ADAPTED WASSERSTEIN DISTANCES AND STABILITY IN MATHEMATICAL FINANCE

J. BACKHOFF-VERAGUAS, D. BARTL, M. BEIGLBÖCK, AND M. EDER

ABSTRACT. Assume that an agent models a financial asset through a measure $\mathbb{Q}$ with the goal to price / hedge some derivative or optimize some expected utility. Even if the model $\mathbb{Q}$ is chosen in the most skilful and sophisticated way, she is left with the possibility that $\mathbb{Q}$ does not provide an *exact* description of reality. This leads us the following question: will the hedge still be somewhat meaningful for models in the proximity of $\mathbb{Q}$?

If we measure proximity with the usual Wasserstein distance (say), the answer is NO. Models which are similar wrt Wasserstein distance may provide dramatically different information on which to base a hedging strategy.

Remarkably, this can be overcome by considering a suitable *adapted* version of the Wasserstein distance which takes the temporal structure of pricing models into account. This adapted Wasserstein distance is most closely related to the nested distance as pioneered by Pflug and Pichler [49, 50, 51]. It allows us to establish Lipschitz properties of hedging strategies for semimartingale models in discrete and continuous time. Notably, these abstract results are sharp already for Brownian motion and European call options.

Keywords: causal optimal transport, nested distance, hedging, utility maximization, stability, weak optimal transport.

1. INTRODUCTION

1.1. **Outline.** Assume that a reference measure $\mathbb{P}$ is used to model the evolution of a financial asset X with the purpose to hedge a financial claim or to maximize some expected utility. We do not expect that the model $\mathbb{P}$ captures reality in an absolutely accurate way. However, supposing that $\mathbb{P}$ is close enough to reality (described by a probability $\mathbb{Q}$) we would still hope that a strategy which is developed for $\mathbb{P}$ leads to reasonable results.

A main goal of this paper is to establish this intuitive idea rigorously based on a new notion of *adapted Wasserstein distance* $\mathcal{AW}_p$ between semimartingale measures. To fix ideas, we provide a first example of the results we are after.

Theorem 1.1. *Let $\mathbb{P}, \mathbb{Q}$ be continuous semimartingale models for the asset price process X, and assume that $C(X)$ denotes an L-Lipschitz payoff of a (pathdependent) derivative C. Assume that a predictable trading strategy $H = (H_t)_{t \in [0,T]}$, $|H| \leq k$ and an initial endowment $m \in \mathbb{R}$ constitute a $\mathbb{P}$-superhedge of $C(X)$, i.e.*

$$C(X) \leq m + (H \bullet X)_T, \quad \mathbb{P}\text{-almost surely.}$$

Then there is a predictable G such that m, G constitute an "almost" $\mathbb{Q}$-superhedge:

$$(1.1) \qquad \mathbb{E}_{\mathbb{Q}}[(C(X) - m - (G \bullet X)_T)^+] \leq 6(k + L) \cdot \mathcal{AW}_1(\mathbb{P}, \mathbb{Q}).$$

While the adapted Wasserstein distance will be defined in abstract terms (see (1.3)), it relates directly to the model parameters for 'simple' models. In particular, if $\mathbb{P}, \mathbb{Q}$ are Brownian models with different volatilities, than the distance between these models is just the difference of these volatilities. Moreover, the bound in (1.1) (as well as further Lipschitz bounds given below) are already sharp in such a simple setting and for C a European call option.

Below we will provide a number of results with similar flavour as Theorem 1.1. E.g. we will provide versions where the hedging error is controlled in terms of risk measures and we will show that a Lipschitz bound of the type (1.1) applies (with bigger constants) if the same trading strategy H is applied in the model $\mathbb{P}$ as well as in the model $\mathbb{Q}$. Importantly, we establish that comparable results of Lipschitz continuity apply to utility maximization.

We emphasize that familiar concepts such as the Lévy-Prokhorov metric or the usual Wasserstein distance do not appear suitable to derive results comparable to Theorem 1.1. E.g. in the vicinity of financial meaningful models there are models with *arbitrarily high arbitrage* even for bounded strategies; similar phenomena appear wrt completeness / incompleteness. Instead we introduce an adapted Wasserstein distance $\mathcal{AW}_p$ which takes the temporal structure of semimartingale models into account. These distances are conceptually closely related to the nested distance as pioneered by Pflug and Pichler [50, 51, 52]; see [1, 27, 18] for first articles which link such a type of distance to finance. We describe these contributions more closely in Section 2 below.

1.2. **Notation and adapted Wasserstein distances.** Throughout we let
$$\Omega := \mathbb{R}^T \quad \text{or} \quad \Omega := C(0, T).$$

The first setting shall be referred to as the discrete time case, and the second as the continuous time case.[1] In the first case we denote by $I = \{1, \ldots, T\}$ the time-index set, and in the second $I = [0, T]$. Throughout the article we will provide definitions and results without specifying which of the two cases we are referring to: This means that the definitions / results apply in both cases. Only occasionally will we consider one case specifically, and in this situation we will state this explicitly.

We interpret Ω as the set of all possible evolutions (in time) of the 1-dimensional asset price. Importantly, mutatis mutandis, all our results (except Lemma 3.3 and Example 3.4) remain true for multi-dimensional asset price processes (corresponding to $\Omega = (\mathbb{R}^d)^T$ / $\Omega = C([0, T], \mathbb{R}^d)$). We chose to go for the 1-dimensional version to simplify notation.

The mappings $X, Y \colon \Omega \to \Omega$ denote the canonical processes (i.e. the identity map), and we make the convention that on $\Omega \times \Omega$ the process X denotes the first coordinate and Y the second one. The spaces Ω and $\Omega \times \Omega$ are endowed with the maximum-norm and the corresponding Borel-σ-field. In continuous time, the space Ω is endowed with the right-continuous filtration generated by X, in discrete time we use the plain filtration generated by X, and in any case $\Omega \times \Omega$ is endowed with the product of these filtrations. The set $\mathrm{Cpl}(\mathbb{P}, \mathbb{Q})$ of couplings between probability measures $\mathbb{P}, \mathbb{Q}$ consists of all probability measures π on $\Omega \times \Omega$ such that $X(\pi) = \mathbb{P}$ and $Y(\pi) = \mathbb{Q}$. A Monge coupling is a coupling that is of the form $\pi = (\mathrm{Id}, T)(\mathbb{P})$ for some Borel mapping $T \colon \Omega \to \Omega$ that transports $\mathbb{P}$ to $\mathbb{Q}$, i.e. satisfies $T(\mathbb{P}) = \mathbb{Q}$. Given a metric d on Ω and $p \geq 1$, the p-Wasserstein distance of $\mathbb{P}, \mathbb{Q}$ is

$$(1.2) \qquad \mathcal{W}_p(\mathbb{P}, \mathbb{Q}) = \inf \left\{ \mathbb{E}_\pi[d(X, Y)^p]^{1/p} : \pi \in \mathrm{Cpl}(\mathbb{P}, \mathbb{Q}) \right\}.$$

In many cases of practical interest the infimum in (1.2) remains unchanged if one minimizes only over Monge couplings, cf. [53].

Before defining the adapted Wasserstein distance between measures $\mathbb{P}$ and $\mathbb{Q}$ on Ω, let us hint why distances related to weak convergence are not suitable for the results we have in mind. Assume for example that we are interested in a utility maximization problem in two periods and that Figure 1 describes the laws $\mathbb{P}, \mathbb{Q}$ of

---

[1]Indeed the arguments in the discrete and the continuous case use the same set of ideas but the presentation is significantly less technical in the discrete case which was an important reason to include the discrete case in the paper.

two traded assets. Clearly they are very close in Wasserstein distance, as follows from considering the obvious Monge coupling induced by $T : \Omega \to \Omega, T(\mathbb{P}) = \mathbb{Q}$ depicted in Figure 1. At the same time, the outcome of utility maximization is certainly very different. Similarly, $\mathbb{P}$ is a martingale measure while $\mathbb{Q}$ allows for arbitrage. The clear reason for that is the different structure of information available at time 1.
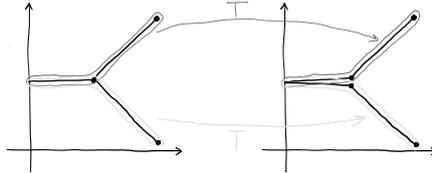
FIGURE 1. Close in Wasserstein, very different for utility maximization.

To exhibit why the Wasserstein distance does not reflect this different structure of information, let us review the transport condition $T(\mathbb{P}) = \mathbb{Q}$. We rephrase it as

$$(1.3) \qquad (T_1(X_1, X_2), T_2(X_1, X_2)) \quad \sim \quad (Y_1, Y_2).$$

While this condition is of course perfectly natural in mass transport, (1.3) almost seems like cheating when viewed from a probabilistic perspective: the map T_1 should not be allowed to consider the future value X_2 in order to determine Y_1. To define an adapted version of the Wasserstein distance, the 'process' $(T_i)_{i=1,2}$ should be taken to be *adapted* in order to account for the different information structures of $\mathbb{P}$ and $\mathbb{Q}$.

Naturally our official definition of adapted Wasserstein distances will not refer to adapted Monge transports but rather to couplings which are 'adapted' in an appropriate sense. Following Lasalle [42], we call such couplings (bi-)causal. Since the definition below may appear a bit technical at first glance, the following may be reassuring: In the discrete time setting and for absolutely continuous measures $\mathbb{P}$, the weak closure of the set of adapted Monge couplings, i.e. $\pi = (\mathrm{Id}, T)(\mathbb{P})$ for T adapted, is precisely the set of all causal couplings, see [39].

Definition 1.2 ((bi-)causal couplings)**.** *For a coupling π of $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(\Omega)$ denote by $\pi(d\omega, d\eta) = \mathbb{P}(d\omega)\pi_\omega(d\eta)$ a regular disintegration w.r.t. $\mathbb{P}$. The set $\mathrm{Cpl}_{\mathrm{C}}(\mathbb{P}, \mathbb{Q})$ of causal couplings consists of all $\pi \in \mathrm{Cpl}(\mathbb{P}, \mathbb{Q})$ such that for all $t \in I$ and $A \in \mathcal{F}_t$*

$$\omega \mapsto \pi_\omega(A) \quad \text{is a.s. } \mathcal{F}_t\text{-measurable.}$$

The set of all bi-causal *couplings $\mathrm{Cpl}_{\mathrm{BC}}(\mathbb{P}, \mathbb{Q})$ consists of all $\pi \in \mathrm{Cpl}_{\mathrm{C}}(\mathbb{P}, \mathbb{Q})$ such that also $S(\pi) \in \mathrm{Cpl}_{\mathrm{C}}(\mathbb{Q}, \mathbb{P})$, where $S : \Omega \times \Omega \to \Omega \times \Omega, S(\omega, \eta) := (\eta, \omega)$.*

In discrete time, a coupling π is causal if and only if

$$\pi\big((Y_1, \ldots, Y_t) \in A | X\big) = \pi\big((Y_1, \ldots, Y_t) \in A | X_1, \ldots X_t\big),$$

for every t and Borel set $A \subseteq \mathbb{R}^t$, that is, at time t, given the past $(X_1, \ldots, X_t)$ of X, the distribution of Y_t does not depend on the future $(X_{t+1}, \ldots, X_N)$ of X.

Replacing couplings by bi-causal couplings in (1.2) one arrives at the nested distance as introduced by Pflug and Pichler [49, 50]. Since our goal is to compare also semimartingale models in continuous time we will work with an adapted Wasserstein distance that is defined slightly differently. (Notably, the two distances are equivalent for probabilities on $\mathbb{R}^N$. We will elaborate in Remark 3.6 below, why the definition in (1.4) is more appropriate for our purposes even in discrete time.)

Denote by $\mathcal{SM}(\Omega)$ the set of all probabilities $\mathbb{P}$ on (the Borel σ-field of) Ω under which the canonical process X is a semimartingale, and for $p \in [1, \infty)$ by $\mathcal{SM}_p(\Omega)$ the subset thereof for which

$$\mathbb{E}_{\mathbb{P}}[[M]_T^{p/2} + |A|_{\text{1-var}}^p] < \infty.$$

Here $X = M + A$ denotes the unique continuous semimartingale decomposition of X under $\mathbb{P}$ into a continuous martingale M starting in zero and a continuous adapted process A of finite variation, $[\cdot]$ is the quadratic variation and $|\cdot|_{\text{1-var}}$ the first variation norm. Of course all adapted processes are semimartingales in the discrete time case.

Definition 1.3 (Adapted Wasserstein distance). *For* $\mathbb{P}, \mathbb{Q} \in \mathcal{SM}_p(\Omega)$, $p \geq 1$ *set*

(1.4)
$$\mathcal{AW}_p(\mathbb{P}, \mathbb{Q}) := \inf \left\{ \mathbb{E}_{\pi} \left[[M^X - M^Y]_T^{p/2} + |A^X - A^Y|_{1-var}^p \right]^{1/p} : \pi \in \text{Cpl}_{\text{BC}}(\mathbb{P}, \mathbb{Q}) \right\},$$

where $X = M^X + A^X, Y = M^Y + A^Y$ *denote the semimartingale decomposition of* X *and* Y *resp.*

It is shown in Lemma 3.1 that $\mathcal{AW}_p$ is well-defined (i.e. that $X - Y$ is a semimartingale under every bi-causal coupling) and in Lemma 3.2 that $\mathcal{AW}_p$ in fact defines a metric.

Remark 1.4. *In the continuous time setup, the adapted Wasserstein distance can also be computed through*

$$\mathcal{AW}_p(\mathbb{P}, \mathbb{Q}) = \inf \left\{ \mathbb{E}_{\pi} \left[[X - Y]_T^{p/2} + MV_T[|X - Y|^p] \right]^{1/p} : \pi \in \text{Cpl}_{\text{BC}}(\mathbb{P}, \mathbb{Q}) \right\}.$$

Here MV *denotes the mean variation, i.e.* $MV_T[Z] = \sup_{\Delta} \sum_{t_j \in \Delta} |\mathbb{E}[Z_{t_{j+1}} - Z_{t_j}|\mathcal{F}_{t_j}]|$, *where the supremum is taken over all finite partitions* Δ *of* $[0, T]$.

In Section 3.1 below we will give explicit formulae for the adapted Wasserstein distance in the case of semi-martingale measures described by simple SDEs.

1.3. **Stability of Superhedging.** For the rest of this article, fix some $k \in \mathbb{R}_+$ and let $\mathcal{H}_k$ be the set of all predictable processes

$$H: \Omega \times I \to [-k, k].$$

For every $p \geq 1$, write b_p for the 'upper' Burkholder-Davis-Gundy (BDG) constant, cf. Remark 3.12 below. In particular it is known that $b_1 \leq 6$ and that $b_2 = 2$.

Our first main result concerns the stability of superhedging and constitutes a stronger version of Theorem 1.1 stated above.

Theorem 1.5. *Let* $\mathbb{P}, \mathbb{Q} \in \mathcal{SM}_1(\Omega)$, $H \in \mathcal{H}_k$ *and let* $C: \Omega \to \mathbb{R}$ *be Lipschitz with constant* L. *Then the hedging error under* $\mathbb{Q}$ *is bounded by the distance of* $\mathbb{P}$ *and* $\mathbb{Q}$ *plus the hedging error under* $\mathbb{P}$ *in the following sense: there exists* $G \in \mathcal{H}_k$ *such that*

(WHI)
$$\mathbb{E}_{\mathbb{Q}}[(C - m - (G \bullet X)_T)^+] \leq \mathbb{E}_{\mathbb{P}}[(C - m - (H \bullet X)_T)^+] \\ + b_1(k + L)\mathcal{AW}_1(\mathbb{P}, \mathbb{Q}).$$

Assume in addition that $H_t: \Omega \to \mathbb{R}$ *is Lipschitz with constant* $\tilde{L}$ *for every* $t \in I$. *Then we can take* $G = H$ *and obtain*

(SHI)
$$\mathbb{E}_{\mathbb{Q}}[(C - m - (G \bullet X)_T)^+] \leq \mathbb{E}_{\mathbb{P}}[(C - m - (H \bullet X)_T)^+] \\ + b_1(k + L)\mathcal{AW}_1(\mathbb{P}, \mathbb{Q}) + \beta \mathcal{AW}_2(\mathbb{P}, \mathbb{Q}),$$

where $\beta := 2\sqrt{2}b_1\tilde{L} \min\{\mathcal{AW}_2(\mathbb{P}, \delta_0), \mathcal{AW}_2(\mathbb{Q}, \delta_0)\}$.

Importantly, it is impossible to transfer a superhedge under $\mathbb{P}$ into a superhedge under $\mathbb{Q}$. This occurs already in a one-period framework and is not a by-product of our definition of adapted Wasserstein distance; see Remark 5.2. A similar reasoning requires to consider only trading strategies bounded by k; see Remark 5.3.

It is worthwhile to compare the relation of the two inequalities (WHI) and (SHI):

(S) In a certain sense the 'strong hedging inequality' (SHI) seems to be the more relevant assertion: after all a trader does not know that the model $\mathbb{Q}$ (rather than the model $\mathbb{P}$) describes reality and hence she might (somewhat stubbornly) stick to the initial plan of hedging her risk according to the strategy H. The inequality (SHI) then allows to quantify the losses due to this model-error.

(W) However, the 'weak hedging inequality' (WHI) also has a particular merit: suppose that a trader W starts with the prior belief that the asset price evolves according to a Black-Scholes model with volatility σ_1 but soon after time 0 realizes that a volatility σ_2 (where $\sigma_2 \neq \sigma_1$) yields a more adequate description of reality. If the witty trader W makes an accurate guess about the correct model and updates her trading strategy accordingly, her losses can be controlled through the tighter bound in (WHI).

In Theorem 4.2 we provide a version of Theorem 1.5, where $(\cdot)^+$ is replaced by a convex, strictly increasing loss function $l\colon \mathbb{R} \to \mathbb{R}_+$.

Another way to gauge the effectiveness of an almost superhedge is by means of risk measures. We postpone the general formulation to Theorem 4.3 and first present a version that appeals to the average value of risk $\mathrm{AVaR}_\alpha^{\mathbb{P}}$. Recall that for a random variable $Z\colon \Omega \to \mathbb{R}$

$$\mathrm{AVaR}_\alpha^{\mathbb{P}}(Z) := \inf_{m \in \mathbb{R}} \mathbb{E}_{\mathbb{P}}[(Z-m)^+/\alpha + m],$$

is the average value at risk at level $\alpha \in (0,1)$ under model $\mathbb{P}$. We then have

Theorem 1.6. *Assume that $C\colon \Omega \to \mathbb{R}$ is Lipschitz with constant L. Then*

$$\left| \inf_{H \in \mathcal{H}_k} \mathrm{AVaR}_\alpha^{\mathbb{P}}(C - (H \bullet X)_T) - \inf_{H \in \mathcal{H}_k} \mathrm{AVaR}_\alpha^{\mathbb{Q}}(C - (H \bullet X)_T) \right| \leq r\,\mathcal{AW}_1(\mathbb{P}, \mathbb{Q}),$$

for $r := b_1(L+k)/\alpha$. If $H \in \mathcal{H}_k$ is such that $H_t\colon \Omega \to [-k,k]$ is Lipschitz with constant $\tilde{L}$ for every $t \in I$ and β is the constant defined in Theorem 1.5, then

$$\left| \mathrm{AVaR}_\alpha^{\mathbb{P}}(C - (H \bullet X)_T) - \mathrm{AVaR}_\alpha^{\mathbb{Q}}(C - (H \bullet X)_T) \right| \leq r\,\mathcal{AW}_1(\mathbb{P}, \mathbb{Q}) + \frac{\beta}{\alpha}\mathcal{AW}_2(\mathbb{P}, \mathbb{Q}).$$

The interpretation of this result is similar to the one of Theorem 1.5: As $\mathrm{AVaR}_\alpha^{\mathbb{P}}(\cdot)$ is translation invariant, one has

$$\inf_{H \in \mathcal{H}_k} \mathrm{AVaR}_\alpha^{\mathbb{P}}(C - (H \bullet X)_T) = \inf\left\{ m \in \mathbb{R} : \begin{array}{l} \text{there is } H \in \mathcal{H}_k \text{ such that} \\ \mathrm{AVaR}_\alpha^{\mathbb{P}}(C - m - (H \bullet X)_T) \leq 0 \end{array} \right\},$$

and the right-hand side constitutes a relaxed version of the superhedging price.

Notably, the explicit calculations of adapted Wasserstein distance given in Section 3.1 imply that Theorem 1.6 (and similarly Theorem 1.5) are sharp

Example 1.7 (Hedging in a Brownian framework). *Consider a European call option $C(X) = (X_T - K)^+$, where for simplicity $K = 0$. Moreover, let $\mathbb{P}^\sigma$ be Wiener measure with constant volatility $\sigma \geq 0$. Then for every $\sigma, \hat{\sigma} \geq 0$, $k \geq 1$, and $\alpha \in (0,1)$ it holds that (we defer the proof of this fact to Section 4)*

$$\left| \inf_{H \in \mathcal{H}_k} \mathrm{AVaR}_\alpha^{\mathbb{P}^\sigma}(C - (H \bullet X)_T) - \inf_{H \in \mathcal{H}_k} \mathrm{AVaR}_\alpha^{\mathbb{P}^{\hat{\sigma}}}(C - (H \bullet X)_T) \right|$$

$$= \left| \mathbb{E}_{\mathbb{P}^\sigma}[C] - \mathbb{E}_{\mathbb{P}^{\hat{\sigma}}}[C] \right| = \frac{1}{\sqrt{2\pi}}T|\sigma - \hat{\sigma}| = \frac{1}{\sqrt{2\pi}}\mathcal{AW}_1(\mathbb{P}^\sigma, \mathbb{P}^{\hat{\sigma}}).$$

This shows that the estimate in Theorem 1.6 is tight (up to constants), in the sense that it is essentially impossible to improve on the probability metric $\mathcal{AW}_1$.

We make the important remark that Glanzer, Pflug, and Pichler [27] use the nested distance to control acceptability prices in discrete time models in a Lipschitz fashion through the nested distance of these models. Specifically, in a discrete one-period framework [27, Proposition 3] and Theorem 1.6 yield almost the same assertion: in this setup, the only difference is that [27, Proposition 3] does not specify a Lipschitz constant and does not assume uniform boundedness of the admissible hedging strategy. (However, the latter seems to be in conflict with our Remark 5.3 below.)

1.4. **Stability of Utility Maximization.** We move on to consider the continuity of utility maximization. Let $U \colon \mathbb{R} \to \mathbb{R}$, be a utility function which is concave, increasing, and denote by U' the left-continuous version of the derivative. We have

Theorem 1.8. *Let $C \colon \Omega \to \mathbb{R}$ be Lipschitz continuous and assume that there exists $c \geq 0$ such that $U'(x) \leq c(1 + |x|^{p-1})$ for all x. Then, for every $R \geq 0$ there exists a constant K depending on k, p, U, C, R such that*

$$\left| \sup_{H \in \mathcal{H}_k} \mathbb{E}_{\mathbb{P}}[U(C + (H \bullet X)_T)] - \sup_{H \in \mathcal{H}_k} \mathbb{E}_{\mathbb{Q}}[U(C + (H \bullet X)_T)] \right| \leq K \cdot \mathcal{AW}_p(\mathbb{P}, \mathbb{Q}),$$

for all $\mathbb{P}$ and $\mathbb{Q}$ with $\mathcal{AW}_p(\mathbb{P}, \delta_0) \leq R$ and $\mathcal{AW}_p(\mathbb{Q}, \delta_0) \leq R$.

The failure of usual Wasserstein distances to guarantee stability of utility maximization is illustrated in Remark 5.1.

1.5. **Structure of the paper.** In Section 2 we briefly review the literature related to this paper. In Section 3 we establish some basic properties of the adapted Wasserstein distance, discuss the choice of cost function and give some examples. Moreover we derive a contraction principle (Theorem 3.10) which relates adapted Wasserstein distance with a 'weak' (in the sense of Gozlan et al [29]) transport distance. This result forms the basis for the proofs of the results mentioned in the introduction, as well as certain extensions of these results, see Section 4. Finally we conclude with some remarks in Section 5.

2. Literature

The articles closest in spirit to ours are [1, 18, 27]. Acciaio, Zalashko and one of the present authors consider in [1] an object related to the adapted Wasserstein distance in continuous time in connection with utility maximization, enlargement of filtrations and optimal stopping. Glanzer, Pflug, and Pichler [27] prove a deviation-inequality for the so-called nested distance in a discrete time framework, and consider acceptability pricing over an ambiguity set given described through the nested distance. Bion-Nadal and Talay [18] study via PDE arguments a continuous-time optimization problem which is related related to the adapted Wasserstein distance.

The concept of causal couplings, and optimal transport over causal couplings, has been recently popularized by Lassalle [42] although precursors can be found in the works [58, 55]. This notion is central to the recent articles [1, 9, 7, 8].

The idea of strengthening weak convergence of measures in order to account for the temporal evolution has some history. Indeed several authors have independently introduced different approaches to address this challenge: The seminal unpublished work by Aldous [2] introduces the notion of extended weak convergence for the study of stability of optimal stopping problems. The principal idea is not to compare the laws of process directly, but rather the laws of the corresponding prediction

processes. Independently, Hellwig [30] introduces the information topology for the stability of equilibrium problems in economics. Roughly, two probability measures on a product of finitely many spaces $X_1 \times \ldots \times X_N$ are considered to be close is for each $t \leq N$ the projections onto the first t coordinates as well as the corresponding conditional (regular) disintegrations are close. Unrelated to these developments Pflug and Pichler [49, 50, 51] have introduced the nested distances for the stability of stochastic programming in discrete time. The nested distance is the obvious role model for the adapted Wasserstein distances considered in this article and (as mentioned above) for a fixed number of time steps and $p \geq 1$, they are obviously equivalent. Yet another idea to account for the temporal evolution of processes would be to symmetrise the causal transport costs $\mathcal{W}_c(\mathbb{P}, \mathbb{Q})$ costs defined by Lassalle [42] by taking the maximum or sum of $\mathcal{W}_c^2(\mathbb{P}, \mathbb{Q})$ and $\mathcal{W}_c^2(\mathbb{Q}, \mathbb{P})$; this was pointed out by Soumik Pal.

In parallel work [6], the four authors of the present article investigate the relations between these concepts in detail. Remarkably, in discrete time *all* of the concepts mentioned above (adapted Wasserstein distances, extended weak convergence, information topology, and nested distances, symmetrised causal transport costs) define the *same topology*.

The question of stability in mathematical finance has been studied from different perspectives over the years. Notably, starting with the articles of Lyons [43] and Avellaneda, Levy, Paras [5] the area of robust finance has mainly focused on extremal models and hedging strategies which dominate the payoff for every model in a specified class. Following the publication of Hobson's seminal article [33] connections with the Skorokhod embedding problem have been a driving force of the field, see the surveys of Hobson [34] and Obłój [45]. Recently this has been complemented by techniques coming from (martingale) optimal transport, early papers which advance this viewpoint include [35, 14, 26, 15, 19, 24, 22, 13]. The literature on 'local' misspecification of volatility in a sense more closely related to the present article appears more spare. El Karoui, Jeanblanc, and Shreve [25] establish in a stochastic volatility framework that if the misspecified volatility dominates the true volatility, then the misspecified price of call options dominates the real price; see also the elegant account of Hobson [36]. More recently, the question of pricing and hedging under uncertainty about the volatility of reference local volatility model is studied by Herrmann, Muhle-Karbe, and Seifried [32] (see also [31]). Less plausible models are penalized through a mean square distance to the volatility of the reference model and the authors obtain explicit formulas for prices and hedging strategies in a limit for small uncertainty aversion. Becherer and Kentia [12] derive worst-case good-deal bounds under model ambiguity which concerns drift as well as volatility. Indeed, discussions with Dirk Becherer motivated us to consider also models with drift in our results on stability of super hedging. The behaviour of the superhedging price in a ball (wrt. various notions of distance) around a reference model is studied in depth by Obłój and Wiesel [46] for a d-dimensional asset and one time period.

A notable implication of this work is that it yields a coherent way to measure model-uncertainty (in the sense of Cont's influential article [23]): Fix a subset M_0 of the set M of all consistent models, i.e. martingale measures which are consistent with benchmark instruments whose price can be observed on the market. Given M_0, the model uncertainty associated to a derivative f can be gauged through

$$\rho_{M_0}(f) := \sup\{\mathbb{E}_{\mathbb{Q}} f : \mathbb{Q} \in M_0\} - \inf\{\mathbb{E}_{\mathbb{Q}} f : \mathbb{Q} \in M_0\}.$$

The worst-case approach typically pursued in robust finance then yields $\rho_{M_0}(f)$ for $M_0 = M$, but it appears equally natural to take M_0 to be an infinitesimal ball around a reference model. This approach is first carried out by Drapeau, Obłój,

Wiesel and one of the present authors [11] in a one period framework. Our results indicate that adapted Wasserstein distance provides a way to extend this to a multi-period setup, we intend to pursue this further in future work.

On a different note, much work has been done regarding the convergence of discrete time models to their continuous time analogues. Due to the vastness of this literature we refer the reader to the book [54] for references. Finally, in more recent times and starting from the works of Kardaras and Žitković, the stability of utility maximization has been studied in [38, 40, 41, 44, 57] among others.

3. The adapted Wasserstein distance

The following Lemma shows that $\mathcal{AW}_p$ is well-defined.

Lemma 3.1. *Let* $\mathbb{P}, \mathbb{Q}$ *be (semi-)martingale measures for* $X, Y \colon \Omega \to \Omega$, *respectively, and let* π *be a bi-causal coupling between* $\mathbb{P}$ *and* $\mathbb{Q}$. *Then* $X, Y, X - Y \colon \Omega \times \Omega \to \Omega$ *are (semi)-martingales w.r.t.* π. *Further, if* $X = M + A$ *denotes the semimartingale decomposition under* $\mathbb{P}$, *then up to evanescence* $M + A$ *is the semimartingale decomposition of* X *under* π.

Proof. Let $X = M + A$ be the semimartingale decomposition under $\mathbb{P}$ and consider M and A as processes on $\Omega \times \Omega$ via $M(\omega, \eta) := M(\omega)$ and $A(\omega, \eta) := A(\omega)$. Further let $\pi = \mathbb{P}(d\omega)\pi_\omega(d\eta)$ be a bi-causal coupling between $\mathbb{P}$ and $\mathbb{Q}$. To show that $X = M + A$ remains the semimartingale decomposition under π, it is enough to show that M is a martingale under π. To that end, let $0 \leq s \leq t$ and let $Z \colon \Omega \times \Omega \to \mathbb{R}$ be $\mathcal{F}_s \otimes \mathcal{F}_s$-measurable and bounded. Then the random variable $Z' \colon \Omega \to \mathbb{R}$ defined by

$$Z'(\omega) := \int Z(\omega, \eta)\, \pi_\omega(d\eta) \quad \text{is } \mathcal{F}_s\text{-measurable}$$

up to $\mathbb{P}$-null sets, and clearly and bounded. Indeed, if $Z(\omega, \eta) = Z^1(\omega)Z^2(\eta)$ for $\mathcal{F}_s$-measurable bounded functions Z^1 and Z^2, then it follows from the definition of bi-causality that Z' is $\mathcal{F}_s$-measurable modulo $\mathbb{P}$; the general statement then follows from a monotone class argument. Therefore

$$\mathbb{E}_\pi[(M_t - M_s)Z] = \int (M_t(\omega) - M_s(\omega)) \int Z(\omega, \eta)\, \pi_\omega(d\eta)\, \mathbb{P}(d\omega)$$
$$= \mathbb{E}_\mathbb{P}[(M_t - M_s)Z']$$
$$= 0,$$

by the martingale property of M under $\mathbb{P}$. This shows that M is a martingale under π and therefore that $X = M + A$ is the semimartingale decomposition under π. $\square$

Lemma 3.2. $\mathcal{AW}_p$ *defines a metric on the set* $\mathcal{SM}_p(\Omega)$.

Proof. It is clear that $\mathcal{AW}_p(\mathbb{P}, \mathbb{Q}) = \mathcal{AW}_p(\mathbb{Q}, \mathbb{P}) \geq 0$ for all $\mathbb{P}, \mathbb{Q} \in \mathcal{SM}_p(\Omega)$. Suppose that $\mathcal{AW}_p(\mathbb{P}, \mathbb{Q}) = 0$. As $\|\cdot\|_\infty \leq |\cdot|_{1\text{-var}}$, it is immediate that if π participates in the minimization problem defining $\mathcal{AW}_p(\mathbb{P}, \mathbb{Q})$, then

$$\mathbb{E}_\pi[\|X - Y\|_\infty^p] \leq 2^{p-1}\mathbb{E}_\pi[\|M\|_\infty^p + |A|_{1\text{-var}}^p]$$
$$\leq 2^{p-1}b_p\mathbb{E}_\pi[[M]_T^{p/2} + |A|_{1\text{-var}}^p]$$

where b_p denotes the BDG constant and we used the BDG inequality for the martingale M. Hence the usual Wasserstein distance between $\mathbb{P}$ and $\mathbb{Q}$ (defined w.r.t. the $\|\cdot\|_\infty$-norm) is dominated from above by $\mathcal{AW}_p(\mathbb{P}, \mathbb{Q})$, and so $\mathbb{P} = \mathbb{Q}$.

We now prove the triangle inequality. Let $\mathbb{P}, \mathbb{Q}, \mathbb{R}$ given. We fix $\varepsilon > 0$ and assume π is bi-causal ε-optimal for $\mathcal{AW}_p(\mathbb{P}, \mathbb{Q})$ and $\tilde{\pi}$ is bi-causal ε-optimal for $\mathcal{AW}_p(\mathbb{Q}, \mathbb{R})$. Let

$$\pi(d\omega, d\eta) = \pi_\eta(d\omega)\, \mathbb{Q}(d\omega) \quad \text{and} \quad \tilde{\pi}(d\eta, d\gamma) = \tilde{\pi}_\eta(d\gamma)\mathbb{Q}(d\eta)$$

be disintegrations, and define $\Pi \in \mathcal{P}(\Omega^3)$ by

$$\Pi(d\omega, d\eta, d\gamma) = \pi_\eta(d\omega)\, \tilde{\pi}_\eta(d\gamma)\, \mathbb{Q}(d\eta).$$

If $\overline{\pi}(d\omega, d\gamma) := \int_\Omega \Pi(d\omega, d\eta, d\gamma)$ is the projection of Π onto the first and third components, then it is clear that the first and second marginals of $\overline{\pi}$ are $\mathbb{P}$ and $\mathbb{R}$ respectively. Moreover, a disintegration of $\overline{\pi} = \overline{\pi}_\omega(d\gamma)\, \mathbb{P}(d\omega)$ is given by

$$\overline{\pi}_\omega(d\gamma) = \int_\Omega \tilde{\pi}_\eta(d\gamma)\, \pi_\omega(d\eta).$$

In particular, for every $A \in \mathcal{F}_t$ we have that $\omega \mapsto \overline{\pi}_\omega(A)$ is $\mathcal{F}_t$-measurable modulo $\mathbb{P}$. The argument for $\overline{\pi} = \overline{\pi}_\gamma(d\omega)\, \mathbb{R}(d\gamma)$ is similar and therefore $\overline{\pi}$ is a bi-causal coupling between $\mathbb{P}$ and $\mathbb{R}$. Finally, it follows as in the proof of Lemma 3.1 that, if $X = M^X + A^X$, $Y = M^Y + A^Y$, and $Z = M^Z + A^Z$ are the semimartingale decompositions under $\mathbb{P}$, $\mathbb{Q}$, and $\mathbb{R}$, then they remain the semimartingale decomposition under Π on Ω^3 endowed with the product filtration.

To finish the proof of the triangle inequality, we observe that

$$\mathcal{AW}_p(\mathbb{P}, \mathbb{R}) \leq \mathbb{E}_{\overline{\pi}}[[M^X - M^Z]_T^{p/2} + |A^X - A^Z|_{1\text{-var}}^p]^{1/p}$$

$$= \mathbb{E}_\Pi\left[[(M^X - M^Y) + (M^Y - M^Z)]_T^{p/2} + \cdots \right.$$

$$\left. \cdots + |(A^X - A^Y) + (A^Y - A^Z)|_{1\text{-var}}^p \right]^{1/p}.$$

The function $M \mapsto \mathbb{E}_\Pi[[M]_T^{p/2}]^{1/p}$ is known to be a norm on the space $\mathcal{M}_p(\Pi)$ of Π-martingales started at zero whose supremum is p-integrable. Likewise $A \mapsto \mathbb{E}_\Pi[|A|_{1\text{-var}}^p]^{1/p}$ is a norm on the space of finite variation processes with p-integrable variation. Hence

$$(M, A) \mapsto \|(M, A)\| := \mathbb{E}_\Pi[[M]_T^{p/2} + |A|_{1\text{-var}}^p]^{1/p}$$

is a norm on the product of these spaces. We conclude the proof for the triangle inequality with

$$\mathcal{AW}_p(\mathbb{P}, \mathbb{R}) \leq \|(M^X - M^Y, A^X - A^Y) + (M^Y - M^Z, A^Y - A^Z)\|$$

$$\leq \|(M^X - M^Y, A^X - A^Y)\| + \|(M^X - M^Y, A^X - A^Y)\|$$

$$= \mathbb{E}_\pi[[M^X - M^Y]_T^{p/2} + |A^X - A^Y|_{1\text{-var}}^p]^{1/p}$$

$$+ \mathbb{E}_{\tilde{\pi}}[[M^Y - M^Z]_T^{p/2} + |A^Y - A^Z|_{1\text{-var}}^p]^{1/p}$$

$$\leq 2\varepsilon + \mathcal{AW}_p(\mathbb{P}, \mathbb{Q}) + \mathcal{AW}_p(\mathbb{Q}, \mathbb{R}),$$

since the semimartingale decomposition of $X - Y$ under π is $(M^X - M^Y) - (A^X - A^Y)$, with an analogous expression for $Y - Z$ under $\tilde{\pi}$.

To conclude the proof, it remains to show that $\mathcal{AW}_p(\mathbb{P}, \mathbb{Q}) < \infty$ for all $\mathbb{P}, \mathbb{Q} \in \mathcal{SM}_p(\Omega)$. By Lemma 3.1, we have $\mathcal{AW}_p(\mathbb{P}, \delta_0) = \mathbb{E}_\mathbb{P}[[M]_T^{p/2} + |A|_{1\text{-var}}^p]^{1/p}$ where $X = M + A$ is the semimartingale decomposition under $\mathbb{P}$. Therefore the triangle inequality implies that $\mathcal{AW}_p$ is real-valued on $\mathcal{SM}_p(\Omega)$. $\qquad \square$

3.1. **On the choice of the cost function and examples.** We start by a simple result which permits to give a closed-form expression of the adapted Wasserstein distance in given continuous-time situations:

Lemma 3.3. *For $i \in \{1, 2\}$ consider the SDEs with bounded progressive coefficients:*

$$(3.1) \qquad\qquad dX_t^i = \mu_i(t, \{X_s^i\}_{s \leq t}) dt + \sigma_i(t, \{X_s^i\}_{s \leq t}) dB_t^i.$$

Assume that each SDE admits a unique strong solution and denote by $\mathbb{P}^{\mu_i, \sigma_i}$ the respective laws. Further assume that

- *μ_1 is a function of time only (namely $\mu_1 \colon [0, T] \to \mathbb{R}$)*

- $\sigma_1, \sigma_2 \geq 0$ and at least one of them is a function of time only.

Then the synchronous coupling (namely $\pi^* =$ joint law of (X^1, X^2), where $B^1 = B^2$ in (3.1)), is optimal in the definition of $\mathcal{AW}_p(\mathbb{P}^{\mu_1, \sigma_1}, \mathbb{P}^{\mu_2, \sigma_2})$.

The discrete time version of the aforementioned synchronous coupling is given by the Knothe-Rosenblatt rearrangement [9], and a variant of the previous result can also be obtained in the discrete time framework.

Proof. Let π be a feasible coupling for $\mathcal{AW}_p(\mathbb{P}^{\mu_1, \sigma_1}, \mathbb{P}^{\mu_2, \sigma_2})$, leading to a finite cost. Naturally for this proof we denote the coordinate process on $\Omega \times \Omega$ by (X^1, X^2). As before we let $X^i = A^i + M^i$ be the unique continuous Doob-Meyer decomposition of X^i under the $\mathbb{P}^{\mu_i, \sigma_i}$-completion of its right-continuous filtration. Observe that $\frac{d}{dt} A^1$ is a.s. deterministic, by the assumption on μ_1, and that the law of $\frac{d}{dt} A^2$ is independent of the coupling π. Both facts can be derived easily from the identity

$$\frac{d}{dt} A_t^i = \lim_{\varepsilon \searrow 0} \frac{\mathbb{E}_\pi \left[X_{t+\varepsilon}^i | \mathcal{F}_t^{X^i} \right] - X_t^i}{\varepsilon},$$

which by Lebesgue differentiation theorem holds $dt \times d\pi$-a.s. As a consequence, the term $\mathbb{E}_\pi[|A^1 - A^2|_{1-var}^p]$ is independent of the coupling π and so we may ignore it and only focus on the term $\mathbb{E}_\pi[[M^1 - M^2]_T^{p/2}]$.

By Doob's martingale representation [37, Theorem 4.2], in a possibly enlarged filtered probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\pi})$ we may represent the martingale (M^1, M^2) by

$$M_t^i = \int_0^t \sigma_{i1} dW + \int_0^t \sigma_{i2} d\hat{W},$$

where $W, \hat{W}$ are independent standard one-dimensional Brownian motions and $\{\sigma_{ik} : i, k \in \{1, 2\}\}$ real-valued processes, both of them adapted in the enlarged filtered space. In the following we will omit the argument $\{X_s^i\}_{s \leq t}$ from σ_i. Necessarily

$$\sigma_i^2 = \frac{d}{dt} [M^i]_t = \sigma_{i1}^2 + \sigma_{i2}^2, \ (dt \times d\tilde{\pi} - a.s.).$$

By Cauchy-Schwarz inequality we deduce that almost surely

$$[M^1, M^2]_T = \int_0^T [\sigma_{11} \sigma_{21} + \sigma_{12} \sigma_{22}] dt \leq \int_0^T \sigma_1 \sigma_2 dt,$$

and accordingly we get the lower bound

$$\mathbb{E}_\pi[[M^1 - M^2]_T^{p/2}] \geq \mathbb{E}_\pi \left[\left(\int_0^T (\sigma_1 - \sigma_2)^2 dt \right)^{p/2} \right].$$

As in the beginning of the proof, the right-hand side does not depend on the coupling π thanks to either σ^i being a function of time only. To conclude observe that for the synchronous coupling π^* we have equality in the above equation.   □

As an easy consequence we have

Example 3.4. *For bounded Lipschitz functions $\mu_1, \mu_2, \sigma_1, \sigma_2$ we denote by $\mathbb{P}^{\mu_i, \sigma_i}$ the law of the diffusion*

$$dX_t^i = \mu_i(t, X_t^i) dt + \sigma_i(t, X_t^i) dB_t.$$

Assume that

- *μ_i is independent of the x-variable, some $i \in \{1, 2\}$, and*
- *σ_k is independent of the x-variable, some $k \in \{1, 2\}$.*

Calling $j \in \{1, 2\}\backslash\{i\}$ and $\ell \in \{1, 2\}\backslash\{k\}$, we have

$$\mathcal{AW}_p(P^{\mu_1, \sigma_1}, P^{\mu_2, \sigma_2})^p = \mathbb{E}\Big[\Big(\int_0^T [\sigma_\ell(t, X_t^\ell) - \sigma_k(t)]^2 dt\Big)^{p/2}\Big]$$

$$+ \mathbb{E}\Big[\Big(\int_0^T |\mu_j(t, X_t^j) - \mu_i(t)|dt\Big)^p\Big].$$

We now illustrate that in general it is not true that the straightforward synchronous coupling of Lemma 3.3 is optimal. As a consequence, we do not expect a closed-form expression for the adapted Wasserstein distance. A discrete-time version of this observation is discussed in [7, Section 7].

Example 3.5. *Consider $d = 1$, $T = 2$, and for each $c \in \mathbb{R}$ introduce*

$$\mu_t^c(\omega) := c 1_{[1,2]}(t) \operatorname{sign}(\omega_1) \quad \text{and} \quad \hat{\mu}_t^c(\omega) := -\mu_t^c(\omega).$$

Assuming that B is a Brownian motion, and for $\sigma \in \mathbb{R}_+$, we introduce the couplings

$$\pi_1 := \operatorname{Law}\Big(\sigma B + \int \mu_t^c(B) dt, \ \sigma B + \int \hat{\mu}_t^c(B) dt\Big),$$

$$\pi_2 := \operatorname{Law}\Big(\sigma B + \int \mu_t^c(B) dt, \ -\sigma B + \int \hat{\mu}_t^c(-B) dt\Big).$$

These couplings share the same marginals and each of them is bi-causal. It is easy to compute

$$\mathbb{E}_{\pi_1}\big[[M]_T^{p/2} + |A|_{1\text{-}var}^p\big] = (2c)^p,$$

$$\mathbb{E}_{\pi_2}\big[[M]_T^{p/2} + |A|_{1\text{-}var}^p\big] = (8\sigma^2)^{p/2}.$$

We conclude that, for each p, there are plenty of pairs (c, σ) such that the "synchronous" coupling π_1 is not optimal between its marginals for the metric $\mathcal{AW}_p$.

Remark 3.6 (The choice of the quadratic variation). *In continuous time, a natural choice for the cost function on Ω is the maximum norm. When the adapted distance is defined w.r.t. this cost function, the BDG-inequalities imply that this yields an equivalent distance on the set of true martingale measures. However, when considering semimartingales, this cost is too coarse. For example, let (ω_n) be a sequence in Ω which converges to zero in maximum norm but for which the first variation tends to infinity. Then $\mathbb{P}_n := \delta_{\omega_n}$ converges to $\mathbb{P} := \delta_0$ (when adapted distance is defined only with maximum norm), however, none of our optimization problems converge (take a strategy $H \in \mathcal{H}_k$ for which $(H(X) \bullet X)_T \approx k|\omega_n|_{1\text{-}var}$ almost surely).*

The above argumentation of course does not work in discrete time, as then all norms are equivalent. However, a revealing reason to stick to "quadratic plus first variation" is that discrete time approximations work well together with quadratic variation, and not with general distances: Consider $\Omega = C([0, 1])$, for deterministic continuous $\sigma \colon [0, 1] \to \mathbb{R}_+$, let $\mathbb{P}^\sigma$ be Wiener measure on Ω with volatility σ. For each N, let $\mathbb{P}_N^\sigma$ be the law of a random walk on $\{0, 1/N, 2/N, \ldots, 1\}$ with independent increments from n/N to $(n+1)/N$ distributed according to $\mathcal{N}(0, \sigma_{n/N}^2/N)$. Then one can compute that

$$\mathcal{AW}_2(\mathbb{P}_N^\sigma, \hat{\mathbb{P}}_N^{\sigma'}) = \Big(\sum_{n=0}^{N-1} \frac{1}{N} |\sigma_{n/N} - \sigma_{n/N}'|^2\Big)^{1/2}$$

$$\to \Big(\int_0^1 |\sigma_t - \sigma_t'|^2 dt\Big)^{1/2} = \mathcal{AW}_2(\mathbb{P}^\sigma, \mathbb{P}^{\sigma'}).$$

That is, the discrete approximation preserves the correct distance. However, this situation changes drastically when $\mathbb{R}^N$ is equipped with say the ℓ^1 norm. Indeed, in this case a computation shows $\mathcal{AW}_2(\mathbb{P}_N^\sigma, \mathbb{P}_N^{\sigma'}) \to \infty$ as $N \to \infty$ whenever $\sigma \neq \sigma'$.

3.2. **Stochastic integrals and a contraction principle.** We present here the two technical results which underlie the proofs of the main theorems in the article. The first one is

Lemma 3.7. *Let $\mathbb{P}, \mathbb{Q} \in \mathcal{SM}(\Omega)$, $H \in \mathcal{H}_k$, and π be a bi-causal coupling between $\mathbb{P}$ and $\mathbb{Q}$. Then there exists a process $G \in \mathcal{H}_k$ such that $G_t(Y) = \mathbb{E}_\pi[H_t(X)|Y]$ for every t, π-almost surely. Moreover, we have $(G(Y) \bullet Y)_T = \mathbb{E}_\pi[(H(X) \bullet Y)_T|Y]$, π-almost surely.*

Proof. In discrete time, write $H = \sum_{t=1}^N H_t 1_{\{t\}}$ for Borel functions $H_t \colon \mathbb{R}^{t-1} \to [-k,k]$. Let $\pi = \pi_\eta(d\omega)\,\mathbb{P}(d\omega)$ be a disintegration and define

$$G'_t(\eta) := \int H_t(\omega)\pi_\eta(d\omega),$$

for every t and $\eta \in \Omega$. By definition of bi-causal coupling G'_t is measurable w.r.t. the $\mathbb{Q}$-completion of $\mathcal{F}_{t-1}$. It remains to pick functions G_t which are $\mathcal{F}_{t-1}$ measurable such $G_t = G'_t$ $\mathbb{Q}$-almost surely. Since $\mathbb{E}_\pi[H_t(X)|Y] = G_t(Y)$ π-almost surely, it is clear that $(G(Y) \bullet Y)_T = \mathbb{E}_\pi[(H(X) \bullet Y)_T|Y]$ π-almost surely.

In continuous time we take G to be the predictable projection of H, under the reference measure π, with respect to the π-completion of the filtration $\{\emptyset, \Omega\} \otimes \mathcal{F}^Y$. By [1, Lemma C.1] the result is π-indistinguishable from a predictable process under the $\mathbb{Q}$-completion of the filtration $\mathcal{F}^Y$. The t-by-t, π-almost sure equality $G_t(Y) = \mathbb{E}_\pi[H_t(X)|Y]$, is then a consequence of the definition of predictable projection. The π-almost sure equality $(G(Y) \bullet Y)_T = \mathbb{E}_\pi[(H \bullet Y)_T|Y]$ is established in Lemma 3.8 below, assuming that $\mathbb{E}_\mathbb{Q}[[Y]_T] < \infty$. The general case follows by localization. $\square$

Lemma 3.8. *In the continuous-time context of Lemma 3.7, assume further that $\mathbb{E}_\mathbb{Q}[[Y]_T] < \infty$. Then we have $(G(Y) \bullet Y)_T = \mathbb{E}_\pi[(H(X) \bullet Y)_T|Y]$, π-almost surely.*

Proof. The statement is true if instead of the stochastic integrals we considered the integrals w.r.t. the finite variation part of Y (either by properties of Riemann-Stieltjes integrals, or directly from the definition of predictable projection). For this reason we may now assume that Y is itself a martingale.

We first take for granted that the following result: if h is bounded and predictable in the filtration of (X, Y), and if g denotes its predictable projection in the filtration of Y under the measure π, then

$$(3.2) \qquad \mathbb{E}_\pi\Big[\int_0^T |g_t|^2 d[Y]_t\Big] \leq \mathbb{E}_\pi\Big[\int_0^T |h_t|^2 d[Y]_t\Big].$$

We know that there exist a sequence (H^n) of predictable simple processes s.t.

$$\lim_{n\to\infty} \mathbb{E}_\pi\Big[\int_0^T |H_t - H_t^n|^2 d[Y]_t\Big].$$

By Itô isometry the stochastic integrals $(H^n \bullet Y)_T$ converge in $L^2(\pi)$ to $(H \bullet Y)_T$. Denoting by G^n the predictable projection of H^n with respect to the Y-filtration, we deduce from (3.2) that

$$\lim_{n\to\infty} \mathbb{E}_\pi\Big[\int_0^T |G_t - G_t^n|^2 d[Y]_t\Big],$$

so again by Itô isometry $(G^n \bullet Y)_T$ converges in $L^2(\pi)$ to $(G \bullet Y)$. The π-almost sure equality $(G^n \bullet Y)_T = \mathbb{E}_\pi[(H^n \bullet Y)_T|Y]$ follows easily by the bi-causality of the coupling π, and by taking L^2 limits the desired conclusion is obtained.

To finish the proof we must establish (3.2). First we observe that

$$\mathbb{E}_\pi\Big[\int_0^T |g_t|^2 d[Y]_t\Big]^{1/2} = \sup_{\substack{f \text{ is } Y\text{-predictable} \\ \|f\| \le 1}} \mathbb{E}_\pi\Big[\int_0^T f_t g_t d[Y]_t\Big]$$

$$= \sup_{\substack{f \text{ is } Y\text{-predictable} \\ \|f\| \le 1}} \mathbb{E}_\pi\Big[\int_0^T f_t h_t d[Y]_t\Big],$$

as follows from predictable projection and upon taking $\|f\|^2 := \mathbb{E}_\pi[\int_0^1 |f_t|^2 d[Y]_t]$. The result is a consequence of the equality

$$\mathbb{E}_\pi\Big[\int_0^T |h_t|^2 d[Y]_t\Big]^{1/2} = \sup_{\substack{f \text{ is } (X,Y)\text{-predictable} \\ \|f\| \le 1}} \mathbb{E}_\pi\Big[\int_0^T f_t h_t d[Y]_t\Big].$$

$\square$

Our next crucial technical result is given in Theorem 3.10 below. But first we need some preparation.

Lemma 3.9. *Let $\mathbb{P}, \mathbb{Q} \in \mathcal{SM}_p(\Omega)$, let π be a bi-causal coupling between $\mathbb{P}$ and $\mathbb{Q}$, let $H \in \mathcal{H}_k$, and write $X - Y = M + A$ for the semimartingale decomposition under π. Then, for every $p \ge 1$, we have*

$$\mathbb{E}_\pi[\|X - Y\|_\infty^p] \le 2^{p-1} b_p \cdot \mathbb{E}_\pi[[M]_T^{p/2} + |A|_{1\text{-}var}^p],$$

$$\mathbb{E}_\pi[|(H(X) \bullet X)_T - (H(X) \bullet Y)_T|^p] \le 2^{p-1} b_p k^p \cdot \mathbb{E}_\pi[[M]_T^{p/2} + |A|_{1\text{-}var}^p],$$

where b_p is the upper constant in the BDG-inequality. If further $H_t \colon \Omega \to \mathbb{R}$ is $\tilde{L}$-Lipschitz continuous for every t, then we have

$$\mathbb{E}_\pi[|(H(X) \bullet X)_T - (H(Y) \bullet Y)_T|^p] \le 2^{2p-2} b_p k^p \cdot \mathbb{E}_\pi[[M]_T^{p/2} + |A|_{1\text{-}var}^p]$$
$$+ \alpha \cdot \mathbb{E}_\pi[[M]_T^p + |A|_{1\text{-}var}^{2p}]^{1/2}$$

where $\alpha = 2^{3p-2} \tilde{L}^p b_p b_{2p}^{1/2} \min\{\mathcal{AW}_{2p}(\mathbb{P}, \delta_0)^p, \mathcal{AW}_{2p}(\mathbb{Q}, \delta_0)^p\}$.

Proof. The elementary inequality $(x + y)^p \le 2^{p-1} x^p + 2^{p-1} y^p$ for $x, y \ge 0$ together with BDG inequality and the fact that $\|\cdot\|_\infty \le |\cdot|_{1\text{-}var}$ imply

$$\mathbb{E}_\pi[\|X - Y\|_\infty^p] \le 2^{p-1} \mathbb{E}_\pi[\|M\|_\infty^p] + 2^{p-1} \mathbb{E}_\pi[|A|_{1\text{-}var}^p]$$
$$\le 2^{p-1} b_p \mathbb{E}_\pi[[M]_T^{p/2} + |A|_{1\text{-}var}^p].$$

This proves the first part. The same arguments imply

$$\mathbb{E}_\pi[|(H(X) \bullet X)_T - (H(X) \bullet Y)_T|^p]$$
$$\le 2^{p-1} \mathbb{E}_\pi[|(H(X) \bullet M)_T|^p] + 2^{p-1} \mathbb{E}_\pi[|(H(X) \bullet A)_T|^p]$$
$$\le 2^{p-1} k^p b_p \mathbb{E}_\pi[[M]_T^{p/2} + |A|_{1\text{-}var}^p]$$

from which the second part follows. To prove the third claim, write

$$\mathbb{E}_\pi[|(H(X) \bullet X)_T - (H(Y) \bullet Y)_T|^p]$$
$$\le 2^{p-1} \mathbb{E}_\pi[|((H(X) - H(Y)) \bullet X)_T|^p] + 2^{p-1} \mathbb{E}_\pi[|(H(Y) \bullet X)_T - (H(Y) \bullet Y)_T|^p].$$

The second term is smaller than $2^{p-1} 2^{p-1} k^p b_p \mathbb{E}_\pi[[M]_T^{p/2} + |A|_{1\text{-}var}^p]$ by the second part. It remains to estimate $\mathbb{E}_\pi[|((H(Y) - H(Y)) \bullet X)_T|^p]$. Write $X = N + B$ for the semimartingale decomposition of X under $\mathbb{P}$. By Lemma 3.1, the semimartingale

decomposition under π is still $X = N + B$. Moreover, the BDG-inequality, the Lipschitz-continuity of H, and Hölder's inequality, imply that

$$\mathbb{E}_\pi[|((H(X) - H(Y)) \bullet X)_T|^p]$$

$$\leq 2^{p-1}\mathbb{E}_\pi[|((H(X) - H(Y)) \bullet N)_T|^p + |((H(X) - H(Y)) \bullet B)_T|^p]$$

$$\leq 2^{p-1}\mathbb{E}_\pi[\|H(X) - H(Y)\|_\infty^p (b_p[N]_T^{p/2} + |B|_{1\text{-var}}^p)]$$

$$\leq 2^{p-1}b_p\tilde{L}^p\mathbb{E}_\pi[\|X - Y\|_\infty^{2p}]^{1/2}\mathbb{E}_\pi[([N]_T^p + |B|_{1\text{-var}}^p)^2]^{1/2}.$$

It now follows from the first part that

$$\mathbb{E}_\pi[\|X - Y\|_\infty^{2p}]^{1/2} \leq (2^{2p-1}b_{2p})^{1/2}\mathbb{E}_\pi[[M]_T^p + |A|_{1\text{-var}}^{2p}]^{1/2}$$

and by Lemma 3.1 we have

$$\mathbb{E}_\pi[([N]_T^{p/2} + |B|_{1\text{-var}}^p)^2]^{1/2} \leq 2^{1/2}\mathcal{A}\mathcal{W}_{2p}(\mathbb{P}, \delta_0)^p.$$

Putting all estimates together and replacing X and Y yields the claim. $\qquad\square$

Denote by $\mathcal{P}_p(\mathbb{R})$ the set of all Borel probability measures μ on $\mathbb{R}$ such that $\int |x|^p \mu(dx) < \infty$. Moreover, let $d_p(\mu, \nu)$ be the usual p-Wasserstein distance, and let d_p^w the weak p-Wasserstein cost, that is,

$$d_p(\mu, \nu) := \inf\left\{\left(\int |x - y|^p \gamma(dx, dy)\right)^{1/p} : \gamma \text{ is a coupling of } \mu \text{ and } \nu\right\},$$

$$d_p^w(\mu, \nu) := \inf\left\{\left(\int \left|x - \int y\, \gamma^x(dy)\right|^p \mu(dx)\right)^{1/p} : \gamma \text{ is a coupling of } \mu \text{ and } \nu\right\}.$$

Here $\gamma = \mu(dx)\gamma^x(dy)$ denotes the disintegration. Note that d_p^w is not symmetric and as a consequence of Jensen's inequality, we always have $d_p^w \leq d_p$. Problems akin to $d_p^w(\mu, \nu)$ go under the name of 'weak optimal transport' and have been recently introduced by Gozlan et al. in [29], but see also [3, 4, 10, 8, 28]. We have

Theorem 3.10 (Contraction). *Let $\mathbb{P}, \mathbb{Q} \in \mathcal{SM}_p(\Omega)$, let π a bi-causal coupling between $\mathbb{P}$ and $\mathbb{Q}$, let $C\colon \Omega \to \mathbb{R}$ be Lipschitz with constant L, and let $H \in \mathcal{H}_k$. Further denote by $X - Y = M + A$ the semimartingale decomposition under π and let $G \in \mathcal{H}_k$ such that $(G(Y) \bullet Y)_T = \mathbb{E}_\pi[(H(X) \bullet Y)_T|Y]$ π-almost surely. Then*

$$(3.3) \qquad d_p^w\Big((C(Y) + (G(Y) \bullet Y)_T)(\mathbb{Q}), (C(X) + (H(X) \bullet X)_T)(\mathbb{P})\Big)$$

$$\leq 2^{(p-1)/p}b_p^{1/p}(k + L) \cdot \mathbb{E}_\pi[[M]_T^{p/2} + |A|_{1\text{-}var}^p]^{1/p}.$$

Now assume in addition that $H_t\colon \Omega \to \mathbb{R}$ is $\tilde{L}$-Lipschitz continuous for every t, then

$$d_p\Big((C(Y) + (H(Y) \bullet Y)_T)(\mathbb{Q}), (C(X) + (H(X) \bullet X)_T)(\mathbb{P})\Big)$$

$$\leq 2^{(3p-3)/p}b_p^{1/p}(k + L)\mathbb{E}_\pi[[M]_T^{p/2} + |A|_{1\text{-}var}^p]^{1/p} + \alpha^{1/p}\mathbb{E}_\pi[[M]_T^p + |A|_{1\text{-}var}^{2p}]^{1/2p},$$

where α is the constant of Lemma 3.9.

Proof. We start by proving the first claim. Let π be as stated, and define $a(X) := C(X) + (H(X) \bullet X)_T$ as well as $b(Y) := C(Y) + (G(Y) \bullet Y)_T$. Now let $\gamma := (b(Y), a(X))(\pi)$ so that γ is trivially a coupling between $b(Y)(\mathbb{Q})$ and $a(X)(\mathbb{P})$. Therefore

$$d_p^w\Big(b(Y)(\mathbb{Q}), a(X)(\mathbb{P})\Big) \leq \mathbb{E}_\pi[|b(Y) - \mathbb{E}_\pi[a(X)|b(Y)]|^p]^{1/p}.$$

By assumption it holds that

$$\mathbb{E}_\pi[(G(Y) \bullet Y)_T - (H(X) \bullet X)_T|Y] = \mathbb{E}_\pi[(H(X) \bullet Y)_T - (H(X) \bullet X)_T|Y].$$

Thus, using the tower property and Jensen's inequality, it follows that

$$\mathbb{E}_\pi[|b(Y) - \mathbb{E}_\pi[a(X)|b(Y)]|^p]^{1/p}$$

$$\leq \mathbb{E}_\pi\big[\,\big|\,\mathbb{E}_\pi[C(Y) - C(X)|Y] + \mathbb{E}_\pi[(G(Y)\bullet Y)_T - (H(X)\bullet X)_T|Y]\,\big|^p\,\big]^{1/p}$$

$$\leq \mathbb{E}_\pi[|C(Y) - C(X)|^p]^{1/p} + \mathbb{E}_\pi\big[\,|(H(X)\bullet Y)_T - (H(X)\bullet X)_T\,|^p\,\big]^{1/p}$$

The claim now follows from the first and second estimates in Lemma 3.9.

In the second case where H is additionally Lipschitz, let $d(X) := C(X) + (H(X)\bullet X)_T$ as well as $e(Y) := C(Y) + (H(Y)\bullet Y)_T$ and $\gamma := (e(Y), d(Y))(\pi)$. Then, similarly as before,

$$d_p\Big(e(Y)(\mathbb{Q})\,, d(X)(\mathbb{P})\Big) \leq \mathbb{E}_\pi[|e(Y) - d(Y)|^p]^{1/p}$$

$$\leq \mathbb{E}_\pi[|C(Y) - C(X)|^p]^{1/p} + \mathbb{E}_\pi[|(H(Y)\bullet Y)_T - (H(X)\bullet X)_T|^p]^{1/p}$$

and the claim follows from the first and third estimates of Lemma 3.9. $\qquad\square$

Remark 3.11. *An evident question is whether an estimate for the usual Wasserstein distance holds true without the (Lipschitz-) continuity assumption on H. Namely if (3.3) holds for d_p instead of d_p^w. The following example shows that this is not true. In a two-period discrete time model ($T = 2$), let*

$$\mathbb{P} := \delta_0 \otimes ((\delta_1 + \delta_{-1})/2) \quad and \quad \mathbb{P}_\varepsilon := ((\delta_\varepsilon + \delta_{-\varepsilon})/2) \otimes ((\delta_1 + \delta_{-1})/2)$$

so that $\mathcal{AW}_p(\mathbb{P}_\varepsilon, \mathbb{P}) \to 0$ as $\varepsilon \to 0$ for every p. Then, set $H_1 := 0$ and $H_2 := 1_{(0,\infty)} - 1_{(-\infty,0)}$. For the projection under any bi-causal coupling between $\mathbb{P}_\varepsilon$ and $\mathbb{P}$ of H onto Y one computes $G_1 = 0$ and $G_2 = 0$. In particular $(G(Y)\bullet Y)_T = 0$ $\mathbb{P}$-almost surely. However, for every $\varepsilon > 0$ one has $\mathbb{P}_\varepsilon((H(X)\bullet X)_T \geq 1 - \varepsilon) \geq 1/4$ which implies that the respective laws cannot converge.

Remark 3.12. *By b_p we denote the smallest real number such that*

$$(3.4) \qquad\qquad \mathbb{E}[\|M\|_\infty^p] \leq b_p \mathbb{E}[[M]^{p/2}]$$

for every martingale M. For $p \geq 2$ it was established by Burkholder [20] that $b_p = p$ but the value of b_p is unknown for $p \in [1, 2)$ according to [47], [48, page 427]. By [16], $b_1 \leq 6$. (The optimal constant in the reverse inequality is known for the trivial case $p = 2$ and for $p = 1$. In the latter instance one obtains $\sqrt{3}$ [21] and $1.2727\ldots$ [56] for continuous martingales, resp.)

4. Proofs of the results stated in the introduction and extensions

Thanks to the previous work, the strategy for the proofs boils down into two parts. In a first step, one forgets about the space Ω and only focuses on continuity of the problem at hand with respect to d_p or d_p^w when image measures on $\mathbb{R}$ are plugged in: e.g. in utility maximization this means to study continuity of $\mu \mapsto \int U(x)\,\mu(dx)$. In a second step, one uses the obtained continuity and the contraction theorem in the previous section.

4.1. **Proof of Theorem 1.5.** We will need the elementary estimate

Lemma 4.1. *Let $\mu, \nu \in \mathcal{P}_1(\mathbb{R})$ and let $f\colon \mathbb{R} \to \mathbb{R}$ be convex and Lipschitz. Then*

$$\int f(x)\mu(dx) - \int f(y)\,\nu(dy) \leq L\,d_1^w(\mu, \nu),$$

where L is Lipschitz constant of f.

Proof. Let γ be a coupling of μ and ν. Applying Jensen's inequality we obtain

$$\int f(x)\,\mu(dx) - \int f(y)\,\nu(dy) = \int f(x) - f(y)\,\gamma(dx, dy)$$

$$= \int \left(f(x) - \int f(y)\,\gamma^x(dy) \right) \mu(dx) \leq \int \left(f(x) - f\left(\int y\,\gamma^x(dy) \right) \right) \mu(dx)$$

$$\leq L \int \left| x - \int y\,\gamma^x(dy) \right| \mu(dx).$$

As γ was arbitrary, this implies the claim. $\square$

In fact there is equality in the previous lemma, as shown in [29, Proposition 3.2].

We now turn to the proof of Theorem 1.5. For $n > 0$ let π be a bi-causal coupling which attains the infimum in the definition of $\mathcal{AW}_1(\mathbb{P}, \mathbb{Q})$ modulo a $1/n$-margin. By Lemma 3.7 there is $G^n \in \mathcal{H}_k$ such that $(G^n(Y) \bullet Y)_T = \mathbb{E}_\pi[(H(X) \bullet Y)_T | Y]$ π-almost surely. Define

$$\mu^n := (C(Y) + (G^n(Y) \bullet Y)_T)(\mathbb{Q}) \quad \text{and} \quad \nu := (C(X) + (H(X) \bullet X)_T)(\mathbb{P}).$$

(Note that $\mu^n, \nu \in \mathcal{P}_1(\mathbb{R})$ as $\mathbb{P}, \mathbb{Q} \in \mathcal{SM}_1(\Omega)$.) By Lemma 4.1 we have

$$\mathbb{E}_{\mathbb{Q}}\left[(C(Y) - m - (G^n(Y) \bullet Y)_T)^+ \right] - \mathbb{E}_{\mathbb{P}}\left[(C(X) - m - (H(X) \bullet X)_T)^+ \right] \leq d_1^w(\mu^n, \nu).$$

From Theorem 3.10 we obtain

$$(4.1) \qquad \begin{aligned} \mathbb{E}_{\mathbb{Q}}[(C(Y) - m - (G^n(Y) \bullet Y)_T))^+] &\leq \mathbb{E}_{\mathbb{P}}[(C(X) - m - (H(X) \bullet X)_T)^+] \\ &\quad + b_1(k + L)\,(\mathcal{AW}_1(\mathbb{P}, \mathbb{Q}) + 1/n). \end{aligned}$$

Assume first that $\mathbb{E}_{\mathbb{Q}}[[Y]_T] < \infty$ and denote by A the finite variation process associated to Y. Then, as (G^n) is uniformly bounded by k, there exists a predictable G and a sequence of forward-convex combinations of (G^n) which converge in $L^2(d\mathbb{Q} \otimes d([Y] + A))$ to G. This, (4.1), and the convexity of $(\cdot)^+$ lead to the desired conclusion. The general case follows by a simple but notationally heavy localization argument.

The proof in case that $G = H$ and H is Lipschitz follows analogously from the second part of Theorem 3.10.

4.2. **Proof of Theorem 1.6.** In a first step notice that for all $\mathbb{P}, \mathbb{P}'$ and random variables Z, Z', it follows as in Lemma 4.1 that

$$\text{AVaR}_\alpha^{\mathbb{P}}(Z) - \text{AVaR}_\alpha^{\mathbb{P}'}(Z') \leq d_1^w(Z(\mathbb{P}), Z'(\mathbb{P}'))/\alpha.$$

Indeed, if γ is a coupling from $\mu := Z(\mathbb{P})$ to $\nu := Z'(\mathbb{P}')$ then

$$\text{AVaR}_\alpha^{\mathbb{P}}(Z) - \text{AVaR}_\alpha^{\mathbb{P}'}(Z')$$

$$= \inf_m \int \frac{1}{\alpha}(x - m)^+ - m\,\mu(dx) - \inf_m \frac{1}{\alpha} \int \int (y - m)^+ \gamma^x(dy) - m\,\mu(dy)$$

$$\leq \sup_m \frac{1}{\alpha} \int (x - m)^+ - (y - m)^+ \,\gamma(dx, dy)$$

$$\leq \sup_m \frac{1}{\alpha} \int (x - m)^+ - \left(\int y\,\gamma^x(dy) - m \right)^+ \mu(dx) \leq \frac{1}{\alpha} \int \left| x - \int y\,\gamma^x(dy) \right| \mu(dx),$$

so minimizing over γ yields the claim.

The rest of the proof now follows the line of argumentation as in the proof for Theorem 1.5. Fix $\mathbb{P}, \mathbb{Q} \in \mathcal{SM}_1(\Omega)$. Assume only for notational simplicity that there exists a bi-causal coupling π which attains the infimum in the definition of $\mathcal{AW}_1(\mathbb{P}, \mathbb{Q})$, and that there exist $H^* \in \mathcal{H}_k$ such that

$$\text{AVaR}_\alpha^{\mathbb{P}}(C(X) - (H^*(X) \bullet X)_T) = \inf_{H \in \mathcal{H}_k} \text{AVaR}_\alpha^{\mathbb{P}}(C(X) - (H(X) \bullet X)_T).$$

By Lemma 3.7 there is $G^* \in \mathcal{H}_k$ such that $(G^*(Y) \bullet Y)_T = \mathbb{E}_\pi[(H^*(X) \bullet Y)_T|Y]$ π-almost surely. Therefore

$$\inf_{G \in \mathcal{H}_k} \text{AVaR}_\alpha^\mathbb{Q}(C(Y) - (G(Y) \bullet Y)_T) - \inf_{H \in \mathcal{H}_k} \text{AVaR}_\alpha^\mathbb{P}(C(X) - (H(X) \bullet X)_T)$$

$$\leq \text{AVaR}_\alpha^\mathbb{Q}(C(Y) - (G^*(Y) \bullet Y)_T) - \text{AVaR}_\alpha^\mathbb{P}(C(X) - (H^*(X) \bullet X)_T)$$

$$\leq \frac{1}{\alpha} d_1^w\Big((C(Y) - (G^*(Y) \bullet Y)_T)(\mathbb{Q}), (C(X) - (H^*(X) \bullet X)_T)(\mathbb{P}) \Big)$$

$$\leq \frac{b_1(k+L)}{\alpha} \mathcal{AW}_1(\mathbb{P}, \mathbb{Q}),$$

where the last inequality is due to Theorem 3.10. Interchanging the role of $\mathbb{P}$ and $\mathbb{Q}$ yields the desired conclusion. The proof for the second estimate follows analogously.

4.3. **Proof of Example 1.7.** First note that $\text{AVaR}_\alpha^\mathbb{P}(Z) \geq \mathbb{E}_\mathbb{P}[Z]$ for every integrable random variable Z. Indeed, this follows from integrating the pointwise inequality $x = x+m-m \leq (x+m)^+/\alpha - m$. Therefore, as the Brownian stochastic integral has expectation zero, we conclude that $\inf_{H \in \mathcal{H}_k} \text{AVaR}_\alpha^\mathbb{P}(C(X) - (H(X) \bullet X)_T) \geq \mathbb{E}_\mathbb{P}[C(X)]$. On the other hand, define

$$f(t, x) := \int c(x+y) \, N(0, \sigma^2(T-t))(dy) \quad \text{for } (t, x) \in [0, T] \times \mathbb{R},$$

where $N(0, \sigma^2(T-t))$ stands for the normal distribution with mean 0 and variance $\sigma^2(T-t)$. Then $C(X) = f(T, X_T)$ and $\mathbb{E}_\mathbb{P}[f(t, X_t)|\mathcal{F}_s] = f(s, X_s)$ for every $0 \leq s \leq t \leq T$. Thus, by Itô's formula and fact that the martingale property implies that the finite variation part vanishes, one has $f(t, X_t) = f(0, 0) + (H^*(X) \cdot X)_T$ for the predictable trading strategy $H_t^* := \partial_x f(t, X_t)$. As further $|H_t^*| \leq 1$ for every t and $f(0, 0) = \sigma/\sqrt{2\pi}$, one has

$$\inf_{H \in \mathcal{H}_1} \text{AVaR}_\alpha^\mathbb{P}(C(X) - (H(X) \bullet X)_T) \leq \text{AVaR}_\alpha^\mathbb{P}(C(X) - (H^*(X) \cdot X)_T) = \frac{\sigma}{\sqrt{2\pi}}.$$

The proof now follows from the explicit formula for the adapted Wasserstein distance derived in Example 3.4 and the fact that $\mathbb{E}_\mathbb{P}[C(X)] = \sigma/\sqrt{2\pi}$.

4.4. **Proof of Theorem 1.8.** Recall that $U'(x) \leq c(1 + |x|^{p-1})$ for all $x \in \mathbb{R}$ and some constant c. Let $\mathbb{P}, \mathbb{Q} \in \mathcal{SM}_p(\Omega)$ be arbitrary and assume only for notational simplicity that there is $H^* \in \mathcal{H}_k$ such that

$$\mathbb{E}_\mathbb{P}[U(C(X) + (H^*(X) \bullet X)_T)] = \sup_{H \in \mathcal{H}_k} \mathbb{E}_\mathbb{P}[U(C(X) + (H(X) \bullet X)_T)]$$

and that there is a bi-causal coupling π coupling between $\mathbb{P}$ and $\mathbb{Q}$ which is optimal for $\mathcal{AW}_p(\mathbb{P}, \mathbb{Q})$. By Lemma 3.7 there is $G^* \in \mathcal{H}_k$ such that $(G^*(Y) \bullet Y)_T = \mathbb{E}_\pi[(H^*(X) \bullet Y)_T|Y]$ π-almost surely. Let

$$\mu := (C(Y) + (G^*(Y) \bullet Y)_T)(\mathbb{Q}) \text{ and } \nu := (C(X) + (H^*(X) \bullet X)_T)(\mathbb{P}),$$

and let γ be an (almost) optimal coupling for $d_p^w(\mu, \nu)$. As U is concave and increasing, we have $U(y) - U(x) \leq U'(\min\{x, y\})|x - y|$. Using Jensen's inequality for the concave function U we have

$$\sup_{H \in \mathcal{H}_k} \mathbb{E}_\mathbb{P}[U(C(X) + (H(X) \bullet X)_T)] - \sup_{G \in \mathcal{H}_k} \mathbb{E}_\mathbb{Q}[U(C(Y) + (G(Y) \bullet Y)_T)]$$

$$\leq \mathbb{E}_\mathbb{P}[U(C(X) + (H^*(X) \bullet X)_T)] - \mathbb{E}_\mathbb{Q}[U(C(Y) + (G^*(Y) \bullet Y)_T)]$$

$$= \int U(y) - U(x) \, \gamma(dx, dy) \leq \int U\Big(\int y \, \gamma^x(dy) \Big) - U(x) \, \mu(dx)$$

$$\leq \Big(\int \Big|U'\Big(\min\Big\{x, \int y \, \gamma^x(dy)\Big\} \Big)\Big|^q \mu(dx) \Big)^{1/q} d_p^w(\mu, \nu),$$

where we used Hölder's inequality in the last line and q denotes the conjugate Hölder exponent of p (that is, $1/p + 1/q = 1$). As $q(p-1) = p$, the growth assumption on U' implies that $|U'(\min\{x,y\})|^q \le c(1 + |x|^p + |y|^p)$ for some (new) constant c. Then, by Lemma 3.9, we have

$$\int \left| U'\left(\min\left\{x, \int y\,\gamma^x(dy)\right\}\right)\right|^q \mu(dx)$$

$$\le c\left(1 + \int |x|^p\,\mu(dx) + \int \left|\int y\,\gamma^x(dy)\right|^p \mu(dy)\right)$$

$$\le c\left(1 + \int |x|^p\,\mu(dx) + \int |y|^p\,\nu(dy)\right)$$

$$\le \tilde{c}\left(1 + \mathcal{AW}_p(\mathbb{Q},\delta_0)^p + \mathcal{AW}_p(\mathbb{P},\delta_0)^p\right) =: e$$

for $e := \tilde{c}(1 + R^p + R^p)$. Exchanging the roles of $\mathbb{P}$ and $\mathbb{Q}$ and using Theorem 3.10 completes the proof.

4.5. **Two generalizations.** The following two results can be proved using almost the same arguments as used in the proofs of Theorem 1.6 and Theorem 1.8. In particular the proofs boil down to establishing convergence for image measures with respect to d_p and give no new insight on adapted Wasserstein distances, so we shall skip them.

Proposition 4.2. *Let $\ell\colon \mathbb{R} \to \mathbb{R}_+$ be a convex and strictly increasing function and let $\delta > 0$. Assume that $p \ge 1$ is such that $\ell'(x) \le c(1 + |x|^{p-1})$ for some constant c. Then, for every Lipschitz continuous function $C\colon \Omega \to \mathbb{R}$, the function*

$$\mathbb{P} \mapsto \inf\left\{m \in \mathbb{R} : \text{there is } H \in \mathcal{H}_k \text{ such that } \mathbb{E}_\mathbb{P}[\ell(C(X) - (H(X) \bullet X)_T - m)] \le \delta\right\}$$

is continuous on $(\mathcal{SM}_p(\Omega), \mathcal{AW}_p)$.

Let ρ be a law-invariant risk measure which we directly view as a functional from $\mathcal{P}_p(\mathbb{R})$ to the reals. For $\mathbb{P} \in \mathcal{SM}_p(\Omega)$ and a random variable $Z\colon \Omega \to \mathbb{R}$ (such that $Z(\mathbb{P}) \in \mathcal{P}_p(\mathbb{R})$) we write $\rho^\mathbb{P}(Z) = \rho(Z(\mathbb{P}))$. A typical example of a law invariant risk measure which satisfies $\rho(\mu) - \rho(\nu) \le L d^w(\mu,\nu)$ for some constant L depending on the p-the moment of μ and ν is the optimized certainty equivalent, introduced to the mathematical finance community in [17]. For a convex, increasing function $\ell\colon \mathbb{R} \to \mathbb{R}$ which is bounded from below and satisfies $\ell(x)/x \to \infty$ as $x \to \infty$, the optimized certainty equivalent is defined via

$$\rho^\mathbb{P}(Z) := \inf_{m \in \mathbb{R}} \left(\mathbb{E}_\mathbb{P}[\ell(Z-m)] + m\right) = \inf_{m \in \mathbb{R}} \left(\int \ell(x-m)\,(Z(\mathbb{P}))(dx) + m\right).$$

If $\ell'(x) \le c(1 + |x|^{p-1})$, then it follows that the infimum over m can be taken in some compact set depending on the p-th moment.s Due to cash additivity of ρ, the following proposition has the same interpretation as Theorem 1.6.

Proposition 4.3. *Assume that $\rho\colon \mathcal{P}_p(\mathbb{R}) \to \mathbb{R}$ satisfies $\rho(\mu) - \rho(\nu) \le L d^w(\mu,\nu)$ for some constant L depending on the p-the moment of μ and ν. Then, for every Lipschitz function $C\colon \Omega \to \mathbb{R}$, the mapping*

$$\mathbb{P} \mapsto \inf_{H \in \mathcal{H}_k} \rho^\mathbb{P}(C(X) - (H(X) \bullet X)_T)$$

is locally Lipschitz continuous on $(\mathcal{SM}_p(\Omega), \mathcal{AW}_p)$.

Finally, let us point out that (though not a convex risk measure) the Value-at-Risk (VaR) would be another natural candidate to study continuity. However, as VaR is not continuous w.r.t. weak convergence, already in a one period model continuity of $\mathbb{P} \mapsto \inf\{m \in \mathbb{R} : \text{there is } H \in \mathcal{H}_k \text{ with } \mathrm{VaR}^\mathbb{P}(C(X) - m - (H(X) \bullet X)_T) \le 0\}$ does not hold.

5. Final remarks

Remark 5.1 (Usual Wasserstein does not work I)**.** *We note that convergence in the usual Wasserstein distance is not sufficient to obtain continuity in any of the problems we study in this paper. Consider a two period market with*

$$\mathbb{P}_n = \frac{1}{4}\Big(\delta_{(1/n,1)} + \delta_{(1/n,0)} + \delta_{(-1/n,0)} + \delta_{(-1/n,-1)} \Big),$$

$$\mathbb{P} = \frac{1}{4}\Big(\delta_{(0,1)} + 2\delta_{(0,0)} + \delta_{(0,-1)} \Big).$$

Then $\mathbb{P}$ and each $\mathbb{P}_n$ satisfy the classical no-arbitrage condition, unlike the situation described in Figure 1. While $\mathbb{P}_n$ converges to $\mathbb{P}$ in usual Wasserstein distance, one can verify that convergence in nested distance does not hold. For example in utility maximization of the trivial claim $C = 0$, we have $\sup_{H \in \mathcal{H}_k} \mathbb{E}_{\mathbb{P}}[U(C(X) + (H(X) \bullet X)_T)] = U(0)$ by Jensen's inequality (as X is a martingale under $\mathbb{P}$). For $\mathbb{P}_n$ taking the strategy H^* consisting of $H_0^* = 0$ and $H_1^*(x) = k\operatorname{sign}(x)$, one gets

$$\sup_{H \in \mathcal{H}_k} \mathbb{E}_{\mathbb{P}_n}[U(C(X) + (H(X) \bullet X)_T)] \geq \mathbb{E}_{\mathbb{P}_n}[U(C(X) + (H^*(X) \bullet X)_T)] \to U(k),$$

showing the lack of continuity.

Remark 5.2 (Usual Wasserstein does not work II)**.** *As explained in the introduction, the objective in Theorem 1.5 can be seen as a relaxed version of the superhedging problem. The reason to consider this relaxation is not a technical simplification but necessary to to obtain continuity without further assumptions. Indeed, the problem of superhedging*

$$\inf\Big\{m \in \mathbb{R} : \text{there is } H \in \mathcal{H}_k \text{ such that } m + (H \bullet X)_T \geq C(X), \mathbb{P}\text{-almost surely}\Big\}$$

is not continuous in $\mathbb{P}$ w.r.t. adapted distance for any $k \in [0, \infty]$. In fact, this already happens in one period, where adapted and the usual Wasserstein distances coincide. Consider and sequence of measures $\mathbb{P}_n$ with full support which converge weakly to a measure $\mathbb{P}$. Then the superhedging price w.r.t. $\mathbb{P}_n$ equals the concave envelope of C, while the superhedging price w.r.t. $\mathbb{P}$ equals the concave envelope of C restricted to the support of $\mathbb{P}$. For a recent paper on this problem in one period, see the work of Obłój and Wiesel [46].

Remark 5.3 (Uniformly bounded strategies are necessary)**.** *Similar as in Remark 5.2 the restriction to trading strategies in $\mathcal{H}_k$ (i.e. uniformly bounded strategies) is also no technical simplification. For example, in a one-period framework, the measures $\mathbb{P}^\varepsilon := (1 - \varepsilon)\delta_{(0,\varepsilon)} + \varepsilon\delta_{(0,-\varepsilon)}$ converges to $\mathbb{P} := \delta_{(0,0)}$ in every (adapted) Wasserstein distance. However, we have for small $\varepsilon > 0$*

$$\inf_{H \in \mathcal{H}_\infty} \operatorname{AVaR}_\alpha^{\mathbb{P}^\varepsilon}((H \bullet X)_T) = -\infty \quad \text{while} \quad \inf_{H \in \mathcal{H}_\infty} \operatorname{AVaR}_\alpha^{\mathbb{P}}((H \bullet X)_T) = 0,$$

where $H \in \mathcal{H}_\infty := \bigcup_{k \in \mathbb{N}} \mathcal{H}_k$ is the set of all bounded trading strategies.

References

[1] B. Acciaio, J. Backhoff-Veraguas, and A. Zalashko. Causal optimal transport and its links to enlargement of filtrations and continuous-time stochastic optimization. *ArXiv e-prints*, 2016.
[2] D. Aldous. Weak convergence and the general theory of stochastic processes. 1981.

[3] A. Alfonsi, J. Corbetta, and B. Jourdain. Sampling of probability measures in the convex order and approximation of Martingale Optimal Transport problems. *ArXiv e-prints*, Sept. 2017.

[4] J.-J. Alibert, G. Bouchitte, and T. Champion. A new class of cost for optimal transport planning. *hal-preprint*, 2018.

[5] M. Avellaneda, A. Levy, and A. Paràs. Pricing and hedging derivative securities in markets with uncertain volatilities. *Appl. Math. Finance*, 2(2):73–88, 1995.

[6] J. Backhoff-Veraguas, D. Bartl, M. Beiglboeck, and M. Eder. All topologies are equal. *manuscript available at request*, 2019.

[7] J. Backhoff-Veraguas, M. Beiglböck, M. Eder, and A. Pichler. Fundamental properties of process distances. *ArXiv e-prints*, 2017.

[8] J. Backhoff-Veraguas, M. Beiglböck, M. Huesmann, and S. Källblad. Martingale Benamou–Brenier: a probabilistic perspective. *ArXiv e-prints*, Aug. 2017.

[9] J. Backhoff-Veraguas, M. Beiglböck, Y. Lin, and A. Zalashko. Causal transport in discrete time and applications. *SIAM Journal on Optimization*, 27(4):2528–2562, 2017.

[10] J. Backhoff-Veraguas, M. Beiglböck, and G. Pammer. Existence, duality, and cyclical monotonicity for weak transport costs. *ArXiv e-prints*, 2018.

[11] D. Bartl, S. Drapeau, J. Obłój, and J. Wiesel. *2019*, private communication.

[12] D. Becherer and K. Kentia. Good deal hedging and valuation under combined uncertainty about drift and volatility. *Probab. Uncertain. Quant. Risk*, 2:Paper No. 13, 40, 2017.

[13] M. Beiglböck, A. Cox, and M. Huesmann. The geometry of multi-marginal Skorokhod Embedding. *preprint available at request*, 2016.

[14] M. Beiglböck, P. Henry-Labordère, and F. Penkner. Model-independent bounds for option prices: A mass transport approach. *Finance Stoch.*, 17(3):477–501, 2013.

[15] M. Beiglböck and N. Juillet. On a problem of optimal transport under marginal martingale constraints. *Ann. Probab.*, 44(1):42–106, 2016.

[16] M. Beiglböck and P. Siorpaes. Pathwise versions of the Burkholder-Davis-Gundy inequality. *Bernoulli*, 21(1):360–373, 2015.

[17] A. Ben-Tal and M. Teboulle. An old-new concept of convex risk measures: The optimized certainty equivalent. *Mathematical Finance*, 17(3):449–476, 2007.

[18] J. Bion-Nadal and D. Talay. On a wasserstein-type distance between solutions to stochastic differential equations. *To appear in Annals of Applied Probability*, 2018.

[19] B. Bouchard and M. Nutz. Arbitrage and duality in nondominated discrete-time models. *The Annals of Applied Probability*, 25(2):823–859, 2015.

[20] D. L. Burkholder. Explorations in martingale theory and its applications. In *École d'Été de Probabilités de Saint-Flour XIX—1989*, volume 1464 of *Lecture Notes in Math.*, pages 1–66. Springer, Berlin, 1991.

[21] D. L. Burkholder. The best constant in the Davis inequality for the expectation of the martingale square function. *Trans. Amer. Math. Soc.*, 354(1):91–105 (electronic), 2002.

[22] L. Campi, I. Laachir, and C. Martini. Change of numeraire in the two-marginals martingale transport problem. *Finance Stoch.*, 21(2):471–486, June 2017.

[23] R. Cont. Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical finance*, 16(3):519–547, 2006.

[24] Y. Dolinsky and H. M. Soner. Martingale optimal transport and robust hedging in continuous time. *Probab. Theory Relat. Fields*, 160(1-2):391–427, 2014.

[25] N. El Karoui, M. Jeanblanc, and S. Shreve. Robustness of the Black and Scholes formula. *Math. Finance*, 8(2):93–126, 1998.

[26] A. Galichon, P. Henry-Labordère, and N. Touzi. A stochastic control approach to no-arbitrage bounds given marginals, with an application to lookback options. *Ann. Appl. Probab.*, 24(1):312–336, 2014.

[27] M. Glanzer, G. Pflug, and A. Pichler. Incorporating statistical model error into the calculation of acceptability prices of contingent claims. *ArXiv e-prints*, 2017.

[28] N. Gozlan, C. Roberto, P.-M. Samson, Y. Shu, and P. Tetali. Characterization of a class of weak transport-entropy inequalities on the line. *Ann. Inst. Henri Poincaré Probab. Stat.*, 54(3):1667–1693, 2018.

[29] N. Gozlan, C. Roberto, P.-M. Samson, and P. Tetali. Kantorovich duality for general transport costs and applications. *J. Funct. Anal.*, 273(11):3327–3405, 2017.

[30] M. F. Hellwig. Sequential decisions under uncertainty and the maximum theorem. *J. Math. Econom.*, 25(4):443–464, 1996.

[31] S. Herrmann and J. Muhle-Karbe. Model uncertainty, recalibration, and the emergence of delta–vega hedging. *Finance and Stochastics*, 21(4):873–930, Oct 2017.

[32] S. Herrmann, J. Muhle-Karbe, and F. T. Seifried. Hedging with small uncertainty aversion. *Finance and Stochastics*, 21(1):1–64, Jan 2017.

[33] D. Hobson. Robust hedging of the lookback option. *Finance and Stochastics*, 2:329–347, 1998.
[34] D. Hobson. The Skorokhod embedding problem and model-independent bounds for option prices. In *Paris-Princeton Lectures on Mathematical Finance 2010*, volume 2003 of *Lecture Notes in Math.*, pages 267–318. Springer, Berlin, 2011.
[35] D. Hobson and A. Neuberger. Robust bounds for forward start options. *Math. Finance*, 22(1):31–56, 2012.
[36] D. G. Hobson. Volatility misspecification, option pricing and superreplication via coupling. *Ann. Appl. Probab.*, 8(1):193–205, 1998.
[37] I. Karatzas and S. Shreve. *Brownian motion and stochastic calculus*, volume 113. Springer Science & Business Media, 2012.
[38] C. Kardaras and G. Žitković. Stability of the utility maximization problem with random endowment in incomplete markets. *Math. Finance*, 21(2):313–333, 2011.
[39] D. Lacker. Dense sets of joint distributions appearing in filtration enlargements, stochastic control, and causal optimal transport. *ArXiv e-prints*, 2018.
[40] K. Larsen. Continuity of utility-maximization with respect to preferences. *Math. Finance*, 19(2):237–250, 2009.
[41] K. Larsen and G. Žitković. Stability of utility-maximization in incomplete markets. *Stochastic Process. Appl.*, 117(11):1642–1662, 2007.
[42] R. Lassalle. Causal transference plans and their Monge-Kantorovich problems. *ArXiv e-prints*, 2013.
[43] T. J. Lyons. Uncertain volatility and the risk-free synthesis of derivatives. *Applied Mathematical Finance*, 2(2):117–133, 1995.
[44] M. Mocha and N. Westray. The stability of the constrained utility maximization problem: a BSDE approach. *SIAM J. Financial Math.*, 4(1):117–150, 2013.
[45] J. Obłój. The Skorokhod embedding problem and its offspring. *Probab. Surv.*, 1:321–390, 2004.
[46] J. Obłój and J. Wiesel. Statistical estimation of superhedging prices. *ArXiv e-prints*, 2018.
[47] A. Osekowski. Sharp maximal inequalities for the martingale square bracket. *Stochastics: An International Journal of Probability and Stochastics Processes*, 82(06):589–605, 2010.
[48] A. Osekowski. *Sharp martingale and semimartingale inequalities*, volume 72 of *Instytut Matematyczny Polskiej Akademii Nauk. Monografie Matematyczne (New Series) [Mathematics Institute of the Polish Academy of Sciences. Mathematical Monographs (New Series)]*. Birkhäuser/Springer Basel AG, Basel, 2012.
[49] G. C. Pflug. Version-independence and nested distributions in multistage stochastic optimization. *SIAM Journal on Optimization*, 20(3):1406–1420, 2009.
[50] G. C. Pflug and A. Pichler. A distance for multistage stochastic optimization models. *SIAM J. Optim.*, 22(1):1–23, 2012.
[51] G. C. Pflug and A. Pichler. *Multistage stochastic optimization*. Springer Series in Operations Research and Financial Engineering. Springer, Cham, 2014.
[52] G. C. Pflug and A. Pichler. From empirical observations to tree models for stochastic optimization: convergence properties. *SIAM J. Optim.*, 26(3):1715–1740, 2016.
[53] A. Pratelli. On the equality between Monge's infimum and Kantorovich's minimum in optimal mass transportation. *Ann. Inst. H. Poincaré Probab. Statist.*, 43(1):1–13, 2007.
[54] J.-L. Prigent. Weak convergence of financial markets. In *Weak Convergence of Financial Markets*, pages 129–265. Springer, 2003.
[55] L. Rüschendorf. The Wasserstein distance and approximation theorems. *Z. Wahrsch. Verw. Gebiete*, 70(1):117–129, 1985.
[56] W. Schachermayer and F. Stebegg. The Sharp Constant for the Burkholder-Davis-Gundy Inequality and Non-Smooth Pasting. *Bernoulli, to appear*, July 2017.
[57] K. Weston. Stability of utility maximization in nonequivalent markets. *Finance and Stochastics*, 20(2):511–541, 2016.
[58] T. Yamada and S. Watanabe. On the uniqueness of solutions of stochastic differential equations. *Journal of Mathematics of Kyoto University*, 11(1):155–167, 1971.