Theorem 4.2](#) in a form suitable for numerical calculation. Using the ordinary functions  $z_i(h, X_{\tau_i}) \equiv -\eta_{\tau_i}(0)[\tau_{i-1} \mapsto h]$  for the price of the zero claim and  $v_i(h, X_{\tau_i}) \equiv \max\{g_{\tau_i}(X_{\tau_i}) + z_i(h, X_{\tau_i}), -\eta_{\tau_i}(-f_i)[\tau_{i-1} \mapsto h]\}$  for the price of the optimally hedged option, both with a current hedging position  $h$ , we get:

$$\begin{aligned} v_{n+1}(h, x) &= \max\{g_{\tau_{n+1}}(x), 0\} + c_{n+1}(h, x) \\ z_{n+1}(h, x) &= c_{n+1}(h, x) \\ \text{Cont}_i(h, x, b, q) &\equiv \frac{1}{\gamma} \ln \mathbb{E}_t \left[ e^{\gamma(b_{i+1}(q, X_{\tau_{i+1}}) - qX_{\tau_{i+1}})} \middle| X_{\tau_i} = x \right] + qx + c_i(h - q, x) \\ q_{\tau_i}^*(h, x) &\equiv \arg \min_{q \in D_{\tau_i}} \text{Cont}_i(h, x, v, q) \end{aligned} \quad (18)$$

$$v_i(h, x) = \max\left\{g_{\tau_i}(x) + z_i(h, x), \text{Cont}_i(h, x, v, q^*(h, x))\right\} \quad (19)$$

$$z_i(h, x) = \min_{q \in D_{\tau_i}} \text{Cont}_i(h, x, z, q)$$

We are going to compare the price obtained under the optimal hedging strategy  $q^*$  with classical delta hedging. Let  $d$  represent the ask price of an agent who rebalances daily to the optimal continuous, zero-transaction costs strategy given by the sum of the option's Black-Scholes delta,  $\Delta_t(x) = \frac{\partial}{\partial x} \text{BS}_t(x)$ , and the utility optimizing strategy from [Eq. \(17\)](#),  $Z_t(x)$ :

$$\begin{aligned} d_{n+1}(h, x) &= v_{n+1}(h, x) \\ d_i(h, x) &= \max\left\{g_{\tau_i}(x) + z_i(h, x), \text{Cont}_i(h, x, d, \Delta_{[\tau_i]}(x) + Z_{[\tau_i]}(x))\right\} \end{aligned} \quad (20)$$

$[\tau_i]$  rounds down  $\tau_i$  to the beginning of the most recent day.

Our numerical program written in C++ is a direct translation of the above equations. The conditional expectations are calculated using a finite difference method with optimal spatial finite difference weights à la [Ito and Toivanen \(2009\)](#) and Crank-Nicolson time-stepping with Rannacher startup ([Rannacher, 1984](#); [Giles and Carter, 2006](#)). The expectation values in  $\text{Cont}_i(h, x, b, q)$  are calculated for all pairs  $(b, q) \in \{v, z, d\} \times D_{\tau_i}$  in parallel, achieving quasi-linear speedup on multi-core CPUs.

We will write normalized prices as  $\bar{v} \equiv v - z$  and  $\bar{d} \equiv d - z$ .

### 5.3. Numerical results

We now present the results of the calculation using the following specifications.

The American put has a strike price  $K = 100$ , volatility  $\sigma = 50\%$  p.a., drift  $\mu = 10\%$  p.a., risk-free rate  $r = 5\%$  p.a., where one year consists of 252 business days. We assume very moderate transaction costs, with a fixed component of  $k_0 = 0.001$  and proportional component of  $k_1 = 0.025\%$ . The writer's coefficient of constant absolute risk aversion is  $\gamma = 0.001$ .

We use 8 hedging and exercise decision times per day, i.e.  $\tau_{i+1} - \tau_i = 1/8$  day. There are 150 allowed hedging positions,  $D_{\tau_i} = \{0, \delta, 2\delta, \dots, 5\}$  with  $\delta = 5/149 \approx 0.0336$ . Using a higher numbers of daily decisions and possible hedging positions does not significantly change the calculated values.

The results comprise three aspects: the writer behavior, the holder behavior and the option price.

The writer's optimal hedging position is given by  $q_t^*(x, h)$  from Eq. (18) and depends on current stock price  $x$  and hedging position  $h$ . Figure 1 plots two examples for fixed values of  $h$ . Instantly after rebalancing to this optimal position,  $h' \equiv q_t^*(x, h)$ , the current stock price  $x$  will lie in one of possibly several plateaus where  $x \mapsto q_t^*(x, h')$  is constant with value  $h'$ . As soon as the stock price leaves this plateau, it is again optimal to rebalance.

The information contained in  $q_t^*$  can be completely described by two corridors, the no-trading and the rebalancing corridor. They are depicted in Figure 2 for two different times  $t$  and have the following interpretation: it is optimal to rebalance to the nearest point of the rebalancing corridor, but only if the stock price leaves the no-trading corridor. The spikes visible in Figure 2 in both corridors occur in the vicinity of the optimal exercise boundary of the holder. An effect, of course only revealed by solving the full optimization problem.

Based on experiments with different values of  $k_0$  and  $k_1$ , we make the following numerical observations. The rebalancing corridor is always fully contained within the no-trading corridor, its width is monotonically increasing in  $k_1$ , the proportional component of the transaction costs, and it collapses for  $k_1 = 0$ . The space between the two corridors exhibits an analogous relationship with  $k_0$ , the fixed component. The delta hedging position lies within the no-trading corridor and for stock prices above the exercise boundary also within the rebalancing corridor. Without transaction costs both corridors collapse to the delta hedging strategy.

The exercise behavior of the holder is characterized by the exercise boundary, which separates the continuation regions from the their complement, the exercise regions. With conservative acceptance the writer is insured against the worst possible or *pessimist* exercise strategy. The corresponding continuation regions consist of states where the maximum in Eq. (19) or Eq. (20) equals the continuation value.

Figure 3 shows the holder's pessimist exercise boundaries against different writers. Against the optimal hedger the continuation region is only slightly larger than in the Black-Scholes case. The most striking finding is that the holder could in fact harm the delta hedger. The boundary against the delta hedger clearly exhibits a daily recurring pattern lining up with her daily rebalancings. As expected, the continuation region is much larger than in the optimal case, which can be explained by the fact that continuing the claim means more hedging costs for the (sub-optimal) delta hedger. This holds true as long as the delta hedger actually rebalances and thus incurs transaction costs. If the delta hedger's current position equals the next delta hedging position, there are no transactions costs and thus at these points (also marked in Figure 3) the extended continuation region stops.

This shows that for a delta hedger the common assumption of a Black-Scholes exercise boundary will result in an underestimation of risk caused by the interplay of American exercise feature, discrete hedging and transaction costs.

Holder and writer behavior are by-products of the main result: the price. Figures 4 and 5 provide two different perspectives on the put's normalized ask prices of an optimally hedging writer,  $\bar{v}$ , and a delta hedging writer,  $\bar{d}$ . Figure 4 shows how the absolute difference  $\bar{v} - \bar{d}$  depends on the current hedging position for fixed stock prices. As is expected, once it is optimal to rebalance, the additional transaction costs of both optimal and delta hedging will be equal, making the difference between the two independent of the current hedging position.

For a fixed current hedging position and varying stock prices the difference between both prices and the Black-Scholes price are reflected in Figure 5. It also contains the relative price reduction to be achieved by optimal hedging, calculated as  $\bar{d}/\bar{v} - 1$ .

Both figures demonstrate that the writer can always offer a more competitive price or make a sure profit by optimal hedging instead of delta hedging. This profit increases with decreasing moneyness. For example, 63 days before expiration at a stock price of 145.3, the optimal-hedging price is  $\bar{v} \approx 0.787$ , which is 10.3% lower than the delta-hedging price.

Last, reducing the risk aversion increases both the optimal-hedging and the delta-hedging price. However, the effect on the delta-hedging price is weaker and consequently, the profit from optimal hedging will be larger for a less risk-averse agent. In the above example, if we change  $\gamma$  to  $\gamma/2 = 0.0005$ ,  $\bar{v}$  will be 17.3% lower than  $\bar{d}$ .


## 6. Conclusion

Applying the methods of Gerer and Dorfleitner (2016) to the problem of hedging options with decisions allows us to derive a general hedging principle in a rigorous but straight forward manner, starting from a small set of clearly stated assumptions. This principle is then further specialized to a formula for realistically hedging American options; a formula that is not conjectured, but formally derived and proved in non-preexisting fashion.

To demonstrate how to turn this completely model-independent formula into actual numbers, we fix a market model, a pricing function and transaction costs and perform numerical calculations. The results of these numerical experiments show that when compared to the delta hedger, the optimal hedger can offer a significantly better price or make a sure profit. Further, they reveal that indeed there is a complex interaction between hedging decisions and the early exercise decisions.

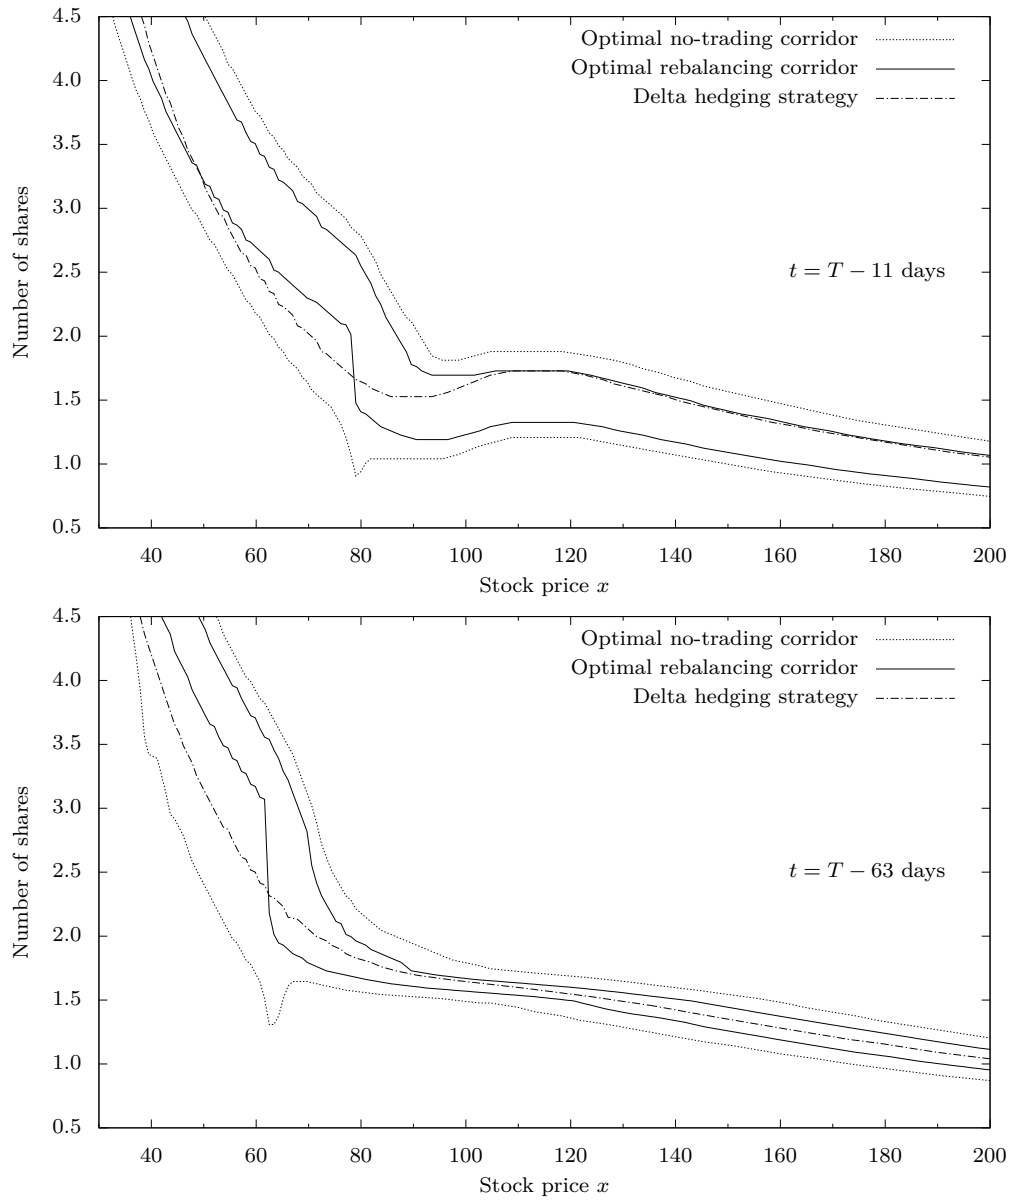
In addition to the conceptual and theoretical advantages demonstrated by our holistic approach to decisions embedded in option contracts, these results prove the usefulness of our method in realistic applications.

We leave it to further research to apply our methods more realistic models than the above example and to overcome the challenge of a numerical implementation of our results for these models.




**Figure 1:** Optimal hedging position at time  $t = T - 11$  days,  $q_t^*(x, h)$ , for different stock prices  $x$  and two different current positions  $h$ . It is optimal not to rebalance if the stock prices stay in regions with  $q_t^*(x, h) = h$ .

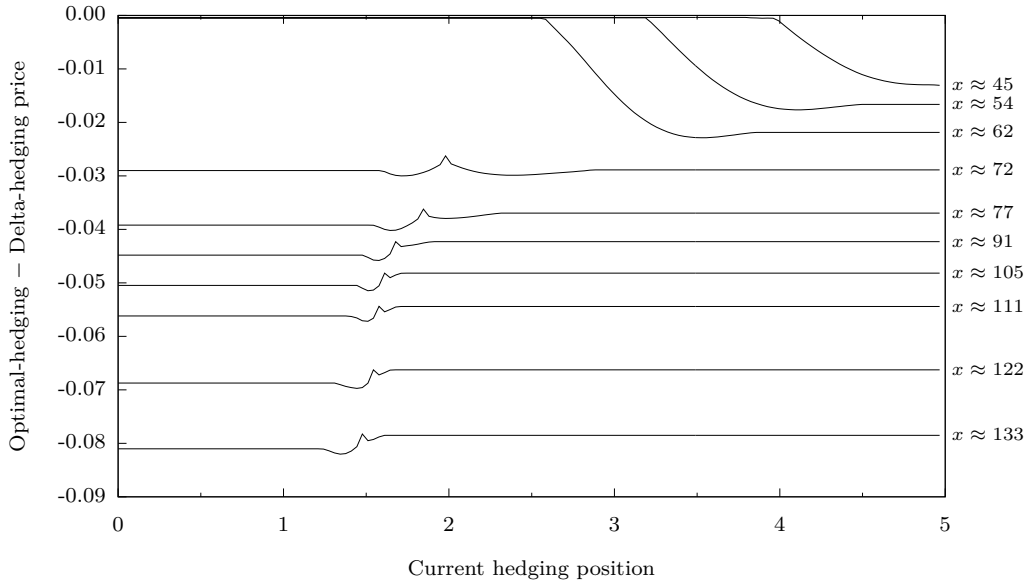





**Figure 2:** The two graphs consolidate the optimal hedging behavior at two different times. Only when the stock price leaves the no-trading corridor, it is optimal to rebalance. The new optimal position is then given by the nearest point of the rebalancing corridor. For comparison the delta hedging strategy,  $\Delta_t(x) + Z_t(x)$ , is also shown.



**Figure 3:** Worst-case exercise behavior of the holder against optimal and delta hedging writer with hedging position  $h \approx 3.02$ , described by boundaries separating exercise and continuation regions. We marked stock prices  $x$  that satisfy  $\Delta_t(x) + Z_t(x) = h$ , i.e. where the delta hedger does not need to re hedge (cf. Section 5.3). For comparison the Black-Scholes exercise boundary is also shown.



**Figure 4:** Price reduction achieved by optimal hedging compared to delta hedging, calculated at  $t = T - 63$  days. Each line corresponds to a fixed stock price  $x$  and varying current hedging positions.



**Figure 5:** Comparison of normalized optimal-hedging, delta-hedging and Black-Scholes price at time  $t = T - 63$  days with current hedging position  $h \approx 4.83$  for different stock prices  $x$ .

# Appendices

## A. Pointwise defined payoffs

**Corollary A.1.** *Given a function  $g : \Omega \times \mathbb{R}^{\mathbb{N}} \rightarrow \mathbb{R}$  and a sequence of times  $(\tau_i)_{i \in \mathbb{N}}$ , define  $f$ :*

$$f(\varphi)(\omega) = g(\omega, (\varphi_{\tau_i}(\omega))_{i \in \mathbb{N}}), \text{ for all } \varphi \in \Phi \text{ and } \omega \in \Omega \quad (21)$$

If  $f(\varphi)$  is  $\mathcal{F}_\infty$ -measurable for all  $\varphi \in \Phi$ , then  $f \in \mathcal{X}_{\{\tau_i\}_{i \in \mathbb{N}}}$ .

*Proof.* As in Definition 2.1, we assume some set  $B \in \mathcal{F}_\infty$  and two decision procedures with  $\psi_{\tau_i} \stackrel{B}{=} \varphi_{\tau_i}$  for all  $i$ . As a direct consequence of Eq. (21) we have  $\bigcap_{i \in \mathbb{N}} \{\varphi_{\tau_i} = \psi_{\tau_i}\} \subseteq \{f(\varphi) = f(\psi)\}$  and applying basic set operations yields:

$$B \setminus \{f(\varphi) = f(\psi)\} \subseteq B \setminus \bigcap_{i \in \mathbb{N}} \{\varphi_{\tau_i} = \psi_{\tau_i}\} = \bigcup_{i \in \mathbb{N}} (B \setminus \{\varphi_{\tau_i} = \psi_{\tau_i}\})$$

From monotonicity and sub-additivity of the probability measure we follow:

$$\mathbb{P}(B \setminus \{f(\varphi) = f(\psi)\}) \leq \sum_{i \in \mathbb{N}} \mathbb{P}(B \setminus \{\varphi_{\tau_i} = \psi_{\tau_i}\})$$

Due to  $\mathbb{P}(\{\varphi_{\tau_i} = \psi_{\tau_i}\} | B) = 1$  for all  $i$  (by assumption) and Corollary C.7.1 both sides of this inequality are zero and thus by the same Corollary  $\mathbb{P}(\{f(\varphi) = f(\psi)\} | B) = 1$ .  $\square$

## B. Proofs

### B.1. Proof of Lemma 3.1

*Proof.* Take any  $\varphi \in \Phi$  and show:

$$\begin{aligned} P[\mathcal{B}_t(\varphi)](f)(\varphi) &= \sup \left\{ x \in L_t^- \mid f[\varphi|_{\langle -\infty, t \rangle}] - x \in \mathcal{B}_t(\varphi) \right\} \\ &= \sup \left\{ x \in L_t^- \mid (f + H_i)[\varphi|_{\langle -\infty, t \rangle}] - x \in \mathcal{A}_t \right\} = \pi_t(f + H_i)(\varphi) \end{aligned}$$

The first and third equation use the definition of  $P[\mathcal{A}]$  (Definition 2.7). The second uses the definition of  $\mathcal{B}_t$  (Eq. (5)).  $\square$

### B.2. Proof of Theorem 3.2

*Proof.* To prove Eq. (7), we start with Lemma 3.1 and apply cash invariance (Definition 2.5, applicable due to  $H_{\tau_{n+1}} = -C_{n+1} \in \mathcal{X}_{\{\tau_n\}}^{\tau_{n+1}}$ ).

To prove Eq. (8), we define  $\Delta_a^b \equiv H_a - H_b$ . By  $H$ 's definition in Eq. (4) and Axiom 1 we can infer for any  $\tau_i < t \leq \tau_{i+1}$  and  $\varphi \in \Phi$ :

$$\begin{aligned} \Delta_{\tau_i}^t(\varphi) &= \varphi_{\tau_i} \cdot (X_t - X_{\tau_i}) - C_i(\varphi) \in L_t \\ \text{and thus } \Delta_{\tau_i}^t &\in \mathcal{X}_{\{\tau_{i-1}, \tau_i\}} \end{aligned} \quad (22)$$

$\mathcal{A}$ 's time consistency (and  $\pi$ 's normalization) with Theorem C.6 implies  $\pi$ 's time consistency. Now we can prove Eq. (8) by eliminating the two decisions at  $\tau_i$  and  $s_i$ :

$$\begin{aligned}
\eta_{\tau_i}(f) &= \pi_{\tau_i}(f + H_{\tau_{i+1}} + \Delta_{\tau_i}^{\tau_{i+1}}) \\
&= \pi_{\tau_i}(\pi_{\tau_{i+1}}(f + H_{\tau_{i+1}}) + \Delta_{\tau_i}^{\tau_{i+1}}) \\
&= \sup_{\varphi \in R_i} \pi_{\tau_i}((p + \Delta_{\tau_i}^{\tau_{i+1}})[\varphi]) \\
&= \sup_{\varphi \in R_i} \pi_{\tau_i}\left(\inf_{\psi \in S_i} \pi_{s_i}((p + \Delta_{\tau_i}^{\tau_{i+1}})[\varphi][\psi])\right) \\
&= \sup_{\varphi \in R_i} \pi_{\tau_i}\left(\inf_{\psi \in S_i} \pi_{s_i}(p[\varphi][\psi] + \varphi_{\tau_i} \cdot (X_{\tau_{i+1}} - X_{\tau_i}))\right) - C_i[\varphi]
\end{aligned}$$

The first equation follows from Lemma 3.1 and  $\Delta$ 's definition. The second uses time consistency (Definition 2.8) to introduce  $\pi_{\tau_{i+1}}$  and Eq. (22) with cash invariance to pull out  $\Delta$ . The third applies Axiom 2 and Theorem C.5 to introduce the supremum. Furthermore it applies Lemma 3.1 in reverse and abbreviates  $p \equiv \eta_{\tau_{i+1}}(f)$ . The fourth uses again time consistency and applies Axiom 2 and Theorem C.4 to introduce  $\inf \pi_{s_i}$ . The fifth expands  $\Delta$  and pulls out  $C_i[\varphi]$  using cash invariance.  $\square$

### B.3. Proof of Lemma 4.1

*Proof "≥".* Follows directly from monotonicity of the supremum<sup>4</sup> and the fact that  $t \mapsto a \in S$  for any  $a \in D_t$  by Definition of  $S$ .  $\square$

*Proof "≤".* Take any  $\varphi \in S$ . We have either  $D_t = \{a_n\}_{n \in \mathbb{N}}$  or  $\Omega = \{\omega_n\}_{n \in \mathbb{N}}$ . In the latter case define  $a_n \equiv \varphi_t(\omega_n)$  ( $\in D_t$  by Definition of  $\Phi_{\mathbf{T}}$  in Eq. (1)). In both cases define  $B_n \equiv \{\varphi_t = a_n\}$  for any  $n$ . It trivially holds  $\bigcup_n B_n = \Omega$ . By definition of  $S$ ,  $\varphi$  is adapted and thus  $B_n \in \mathcal{F}_t$ . By Definition 2.1 of  $\mathcal{X}$ , we have  $f[\varphi] \stackrel{B_n}{=} f[t \mapsto a_n]$  and thus from locality (by Corollary C.3) we can follow:

$$\pi(f[\varphi]) \stackrel{B_n}{=} \pi(f[t \mapsto a_n]) \leq \sup_{a \in D_t} \pi(f[t \mapsto a])$$

$\bigcup_n B_n = \Omega$  together with Corollary C.7.2 yields

$$\pi(f[\varphi]) \leq \sup_{a \in D_t} \pi(f[t \mapsto a]).$$

The assertion now follows from the least upper bound property of the supremum.  $\square$

### B.4. Proof of Theorem 4.2

*Proof.* Eq. (14) follows from Theorem 3.2 Eq. (7) due to  $f_{n+1} = 0$  and  $\pi$ 's normalization.

Now prove Eq. (15). With  $p$ 's definition and Theorem 3.2 Eq. (8) we get:

$$\begin{aligned}
p_i &= - \sup_{\varphi \in R_i} \pi_{\tau_i}\left(\inf_{\psi \in S_i} \pi_{s_i}(\eta_{\tau_{i+1}}(-f_i)[\varphi][\psi] + \varphi_{\tau_i} \cdot (X_{\tau_{i+1}} - X_{\tau_i}))\right) - C_i[\varphi] \\
&= \inf_{\varphi \in R_i} -\pi_{\tau_i}^0\left(\inf_{\psi \in S_i} \eta_{\tau_{i+1}}(-f_i)[\varphi][\psi] + \varphi_{\tau_i} \cdot (X_{\tau_{i+1}} - X_{\tau_i}))\right) + C_i[\varphi]
\end{aligned}$$

---

<sup>4</sup> $\sup A \leq \sup B$ , if  $A \subseteq B$

The last equation uses the inf / sup-duality and eliminates  $\pi_{s_i}$  using Eq. (13) and  $\pi$ 's cash invariance and normalization. Furthermore, as the argument of the outer pricing function does not depend on any decision, we can replace the outer pricing function by its version without decisions (according to Lemma C.1). It remains to show:

$$\inf_{\psi \in S_i} \eta_{\tau_{i+1}}(-f_i)[\psi] = \min\{\eta_{\tau_{i+1}}(-g_{s_i}), \eta_{\tau_{i+1}}(-f_{i+1})\} = -\max\{g_{\tau_{i+1}} - \eta_{\tau_{i+1}}(0), p_{i+1}\}$$

The first equation uses Eq. (11) after an application of Lemma 4.1, possible due to the explicit definition of  $S$  in Eq. (12) and the fact that  $D_{s_i}$  from Eq. (9) is finite. The second equation employs Eq. (13),  $\eta$ 's cash invariance,  $p$ 's definition and the min / max-duality.  $\square$

### C. Results from Gerer and Dorfleitner (2016, referenced hereafter as GD16)

**Lemma C.1** (Lemma A.9 in GD16). *Take a  $\mathcal{A}^0$  with dual  $\pi^0$  and some  $\mathcal{A}$ , with dual  $\pi$ , that satisfies  $\mathcal{A} \cap \mathcal{X}_\emptyset = \mathcal{A}^0$ . Taking into account Remark 3, it holds  $\pi_t(f) = \pi_t^0(f)$  for any payoff with no decisions at or after time  $t$ .*

**Definition C.2** (Locality, Corollary 3.3 in GD16). *A  $t$ -pricing function  $\pi$  is called local, if*

$$\pi(f) \stackrel{B}{=} \pi(g), \text{ for all } B \in \mathcal{F}_t \text{ and } f \stackrel{B}{=} g$$

**Corollary C.3** (Corollary 3.3 in GD16). *Cash invariance implies locality.*

**Theorem C.4** (Theorem 4.1 in GD16). *If  $\mathcal{A}$  is a proper  $t$ -acceptance set with price  $\pi$ , then  $\mathcal{A}^{\forall S}$  also is a proper  $t$ -acceptance set and its pricing function is given by*

$$P[\mathcal{A}^{\forall S}](f) = \inf_{\varphi \in S} \pi(f[\varphi]) \text{ for all } f \in \mathcal{X}_{[t, \infty)}$$

*The agent's price for any actual decisions procedure followed by the counterparty is always higher than this price. The agent can keep the difference in the form of his or her own profit. However, the counterparty can make this profit arbitrarily small (if the infimum is finite).*

**Theorem C.5** (Theorem 4.2 in GD16). *If  $S$  is  $t$ -compatible and  $\mathcal{A}$  is a proper  $t$ -acceptance set with price  $\pi$ , then  $\mathcal{A}^{\exists S}$  is also a proper  $t$ -acceptance set and its pricing function is given by:*

$$P[\mathcal{A}^{\exists S}](f) = \sup_{\varphi \in S} \pi(f[\varphi]) \text{ for all } f \in \mathcal{X}_{[t, \infty)}$$

*While this price can be higher than the price for an actually realized decision procedure by the agent, he or she can make this loss arbitrarily small (if the supremum is finite).*

**Theorem C.6** (Theorem 5.1 in GD16). *An acceptance family is time consistent (cf. Definition 2.8) if and only if its dual pricing family is time consistent.*

**Corollary C.7** (Properties of conditionally almost sure, Corollary B.1 in GD16). *Let  $\mathbf{D}$  be a countable collection of sets with positive probability and  $A, B \in \mathbf{D}$ .*

- (1)  $\mathbb{P}(B \setminus C) = 0 \iff \mathbb{P}(C|B) = 1$
- (2)  $(\forall A \in \mathbf{D} : \mathbb{P}(C|A) = 1) \implies \mathbb{P}(C|\bigcup \mathbf{D}) = 1$
- (3)  $\mathbb{P}(B|A) = \mathbb{P}(C|B) = 1 \implies \mathbb{P}(C|A) = 1$

## References

- Ahn, H. and P. Wilmott (2009). A note on hedging: restricted but optimal delta hedging, mean, variance, jumps, stochastic volatility, and costs. *Wilmott Journal* 1(3), 121–131.
- Becherer, D. (2003, aug). Rational hedging and valuation of integrated risks under constant absolute risk aversion. *Insurance: Mathematics and Economics* 33(1), 1–28.
- Bouchard, B. and E. Temam (2005). On the hedging of American options in discrete time markets with proportional transaction costs. *Electronic Journal of Probability* 10(2000), 746–760.
- Chalasani, P. and S. Jha (2001, jan). Randomized stopping times and American option pricing with transaction costs. *Mathematical Finance* 11(1), 33–77.
- Cheridito, P. and M. Kupper (2009). Recursiveness of indifference prices and translation-invariant preferences. *Mathematics and Financial Economics* 2(3), 173–188.
- Cheridito, P. and M. Stadjje (2009). Time-inconsistency of VaR and time-consistent alternatives. *Finance Research Letters* 6(1), 40–46.
- Coleman, T. F., D. Levchenkov, and Y. Li (2007, nov). Discrete hedging of American-type options using local risk minimization. *Journal of Banking & Finance* 31(11), 3398–3419.
- Constantinides, G. M. and S. Perrakis (2007, jan). Stochastic dominance bounds on American option prices in markets with frictions. *Review of Finance* 11(1), 71–115.
- Constantinides, G. M. and T. Zariphopoulou (2001). Bounds on derivative prices in an intertemporal setting with proportional transaction costs and multiple securities. *Mathematical Finance* 11(3), 331–346.
- Davis, M. H. A., V. G. Panas, and T. Zariphopoulou (1993, mar). European option pricing with transaction costs. *SIAM Journal on Control and Optimization* 31(2), 470–493.
- De Vallière, D., E. Denis, and Y. Kabanov (2008, sep). Hedging of American options under transaction costs. *Finance and Stochastics* 13(1), 105–119.
- Gerer, J. and G. Dorfleitner (2016). Time consistent pricing of options with embedded decisions. *unpublished working paper*.
- Giles, M. B. and R. Carter (2006). Convergence analysis of Crank–Nicolson and Rannacher time-marching. *The Journal of Computational Finance* 9(4), 89–112.
- Gobet, E. and N. Landon (2014, aug). Almost sure optimal hedging strategy. *Ann. Appl. Probab.* 24(4), 1652–1690.
- Henderson, V. and D. G. Hobson (2002, may). Substitute hedging. *RISK* 15(5), 71–75.
- Ito, K. and J. Toivanen (2009, jan). Lagrange multiplier approach with optimized finite difference stencils for pricing American options under stochastic volatility. *SIAM J. Sci. Comput.* 31(4), 2646–2664.
- Rannacher, R. (1984). Finite element solution of diffusion problems with irregular data. *Numerische Mathematik* 43(2), 309–327.
- Roux, A. and T. Zastawniak (2008, aug). American Options under Proportional Transaction Costs: Pricing, Hedging and Stopping Algorithms for Long and Short Positions. *Acta Applicandae Mathematicae* 106(2), 199–228.