

A strong duality principle for equivalence couplings and total variation*

Adam Quinn Jaffe[†]

Abstract

We introduce and study a notion of duality for two classes of optimization problems commonly occurring in probability theory. That is, on an abstract measurable space (Ω, \mathcal{F}) , we consider pairs (E, \mathcal{G}) where E is an equivalence relation on Ω and \mathcal{G} is a sub- σ -algebra of \mathcal{F} ; we say that (E, \mathcal{G}) satisfies “strong duality” if E is $(\mathcal{F} \otimes \mathcal{F})$ -measurable and if for all probability measures \mathbb{P}, \mathbb{P}' on (Ω, \mathcal{F}) we have

$$\max_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)| = \min_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(E)),$$

where $\Pi(\mathbb{P}, \mathbb{P}')$ denotes the space of couplings of \mathbb{P} and \mathbb{P}' , and where “max” and “min” assert that the supremum and infimum are in fact achieved. The results herein give wide sufficient conditions for strong duality to hold, thereby extending a form of Kantorovich duality to a class of cost functions which are irregular from the point of view of topology but regular from the point of view of descriptive set theory. The given conditions recover or strengthen classical results, and they have novel consequences in stochastic calculus, point process theory, and random sequence simulation.

Keywords: coupling; equivalence relation; total variation; duality; optimal transport; Kantorovich duality; Borel equivalence relation; smoothness; hypersmoothness.

MSC2020 subject classifications: 60A10; 03E15; 28A35; 90C46.

Submitted to EJP on December 6, 2022, final version accepted on September 7, 2023.

Supersedes arXiv:2207.14239.

1 Introduction

The objects of interest in this paper are two optimization problems commonly occurring in probability theory. To state them, we consider two probability measures \mathbb{P} and \mathbb{P}' on an abstract measurable space (Ω, \mathcal{F}) . For an equivalence relation E on Ω (which, when viewed as a subset of $\Omega \times \Omega$, is $(\mathcal{F} \otimes \mathcal{F})$ -measurable) we often aim to solve

$$\begin{aligned} &\text{minimize} && 1 - \tilde{\mathbb{P}}(E) \\ &\text{over} && \text{all couplings } \tilde{\mathbb{P}} \text{ of } \mathbb{P}, \mathbb{P}', \end{aligned} \tag{EC}$$

*This material is based upon work for which the author was supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE 1752814.

[†]University of California, Berkeley, United States of America. E-mail: aqjaffe@berkeley.edu

which we refer to as the *equivalence coupling problem for E* or simply the *E-coupling problem*. Alternatively, for a sub- σ -algebra \mathcal{G} of \mathcal{F} , we might aim to solve

$$\begin{aligned} & \text{maximize} && |\mathbb{P}(A) - \mathbb{P}'(A)| \\ & \text{over} && A \in \mathcal{G}, \end{aligned} \tag{TV}$$

which we refer to as *total variation problem for G* or simply the *G-total variation problem*. While one is typically interested in one of (EC) or (TV) for particular applications, our goal in this paper is to show that, in great generality, they are in fact equivalent.

1.1 Review of prior work

There already exist a few particular instances of strong duality which motivate the potential success of such a general theory; the most classical example will certainly be familiar to the reader. It states that, if Ω is Polish space with $\mathcal{B}(\Omega)$ is its Borel σ -algebra, and if $\Delta = \{(\omega, \omega) \in \Omega \times \Omega : \omega \in \Omega\}$ denotes the diagonal in $\Omega \times \Omega$, we have

$$\max_{A \in \mathcal{B}(\Omega)} |\mathbb{P}(A) - \mathbb{P}'(A)| = \min_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(\Delta)) \tag{1.1}$$

for all Borel probability measures \mathbb{P}, \mathbb{P}' on Ω , where $\Pi(\mathbb{P}, \mathbb{P}')$ denotes the space of all couplings of \mathbb{P} and \mathbb{P}' . This result has certainly been known for a long time, at least for countable sets Ω , so its exact source is difficult to track down [50, Chapter I.7].

A second known example concerns the space of binary sequences $\Omega := \{0, 1\}^{\mathbb{N}}$ with the Borel σ -algebra of its product topology, and E_0 the equivalence relation of eventual equality. (Actually, $\{0, 1\}$ can be replaced with any Polish space.) A coupling of two Borel probability measures \mathbb{P}, \mathbb{P}' on Ω is called *successful* if E_0 has probability one, that is, if the two random sequences are eventually equal almost surely. It was shown in a series of works [58, 27, 70, 35], culminating in [34], that the existence of a successful coupling of two Borel probability measures \mathbb{P}, \mathbb{P}' on Ω is closely related to \mathbb{P} and \mathbb{P}' assigning the same probability to all elements of the tail σ -algebra \mathcal{T} . More precisely [34, Theorem 2.1], one has

$$\max_{A \in \mathcal{T}} |\mathbb{P}(A) - \mathbb{P}'(A)| = 0 \quad \text{if and only if} \quad \min_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(E_0)) = 0, \tag{1.2}$$

for all Borel probability measures \mathbb{P}, \mathbb{P}' on Ω .

A third example, from ergodic theory, shows that eventual equality and the tail σ -algebra in (1.2) can be replaced with the analogous objects for the notion of shift-invariance. Adopting the notation of the preceding paragraph, write \mathcal{I} for the shift-invariant σ -algebra and E_S for the equivalence relation of equality modulo shifts. A coupling of two probability measures \mathbb{P}, \mathbb{P}' on Ω is called a *successful shift-coupling* if E_S has probability one, that is, if the random sequences are almost surely equal modulo a random shift. Then, it was shown [2] that one has

$$\max_{A \in \mathcal{I}} |\mathbb{P}(A) - \mathbb{P}'(A)| = 0 \quad \text{if and only if} \quad \min_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(E_S)) = 0, \tag{1.3}$$

for \mathbb{P}, \mathbb{P}' any two Borel probability measures on Ω .

Subsequent work also saw many generalizations of (1.3). For example, in [29] it is shown that a similar statement holds for measurable actions of very many groups and semigroups on Polish spaces. There is also [65] which generalizes this fact to the setting of cardinal algebras. We also have [46, Theorem 1] which establishes an analogous statement for unimodular random networks, where the shift operation is replaced by a root-change operation; this example is particularly interesting since the underlying measurable space is not standard Borel [46, Proposition 1].

The most general formulation of an existing result in the spirit of strong duality that we are aware of is [59, Theorem 1] which proves, as a consequence of a much more general result on constrained versions of the Skorokhod representation theorem, the following: If two Borel probability measures on a metric space S agree on the sub- σ -algebra generated by a measurable function $g : S \rightarrow T$ for a sufficiently nice space T , then they can be coupled together so that their images under g are almost surely equal. This is a sort of strong duality statement for equivalence relations on S that can be represented as pullbacks of Δ in T by g ; however, the regularity conditions required on T mean this result does not recover strong duality for many examples of interest (for example, the pairs (E_0, T) and (E_S, \mathcal{I}) above).

1.2 Relation to Kantorovich duality

The theory of Monge-Kantorovich optimal transport and Kantorovich duality can also be seen as an important precedent for our work; see [74, 73] for the standard background, and [61] for a more comprehensive treatment which is better suited to our setting. Since (EC) is exactly a Monge-Kantorovich optimal transport problem with cost $c := 1 - \mathbb{1}_E$, it is natural to try to prove strong duality as a consequence of some sufficiently general theorem in the vast literature on Kantorovich duality [74, 73, 45, 62, 64, 5, 6]. In this subsection we address the limitations of such an approach.

First, we consider the case that Ω is Polish and E is closed in $\Omega \times \Omega$. Then the cost function $c = 1 - \mathbb{1}_E$ is lower-semicontinuous, so a standard form of Kantorovich duality [74, Theorem 5.10] gives

$$\max_{f, f'} \left(\int_{\Omega} f d\mathbb{P} + \int_{\Omega} f' d\mathbb{P}' \right) = \min_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(E)),$$

where the right side is achieved and the left side is taken over, say, all bounded measurable $f, f' : \Omega \rightarrow \mathbb{R}$ with $f \oplus f' \leq 1 - \mathbb{1}_E$. In fact, it is standard that one can equivalently take the maximum over function classes which are assumed to have further structure, hence massaging the left side in the following ways: first, one can assume that f is \mathbb{P} -integrable and c -convex; second, one can assume that f' is equal to f^c , the c -transform of f ; third, one can use that c satisfies the triangle inequality (this follows from E being transitive) to get that $f^c = -f$; fourth, one can use that c satisfies the bounds $0 \leq c \leq 1$ to assume that f satisfies the bounds $0 \leq f \leq 1$; fifth, one can use that c is symmetric (this follows from E being reflexive) to replace the parentheses with absolute values. At this point some straightforward additional analysis shows that one can take f to be of the form $\mathbb{1}_A$ for A ranging over a suitable sub- σ -algebra \mathcal{G} of \mathcal{F} . (Alternatively, one can begin by directly applying a form of Kantorovich duality which is specifically tailored to cost functions which only take on values in $\{0, 1\}$, like [73, Theorem 1.27].) Thus, Kantorovich duality implies that closed equivalence relations satisfy our strong duality.

However, for the applications most of interest in probability theory, one needs to consider the case that Ω is Polish and E is F_{σ} in $\Omega \times \Omega$. Since $c = 1 - \mathbb{1}_E$ is no longer guaranteed to be lower semicontinuous, the standard form of Kantorovich duality does not apply. Nonetheless, one has [61, Corollary 2.3.9] for all measurable costs, which guarantees

$$\sup_{f, f'} \left(\int_{\Omega} f d\mathbb{P} + \int_{\Omega} f' d\mathbb{P}' \right) = \inf_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(E)),$$

where again the left side is taken over all bounded measurable $f, f' : \Omega \rightarrow \mathbb{R}$ with $f \oplus f' \leq 1 - \mathbb{1}_E$. While this may give the impression that strong duality is within reach, there are two crucial ways in which we become stuck: For one, there are not sufficient results in the literature which guarantee that the left side can be massaged, as in the

preceding paragraph, into the form of a total variation norm; indeed, steps three, four, and five can be executed similarly, but steps one and two are hard to rigorously justify, due to the fact that classical approximation arguments break down when c lacks the appropriate topological regularity. For another, which we explain further in the next paragraph, the existence of primal minimizers is not guaranteed. Thus, existing forms of Kantorovich duality do not appear to imply that F_σ equivalence relations satisfy our strong duality for a suitable sub- σ -algebra \mathcal{G} of \mathcal{F} .

In our opinion, the question of the existence of primal minimizers has received surprisingly little attention. In the classical setting where Ω is a Polish space and the cost c is lower semi-continuous, existence of course follows immediately from topological considerations. Yet, many authors do not attempt to expand the scope of existence results; for instance, in [6] it is stated: “If c fails to be lower semi-continuous, there is little reason why a primal optimizer should exist”. Nonetheless, there are, as far as we know, exactly two known results outside the classical setting: The first is [61, Remark 2.12(a)] in which it is shown that primal minimizers always exist if one works in the world of finitely-additive probability measures, but this is not useful to us since we are interested in the more standard setting of countably-additive probability measures. The second is [61, Theorem 2.3.10] (originally proven in [45]) in which it is shown that primal minimizers exist for costs which are limits of regular costs, with respect to a certain metric; unfortunately, this condition is still not general enough to cover our setting of interest.

For all these reasons, the state-of-the-art technology on Kantorovich duality does not appear to imply our notion of strong duality in a sufficiently general setting. Nonetheless, our strong duality can certainly be seen as a form of Kantorovich duality, at least in a formal sense. In this way, our positive results can be seen as rather surprising in that they establish a form of Kantorovich duality for highly irregular costs. From the point of view of optimal transport, our results can be succinctly summarized as follows: One can exchange topological regularity (the cost function is lower semi-continuous) for descriptive set-theoretic regularity (the cost function is supported on an equivalence relation) while maintaining a form of Kantorovich duality.

1.3 Statement of main results

Having established these precedents for duality between (EC) and (TV), we now state the main results of the paper. Throughout, let (Ω, \mathcal{F}) be an abstract measurable space.

Write $\mathcal{P}(\Omega, \mathcal{F})$ for the set of probability measures on (Ω, \mathcal{F}) . For $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$, write $\Pi(\mathbb{P}, \mathbb{P}')$ for the space of all couplings of \mathbb{P}, \mathbb{P}' , that is, all probability measures on $(\Omega \times \Omega, \mathcal{F} \otimes \mathcal{F})$ with marginals given by \mathbb{P} and \mathbb{P}' , respectively. Then, we consider pairs (E, \mathcal{G}) , where E is an equivalence relation on Ω and where \mathcal{G} is a sub- σ -algebra of \mathcal{F} . We say that E is *measurable* if $E \in \mathcal{F} \otimes \mathcal{F}$ when viewed as a subset of $\Omega \times \Omega$, and we say that a pair (E, \mathcal{G}) satisfies *strong duality* if E is measurable and if we have

$$\max_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)| = \min_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(E)) \tag{1.4}$$

for all $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$; the appearance of “max” and “min” assert that the supremum and infimum are both achieved.

Our main results, as is the common practice in mathematical optimization, come in the form of sufficient conditions for strong duality to hold. To state them, we need to introduce an important simplification. For an equivalence relation E on Ω , define the sub- σ -algebra

$$E^* := \{A \in \mathcal{F} : \forall(\omega, \omega') \in E : (\omega \in A \Leftrightarrow \omega' \in A)\}$$

of \mathcal{F} , and, for sub- σ -algebra \mathcal{G} of \mathcal{F} , define the equivalence relation

$$\mathcal{G}^* := \{(\omega, \omega') \in \Omega \times \Omega : \forall A \in \mathcal{G} : (\omega \in A \Leftrightarrow \omega' \in A)\}$$

on Ω . (An equivalent definition of \mathcal{G}^* which has been extensively studied in classical works on the existence of regular conditional distributions [11, 8, 37] is $\mathcal{G}^* = \bigcup_{H \text{ is a } \mathcal{G}\text{-atom}} (H \times H)$, where by a \mathcal{G} -atom we mean an intersection of all events in \mathcal{G} containing a given point.) It turns out (Lemma 3.3) that $*$ is an antitone Galois correspondence which plays a central role in the notion of strong duality.

With the correspondence in hand, we can state our first main result:

Theorem 3.13. If (Ω, \mathcal{F}) is a standard Borel space and if a pair (E, \mathcal{G}) satisfies $E \subseteq \mathcal{G}^*$, $\mathcal{G} \subseteq E^*$, and $E \in \mathcal{G} \otimes \mathcal{G}$, then (E, \mathcal{G}) satisfies strong duality.

While the conditions $E \subseteq \mathcal{G}^*$ and $\mathcal{G} \subseteq E^*$ are usually trivial to verify, the condition $E \in \mathcal{G} \otimes \mathcal{G}$ is more complicated. In fact, while this condition appears to be rather general, there are important examples of pairs satisfying strong duality for which this result does not apply; notably, the tail equivalence relation and the tail σ -algebra (E_0, \mathcal{T}) given above. This suggests that strong duality is a more general phenomenon, and this leads us to our second main result:

Theorem 3.14. If (Ω, \mathcal{F}) is a standard Borel space and if equivalence relations $E_1 \subseteq E_2 \subseteq \dots$ on Ω and sub- σ -algebras $\mathcal{G}_1 \supseteq \mathcal{G}_2 \supseteq \dots$ of \mathcal{F} are such that for each $n \in \mathbb{N}$ the pair (E_n, \mathcal{G}_n) satisfies strong duality, then the pair $(\bigcup_{n \in \mathbb{N}} E_n, \bigcap_{n \in \mathbb{N}} \mathcal{G}_n)$ satisfies strong duality.

It is now useful to make a few remarks about these main theorems. First, we note that Theorem 3.13 and Theorem 3.14 are together powerful enough to establish strong dualizability for most “reasonable” pairs of equivalence relations and sub- σ -algebras appearing in probability theory. To illustrate this, we investigate three major applications (in stochastic calculus, point process theory, and random sequence simulation) in Section 2. Additionally, we show (Proposition 3.9) that many existing results in ergodic theory can be immediately upgraded into results about strong duality. For the remaining remarks, let (Ω, \mathcal{F}) denote a standard Borel space, (E, \mathcal{G}) a pair satisfying strong duality, and $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$ any probability measures.

Second, it is crucial to emphasize that the statements and proofs of our main results are purely measure-theoretic. Even though expanding the generality to this setting forces one to leave behind many useful topological tools, our proofs reveal that only a few non-elementary results (notably, the existence of regular conditional probabilities, the Hahn-Jordan decomposition for finite signed measures, and some basic aspects of Hilbert spaces) are needed. Keeping this generality in mind, it is highly non-trivial that optimizers can even be guaranteed to exist; since compactness arguments are not available, we must instead directly construct the putative optimizers and show that they have the desired properties. Thus, it is exactly this measure-theoretic perspective that allows us to push beyond the reach of standard Kantorovich duality.

Third, we comment on the structure of solutions to (EC) and (TV). Indeed, for (EC), the objective $\tilde{\mathbb{P}} \mapsto 1 - \tilde{\mathbb{P}}(E)$ is affine and the feasible region $\Pi(\mathbb{P}, \mathbb{P}')$ is convex, so it follows that its solution set is convex and, by strong duality, non-empty; in fact, we have an exact characterization of minimizers (Proposition 3.12), which largely parallels a classical characterization of primal minimizers for (1.1). For (TV), it is known (Lemma 3.11) that maximizers always exist, even without the assumption of strong duality; however, there does not appear to be a useful characterization of maximizers. We also remark that, while our results give a very wide guarantee of the existence of solutions, concrete settings of interest will generally require further analysis in order to show that there exist solutions with desirable properties.

We now consider the perspective that one of (EC) and (TV) is a given *primal* problem and that one seeks a suitable *dual* problem of the opposite form. As we outline in this last part, both perspectives are possible in great generality, although the situation is not completely symmetric.

On the one hand, we consider (Subsection 3.6) the setting that an equivalence relation E on Ω is given, and that one seeks a suitable sub- σ -algebra \mathcal{G} of \mathcal{F} such that (E, \mathcal{G}) satisfies strong duality. In this setting, it can be shown (Proposition 3.15) that (E, E^*) satisfies strong duality as soon as (E, \mathcal{G}) satisfies strong duality for some sub- σ -algebra \mathcal{G} of \mathcal{F} . Since E^* is thus a canonical choice of \mathcal{G} , we can say that E itself is *strongly dualizable* whenever the pair (E, E^*) satisfies strong duality. Let us also recall a standard term from descriptive set theory, that an equivalence relation E on a standard Borel space (Ω, \mathcal{F}) is called *smooth* if there exists a standard Borel space (X, \mathcal{X}) and a measurable map $\phi : (\Omega, \mathcal{F}) \rightarrow (X, \mathcal{X})$ such that for all $\omega, \omega' \in \Omega$ we have $(\omega, \omega') \in E$ if and only if $\phi(\omega) = \phi(\omega')$. Then we have the following:

Corollary 3.18. On a standard Borel space, any equivalence relation that can be written as a countable increasing union of smooth equivalence relations is strongly dualizable.

An equivalence relation that can be written as a countable increasing union of smooth equivalence relations is typically called *hypersmooth* in descriptive set theory. Indeed, many equivalence relations encountered in probability theory are hypersmooth.

On the other hand, we consider (Subsection 3.7) the setting that a sub- σ -algebra \mathcal{G} of \mathcal{F} is given, and that one seeks a suitable equivalence relation E on Ω such that (E, \mathcal{G}) satisfies strong duality. As before we still say that \mathcal{G} itself is *strongly dualizable* whenever the pair $(\mathcal{G}^*, \mathcal{G})$ satisfies strong duality. This setting is slightly more complicated than the previous, since \mathcal{G} being strongly dualizable need not be equivalent to the existence of some E such that (E, \mathcal{G}) satisfies strong duality, and since there exist (Example 3.20) sub- σ -algebras that are not strongly dualizable. Nonetheless, we have the following:

Corollary 3.23. On a standard Borel space, any sub- σ -algebra that can be written as a countable decreasing intersection of countably-generated sub- σ -algebras is strongly dualizable.

A sub- σ -algebra that can be written as a countable decreasing intersection of countably-generated sub- σ -algebras is sometimes called a *tail* σ -algebra in probability theory [22]. (Note that we say “a tail σ -algebra” rather than “the tail σ -algebra”, in order to distinguish the general case from the particular example of \mathcal{T} .) Indeed, many sub- σ -algebras encountered in probability theory are tail σ -algebras.

Finally, we note that, in addition to recovering some known results, our main theorems have novel consequences in a few probabilistic settings. Most importantly, we obtain a characterization of the solution to the so-called “Brownian germ coupling problem” for one-dimensional diffusions (Theorem 2.3) which complements recent work in [25, 75]. We also obtain (Corollary 2.6) some partial results on couplings of point processes; these results are simple (and perhaps already known) in the setting of Poisson point processes, but our robust method of proof allows us to conclude the same about the more complicated setting of determinantal point processes. Our work also yields (Corollary 2.10) a non-constructive proof of the existence of some interesting randomized sorting algorithms, but many questions remain.

1.4 Future work

We now outline a few open questions related to this work.

A main open question remaining at the end of this work concerns the generality of strong duality. That is, can we find necessary and sufficient conditions for a pair (E, \mathcal{G}) to satisfy strong duality? While the sufficient conditions of this paper are already powerful

enough to prove strong duality for most applications of interest to probabilists, we have two motivations for posing the question more generally.

The first motivation concerns the question of strong dualizability for equivalence relations. In this case, we have shown that all hypersmooth equivalence relations are strongly dualizable, and also (Corollary 3.19) that all countable Borel equivalence relations are strongly dualizable. Additionally, we were unable, despite trying for quite some time, to come up with a single example of a measurable equivalence relation which is not strongly dualizable. Thus, from the perspective of descriptive set theory, the following is natural:

Conjecture 1.1. *On a standard Borel space, all Borel equivalence relations are strongly dualizable.*

The second motivation concerns the question of strong dualizability for sub- σ -algebras. Indeed, in the general theory of couplings, one has Thorisson’s “Working Hypothesis” from [70] that “Each meaningful distributional relation should have a coupling counterpart”. The present results extend the existing work of [59] in developing the theme that, if two probability measures agree on a given sub- σ -algebra, then they can be coupled to be almost surely equivalent under a suitable equivalence relation. While there indeed exist sub- σ -algebras which are not strongly dualizable, we believe the present work shows that the term “meaningful” in fact excludes very few cases.

Another possible avenue for future work concerns a notion of duality which lies in between weak duality and strong duality; that is, that a pair (E, \mathcal{G}) has the property that we have

$$\max_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)| = \inf_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(E))$$

for all probability measures $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$. (Recall that the supremum on the left is always achieved, by Lemma 3.11.) While this property is likely interesting in its own right, we did not pursue it in the present work, since most interesting applications to probability require the infimum to be realized.

Additionally, several open questions arise when one interprets the present results in the setting of optimal transport theory. That is, we can consider slight modifications to the primal problem (EC) and ask whether there exists a suitable modification to the dual problem (TV) such that one maintains strong duality. For one, we can fix a family of, say, bounded measurable functions $\{f_i\}_{i \in I}$ from $\Omega \times \Omega$ to \mathbb{R} and consider the problem

$$\begin{aligned} &\text{minimize} && 1 - \tilde{\mathbb{P}}(E) \\ &\text{over} && \text{all couplings } \tilde{\mathbb{P}} \text{ of } \mathbb{P}, \mathbb{P}' \\ &\text{subject to} && \int_{\Omega \times \Omega} f_i d\tilde{\mathbb{P}} = 0 \text{ for all } i \in I. \end{aligned}$$

This constraint encodes, for example, the so-called problems of *martingale optimal transport* [4, 7] and *optimal transport under symmetry* [51]. (See [76].) Another possibility is to consider, for a parameter $\varepsilon > 0$, the regularized problem

$$\begin{aligned} &\text{minimize} && 1 - \tilde{\mathbb{P}}(E) + \varepsilon H(\tilde{\mathbb{P}} | \mathbb{P} \otimes \mathbb{P}') \\ &\text{over} && \text{all couplings } \tilde{\mathbb{P}} \text{ of } \mathbb{P}, \mathbb{P}' \end{aligned}$$

where $H(\tilde{\mathbb{P}} | \mathbb{P} \otimes \mathbb{P}')$ denotes the relative entropy. This is the analog of the so-called problem of *entropically-regularized optimal transport* which has been extensively studied in many recent works [54, 30].

Lastly, we pose the question of whether one can use Kantorovich duality alone to prove strong duality for a sufficiently rich class of examples. As we have seen in Subsection 1.2,

this is not straightforward from existing literature, but it still seems possible to sharpen various approximation arguments in order to conclude. Interestingly, it appears that arguments based on Kantorovich duality will lead to slightly different results than the arguments herein. For example, we believe that some clever applications of Kantorovich duality should allow one to prove:

Conjecture 1.2. *On a Polish space, all F_σ equivalence relations are strongly dualizable.*

Strictly speaking, this result is not implied by Corollary 3.18, since there exist equivalence relations which can be written as a countable increasing union of closed sets but not as an countable increasing union of closed equivalence relations (for instance, the shift equivalence relation E_S , and also many orbit equivalence relations). Yet, for most applications, both Corollary 3.18 and Conjecture 1.2 apply.

2 Applications

In this section we apply our main results to three different areas of probability theory, and we hope that our notion of strong duality and our sufficient conditions will inspire similar work in the future. Throughout this section, if Ω is a Polish space, we write $\mathcal{P}(\Omega)$ for the space of Borel probability measures on Ω .

2.1 Stochastic calculus

Our first application concerns stochastic calculus, so for this subsection we assume that the reader has familiarity with the basic theory, as outlined in, say, [63]. For the setting, let us write $\Omega := C_0([0, \infty), \mathbb{R})$ for the usual Wiener space of continuous real-valued functions vanishing at zero, endowed with the topology of uniform convergence on compact sets. Also write W for the Wiener measure on Ω , that is, the law of a standard Brownian motion.

A common question in the theory of stochastic processes is that of determining when two stochastic processes behave similarly at small time scales. For example, it is known that any process which is locally absolutely continuous with respect to Brownian motion must have the local limit of a Brownian motion [21, Lemma 4.3] and that “the compound Poisson component [of the Lévy-Itô decomposition] does not contribute to the initial sample path behaviour of a Lévy process” [9, p. 15]. Another well-studied notion of small-time similarity (which is especially important from the point of view of mathematical finance) is that of separating times [72, 24]: For each pair of stochastic processes started at the same point, there exists a unique stopping time such that their laws are mutually absolutely continuous before this time and mutually singular after it; in fact, exact forms of the separating times are known for pairs of processes such as Bessel processes, Lévy processes, and general diffusions.

Our object of interest in this application is a somewhat different notion of small-time similarity. For the sake of simplicity, we focus on the case of comparing an arbitrary continuous process to a Brownian motion, in the following sense:

Definition 2.1. *A probability measure $\mathbb{P} \in \mathcal{P}(\Omega)$ is said to have the Brownian germ coupling property (Brownian GCP) if one can construct a probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ on which are defined*

- (i) *a standard Brownian motion $B = \{B_t\}_{t \geq 0}$,*
- (ii) *a stochastic process $X = \{X_t\}_{t \geq 0}$ with law \mathbb{P} , and*
- (iii) *a random time T with $\tilde{\mathbb{P}}(T > 0) = 1$,*

such that $\tilde{\mathbb{P}}(X_t = B_t \text{ for all } 0 \leq t \leq T) = 1$. A stochastic process is said to have the Brownian GCP if its law has the Brownian GCP.

In words, a stochastic process with the Brownian GCP is a process which can be coupled to almost surely travel alongside a Brownian motion for some positive amount of time. This notion of local Brownianity appears rather strong, and, apart from trivialities, it is not at all obvious how to construct stochastic processes satisfying this property. Thus, we were greatly intrigued by the results of [25], which show that all Brownian motions with drift satisfy the Brownian GCP. This suggests that the class of stochastic processes satisfying the Brownian GCP may actually be quite rich; our present goal is to understand exactly how rich it is.

Our major contribution is to show that the duality theory presents a robust approach to this problem. Towards this end, if we write $\omega = \{\omega_t\}_{t \geq 0}$ for the canonical coordinate process in Ω and if we define the usual germ σ -algebra via

$$\mathcal{F}_{0+} := \bigcap_{t>0} \sigma(\omega_s : 0 \leq s \leq t),$$

then we are led to the following fundamental result:

Proposition 2.2. *A probability measure $\mathbb{P} \in \mathcal{P}(\Omega)$ has the Brownian GCP if and only if we have $\mathbb{P}(A) = \mathbb{W}(A)$ for all $A \in \mathcal{F}_{0+}$.*

Proof. It can be easily checked that the equivalence relation

$$E_{0+} := \bigcup_{t>0} \{(\omega, \omega') \in \Omega \times \Omega : \{\omega_s\}_{0 \leq s \leq t} = \{\omega'_s\}_{0 \leq s \leq t}\}$$

is hypersmooth and satisfies $E_{0+}^* = \mathcal{F}_{0+}$, hence by Corollary 3.18 that $(E_{0+}, \mathcal{F}_{0+})$ satisfies strong duality. Thus, we claim that a probability measure $\mathbb{P} \in \mathcal{P}(\Omega)$ has the Brownian GCP if and only if there exists some coupling $\tilde{\mathbb{P}} \in \Pi(\mathbb{W}, \mathbb{P})$ with $\tilde{\mathbb{P}}(E_{0+}) = 1$. For one direction, observe that, if \mathbb{P} satisfies the GCP, then the joint law $\tilde{\mathbb{P}}$ of (B, X) is exactly a coupling $\tilde{\mathbb{P}} \in \Pi(\mathbb{W}, \mathbb{P})$ with $\tilde{\mathbb{P}}(E_{0+}) = 1$. For the other direction, suppose that $\tilde{\mathbb{P}} \in \Pi(\mathbb{W}, \mathbb{P})$ is some coupling with $\tilde{\mathbb{P}}(E_{0+}) = 1$. We claim that the desired space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ is given by $\tilde{\Omega} := \Omega \times \Omega$, and $\tilde{\mathcal{F}} := \mathcal{B}(\Omega) \otimes \mathcal{B}(\Omega)$, and $\tilde{\mathbb{P}}$ as given. Indeed, one can define $T := \sup\{t \geq 0 : \omega_s = \omega'_s \text{ for all } 0 \leq s \leq t\}$, which is clearly a measurable function of the pair (ω, ω') . Thus, the result holds by the definition of strong duality. \square

On the one hand, this provides an authoritative answer to the question of which stochastic processes have the Brownian GCP. On the other hand, it remains to show that this equivalent condition is easy to verify in some concrete cases. In fact, the value of Proposition 2.2 is that the classical tools of stochastic calculus have very much to say about the germ σ -algebra \mathcal{F}_{0+} , and thus many soft arguments suddenly become available to us.

As an illustration of this, we have a simple and comprehensive characterization of the Brownian GCP for the case of time-homogeneous one-dimensional diffusions within the realm of classical conditions ensuring the existence of strong solutions to a given SDE [63, Chapter IX, Theorem 2.1]. Since the Brownian GCP is spatially translation-invariant and temporally local, we lose no generality in restricting our attention to diffusions started at the origin and viewed on a finite time interval.

Theorem 2.3. *If $\mu : \mathbb{R} \rightarrow \mathbb{R}$ and $\sigma : \mathbb{R} \rightarrow [0, \infty)$ are Lipschitz continuous, then the strong solution $X = \{X_t\}_{0 \leq t \leq 1}$ of the SDE*

$$\begin{cases} dX_t = \mu(X_t)dt + \sigma(X_t)dB_t & \text{for } 0 \leq t \leq 1 \\ X_0 = 0, \end{cases} \tag{2.1}$$

has the Brownian GCP if and only if $\sigma \equiv 1$ on some neighborhood of 0.

Proof. For one direction, let $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ be a probability space witnessing the fact that X satisfies the Brownian GCP. Then consider the event

$$A := \bigcup_{n \in \mathbb{N}} \{\omega \in \Omega : \langle \omega, \omega \rangle_t = t \text{ for all } [0, 2^{-n}]\}$$

where $\langle \omega, \omega \rangle$ denotes the Itô quadratic variation of ω . Since $A \in \mathcal{F}_{0+}$ and $\tilde{\mathbb{P}}(B \in A) = 1$, Proposition 2.2 implies that we must have $\tilde{\mathbb{P}}(X \in A) = 1$. Now note that the quadratic variation of X is given by $\langle X, X \rangle_t = \int_0^t (\sigma(X_s))^2 ds$, so we conclude that there almost surely exists some $N \in \mathbb{N}$ with $\sigma(X_s) = 1$ for all $0 \leq s \leq 2^{-N}$. In particular, we conclude $\sigma(0) = 1$.

Next we define the random times $\tau^- := \inf\{t > 0 : X_t < 0\}$ and $\tau^+ := \inf\{t > 0 : X_t > 0\}$ on $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$, which are stopping times with respect to the natural filtration of X . We recall that the standard small-time approximation for diffusions with Lipschitz coefficients guarantees that X_t/\sqrt{t} converges in distribution as $t \rightarrow 0$ to a Gaussian random variable with mean 0 and variance $\sigma(0) = 1$, hence

$$\begin{aligned} \tilde{\mathbb{P}}(\tau^- = 0) &= \lim_{t \rightarrow 0} \tilde{\mathbb{P}}(X_s < 0 \text{ for some } 0 \leq s \leq t) \\ &\geq \liminf_{t \rightarrow \infty} \tilde{\mathbb{P}}(X_t < 0) \\ &= \liminf_{t \rightarrow \infty} \tilde{\mathbb{P}}\left(\frac{X_t}{\sqrt{t}} < 0\right) = \frac{1}{2} > 0. \end{aligned}$$

By Blumenthal’s zero-one law applied to the strong Markov process X , we see that $\tilde{\mathbb{P}}(\tau^- = 0) > 0$ implies $\tilde{\mathbb{P}}(\tau^- = 0) = 1$. The same argument applies to show $\tilde{\mathbb{P}}(\tau^+ = 0) = 1$.

Finally, we put all the pieces together. By the almost sure continuity of X and the fact that $\tilde{\mathbb{P}}(\tau^- = 0) = \tilde{\mathbb{P}}(\tau^+ = 0) = 1$, we conclude that $\{X_s : 0 \leq s \leq T\}$ contains an open neighborhood of 0 for any random variable T which is almost surely strictly positive. Since we already know that there exists an \mathbb{N} -valued random variable N with $\sigma(X_s) = 1$ for all $0 \leq s \leq 2^{-N}$, we conclude that we have $\sigma \equiv 1$ on some neighborhood of 0.

For the other direction, suppose that $\sigma \equiv 1$ on some neighborhood U of 0. Then get reals $a, b \in \mathbb{R}$ with $a < b$ and $0 \in (a, b) \subseteq [a, b] \subseteq U$. Since (2.1) admits strong solutions, we can construct $X = \{X_t\}_{t \geq 0}$ on the probability space $(\Omega, \mathcal{B}(\Omega), \mathbb{W})$ with its natural filtration $\{\mathcal{F}_t\}_{t \geq 0}$ defined via $\mathcal{F}_t := \sigma(\omega_s : 0 \leq s \leq t)$ for $t \geq 0$. We write \mathbb{E} for the expectation on this space.

To begin, we claim that the exit time $\tau_{a,b}^X := \inf\{t > 0 : X_t \notin [a, b]\}$ has finite exponential moments of all orders. To see this, define $m := \max_{-a \leq x \leq b} \mu(x)$, which is finite by the continuity of μ on $[a, b]$. Since we have $\mu(X_t) \leq m$ for all $0 \leq t \leq \tau_{a,b}^X$ almost surely, it follows that one can construct $B^m = \{B_t^m\}_{t \geq 0}$ a Brownian motion with drift m on the same probability space in such a way that we have $X_t \leq B_t^m$ for all $0 \leq t \leq \tau_{a,b}^X$ almost surely. Then define the stopping time $\tau_a^{B^m} := \inf\{t > 0 : B_t^m \leq a\}$, and note that we have $\tau_{a,b}^X \leq \tau_a^{B^m}$ almost surely. It is known that $\tau_a^{B^m}$ has finite exponential moments of all orders, so the same must be true of $\tau_{a,b}^X$.

In particular, we have shown

$$\mathbb{E} \left[\exp \left(\frac{1}{2} \int_0^{\tau_{a,b}^X} (\mu(X_s))^2 ds \right) \right] \leq \mathbb{E} \left[\exp \left(\frac{m^2}{2} \tau_{a,b}^X \right) \right] < \infty.$$

This stopping-time version of Novikov’s condition implies a stopping-time version of Girsanov’s theorem which states that the law of X is mutually absolutely continuous with the law of $X' = \{X'_t\}_{t \geq 0}$, the strong solution to the SDE

$$\begin{cases} dX'_t = \sigma(X'_t) dB_t & \text{for } 0 \leq t \leq 1 \\ X'_0 = 0, \end{cases}$$

when restricted to the stopped σ -algebra $\mathcal{F}_{\tau_{a,b}} \subseteq \mathcal{B}(\Omega)$, for $\tau_{a,b} := \inf\{t > 0 : \omega_t \notin [a, b]\}$. Importantly, note $\mathcal{F}_{0+} \subseteq \mathcal{F}_{\tau_{a,b}}$ by the continuity of sample paths. Since Blumenthal's zero-one law implies that all events in \mathcal{F}_{0+} have probability in $\{0, 1\}$ under the laws $\mathbb{W}(X \in \cdot)$ and $\mathbb{W}(X' \in \cdot)$, we conclude $\mathbb{W}(X \in A) = \mathbb{W}(X' \in A)$ for all $A \in \mathcal{F}_{0+}$.

Finally, observe that $\sigma \equiv 1$ on U implies that X' is a standard Brownian motion up to the stopping time $\tau_U := \inf\{t > 0 : X' \notin U\}$. Since $\tau_U > 0$ by the continuity of sample paths, this means X' has the Brownian GCP. Now by Proposition 2.2 we get $\mathbb{W}(X' \in A) = \mathbb{W}(A)$ for all $A \in \mathcal{F}_{0+}$, and, consequently, $\mathbb{W}(X \in A) = \mathbb{W}(A)$ for all $A \in \mathcal{F}_{0+}$. Therefore, one last application of Proposition 2.2 implies that X has the Brownian GCP. \square

Observe that Theorem 2.3 recovers the result of [25], since Brownian motion with drift is certainly included as a special case. However, the generality comes at a cost, since our result is essentially non-constructive. In contrast, the work in [25] is satisfyingly concrete, being based on William's excursion theory for Brownian motions with drift. Thus, we would be greatly interested, for certain diffusions, in finding concrete couplings alongside a Brownian motion such that they agree for some positive random initial amount of time. For instance, is this possible for a standard Brownian motion and a standard Ornstein-Uhlenbeck process?

2.2 Point process theory

Our second application concerns point processes, for which [41] is the standard reference. To set things up, fix $k \in \mathbb{N}$ with $k \geq 1$, and let Ω denote the space of non-negative, locally finite, integer-valued, simple Borel measures on \mathbb{R}^k . It is known [60, Proposition 1.11] that Ω is a standard Borel space when endowed with the Borel σ -algebra of the weak topology. (In fact, \mathbb{R}^k can be replaced with any locally compact Polish space). An element of Ω is often called a *configuration*, and it can equivalently be regarded as a locally finite countable subset of \mathbb{R}^k . Each $\mathbb{P} \in \mathcal{P}(\Omega)$ is called (the law of) a *point process*.

The Poisson point processes are of course the simplest family of point processes, although recent attention has mostly turned to point processes exhibiting strong spatial correlations. Recent work has significantly improved understanding of such aspects as Palm conditioning [31, 17, 57], number rigidity [33, 32], connections to infinite-dimensional stochastic differential equations [55, 38], and more. Much of this progress concerns the so-called determinantal point processes (including those arising from the Gamma, Airy, Bessel, and Bergman kernels) for which many analytic and algebraic properties form essential tools; however, there have also been significant advances in the more general setting, notably including the point process of zeros of the standard planar Gaussian analytic function.

The statements of our results require some notation and terminology. First, a finite collection of distinct points $\mathbf{x} \in (\mathbb{R}^k)^\ell$ will be called *point set* and $\ell \in \mathbb{N}$ will be called its *length*, denoted $|\mathbf{x}|$; here we use the convention that \emptyset is the unique point set with $|\emptyset| = 0$. Next, for a point process $\mathbb{P} \in \mathcal{P}(\Omega)$ and a point set $\mathbf{x} = (x_1, \dots, x_\ell) \in (\mathbb{R}^k)^\ell$, we write $\mathbb{P}_{\mathbf{x}}$ for the *reduced Palm measure* of \mathbb{P} under \mathbf{x} ; roughly speaking, the point process $\mathbb{P}_{\mathbf{x}}$ is the result of first conditioning \mathbb{P} to contain the points $x_1, \dots, x_\ell \in \mathbb{R}^k$ and then deleting them from the resulting configuration. Then we introduce the following:

Definition 2.4. *With respect to a point process $\mathbb{P} \in \mathcal{P}(\Omega)$, two points sets \mathbf{x} and \mathbf{x}' are said to have similar potential if there exists a coupling $\omega_{\mathbf{x}}$ and $\omega_{\mathbf{x}'}$ of $\mathbb{P}_{\mathbf{x}}$ and $\mathbb{P}_{\mathbf{x}'}$, respectively, on a probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ such that we have $\tilde{\mathbb{P}}(\|\omega_{\mathbf{x}} - \omega_{\mathbf{x}'}\|_{\text{TV}} < \infty) = 1$.*

Let us now provide some interpretation for this notion and this terminology. Indeed, consider taking a sample ω from \mathbb{P} and transforming it into a sample $\omega_{\mathbf{x}}$ from $\mathbb{P}_{\mathbf{x}}$. One

of course has to modify some subset of the points of ω to achieve this, but the amount of total modification necessary depends on \mathbf{x} ; if it is extremely rare that \mathbf{x} appears in ω then many points will need to be moved, and if it is extremely common that \mathbf{x} appears in ω then not so many points will need to be moved. Now let us consider the effect of conditionings on two different point sets, \mathbf{x} and \mathbf{x}' . If \mathbf{x} is much rarer than \mathbf{x}' , then many points need to be moved in order to get from ω to $\omega_{\mathbf{x}}$ while not so many points need to be moved in order to get from ω to $\omega_{\mathbf{x}'}$. Since the condition $\|\omega_{\mathbf{x}} - \omega_{\mathbf{x}'}\|_{\text{TV}} < \infty$ exactly means that the configurations $\omega_{\mathbf{x}}$ and $\omega_{\mathbf{x}'}$ differ only in finitely many points, this means that \mathbf{x} and \mathbf{x}' have a similar probability of appearing in ω .

Alternatively, one can view conditioning on containing \mathbf{x} as adding some amount of potential energy to the configuration ω ; if the amount of energy added is very high the system is likely to find a completely different equilibrium, and if the amount of energy added is very low the system is likely to stay close to its current equilibrium. Thus, the condition $\|\omega_{\mathbf{x}} - \omega_{\mathbf{x}'}\|_{\text{TV}} < \infty$ can be interpreted as saying that \mathbf{x} and \mathbf{x}' add a comparable amount of potential energy to the system ω . In fact, the idea to consider the energy added to a point process by (Palm) conditioning is not a new one; this is usually made quantitative via the analysis of so-called *logarithmic derivatives* [16], but only for the case of (Palm) conditioning on single points.

The contribution of this subsection is to show that the similarity of potentials is closely related, via our duality theory, to many aspects of the existing literature. In order to do this, we need to introduce some notation. For a compact set $B \subseteq \mathbb{R}^k$, let us write $\mathcal{F}_B := \sigma(\omega(U \setminus B) : U \subseteq \mathbb{R}^k \text{ open})$ for the σ -algebra generated by the configuration ω outside B . Then we define the *tail σ -algebra*

$$\mathcal{T} := \bigcap_{\substack{B \subseteq \mathbb{R}^k \\ B \text{ compact}}} \mathcal{F}_B,$$

and we prove the following:

Proposition 2.5. *Two point sets \mathbf{x} and \mathbf{x}' have similar potential with respect to a point process $\mathbb{P} \in \mathcal{P}(\Omega)$ if and only if $\mathbb{P}_{\mathbf{x}}(A) = \mathbb{P}_{\mathbf{x}'}(A)$ for all $A \in \mathcal{T}$.*

Proof. It can be easily shown that the equivalence relation

$$E := \bigcup_{\substack{B \subseteq \mathbb{R}^k \\ B \text{ compact}}} \{(\omega, \omega') \in \Omega \times \Omega : \omega(U \setminus B) = \omega'(U \setminus B) \text{ for all } U \subseteq \mathbb{R}^k \text{ open}\}$$

is hypersmooth and satisfies $E^* = \mathcal{T}$, hence by Corollary 3.18 that (E, \mathcal{T}) satisfies strong duality. Moreover, we observe that two point sets \mathbf{x} and \mathbf{x}' have similar potential with respect to a point process \mathbb{P} if and only if there exists a coupling $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}_{\mathbf{x}}, \mathbb{P}_{\mathbf{x}'})$ with $\tilde{\mathbb{P}}(E) = 1$. Therefore, the result follows from the definition of strong duality. \square

To push this one step further, we provide the following consequence for determinantal point processes.

Corollary 2.6. *If \mathbb{P} is a determinantal point process and the point sets \mathbf{x} and \mathbf{x}' are such that the reduced Palm measures $\mathbb{P}_{\mathbf{x}}$ and $\mathbb{P}_{\mathbf{x}'}$ are mutually absolutely continuous, then \mathbf{x} and \mathbf{x}' have similar potential with respect to \mathbb{P} .*

Proof. By the Shirai-Takahashi theorem [66, Theorem 1.7], reduced Palm measures of determinantal point processes are themselves determinantal point processes. Moreover,

we have the the Osada-Osada-Lyons zero-one law¹ which states that \mathcal{T} is trivial under every determinantal point processes. Consequently, \mathbb{P}_x and $\mathbb{P}_{x'}$ being mutually absolutely continuous implies that we must have $\mathbb{P}_x(A) = \mathbb{P}_{x'}(A)$ for all $A \in \mathcal{T}$. Thus, the proof is completed by applying Proposition 2.5. \square

To close this subsection, we explore some consequences of combining Corollary 2.6 with existing literature. First, we note that many point processes of interest have the property that the reduced Palm measures \mathbb{P}_x and $\mathbb{P}_{x'}$ are mutually absolutely if and only if the point sets x and x' satisfy $|x| = |x'|$: indeed, this is the case for many determinantal point processes on \mathbb{R} with integrable kernels [15, Theorem 1.4], for the Ginibre point process [57, Theorem 1.1], and for certain determinantal point processes on \mathbb{C} associated to Hilbert spaces of entire functions [18, Theorem 1.1]. Consequently, all of these point processes have the property that any two point sets have similar potential whenever they have the same length. For a somewhat different example, we recall that certain determinantal point processes on the unit disc $\mathbb{D} \subseteq \mathbb{C}$ associated to Hilbert spaces of holomorphic functions have the property that the reduced Palm measures of any two point sets are mutually absolutely continuous [18, Theorem 1.4]. Thus, these determinantal point processes have the property that all points sets have similar potential to each other, under no assumption on their lengths.

2.3 Random sequence simulation

Our final application concerns simulating random sequences with complicated dependency structures. For this section, let X be a Polish space, and set $\Omega := X^{\mathbb{N}}$ with the Borel σ -algebra of the product topology on Ω . We write $\omega = \{\omega_n\}_{n \in \mathbb{N}}$ for an arbitrary element of Ω .

Our motivation is to try to understand which laws of random sequences can be simulated by applying a random sorting algorithm to an independent identically distributed sequence. Towards making this precise, let us define some objects of interest. First, write S for the group of all permutations of \mathbb{N} (that is, all bijections $\sigma : \mathbb{N} \rightarrow \mathbb{N}$), and let us endow S with the topology of pointwise convergence. Then for $n \in \mathbb{N}$, write Π_n for the subgroup of all permutations $\sigma \in S$ with $\sigma(i) = i$ for $i > n$; these are the permutations that fix all elements except possibly those in $\{1, \dots, n\}$. We also define $\Pi_\infty := \bigcup_{n \in \mathbb{N}} \Pi_n \subseteq S$; it is known that the topology inherited from S makes Π_∞ into a Polish group [44, Theorem 6.26]. We also let Π_∞ act on Ω in the natural way, by setting $\sigma \cdot \{\omega_n\}_{n \in \mathbb{N}} := \{\omega_{\sigma^{-1}(n)}\}_{n \in \mathbb{N}}$. In particular, observe that the action is jointly continuous when viewed as a map $a : \Pi_\infty \times \Omega \rightarrow \Omega$. When we wish to emphasize that one argument is fixed, we will equivalently write $a(\sigma, \omega) =: a_\omega(\sigma) =: a_\sigma(\omega)$ for $\omega \in \Omega$ and $\sigma \in \Pi_\infty$.

Now we introduce our main object of interest.

Definition 2.7. For $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega)$, an algorithmic reassortment from \mathbb{P} to \mathbb{P}' is a transition kernel $K : \Omega \times \mathcal{B}(\Pi_\infty) \rightarrow [0, 1]$ such that we have

$$\mathbb{P}'(A) = \int_{\Omega} K(\omega, a_\omega^{-1}(A)) d\mathbb{P}(\omega)$$

for all $A \in \mathcal{B}(\Omega)$; we say that \mathbb{P}' is an algorithmic reassortment of \mathbb{P} if there exists an algorithmic reassortment from \mathbb{P} to \mathbb{P}' .

¹There appears to be some ambiguity about whose names should appear with this result: The zero-one law was first proven for discrete spaces and conjectured for more general spaces by Lyons in [52]. Subsequently, it was proven more generally (under very mild assumptions) by Osada and Osada in [56]. However, the proof in [56] contained a small error which was later repaired by Lyons in [53]. Also see [19, Theorem 1.7] by Bufetov, Qiu, and Shamov for closely-related concurrent work.

An algorithmic reassortment from \mathbb{P} to \mathbb{P}' exactly corresponds to an algorithm described at the beginning of this subsection: For each $\omega \in \Omega$, we have $K(\omega, \cdot)$ which represents the required (possibly-random) permutation bringing ω to ω' . Since a random element from $K(\omega, \cdot)$ fixes all but finitely elements of \mathbb{N} almost surely, it can be regarded as (in fact, decomposed into) a finite sequence of transpositions on \mathbb{N} , and this means that an algorithm represented by an algorithmic reassortment K must terminate in a finite number of swaps. However, such algorithms are not required to terminate in a finite number of reads; they may need to read the entire infinite sequence ω before deciding which swaps to apply.

Of course, it will turn out to be the case that algorithmic reassortment can be studied via our duality theory. To do this, let us define the usual *exchangeable σ -algebra* via

$$\mathcal{E} := \{A \in \mathcal{B}(\Omega) : a_\sigma^{-1}(A) = A \text{ for all } \sigma \in \Pi_\infty\}.$$

Then we have the following:

Theorem 2.8. *For probability measures $\mathbb{P}, \mathbb{P}' \in \mathcal{B}(\Omega)$, we have that \mathbb{P}' is an algorithmic reassortment of \mathbb{P} if and only if $\mathbb{P}(A) = \mathbb{P}'(A)$ for all $A \in \mathcal{E}$.*

Proof. One can easily check that the equivalence relation

$$E := \{(\omega, \omega') \in \Omega \times \Omega : \sigma \cdot \omega = \omega' \text{ for some } \sigma \in \Pi_\infty\},$$

is hypersmooth and satisfies $E^* = \mathcal{E}$, so Corollary 3.18 implies that (E, \mathcal{E}) satisfies strong duality. (Alternatively, one can directly apply Corollary 3.19.)

By the definition of strong duality, it therefore suffices to show that \mathbb{P}' being an algorithmic reassortment of \mathbb{P} is equivalent to the existence of some $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ with $\tilde{\mathbb{P}}(E) = 1$. For the first direction, suppose that \mathbb{P}' is an algorithmic reassortment of \mathbb{P} , and let K be the guaranteed kernel. Now construct a probability space on which are defined ω a random element of Ω with law \mathbb{P} and σ a random element of Π_∞ whose law conditional on ω is $K(\omega, \cdot)$. Then, setting, $\omega' := \sigma \cdot \omega$, it is easy to see that the joint law $\tilde{\mathbb{P}}$ of (ω, ω') satisfies $\tilde{\mathbb{P}}(E) = 1$. For the second direction, we suppose that there exists $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ with $\tilde{\mathbb{P}}(E) = 1$. Now we claim that there exists a measurable map $\psi : E \rightarrow \Pi_\infty$ with $\psi(\omega, \omega') \cdot \omega = \omega'$ for all $(\omega, \omega') \in E$, where E is given the Borel σ -algebra of the topology inherited from $\Omega \times \Omega$; indeed this can be verified easily from the Kuratowski and Ryll-Nardzewski measurable selection theorem, but we omit the details. (Note, however, that E is not Polish and that $(E, \mathcal{B}(E))$ is not standard Borel.) It follows that the probability space $(\Omega \times \Omega, \mathcal{B}(\Omega) \otimes \mathcal{B}(\Omega), \tilde{\mathbb{P}})$ supports a random variable $\sigma : \Omega \times \Omega \rightarrow \Pi_\infty$ defined for $(\omega, \omega') \in \Omega \times \Omega$ via

$$\sigma(\omega, \omega') := \begin{cases} \psi(\omega, \omega'), & \text{if } (\omega, \omega') \in E, \\ \sigma_0, & \text{if } (\omega, \omega') \notin E, \end{cases}$$

where $\sigma_0 \in \Pi_\infty$ is any arbitrary fixed element. Since Ω and Π are both Polish, there exists a regular conditional distribution $K : \Omega \times \mathcal{B}(\Pi_\infty) \rightarrow [0, 1]$ of σ with respect to ω , and this completes the proof. \square

Remark 2.9. The preceding result shows that \mathbb{P} is an algorithmic reassortment of \mathbb{P}' if and only if \mathbb{P}' is an algorithmic reassortment of \mathbb{P} , which is not obvious from the definition.

Our main application of this characterization is as follows:

Corollary 2.10. *If $\mathbb{P} = \bigotimes_{n \in \mathbb{N}} \mu$ for some $\mu \in \mathcal{P}(X)$ and $V : \Omega \rightarrow \mathbb{R}$ is measurable, non-negative, and has $0 < \int_\Omega V d\mathbb{P} < \infty$, then $\mathbb{P}' \in \mathcal{P}(\Omega)$ defined via*

$$\frac{d\mathbb{P}'}{d\mathbb{P}}(\omega) := \frac{V(\omega)}{\int_\Omega V d\mathbb{P}}$$

is an algorithmic reassortment of \mathbb{P} .

Proof. By the Hewitt-Savage zero-one law, we have $\mathbb{P}(A) \in \{0, 1\}$ for all $A \in \mathcal{E}$. Moreover, \mathbb{P}' and \mathbb{P} are mutually absolutely continuous by construction. Therefore, $\mathbb{P}'(A) = \mathbb{P}(A)$ for all $A \in \mathcal{E}$, so the result follows from Theorem 2.8. \square

There are still a number of interesting questions in this direction. For instance, while we have given a rather abstract existence proof of such algorithms, we would be interested in understanding whether they can be constructed, in some generality, in a more concrete way. Moreover, we would be interested in understanding when such algorithms can be guaranteed to require only finitely many reads, in addition to the existing guarantee of only finitely many writes.

3 Proofs of main results

In this section we develop the abstract duality theory as outlined in the introduction. Throughout this section, (Ω, \mathcal{F}) denotes a fixed measurable space.

3.1 Preliminaries

In this subsection we review some notation, definitions, and results that will be needed throughout the paper.

The first concepts concern the elementary notions of relations. By a *relation* R on Ω we mean any subset $R \subseteq \Omega \times \Omega$. By an *equivalence relation* E on Ω we mean a relation Ω which is reflexive, symmetric, and transitive. There is a well-known correspondence between equivalence relations on Ω and partitions of Ω , where an equivalence relation gives rise to a partition by dividing the space into equivalence classes, and where a partition gives rise to an equivalence relation by declaring two points to be equivalent if and only if they occur in the same element of the partition.

Next we consider basic measure theory, as found in, say [40] or [13]. For this part, we let (S, \mathcal{S}) denote an abstract measurable space; the results herein will typically be applied when (S, \mathcal{S}) is taken to be (Ω, \mathcal{F}) or $(\Omega \times \Omega, \mathcal{F} \otimes \mathcal{F})$, but some other cases will also be used. For example, write \mathbb{R} for the set of real numbers and $\mathcal{B}(\mathbb{R})$ for the Borel σ -algebra of its standard topology. We write bS for the space of bounded measurable functions from (S, \mathcal{S}) to $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$. The Cartesian product of S with itself is denoted $S \times S$, and the product σ -algebra of S with itself, that is the σ -algebra generated by the $S \times S$, is denoted $\mathcal{S} \otimes \mathcal{S}$. We write $\pi, \pi' : S \times S \rightarrow S$ for the projection maps onto the first and second coordinates, respectively.

Now, we write $\mathcal{M}(S, \mathcal{S})$ for the space of finite signed measures on (S, \mathcal{S}) . We endow $\mathcal{M}(S, \mathcal{S})$ with the partial order \leq where $\mu, \mu' \in \mathcal{M}(S, \mathcal{S})$ have $\mu \leq \mu'$ if and only if we have $\mu(A) \leq \mu'(A)$ for all $A \in \mathcal{S}$. For $\mu, \mu' \in \mathcal{M}(S, \mathcal{S})$ there exists [40, Corollary 2.9] a unique element of $\mathcal{M}(S, \mathcal{S})$ which is \leq -maximal among all elements which are \leq -bounded above by both μ and μ' , and we denote this by $\mu \wedge \mu'$. In fact, if $P, N \in \mathcal{S}$ are respectively the positive part and negative part of the Jordan decomposition of the signed measure $\mu - \mu'$, then we have $\mu \wedge \mu' = \mu(\cdot \cap N) + \mu'(\cdot \cap P)$.

We write $\mathcal{P}(S, \mathcal{S})$ for the space of probability measures on (S, \mathcal{S}) . By a *sub-probability measure* on (S, \mathcal{S}) we mean an element $\mu \in \mathcal{M}(S, \mathcal{S})$ with $0 \leq \mu(A) \leq 1$ for all $A \in \mathcal{S}$, and we write $\mathcal{P}_s(S, \mathcal{S}) \supseteq \mathcal{P}(S, \mathcal{S})$ for the space of all sub-probability measures on (S, \mathcal{S}) . For any $\mu, \mu' \in \mathcal{P}_s(S, \mathcal{S})$, a *coupling* of μ and μ' is an element $\nu \in \mathcal{P}_s(S \times S, \mathcal{S} \otimes \mathcal{S})$ satisfying $\nu \circ \pi^{-1} = \mu$ and $\nu \circ (\pi')^{-1} = \mu'$, and a *sub-coupling* of μ and μ' is an element $\nu \in \mathcal{P}_s(S \times S, \mathcal{S} \otimes \mathcal{S})$ satisfying $\nu \circ \pi^{-1} \leq \mu$ and $\nu \circ (\pi')^{-1} \leq \mu'$. The spaces of couplings and sub-couplings of $\mu, \mu' \in \mathcal{P}_s(S, \mathcal{S})$ are denoted $\Pi(\mu, \mu')$ and $\Pi_s(\mu, \mu')$, respectively; observe that $\Pi(\mu, \mu')$ is empty unless $\mu(S) = \mu'(S)$.

Finally, we review some concepts related to Polish spaces; these ideas can be found in standard sources on measure theory like [13], or in more specialized treatments like [42]. By a *Polish space* we mean a separable topological space (X, τ) such that there exists a complete metric d on X which generates the topology τ . We write $\mathcal{B}(\tau)$ for the *Borel σ -algebra* of a Polish space, that is, the σ -algebra generated by the open sets. By a *standard Borel space* we mean a measurable space (X, \mathcal{X}) such that there exists a topology τ on X making (X, τ) into a Polish space and such that we have $\mathcal{B}(\tau) = \mathcal{X}$.

Let us also review our central definition; let (Ω, \mathcal{F}) denote an abstract measurable space, and consider any pair (E, \mathcal{G}) where E is a relation on Ω and \mathcal{G} is a subset of \mathcal{F} . We say that E is *measurable* if $E \in \mathcal{F} \otimes \mathcal{F}$. We say that (E, \mathcal{G}) satisfies *strong duality* if E is measurable and if we have

$$\max_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)| = \min_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(E))$$

for all $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$. An event $A \in \mathcal{G}$ achieving the supremum on the left is called *\mathcal{G} -optimal* for \mathbb{P}, \mathbb{P}' and a coupling $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ achieving the minimum on the right side is called *E -optimal* for \mathbb{P}, \mathbb{P}' ; a coupling $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ satisfying $1 - \tilde{\mathbb{P}}(A) = \sup_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)|$ will be called *(E, \mathcal{G}) -optimal*, and note that E -optimality is equivalent to (E, \mathcal{G}) -optimality if (E, \mathcal{G}) satisfies strong duality. Further, a coupling satisfying $\tilde{\mathbb{P}}(E) = 1$ is called *E -successful*. It turns out that it will also be useful along the way to consider two further notions of duality; we briefly introduce them now, and in the following subsections we study them more carefully.

First, let us say that a pair (E, \mathcal{G}) satisfies *weak duality* if E is measurable and if we have

$$\max_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)| \leq \inf_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(E)) \tag{3.1}$$

for all $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$. Let us emphasize that the infimum need not be achieved in that E -optimal couplings are not required to exist (and that we will later show that the supremum is always achieved.) Moreover, observe that strong duality obviously implies weak duality.

Second, let us say that a pair (E, \mathcal{G}) satisfies *quasi-strong duality* if, for all $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$, the following are equivalent:

- (i) For all $A \in \mathcal{G}$, we have $\mathbb{P}(A) = \mathbb{P}'(A)$.
- (ii) There exists a $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ and a $N \in \mathcal{F} \otimes \mathcal{F}$ with $\tilde{\mathbb{P}}(N) = 0$ and $(\Omega \times \Omega) \setminus E \subseteq N$.

Crucially, observe that the measurability of E is not required in order for (E, \mathcal{G}) to satisfy quasi-strong duality; if E happens to be measurable, then (ii) above is simply equivalent to the existence of an E -successful coupling. The primary desire to generalize beyond measurable relations is to be able to say something about analytic equivalence relations on standard Borel spaces, which are a common object of study in ergodic theory [29].

3.2 Weak duality

In this subsection we study weak duality as a stepping stone to strong duality. The results herein provide some reductions which simplify our later work.

To begin, we state a fundamental set-theoretic correspondence between equivalence relations and sub- σ -algebras.

Definition 3.1. For any relation E on Ω , define the sub- σ -algebra

$$E^* := \{A \in \mathcal{F} : \forall (\omega, \omega') \in E : (\omega \in A \Leftrightarrow \omega' \in A)\}$$

of \mathcal{F} , and, for any subset \mathcal{G} of \mathcal{F} , define the equivalence relation

$$\mathcal{G}^* := \{(\omega, \omega') \in \Omega \times \Omega : \forall A \in \mathcal{G} : (\omega \in A \Leftrightarrow \omega' \in A)\}.$$

on Ω .

We point out that the collection E^* is a σ -algebra even when E is not an equivalence relation, and likewise the relation \mathcal{G}^* is an equivalence relation even when \mathcal{G} is not a σ -algebra. Although these operations will typically be applied when E and \mathcal{G} are, respectively, an equivalence relation and a σ -algebra, the added generality will be useful in some cases. We also have the following:

Remark 3.2. The map $E \mapsto E^*$ depends on the ambient σ -algebra \mathcal{F} with which Ω is endowed, while the map $\mathcal{G} \mapsto \mathcal{G}^*$ does not. We hope it causes no confusion that this dependence is not emphasized in the notation.

Another important collection of remarks concerns the algebraic structure of this correspondence. While such algebraic terminology does not immediately pay dividends for our work, we believe that it is nonetheless useful to highlight the perspective; we direct the reader to [12, Section 1.6] for further information on abstract order theory.

Lemma 3.3. For E a relation on Ω and \mathcal{G} a subset of \mathcal{F} , we have

- (i) $E^{**} \supseteq E$ and $\mathcal{G}^{**} \supseteq \mathcal{G}$, and
- (ii) $E \subseteq \mathcal{G}^*$ if and only if $\mathcal{G} \subseteq E^*$.

In other words, the correspondence given by $*$ is an antitone Galois correspondence.

As is the case in many antitone Galois correspondences, it is instructive to understand the operation arising by applying the correspondence twice. (In general, the resulting operation is not the identity.)

Lemma 3.4. For E a relation on Ω we have $E^{**} = E$ if and only if E is an equivalence relation.

Proof. First, suppose that E is an equivalence relation, and let us show that $E^{**} = E$. Since we have $E \subseteq E^{**}$ by Lemma 3.3, it suffices to show $E^{**} \subseteq E$. Indeed, take any $(\omega, \omega') \in E^{**}$ and recall that for any $A \in E^*$ that $\omega \in A$ if and only if $\omega' \in A$. Now combine $E \in \mathcal{F} \otimes \mathcal{F}$ with Fubini's theorem to get $[\omega]_E := \{\omega'' \in \Omega : (\omega, \omega'') \in E\} \in \mathcal{F}$, then use the symmetry and transitivity of E to get $[\omega]_E \in E^*$. By the reflexivity of E we have $\omega \in [\omega]_E$, hence $(\omega, \omega') \in E^{**}$ gives $\omega' \in [\omega]_E$. This is exactly $(\omega, \omega') \in E$, hence $E^{**} \subseteq E$, as needed. Second, note that $E^{**} = E$ implies that E is an equivalence relation, since E^{**} is always an equivalence relation. \square

As we see in the following, it is not true that \mathcal{G} being a σ -algebra implies $\mathcal{G}^{**} = \mathcal{G}$.

Example 3.5. Let Ω be an uncountable set, let \mathcal{F} be the σ -algebra of all subsets of Ω , and let \mathcal{G} be the σ -algebra of all sets which are countable or whose complements are countable. Then $\mathcal{G}^* = \Delta$ and $\mathcal{G}^{**} = \mathcal{F} \supsetneq \mathcal{G}$.

Having clarified some aspects of the Galois correspondence, our next goal is to show that it is indeed useful in studying weak and strong duality. As a first indication of this, let us show how the Galois correspondence leads to a necessary and sufficient condition for weak duality.

Proposition 3.6. For E a relation on Ω and \mathcal{G} a subset of \mathcal{F} with $\Omega \in \mathcal{G}$, the pair (E, \mathcal{G}) satisfies weak duality if and only if E is measurable and $E \subseteq \mathcal{G}^*$ and (equivalently, or) $\mathcal{G} \subseteq E^*$.

Proof. First suppose that (E, \mathcal{G}) satisfies weak duality. By definition, E is measurable, so we only need to show $\mathcal{G} \subseteq E^*$ and $E \subseteq \mathcal{G}^*$. These are equivalent by part (ii) of Lemma 3.3, so it suffices to show $\mathcal{G} \subseteq E^*$. Indeed, take any $(\omega, \omega') \in E$, and note that by setting $\mathbb{P} := \delta_\omega$ and $\mathbb{P}' := \delta_{\omega'}$, we have

$$\sup_{A \in \mathcal{G}} |\mathbb{1}\{\omega \in A\} - \mathbb{1}\{\omega' \in A\}| \leq 1 - \mathbb{1}\{(\omega, \omega') \in E\} = 0.$$

This says that for any $A \in \mathcal{G}$ we have $\omega \in A$ if and only if $\omega' \in A$. In other words, $\mathcal{G} \subseteq E^*$.

Conversely, suppose that E is measurable and that $\mathcal{G} \subseteq E^*$ and $E \subseteq \mathcal{G}^*$. Note that $E \subseteq \mathcal{G}^*$ implies that for any $A \in \mathcal{G}$ we have

$$\{(\omega, \omega') \in \Omega \times \Omega : \omega \in A\} \cap E = \{(\omega, \omega') \in \Omega \times \Omega : \omega' \in A\} \cap E.$$

Thus, for any $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$ and $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$, we can bound:

$$\begin{aligned} \mathbb{P}(A) - \mathbb{P}'(A) &= \tilde{\mathbb{P}}(\omega \in A) - \tilde{\mathbb{P}}(\omega' \in A) = \tilde{\mathbb{P}}(\{\omega \in A\} \setminus E) \\ &\quad - \tilde{\mathbb{P}}(\{\omega' \in A\} \setminus E) \\ &\leq \tilde{\mathbb{P}}(\{\omega \in A\} \setminus E). \end{aligned}$$

Now take the supremum over $A \in \mathcal{G}$ and use that $\Omega \in \mathcal{G}$, then take the infimum over $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ to conclude. \square

As we have seen in Remark 3.2 and Example 3.5, the roles of E and \mathcal{G} in our duality theory are not completely symmetric. While in the bulk of the paper we typically aim to state our results in the most symmetric form possible, we will address, in Subsection 3.6 and Subsection 3.7, the particularities associated with taking either of E or \mathcal{G} as given.

Corollary 3.7. *For any measurable relation E on Ω , the pair (E, E^*) satisfies weak duality. For any $\mathcal{G} \subseteq \mathcal{F}$ such that \mathcal{G}^* is measurable and $\Omega \in \mathcal{G}$, the pair $(\mathcal{G}^*, \mathcal{G})$ satisfies weak duality.*

Proof. By part (i) of Lemma 3.3, we have $E \subseteq E^{**}$ and $\mathcal{G} \subseteq \mathcal{G}^{**}$, and we of course also have $\Omega \in E^*$. Thus, the result follows from Proposition 3.6. \square

3.3 Quasi-strong duality

In this brief subsection we study quasi-strong duality as it relates to strong duality. We begin with a simple but useful auxiliary result.

Lemma 3.8. *For any $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$, if $\tilde{\mathbb{M}} \in \Pi_s(\mathbb{P}, \mathbb{P}')$ is a sub-coupling of \mathbb{P} and \mathbb{P}' , then there exists a coupling $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ with $\tilde{\mathbb{M}} \leq \tilde{\mathbb{P}}$.*

Proof. By assumption, $\mathbb{Q} := \mathbb{P} - \tilde{\mathbb{M}} \circ \pi^{-1}$ and $\mathbb{Q}' := \mathbb{P}' - \tilde{\mathbb{M}} \circ (\pi')^{-1}$ are sub-probability measures on (Ω, \mathcal{F}) . Of course, they have the same total mass, which we denote $\gamma \in [0, 1]$. More explicitly, we have

$$\gamma = \mathbb{Q}(\Omega) = \mathbb{P}(\Omega) - (\tilde{\mathbb{M}} \circ \pi^{-1})(\Omega) = 1 - \tilde{\mathbb{M}}(\Omega \times \Omega).$$

The remainder of the proof proceeds in two cases.

If $\gamma = 0$, we take $\tilde{\mathbb{P}} := \tilde{\mathbb{M}}$. To see that it is a probability measure, compute

$$\begin{aligned} 0 = \gamma = \mathbb{Q}(\Omega) &= \mathbb{P}(\Omega) - (\tilde{\mathbb{M}} \circ \pi^{-1})(\Omega) \\ &= 1 - \tilde{\mathbb{M}}(\Omega \times \Omega) \\ &= 1 - \tilde{\mathbb{P}}(\Omega \times \Omega). \end{aligned}$$

To see that it has the correct marginals, observe that $\tilde{\mathbb{P}} \circ \pi^{-1} \leq \mathbb{P}$ with $\tilde{\mathbb{P}}$ a probability measure implies that $\tilde{\mathbb{P}} \circ \pi^{-1} = \mathbb{P}$; likewise for $\tilde{\mathbb{P}} \circ (\pi')^{-1}$.

If $\gamma > 0$, we take

$$\tilde{\mathbb{P}} := \tilde{\mathbb{M}} + \frac{\mathbb{Q} \otimes \mathbb{Q}'}{\gamma}.$$

To see that it is a probability measure, compute

$$\tilde{\mathbb{P}}(\Omega \times \Omega) = \tilde{\mathbb{M}}(\Omega \times \Omega) + \frac{\gamma^2}{\gamma} = \tilde{\mathbb{M}}(\Omega \times \Omega) + \gamma = 1.$$

To see that it has the correct marginals, compute

$$\tilde{\mathbb{P}} \circ \pi^{-1} = \tilde{\mathbb{M}} \circ \pi^{-1} + \frac{\mathbb{Q}'(\Omega)}{\gamma} \mathbb{Q} = \tilde{\mathbb{M}} \circ \pi^{-1} + \mathbb{Q} = \mathbb{P},$$

and likewise for $\tilde{\mathbb{P}} \circ (\pi')^{-1}$. □

From this we get the following.

Proposition 3.9. *On a standard Borel space (Ω, \mathcal{F}) , a pair (E, \mathcal{G}) satisfies strong duality if and only if E is measurable and (E, \mathcal{G}) satisfies quasi-strong duality.*

Proof. Let E be an equivalence relation on Ω and \mathcal{G} a sub- σ -algebra of \mathcal{F} . It is clear that (E, \mathcal{G}) satisfying strong duality implies that it satisfies quasi-strong duality. Conversely, suppose that E is measurable and that (E, \mathcal{G}) satisfies quasi-strong duality. To see that (E, \mathcal{G}) satisfies strong duality, let us show that an (E, \mathcal{G}) -optimal coupling exists for all $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$. To do this, consider $\mu := \mathbb{P}|_{\mathcal{G}}$ and $\mu' := \mathbb{P}'|_{\mathcal{G}}$ as probability measures on the measurable space (Ω, \mathcal{G}) , and set $\nu := \mu \wedge \mu'$. (Note that ν is not in general equal to $(\mathbb{P} \wedge \mathbb{P}')|_{\mathcal{G}}$.) Next, write $K : \Omega \times \mathcal{F} \rightarrow [0, 1]$ for a regular conditional distribution of \mathbb{P} with respect to \mathcal{G} , and $K' : \Omega \times \mathcal{F} \rightarrow [0, 1]$ for a regular conditional distribution of \mathbb{P}' with respect to \mathcal{G} . Finally, define sub-probability measures $\mathbb{M}, \mathbb{M}' \in \mathcal{P}_s(\Omega, \mathcal{F})$ via $\mathbb{M}(A) := \int_{\Omega} K(\omega, A) d\nu(\omega)$ and $\mathbb{M}'(A) := \int_{\Omega} K'(\omega, A) d\nu(\omega)$ for all $A \in \mathcal{F}$.

By construction we have $\mathbb{M}(\Omega) = \mathbb{M}'(\Omega) = \nu(\Omega)$. Moreover, for $A \in \mathcal{G}$ we have $K(\omega, A) = \mathbb{1}_A(\omega)$ holding \mathbb{P} -almost surely hence μ -almost surely, and $\nu \ll \mu$ implies that it further holds ν -almost surely; similarly, for $A \in \mathcal{G}$ we have $K'(\omega, A) = \mathbb{1}_A(\omega)$ holding ν -almost surely. In particular, for $A \in \mathcal{G}$ we have

$$\mathbb{M}(A) = \int_{\Omega} K(\omega, A) d\nu(\omega) = \nu(A) = \int_{\Omega} K'(\omega, A) d\nu(\omega) = \mathbb{M}'(A).$$

This means we can apply quasi-strong duality (to a suitable rescaling of \mathbb{M} and \mathbb{M}') to get that there must exist some $\tilde{\mathbb{M}} \in \Pi(\mathbb{M}, \mathbb{M}')$ and some $N \in \mathcal{F} \otimes \mathcal{F}$ with $\tilde{\mathbb{M}}(N) = 0$ and $(\Omega \times \Omega) \setminus E \subseteq N$.

Since E is measurable, this implies

$$\tilde{\mathbb{M}}(\Omega \times \Omega) - \tilde{\mathbb{M}}(E) = \tilde{\mathbb{M}}((\Omega \times \Omega) \setminus E) \leq \tilde{\mathbb{M}}(N) = 0$$

hence $\tilde{\mathbb{M}}(E) = \tilde{\mathbb{M}}(\Omega \times \Omega) = \nu(\Omega)$. In fact, for $A \in \mathcal{F}$ we have

$$\begin{aligned} (\tilde{\mathbb{M}} \circ \pi^{-1})(A) &= \mathbb{M}(A) = \int_{\Omega} K(\omega, A) d\nu(\omega) \\ &\leq \int_{\Omega} K(\omega, A) d\mu(\omega) \\ &= \int_{\Omega} K(\omega, A) d\mathbb{P}(\omega) = \mathbb{P}(A), \end{aligned}$$

hence $\tilde{M} \circ \pi^{-1} \leq \mathbb{P}$. Likewise we have $\tilde{M} \circ (\pi')^{-1} \leq \mathbb{P}'$. This gives $\tilde{M} \in \Pi_s(\mathbb{P}, \mathbb{P}')$, so we can apply Lemma 3.8 to get some $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ with $\tilde{M} \leq \tilde{\mathbb{P}}$.

Finally, let $P, N \in \mathcal{G}$ denote the positive and negative part of the Jordan decomposition of the signed measure $\mu - \mu'$ on (Ω, \mathcal{G}) . It is well-known [40, Corollary 2.9] that we have

$$\nu(A) = \mu(A) - \mu(A \cap P) + \mu'(A \cap P)$$

for all $A \in \mathcal{G}$, and thus

$$\begin{aligned} \nu(\Omega) &= 1 - \mu(P) + \mu'(P) \\ &\geq 1 - \sup_{A \in \mathcal{G}} |\mu(A) - \mu'(A)| \\ &= 1 - \sup_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)|. \end{aligned}$$

Combining this all, we have shown

$$1 - \tilde{\mathbb{P}}(E) \leq 1 - \tilde{M}(E) = 1 - \nu(\Omega) \leq \sup_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)|.$$

Since Proposition 3.6 establishes the reverse inequality, we have shown that $\tilde{\mathbb{P}}$ is (E, \mathcal{G}) -optimal, hence that (E, \mathcal{G}) satisfies strong duality. \square

The value of Proposition 3.9 is twofold. First, it allows us to simplify the proofs of some later results by demonstrating quasi-strong duality in place of strong duality. Second, it allows us to immediately upgrade many results from the existing literature, which primarily come in the form of quasi-strong duality, into statements of strong duality. For example, it is known (see [29, Proposition 3.1] and the remarks thereafter) that if G countable group acting measurably on a standard Borel space (Ω, \mathcal{F}) , then the orbit equivalence relation

$$E_G := \{(\omega, \omega') \in \Omega \times \Omega : g \cdot \omega = g' \cdot \omega' \text{ for some } g \in G\}$$

is measurable, the invariant σ -algebra

$$\mathcal{I}_G := \{A \in \mathcal{F} : g^{-1}(A) = A \text{ for all } g \in G\}$$

satisfies $E_G^* = \mathcal{I}_G$, and the pair (E_G, \mathcal{I}_G) satisfies quasi-strong duality. Thus we conclude:

Theorem 3.10. *If G is a countable group acting measurably on a standard Borel space, then (E_G, \mathcal{I}_G) satisfies strong duality.*

The result of [29] in fact guarantees quasi-strong duality for a more general collection of semigroups acting measurably on a standard Borel space, and thus we get an analogous statement of strong duality for free. However, the precise sufficient conditions of [29] are rather cumbersome to state, so we omit them for brevity's sake. Suffice it to say that quasi-strong duality is known for most orbit equivalence relations, so our results establish strong duality for most measurable orbit equivalence relations.

3.4 Optimality conditions

Before we get to the next subsection in which we prove our main results of interest on strong duality, we address a manifestation of a common principle in mathematical optimization, that identifying a useful duality between classes of optimization problems can lead one to optimality conditions.

Along the way, we will need an important intermediate result. It establishes that (TV) always admits a maximizer, and that the maximization problem is equivalent to a suitable convex relaxation. Since this result is well-known, we omit its proof.

Lemma 3.11. For any sub- σ -algebra \mathcal{G} of \mathcal{F} and any $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$, we have

$$\max_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)| = \max_{\substack{f \in b\mathcal{G} \\ 0 \leq f \leq 1}} \left| \int_{\Omega} f d\mathbb{P} - \int_{\Omega} f d\mathbb{P}' \right|.$$

Now we turn to the question of characterizing optimal solutions to the problems (EC) and (TV). We primarily pursue the study minimizers of (EC), which, as we have indicated, might not always exist. Our main result in this direction is an exact characterization of minimizing couplings.

Proposition 3.12. For a pair (E, \mathcal{G}) satisfying strong duality and any $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$, a coupling $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ is (E, \mathcal{G}) -optimal if and only if $\tilde{\mathbb{P}}(\cdot \setminus E) \circ \pi^{-1}$ and $\tilde{\mathbb{P}}(\cdot \setminus E) \circ (\pi')^{-1}$ are mutually singular on (Ω, \mathcal{G}) .

Proof. For the “if” direction, we suppose that $\tilde{\mathbb{P}}(\cdot \setminus E) \circ \pi^{-1}$ and $\tilde{\mathbb{P}}(\cdot \setminus E) \circ (\pi')^{-1}$ are mutually singular on (Ω, \mathcal{G}) , that is, that there exists $A \in \mathcal{G}$ such that we have $(\tilde{\mathbb{P}}(\cdot \setminus E) \circ \pi^{-1})(A) = 1 - \tilde{\mathbb{P}}(E)$ and $(\tilde{\mathbb{P}}(\cdot \setminus E) \circ (\pi')^{-1})(A) = 0$. By Corollary 3.7, it now suffices to show that we have $1 - \tilde{\mathbb{P}}(E) \leq |\mathbb{P}(A) - \mathbb{P}'(A)|$. Indeed, we simply bound:

$$\begin{aligned} 1 - \tilde{\mathbb{P}}(E) &= \tilde{\mathbb{P}}(\{(\omega, \omega') \in \Omega \times \Omega : \mathbb{1}_A(\omega) = 1, \mathbb{1}_A(\omega') = 0\} \setminus E) \\ &= \int_{\Omega \times \Omega} (\mathbb{1}_A(\omega) - \mathbb{1}_A(\omega')) d(\tilde{\mathbb{P}}(\cdot \setminus E))(\omega, \omega') \\ &= \int_{\Omega \times \Omega} (\mathbb{1}_A(\omega) - \mathbb{1}_A(\omega')) d\tilde{\mathbb{P}}(\omega, \omega') = \mathbb{P}(A) - \mathbb{P}'(A) \leq |\mathbb{P}(A) - \mathbb{P}'(A)|. \end{aligned}$$

For the “only if” direction, we suppose that $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ is E -optimal for \mathbb{P}, \mathbb{P}' . By Lemma 3.11, there is some event $A \in \mathcal{G}$ that is \mathcal{G} -optimal for \mathbb{P}, \mathbb{P}' . Note that $A \in \mathcal{G}$ is equivalent to having $|\mathbb{1}_A(\omega) - \mathbb{1}_A(\omega')| = 1 - \mathbb{1}_E(\omega, \omega')$ for all $\omega, \omega' \in \Omega$. At the same time, we can compute:

$$\begin{aligned} |\mathbb{P}(A) - \mathbb{P}'(A)| &= \left| \int_{\Omega \times \Omega} (\mathbb{1}_A(\omega) - \mathbb{1}_A(\omega')) d\tilde{\mathbb{P}}(\omega, \omega') \right| \\ &\leq \int_{\Omega \times \Omega} |\mathbb{1}_A(\omega) - \mathbb{1}_A(\omega')| d\tilde{\mathbb{P}}(\omega, \omega') \\ &= \int_{\Omega \times \Omega} 1 - \mathbb{1}_E(\omega, \omega') d\tilde{\mathbb{P}}(\omega, \omega') = 1 - \tilde{\mathbb{P}}(E). \end{aligned}$$

However, note that $\tilde{\mathbb{P}}$ being E -optimal and A being \mathcal{G} -optimal combine with strong duality to show that $|\mathbb{P}(A) - \mathbb{P}'(A)| = 1 - \tilde{\mathbb{P}}(E)$. Consequently, the inequality above is an equality, and we have

$$\left| \int_{\Omega \times \Omega} (\mathbb{1}_A(\omega) - \mathbb{1}_A(\omega')) d\tilde{\mathbb{P}}(\omega, \omega') \right| = \int_{\Omega \times \Omega} |\mathbb{1}_A(\omega) - \mathbb{1}_A(\omega')| d\tilde{\mathbb{P}}(\omega, \omega').$$

This implies that we have either

$$\tilde{\mathbb{P}}(\{(\omega, \omega') \in \Omega \times \Omega : \mathbb{1}_A(\omega) \geq \mathbb{1}_A(\omega')\}) = 1$$

or

$$\tilde{\mathbb{P}}(\{(\omega, \omega') \in \Omega \times \Omega : \mathbb{1}_A(\omega) \leq \mathbb{1}_A(\omega')\}) = 1.$$

In the first case, we use $|\mathbb{1}_A(\omega) - \mathbb{1}_A(\omega')| = 1$ for $(\omega, \omega') \in (\Omega \times \Omega) \setminus E$ to get

$$\tilde{\mathbb{P}}(\{(\omega, \omega') \in \Omega \times \Omega : \mathbb{1}_A(\omega) = 1, \mathbb{1}_A(\omega') = 0\} \setminus E) = 1 - \tilde{\mathbb{P}}(E)$$

hence $(\tilde{\mathbb{P}}(\cdot \setminus E) \circ \pi^{-1})(A) = 1 - \tilde{\mathbb{P}}(E)$ and $(\tilde{\mathbb{P}}(\cdot \setminus E) \circ (\pi')^{-1})(A) = 0$. In the second case, the same argument shows

$$\tilde{\mathbb{P}}(\{(\omega, \omega') \in \Omega \times \Omega : \mathbb{1}_A(\omega) = 0, \mathbb{1}_A(\omega') = 1\} \setminus E) = 1 - \tilde{\mathbb{P}}(E)$$

hence $(\tilde{\mathbb{P}}(\cdot \setminus E) \circ \pi^{-1})(A) = 0$ and $(\tilde{\mathbb{P}}(\cdot \setminus E) \circ (\pi')^{-1})(A) = 1 - \tilde{\mathbb{P}}(E)$. In either case, we have shown that $\tilde{\mathbb{P}}(\cdot \setminus E) \circ \pi^{-1}$ and $\tilde{\mathbb{P}}(\cdot \setminus E) \circ (\pi')^{-1}$ are mutually singular on (Ω, \mathcal{G}) , as claimed. \square

We believe, but are unable to prove, that the previous condition for minimality should be necessary and sufficient in great generality. For one indication of this, we recall the result [3, Theorem 1]: For a measurable cost function on a standard Borel space, a coupling is minimal for the Monge-Kantorovich optimal transport problem if and only if it is concentrated on a cyclically monotone set. Thus, our condition is necessary and sufficient if one can show that a set $S \subseteq \Omega \times \Omega$ is $(1 - \mathbb{1}_E)$ -cyclically monotone if and only if there exists an event $A \in \mathcal{G}$ satisfying $S \setminus E \subseteq (A \times \Omega) \cup (\Omega \times (\Omega \setminus A))$. We believe this is an interesting avenue for future work.

One can certainly also use this duality to characterize maximizers of (TV), but the resulting condition is not so useful: An event $A \in \mathcal{G}$ is \mathcal{G} -optimal if and only if there exists a coupling $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ such that A witnesses the mutual singularity of $\tilde{\mathbb{P}}(\cdot \setminus E) \circ \pi^{-1}$ and $\tilde{\mathbb{P}}(\cdot \setminus E) \circ (\pi')^{-1}$. We believe it would be interesting to see if this condition can be re-interpreted in a manner which does not reference a quantifier over all possible couplings $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$.

3.5 Strong duality

In this subsection, we finally devote our attention to strong duality. In particular, we state and prove the sufficient conditions of our main results.

Theorem 3.13. *If (Ω, \mathcal{F}) is a standard Borel space and if a pair (E, \mathcal{G}) satisfies $E \subseteq \mathcal{G}^*$ and (equivalently, or) $\mathcal{G} \subseteq E^*$ as well as $E \in \mathcal{G} \otimes \mathcal{G}$, then (E, \mathcal{G}) satisfies strong duality.*

Proof. We see that $\mathcal{G} \subseteq \mathcal{F}$ implies that E is measurable, so, by Proposition 3.9, it suffices to show that (E, \mathcal{G}) satisfies quasi-strong duality. Further, note that Proposition 3.6 guarantees that (ii) implies (i) in the definition of quasi-strong duality, so it only remains to show that (i) implies (ii). That is, for any $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$ satisfying $\mathbb{P}(A) = \mathbb{P}'(A)$ for all $A \in \mathcal{G}$, we must construct an E -successful coupling.

Our construction is as follows. First, define $\nu := \mathbb{P}|_{\mathcal{G}}$ as a probability measure on (Ω, \mathcal{G}) , and note by assumption that we also have $\nu = \mathbb{P}'|_{\mathcal{G}}$. Second, we use the fact that (Ω, \mathcal{F}) is standard Borel to get [13, Corollary 10.4.6] a regular conditional distribution of \mathbb{P} with respect to \mathcal{G} denoted $K : \Omega \times \mathcal{F} \rightarrow [0, 1]$ as well as a regular conditional probability of \mathbb{P}' with respect to \mathcal{G} denoted $K' : \Omega \times \mathcal{F} \rightarrow [0, 1]$. Next, we define the set-function $\tilde{\mathbb{P}} : \mathcal{F} \times \mathcal{F} \rightarrow [0, 1]$ via

$$\tilde{\mathbb{P}}(A \times A') := \int_{\Omega} K(\omega, A)K'(\omega, A') d\nu(\omega)$$

for $A, A' \in \mathcal{F}$. The last step of the construction is to show that $\tilde{\mathbb{P}}$ extends uniquely to a probability measure on $(\Omega \times \Omega, \mathcal{F} \otimes \mathcal{F})$, which (by a slight abuse of notation) we also denote by $\tilde{\mathbb{P}}$. This of course follows from Carathéodory's extension theorem [40, Theorem 2.5] if we can show that $\tilde{\mathbb{P}}$ is countably additive on the semi-ring $\mathcal{F} \times \mathcal{F}$, so suppose that we have $A \times A' = \bigcup_{n \in \mathbb{N}} (A_n \times A'_n)$ for $A, A', \{A_n\}_{n \in \mathbb{N}}$, and $\{A'_n\}_{n \in \mathbb{N}}$ in \mathcal{F} such that $\{A_n \times A'_n\}_{n \in \mathbb{N}}$ are disjoint. This implies $\mathbb{1}_A \otimes \mathbb{1}_{A'} = \sum_{n \in \mathbb{N}} (\mathbb{1}_{A_n} \otimes \mathbb{1}_{A'_n})$, so for

a fixed $\omega \in \Omega$ we can take the probability of both sides under the product measure $K(\omega, \cdot) \otimes K'(\omega, \cdot)$ and we get

$$K(\omega, A)K'(\omega, A') = \sum_{n \in \mathbb{N}} K(\omega, A_n)K'(\omega, A'_n).$$

Now integrate both sides with respect to ν , and use monotone convergence to get

$$\begin{aligned} \tilde{\mathbb{P}}(A \times A') &= \int_{\Omega} K(\omega, A)K'(\omega, A') d\nu(\omega) \\ &= \int_{\Omega} \sum_{n \in \mathbb{N}} K(\omega, A_n)K'(\omega, A'_n) d\nu(\omega) \\ &= \sum_{n \in \mathbb{N}} \int_{\Omega} K(\omega, A_n)K'(\omega, A'_n) d\nu(\omega) = \sum_{n \in \mathbb{N}} \tilde{\mathbb{P}}(A_n \times A'_n). \end{aligned}$$

This is as desired, and thus completes the construction. To complete the proof, we need to verify two further properties about $\tilde{\mathbb{P}}$.

First, we claim that $\tilde{\mathbb{P}}$ is a coupling of \mathbb{P} and \mathbb{P}' . Indeed, for any $A \in \mathcal{F}$:

$$\begin{aligned} (\tilde{\mathbb{P}} \circ \pi^{-1})(A) &= \tilde{\mathbb{P}}(A \times \Omega) \\ &= \int_{\Omega} K(\omega, A)K'(\omega, \Omega) d\nu(\omega) \\ &= \int_{\Omega} K(\omega, A) d\nu(\omega) = \int_{\Omega} K(\omega, A) d\mathbb{P}(\omega) = \mathbb{P}(A), \end{aligned}$$

hence $\tilde{\mathbb{P}} \circ \pi^{-1} = \mathbb{P}$. The same calculation shows $\tilde{\mathbb{P}} \circ (\pi')^{-1} = \mathbb{P}'$.

Second, we claim that $\tilde{\mathbb{P}}(E) = 1$. To do this, write $f : (\Omega, \mathcal{G}) \rightarrow (\Omega \times \Omega, \mathcal{G} \otimes \mathcal{G})$ for the measurable map $f(\omega) := (\omega, \omega)$, and let us show that we have $\tilde{\mathbb{P}}(S) = \nu(f^{-1}(S))$ for all $S \in \mathcal{G} \otimes \mathcal{G}$. Indeed, it is straightforward to show that for all $A, A' \in \mathcal{G}$ we have

$$\begin{aligned} \tilde{\mathbb{P}}(A \times A') &= \int_{\Omega} K(\omega, A)K'(\omega, A') d\nu(\omega) \\ &= \int_{\Omega} \mathbb{1}_A(\omega)\mathbb{1}_{A'}(\omega) d\nu(\omega) = \nu(A \cap A') = \nu(f^{-1}(A \times A')). \end{aligned}$$

In the second equality, we used that $K(\omega, A) = \mathbb{1}_A(\omega)$ holds \mathbb{P} -almost surely hence ν -almost surely, and that $K'(\omega, A') = \mathbb{1}_{A'}(\omega)$ holds \mathbb{P}' -almost surely hence ν -almost surely. This shows that the probability measures $\tilde{\mathbb{P}}$ and $\nu \circ f^{-1}$ agree on the π -system $\mathcal{G} \times \mathcal{G}$, so it follows that they agree on $\mathcal{G} \otimes \mathcal{G}$. Finally, we use $E \in \mathcal{G} \otimes \mathcal{G}$ to compute:

$$\tilde{\mathbb{P}}(E) = \nu(f^{-1}(E)) \geq \nu(f^{-1}(\Delta)) = \nu(\Omega) = 1.$$

This finishes the proof. □

One may be tempted to think that Theorem 3.13 is powerful enough to establish strong duality most cases of interest in probability theory. However, many examples of (E, \mathcal{G}) of interest do not satisfy $E \in \mathcal{G} \otimes \mathcal{G}$: importantly, the tail equivalence relation E_0 and the tail σ -algebra \mathcal{T} are such that (E_0, \mathcal{T}) satisfies strong duality (to see this, combine [34] with Proposition 3.9), but $E_0 \notin \mathcal{T} \otimes \mathcal{T}$. (The fact that $E_0 \notin \mathcal{T} \otimes \mathcal{T}$ is closely related to the Glimm-Effros dichotomy [43, Theorem 6.5], but it can also be shown directly via elementary considerations.) Thus, our goal is to widen our sufficient conditions to include (E_0, \mathcal{T}) . Towards filling this gap, we establish our next main result:

Theorem 3.14. *On a standard Borel space (Ω, \mathcal{F}) , if equivalence relations $E_1 \subseteq E_2 \subseteq \dots$ on Ω and sub- σ -algebras $\mathcal{G}_1 \supseteq \mathcal{G}_2 \supseteq \dots$ of \mathcal{F} are such that for each $n \in \mathbb{N}$ the pair (E_n, \mathcal{G}_n) satisfies strong duality, then the pair $(\bigcup_{n \in \mathbb{N}} E_n, \bigcap_{n \in \mathbb{N}} \mathcal{G}_n)$ satisfies strong duality.*

Proof. For convenience, write $E := \bigcup_{n \in \mathbb{N}} E_n$ and $\mathcal{G} := \bigcap_{n \in \mathbb{N}} \mathcal{G}_n$. Since E_n is measurable for all $n \in \mathbb{N}$ it follows that E is measurable. Hence, it suffices by Proposition 3.9 to show that (E, \mathcal{G}) satisfies quasi-strong duality. To see that (ii) implies (i) in the definition of quasi-strong duality, use Proposition 3.6 to see that, for each $n \in \mathbb{N}$, the pair (E_n, \mathcal{G}_n) being strongly hence weakly dual implies $E_n \subseteq \mathcal{G}_n^*$ for all $n \in \mathbb{N}$. Now the antimonicity of the Galois correspondence gives $E_n \subseteq \mathcal{G}_n^* \subseteq \mathcal{G}^*$ for all $n \in \mathbb{N}$ hence $E \subseteq \mathcal{G}^*$. Thus, Proposition 3.6, shows that (E, \mathcal{G}) satisfies weak duality, which implies that (ii) implies (i) in the definition of quasi-strong duality. To see that (i) implies (ii), we take any $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$ with $\mathbb{P}(A) = \mathbb{P}'(A)$ for all $A \in \mathcal{G}$, and we show how to construct an (E, \mathcal{G}) -optimal coupling.

To begin, set $\mathbb{P}_0 := \mathbb{P}$ and $\mathbb{P}'_0 := \mathbb{P}'$. Then, inductively for $n \in \mathbb{N}$, use the strong duality of (E_n, \mathcal{G}_n) to let $\tilde{\mathbb{P}}_n$ be an (E_n, \mathcal{G}_n) -optimal coupling of \mathbb{P}_n and \mathbb{P}'_n , and set $\mathbb{P}_{n+1} := \mathbb{P}_n - \tilde{\mathbb{P}}_n(\cdot \cap E_n) \circ \pi^{-1}$ and $\mathbb{P}'_{n+1} := \mathbb{P}'_n - \tilde{\mathbb{P}}_n(\cdot \cap E_n) \circ (\pi')^{-1}$. To ensure that this construction is well-defined, we must verify that \mathbb{P}_n and \mathbb{P}'_n are, for each $n \in \mathbb{N}$, sub-probability measures with the same total mass. Indeed, we claim by induction on $n \in \mathbb{N}$ that

$$\mathbb{P}_n(\Omega) = \mathbb{P}'_n(\Omega) = 1 - \sum_{m=0}^{n-1} \tilde{\mathbb{P}}_m(E_m).$$

The base case $n = 0$ holds because \mathbb{P}_0 and \mathbb{P}'_0 are probability measures, and the inductive step for $n \in \mathbb{N}$ follows from combining

$$\begin{aligned} \mathbb{P}_{n+1}(\Omega) &= \mathbb{P}_n(\Omega) - (\tilde{\mathbb{P}}_n(\cdot \cap E_n) \circ \pi^{-1})(\Omega) \\ &= \mathbb{P}_n(\Omega) - \tilde{\mathbb{P}}_n(E_n) \\ &= 1 - \sum_{m=0}^{n-1} \tilde{\mathbb{P}}_m(E_m) - \tilde{\mathbb{P}}_n(E_n) = 1 - \sum_{m=0}^n \tilde{\mathbb{P}}_m(E_m) \end{aligned}$$

with the analogous calculation for $\mathbb{P}'_{n+1}(\Omega)$.

Next, fix $n \in \mathbb{N}$, and recall that by construction we have $\mathbb{P}_{n+1} = \mathbb{P} - \sum_{m=0}^n \tilde{\mathbb{P}}_m(\cdot \cap E_m) \circ \pi^{-1}$ and $\mathbb{P}'_{n+1} = \mathbb{P}' - \sum_{m=0}^n \tilde{\mathbb{P}}_m(\cdot \cap E_m) \circ (\pi')^{-1}$. Thus, combining the $(E_{n+1}, \mathcal{G}_{n+1})$ -optimality of $\tilde{\mathbb{P}}_{n+1}$ with the triangle inequality, we get

$$\begin{aligned} &1 - \sum_{m=0}^{n+1} \tilde{\mathbb{P}}_m(E_m) \\ &= 1 - \sum_{m=0}^n \tilde{\mathbb{P}}_m(E_m) - \tilde{\mathbb{P}}_{n+1}(E_{n+1}) \\ &= \sup_{A \in \mathcal{G}_{n+1}} |\mathbb{P}_{n+1}(A) - \mathbb{P}'_{n+1}(A)| \\ &\leq \sup_{A \in \mathcal{G}_{n+1}} |\mathbb{P}(A) - \mathbb{P}'(A)| \\ &\quad + \sum_{m=0}^n \sup_{A \in \mathcal{G}_{n+1}} |(\tilde{\mathbb{P}}_m(\cdot \cap E_m) \circ \pi^{-1})(A) - (\tilde{\mathbb{P}}_m(\cdot \cap E_m) \circ (\pi')^{-1})(A)|. \end{aligned}$$

We claim that the sum in the last line above is equal to zero. In fact, we claim that all summands are equal to zero, in that for all $m = 0, 1, \dots, n$ we have

$$\sup_{A \in \mathcal{G}_{n+1}} |(\tilde{\mathbb{P}}_m(\cdot \cap E_m) \circ \pi^{-1})(A) - (\tilde{\mathbb{P}}_m(\cdot \cap E_m) \circ (\pi')^{-1})(A)| = 0. \tag{3.2}$$

Since

$$(\tilde{\mathbb{P}}_m(\cdot \cap E_m) \circ \pi^{-1})(\Omega) = (\tilde{\mathbb{P}}_m(\cdot \cap E_m) \circ (\pi')^{-1})(\Omega) = \tilde{\mathbb{P}}_m(E_m),$$

it follows by normalizing that (3.2) holds if and only if we have

$$\sup_{A \in \mathcal{G}_{n+1}} |(\tilde{\mathbb{P}}_m(\cdot | E_m) \circ \pi^{-1})(A) - (\tilde{\mathbb{P}}_m(\cdot | E_m) \circ (\pi')^{-1})(A)| = 0. \tag{3.3}$$

By the strong duality of $(E_{n+1}, \mathcal{G}_{n+1})$ we know that (3.3) holds if and only if there exists an E_{n+1} -successful coupling of $\tilde{\mathbb{P}}_m(\cdot | E_m) \circ \pi^{-1}$ and $\tilde{\mathbb{P}}_m(\cdot | E_m) \circ (\pi')^{-1}$. Of course, the coupling $\tilde{\mathbb{P}}_m(\cdot | E_m)$ is exactly what is needed, since $E_m \subseteq E_{n+1}$ implies $\tilde{\mathbb{P}}_m(E_{n+1} | E_m) = 1$. Thus, we have shown

$$1 - \sum_{m=0}^{n+1} \tilde{\mathbb{P}}_m(E_m) \leq \sup_{A \in \mathcal{G}_{n+1}} |\mathbb{P}(A) - \mathbb{P}'(A)| \tag{3.4}$$

for all $n \in \mathbb{N}$.

Now let us get for each $n \in \mathbb{N}$ some $A_n \in \mathcal{G}_n$ with

$$|\mathbb{P}(A_n) - \mathbb{P}'(A_n)| \geq \sup_{A \in \mathcal{G}_n} |\mathbb{P}(A) - \mathbb{P}'(A)| - \frac{1}{2^n}$$

Now consider the Hilbert space $L^2(\Omega, \mathcal{F}, \frac{1}{2}(\mathbb{P} + \mathbb{P}'))$, in which $\{\mathbb{1}_{A_n}\}_{n \in \mathbb{N}}$ form a norm-bounded sequence. By the Banach-Alaoglu theorem, there exists a subsequence $\{n_j\}_{j \in \mathbb{N}}$ and some $f \in L^2(\Omega, \mathcal{F}, \frac{1}{2}(\mathbb{P} + \mathbb{P}'))$ with $\mathbb{1}_{A_{n_j}} \rightarrow f$ weakly. Since $L^2(\Omega, \mathcal{G}_n, \frac{1}{2}(\mathbb{P} + \mathbb{P}'))$ is a strongly closed subspace of $L^2(\Omega, \mathcal{F}, \frac{1}{2}(\mathbb{P} + \mathbb{P}'))$, it follows that it is also weakly closed. This implies that $f \in L^2(\Omega, \mathcal{G}_n, \frac{1}{2}(\mathbb{P} + \mathbb{P}'))$ for all $n \in \mathbb{N}$, hence $f \in L^2(\Omega, \mathcal{G}, \frac{1}{2}(\mathbb{P} + \mathbb{P}'))$. Also, we have

$$0 \leq \lim_{j \rightarrow \infty} \int_{\Omega} \mathbb{1}_{A_{n_j}} \mathbb{1}_{\{f \leq 0\}} d\left(\frac{\mathbb{P} + \mathbb{P}'}{2}\right) = \int_{\Omega} f \mathbb{1}_{\{f \leq 0\}} d\left(\frac{\mathbb{P} + \mathbb{P}'}{2}\right) \leq 0$$

which shows that $f \geq 0$ holds \mathbb{P} - and \mathbb{P}' -almost surely; a similar argument shows that $f \leq 1$ holds \mathbb{P} - and \mathbb{P}' -almost surely. Putting this all together, we conclude that there exists a function $g : \Omega \rightarrow \mathbb{R}$ which is \mathcal{G} -measurable and satisfies $0 \leq g \leq 1$ such that $f = g$ both \mathbb{P} - and \mathbb{P}' -almost surely. Consequently, Lemma 3.11 gives:

$$\begin{aligned} \liminf_{n \rightarrow \infty} \sup_{A \in \mathcal{G}_n} |\mathbb{P}(A) - \mathbb{P}'(A)| &\leq \lim_{j \rightarrow \infty} |\mathbb{P}(A_{n_j}) - \mathbb{P}'(A_{n_j})| \\ &= \left| \int_{\Omega} f d\mathbb{P} - \int_{\Omega} f d\mathbb{P}' \right| \\ &= \left| \int_{\Omega} g d\mathbb{P} - \int_{\Omega} g d\mathbb{P}' \right| \\ &\leq \sup_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)|. \end{aligned}$$

Therefore, we conclude

$$1 - \sum_{n \in \mathbb{N}} \tilde{\mathbb{P}}_n(E_n) \leq \sup_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)|, \tag{3.5}$$

by taking $n \rightarrow \infty$ in (3.4).

Now we have all of the ingredients to construct our coupling. First, set

$$\tilde{\mathbb{M}} := \sum_{n \in \mathbb{N}} \tilde{\mathbb{P}}_n(\cdot \cap E_n),$$

which is evidently a sub-probability measure on $(\Omega \times \Omega, \mathcal{F} \otimes \mathcal{F})$. Next, we claim that \tilde{M} is a sub-coupling of \mathbb{P} and \mathbb{P}' . Indeed, for any $A \in \mathcal{F}$ and $n \in \mathbb{N}$:

$$\begin{aligned} \sum_{m=0}^n (\tilde{\mathbb{P}}_m(\cdot \cap E_m) \circ \pi^{-1})(A) &= \sum_{m=0}^n (\mathbb{P}_{m+1}(A) - \mathbb{P}_{m+1}(A)) \\ &= \mathbb{P}(A) - \mathbb{P}_{n+1}(A) \\ &\leq \mathbb{P}(A). \end{aligned}$$

Thus, taking $n \rightarrow \infty$ we get

$$(\tilde{M} \circ \pi^{-1})(A) = \sum_{n \in \mathbb{N}} (\tilde{\mathbb{P}}_n(\cdot \cap E_n) \circ \pi^{-1})(A) \leq \mathbb{P}(A)$$

as claimed. We also get $\tilde{M} \circ \pi^{-1} \leq \mathbb{P}$ by the same calculation. Lastly, we apply Lemma 3.8 to get some $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ with $\tilde{M} \leq \tilde{\mathbb{P}}$.

It only remains to show that $\tilde{\mathbb{P}}$ is (E, \mathcal{G}) -optimal. Indeed, note that for all $n \in \mathbb{N}$:

$$\tilde{M}(E_{n+1}) \geq \sum_{m=0}^{n+1} \tilde{\mathbb{P}}_m(E_m \cap E_{n+1}) = \sum_{m=0}^{n+1} \tilde{\mathbb{P}}_m(E_m).$$

Thus, taking $n \rightarrow \infty$ gives

$$\tilde{M}(E) \geq \sum_{n \in \mathbb{N}} \tilde{\mathbb{P}}_n(E_n).$$

Finally, by applying $\tilde{M} \leq \tilde{\mathbb{P}}$ and (3.5), we find

$$1 - \tilde{\mathbb{P}}(E) \leq 1 - \tilde{M}(E) \leq 1 - \sum_{n \in \mathbb{N}} \tilde{\mathbb{P}}_n(E_n) \leq \sup_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)|.$$

This shows that $\tilde{\mathbb{P}}$ is (E, \mathcal{G}) -optimal, so the result is proved. □

3.6 Given an equivalence relation

In this subsection we consider the setting that E is a given equivalence relation on Ω and that one wants to find a suitable sub- σ -algebra \mathcal{G} of \mathcal{F} such that (E, \mathcal{G}) satisfies strong duality. Indeed, we show that this is possible in great generality, and such results will follow readily by specializing the abstract results of the previous subsection.

The first result shows that E^* is a canonical choice of \mathcal{G} :

Proposition 3.15. *For E a relation on Ω and \mathcal{G} a subset of \mathcal{F} with $\Omega \in \mathcal{G}$, if the pair (E, \mathcal{G}) satisfies strong duality, then (E, E^*) satisfies strong duality.*

Proof. By Proposition 3.6, the weak duality of (E, \mathcal{G}) implies $\mathcal{G} \subseteq E^*$. Thus, for any $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$ we can use the strong duality of (E, \mathcal{G}) to bound:

$$\sup_{A \in E^*} |\mathbb{P}(A) - \mathbb{P}'(A)| \geq \sup_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)| = \min_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(E)).$$

At the same time, Corollary 3.7 gives the bound:

$$\sup_{A \in E^*} |\mathbb{P}(A) - \mathbb{P}'(A)| \leq \inf_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(E)).$$

Combining these completes the proof. □

Motivated by this result, a measurable equivalence relation E on Ω has (E, E^*) satisfying strong duality if and only if there exists a sub- σ -algebra \mathcal{G} of \mathcal{F} such that (E, \mathcal{G}) satisfies strong duality; in this case we say that E is *strongly dualizable*. Let us also say that E is *weakly dualizable* if (E, E^*) satisfies weak duality, and that it is *quasi-strongly dualizable* if (E, E^*) satisfies quasi-strong duality.

Next, we require some basic definitions. Let us say that an equivalence relation E on a standard Borel space (Ω, \mathcal{F}) is *smooth* if there exists a standard Borel space (X, \mathcal{X}) and some measurable map $\phi : (\Omega, \mathcal{F}) \rightarrow (X, \mathcal{X})$ such that $\omega, \omega' \in \Omega$ have $(\omega, \omega') \in E$ if and only if $\phi(\omega) = \phi(\omega')$. Roughly speaking, a smooth equivalence relation is one that is “generated by” equality of Polish-space valued random variables. Then, we have the following equivalent characterizations of smoothness:

Lemma 3.16. *For an equivalence relation E on Ω , the following are equivalent:*

- (i) E is smooth.
- (ii) E^* is countably-generated.
- (iii) $E \in E^* \otimes E^*$.

Remark 3.17. Although it is not in our notation, the properties (i), (ii), and (iii) all depend on the ambient σ -algebra \mathcal{F} with which Ω is endowed. (cf. Remark 3.2)

Proof. It is classical that (i) and (ii) are equivalent (see [67, Exercise 5.1.10]), so it suffices to prove that (i) and (iii) are equivalent. To show (i) implies (iii), suppose E is smooth so there exists a standard Borel space (X, \mathcal{X}) and a measurable map $\phi : \Omega \rightarrow X$ such that $\omega, \omega' \in \Omega$ have $(\omega, \omega') \in E$ if and only if $\phi(\omega) = \phi(\omega')$. It readily follows that $\phi : (\Omega, E^*) \rightarrow (X, \mathcal{X})$ is measurable, hence also that $f : (\Omega \times \Omega, E^* \otimes E^*) \rightarrow (X \times X, \mathcal{X} \otimes \mathcal{X})$ defined via $f(\omega, \omega') := (\phi(\omega), \phi(\omega'))$ is measurable. Finally, observe that we have $\Delta \in \mathcal{X} \otimes \mathcal{X}$ since (X, \mathcal{X}) is standard Borel, hence $E = f^{-1}(\Delta) \in E^* \otimes E^*$. Thus, (iii) holds.

It requires a bit more work to show (iii) implies (i), so suppose $E \in E^* \otimes E^*$. It is classical that $\Sigma_1 := \{B \in E^* \otimes E^* : \text{there exist } A_1, A_2, \dots \in E^* \times E^* \text{ with } B \in \sigma(A_m \times A_n : m, n \in \mathbb{N})\}$ is a σ -algebra containing $E^* \times E^*$, hence $E^* \otimes E^* \subseteq \Sigma_1$. In particular, there exist $A_1, A_2, \dots \in E^* \times E^*$ with $E \in \sigma(A_m \times A_n : m, n \in \mathbb{N})$.

Next, we aim to show that the equivalence classes $[\omega]_E := \{\omega' \in \Omega : (\omega, \omega') \in E\}$ are in $\sigma(A_n : n \in \mathbb{N})$ for all $\omega \in \Omega$. To do this, take arbitrary $\omega \in \Omega$ and define $\Sigma_2(\omega) := \{B \in E^* \otimes E^* : ([\omega]_E \times [\omega]_E) \cap B \in \sigma(A_m \times A_n : m, n \in \mathbb{N})\}$, which is easily seen to be a σ -algebra. Moreover, observe that for all $m, n \in \mathbb{N}$ we have

$$([\omega]_E \times [\omega]_E) \cap (A_m \times A_n) = \begin{cases} A_m \times A_n, & \text{if } \omega \in A_m \cap A_n, \\ \emptyset, & \text{if } \omega \notin A_m \cap A_n, \end{cases}$$

since $A_m, A_n \in E^*$. This implies $\sigma(A_n \times A_m : m, n \in \mathbb{N}) \subseteq \Sigma_2(\omega)$, hence $E \in \Sigma_2(\omega)$. Consequently, $[\omega]_E \times [\omega]_E = ([\omega]_E \times [\omega]_E) \cap E \in \sigma(A_m \times A_n : m, n \in \mathbb{N}) \subseteq \sigma(A_n : n \in \mathbb{N}) \otimes \sigma(A_n : n \in \mathbb{N})$. Now Fubini’s theorem gives $[\omega]_E \in \sigma(A_n : n \in \mathbb{N})$, as claimed.

Moving on, we claim that, for all $\omega, \omega' \in \Omega$ with $(\omega, \omega') \notin E$, there exists $n \in \mathbb{N}$ such that we have either $[\omega]_E \subseteq A_n$ and $[\omega']_E \cap A_n = \emptyset$ or $[\omega]_E \subseteq \Omega \setminus A_n$ and $[\omega']_E \subseteq A_n$. If this is not true, then, recalling $\{A_n\}_{n \in \mathbb{N}} \subseteq E^*$, there must exist $\omega, \omega' \in \Omega$ with $(\omega, \omega') \notin E$ such that for all $n \in \mathbb{N}$ we have $[\omega]_E, [\omega']_E \subseteq A_n$ or $[\omega]_E, [\omega']_E \subseteq \Omega \setminus A_n$. But $\Sigma_3(\omega, \omega') := \{A \in E^* : [\omega]_E, [\omega']_E \subseteq A \text{ or } [\omega]_E, [\omega']_E \subseteq \Omega \setminus A\}$ is a σ -algebra, so $\sigma(A_n : n \in \mathbb{N}) \subseteq \Sigma_3(\omega, \omega')$. This contradicts the conclusion of the previous paragraph.

Finally, we define the function $\phi : (\Omega, \mathcal{F}) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ via the summation $\phi(\omega) := \sum_{n \in \mathbb{N}} 3^{-n} \mathbb{1}_{A_n}(\omega)$ for all $\omega \in \Omega$, which is clearly measurable. Also, the previous paragraph shows that $\omega, \omega' \in \Omega$ have $(\omega, \omega') \in E$ if and only if $\phi(\omega) = \phi(\omega')$. Therefore, E is smooth, so (i) holds. This finishes the proof. \square

This characterization leads us to our main result:

Corollary 3.18. *On a standard Borel space, any equivalence relation that can be written as a countable increasing union of smooth equivalence relations is strongly dualizable.*

Proof. Suppose that $E_1 \subseteq E_2 \subseteq \dots$ are smooth equivalence relations and $E = \bigcup_{n \in \mathbb{N}} E_n$. For each $n \in \mathbb{N}$, we have $E_n \in E_n^* \otimes E_n^*$ by Lemma 3.16 and $E_n \subseteq E_n^{**}$ by part (i) of Lemma 3.3; thus, Theorem 3.13 implies that (E_n, E_n^*) satisfies strong duality for all $n \in \mathbb{N}$. Consequently, Theorem 3.14 implies that $(E, \bigcap_{n \in \mathbb{N}} E_n^*)$ satisfies strong duality, hence E is strongly dualizable. \square

In descriptive set theory, an equivalence relation that can be written as a countable increasing union of smooth equivalence relations is called a *hypersmooth* equivalence relation.

We also prove the following result, which is interesting from the point of view of descriptive set theory, Recall that a *Borel equivalence relation* is a measurable equivalence relation on a standard Borel space, and that a *countable Borel equivalence relation* is a Borel equivalence relation whose equivalence classes are all countable.

Corollary 3.19. *On a standard Borel space, all countable Borel equivalence relations are strongly dualizable.*

Proof. The Feldman-Moore theorem [43, Theorem 2.3] states that, for any countable Borel equivalence relation E on a standard Borel space (Ω, \mathcal{F}) , there exists a countable group G acting measurably on Ω such that $E = E_G$. Therefore, Theorem 3.10 guarantees that E is strongly dualizable. \square

3.7 Given a sub- σ -algebra

In this subsection we consider the setting that \mathcal{G} is a given sub- σ -algebra of \mathcal{F} and that one wants to find a suitable equivalence relation E on Ω such that (E, \mathcal{G}) satisfies strong duality. As before, we show that this is possible in great generality, as a consequence of our general results.

We begin, however, with a negative result which shows that the task in this setting can be impossible:

Example 3.20. Take $\Omega := [0, 1]$ with \mathcal{F} its Borel σ -algebra, and let \mathcal{G} be the σ -algebra of all subsets of Ω which are countable or whose complements are countable. Now suppose that E is an equivalence relation on Ω such that the pair (E, \mathcal{G}) satisfies strong duality. By Proposition 3.6, we have $\mathcal{G}^* \supseteq E$. Since \mathcal{G} separates points, we also have $\mathcal{G}^* = \Delta$. Thus, these considerations give $E \subseteq \mathcal{G}^* = \Delta$. Since E is an equivalence relation, it also satisfies $E \supseteq \Delta$, hence we have $E = \Delta$. Then, let λ denote the Lebesgue measure on (Ω, \mathcal{F}) , and define probability measures \mathbb{P} and \mathbb{P}' via

$$\frac{d\mathbb{P}}{d\lambda}(x) = 2x \quad \text{and} \quad \frac{d\mathbb{P}'}{d\lambda}(x) = 2(1 - x)$$

for $x \in \Omega$. It is easy to see that we have

$$\sup_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)| = 0.$$

However, since we have $1 - \mathbb{1}_\Delta(x, x') \geq x - x'$ for all $x, x' \in \Omega$ it also follows that

$$\begin{aligned} \min_{\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')} (1 - \tilde{\mathbb{P}}(\Delta)) &\geq \min_{\tilde{\mathbb{P}}} \left(\int_{\Omega} x d\mathbb{P}(x) - \int_{\Omega} x' d\mathbb{P}'(x') \right) \\ &= \int_0^1 2x^2 d\lambda(x) - \int_0^1 2x'(1 - x') d\lambda(x') = \frac{2}{3} - \frac{1}{3} = \frac{1}{3} > 0. \end{aligned}$$

Since this contradicts the assumption that (E, \mathcal{G}) satisfies strong duality, there can be no equivalence relation E on Ω for which the pair (E, \mathcal{G}) satisfies strong duality.

Despite this difficulty, we will develop some wide sufficient conditions for a positive answer, and these will be general enough to cover most cases of interest to probabilists.

Next we show that \mathcal{G}^* is a (somewhat) canonical choice of E , although observe that it can be hard to directly check for the measurability of \mathcal{G}^* given \mathcal{G} :

Proposition 3.21. *For \mathcal{G} a sub- σ -algebra of \mathcal{F} and E an equivalence relation on Ω , if the pair (E, \mathcal{G}) satisfies strong duality and \mathcal{G}^* is measurable, then $(\mathcal{G}^*, \mathcal{G})$ satisfies strong duality.*

Proof. By part (i) of Lemma 3.3 we have $\mathcal{G} \subseteq \mathcal{G}^{**}$, so Proposition 3.6 implies that $(\mathcal{G}^*, \mathcal{G})$ satisfies weak duality. To see that $(\mathcal{G}^*, \mathcal{G})$ in fact satisfies strong duality, take arbitrary $\mathbb{P}, \mathbb{P}' \in \mathcal{P}(\Omega, \mathcal{F})$, and use the strong duality of (E, \mathcal{G}) to get $\tilde{\mathbb{P}} \in \Pi(\mathbb{P}, \mathbb{P}')$ with

$$\sup_{A \in \mathcal{G}} |\mathbb{P}(A) - \mathbb{P}'(A)| = 1 - \tilde{\mathbb{P}}(E).$$

Now note that (E, \mathcal{G}) satisfying strong duality implies that it satisfies weak duality, whence $E \subseteq \mathcal{G}^*$ by Proposition 3.6. Consequently, $1 - \tilde{\mathbb{P}}(E) \geq 1 - \tilde{\mathbb{P}}(\mathcal{G}^*)$, and this completes the proof. \square

As before, let us say that a sub- σ -algebra \mathcal{G} of \mathcal{F} is *strongly dualizable* if $(\mathcal{G}^*, \mathcal{G})$ satisfies strong duality. In contrast to the previous subsection, this need not be equivalent to the existence of some measurable equivalence relation E on Ω such that (E, \mathcal{G}) satisfies strong duality, since \mathcal{G}^* can fail to be measurable. Let us also say that \mathcal{G} is *weakly dualizable* or *quasi-strongly dualizable* if $(\mathcal{G}^*, \mathcal{G})$ satisfies weak duality or quasi-strong duality, respectively.

Now we turn to another important result.

Lemma 3.22. *If a sub- σ -algebra \mathcal{G} of \mathcal{F} is countably-generated, then $\mathcal{G}^* \in \mathcal{G} \otimes \mathcal{G}$.*

Proof. Suppose that there exists a sequence $\{A_n\}_{n \in \mathbb{N}}$ in \mathcal{G} satisfying $\mathcal{G} = \sigma(A_n : n \in \mathbb{N})$. We claim in this case that

$$\mathcal{G}^* = \{(\omega, \omega') \in \Omega \times \Omega : \forall n \in \mathbb{N} (\omega \in A_n \Leftrightarrow \omega' \in A_n)\}.$$

It is obvious that the right side contains \mathcal{G}^* , so it only remains to show that the right side is contained in \mathcal{G}^* . Indeed, suppose that (ω, ω') is an element of the right side. Then $\Sigma(\omega, \omega') := \{A \in \mathcal{G} : \omega \in A \Leftrightarrow \omega' \in A\}$ is a σ -algebra containing $\{A_n\}_{n \in \mathbb{N}}$, so it must contain $\sigma(A_n : n \in \mathbb{N}) = \mathcal{G}$, and this yields $(\omega, \omega') \in \mathcal{G}^*$. Thus, we can write the expression above as

$$\begin{aligned} \mathcal{G}^* = & \bigcap_{n \in \mathbb{N}} (((\Omega \setminus A_n) \times \Omega) \cup (\Omega \times A_n)) \\ & \cap \bigcap_{n \in \mathbb{N}} ((\Omega \times (\Omega \setminus A_n)) \cup (A_n \times \Omega)). \end{aligned}$$

Since each set on the right side is in \mathcal{G} , it follows that $\mathcal{G}^* \in \mathcal{G} \otimes \mathcal{G}$. \square

Finally, we come to the main result of this subsection:

Corollary 3.23. *On a standard Borel space, any sub- σ -algebra that can be written as a countable decreasing intersection of countably-generated sub- σ -algebras is strongly dualizable.*

Proof. Suppose that $\mathcal{G}_1 \supseteq \mathcal{G}_2 \supseteq \dots$ are countably-generated sub- σ -algebras and $\mathcal{G} = \bigcap_{n \in \mathbb{N}} \mathcal{G}_n$. For each $n \in \mathbb{N}$, we have $\mathcal{G}_n^* \in \mathcal{G}_n \otimes \mathcal{G}_n$ by Lemma 3.22 and $\mathcal{G}_n \subseteq \mathcal{G}_n^{**}$ by part (i) of Lemma 3.3; thus, Theorem 3.13 implies that $(\mathcal{G}_n^*, \mathcal{G}_n)$ satisfies strong duality for all $n \in \mathbb{N}$. Consequently, Theorem 3.14 implies that $(\bigcup_{n \in \mathbb{N}} \mathcal{G}_n^*, \mathcal{G})$ satisfies strong duality. Thus, it only remains to show that $\bigcup_{n \in \mathbb{N}} \mathcal{G}_n^* = \mathcal{G}^*$. To do this, note that we readily have $\bigcup_{n \in \mathbb{N}} \mathcal{G}_n^* \subseteq (\bigcap_{n \in \mathbb{N}} \mathcal{G}_n)^*$ by the antimonotonicity of the Galois correspondence. For the converse, suppose that (ω, ω') is not in the left side. Then for each $n \in \mathbb{N}$ there exists $A_n \in \mathcal{G}_n$ with $\omega \in A_n$ and $\omega' \notin A_n$. It follows that the set $A := \bigcap_{m \in \mathbb{N}} \bigcup_{n \geq m} A_n$ has $A \in \bigcap_{n \in \mathbb{N}} \mathcal{G}_n$ and $\omega \in A$ and $\omega' \notin A$, whence (ω, ω') is not in the right side. This shows $\bigcup_{n \in \mathbb{N}} \mathcal{G}_n^* = (\bigcap_{n \in \mathbb{N}} \mathcal{G}_n)^* = \mathcal{G}^*$, so \mathcal{G} is strongly dualizable. \square

Some authors refer to a sub- σ -algebra that can be written as a countable decreasing intersection of countably-generated sub- σ -algebras a *tail σ -algebra*.

Let us conclude by giving an equivalent definition of \mathcal{G}^* which can be desirable in some situations. We write $\mathcal{G}(\omega) := \bigcap \{A \in \mathcal{G} : \omega \in A\}$ for $\omega \in \Omega$ as the intersection of all events in \mathcal{G} which contain $\omega \in \Omega$; this is called the *atom* of \mathcal{G} at ω , and we write $\text{at}(\mathcal{G}) = \{\mathcal{G}(\omega) : \omega \in \Omega\}$ for the collection of all atoms of \mathcal{G} . Then it is easy to show that we have

$$\mathcal{G}^* := \bigcup_{H \in \text{at}(\mathcal{G})} H \times H.$$

Note that, in general, atoms of \mathcal{G} need not be \mathcal{G} -measurable and, as always, that \mathcal{G}^* need not be $(\mathcal{F} \otimes \mathcal{F})$ -measurable.

The difficulties of this subsection, compared the previous subsection, are largely due to the fact that \mathcal{G}^* need not be measurable. Thus, in light of Proposition 3.9, the study of quasi-strong duality may be more natural, compared to the study of strong duality, for sub- σ -algebras \mathcal{G} .

References

- [1] Allan F. Abrahamse. The tail field of a Markov chain. *Ann. Math. Statist.*, 40:127–136, 1969. MR0248913
- [2] David J. Aldous and Hermann Thorisson. Shift-coupling. *Stochastic Process. Appl.*, 44(1):1–14, 1993. MR1198659
- [3] Mathias Beiglböck, Martin Goldstern, Gabriel Maresch, and Walter Schachermayer. Optimal and better transport plans. *J. Funct. Anal.*, 256(6):1907–1927, 2009. MR2498564
- [4] Mathias Beiglböck, Pierre Henry-Labordère, and Friedrich Penkner. Model-independent bounds for option prices—a mass transport approach. *Finance Stoch.*, 17(3):477–501, 2013. MR3066985
- [5] Mathias Beiglböck, Christian Léonard, and Walter Schachermayer. A general duality theorem for the Monge-Kantorovich transport problem. *Studia Math.*, 209(2):151–167, 2012. MR2943840
- [6] Mathias Beiglböck, Christian Léonard, and Walter Schachermayer. On the duality theory for the Monge-Kantorovich transport problem. In *Optimal Transportation*, volume 413 of *London Math. Soc. Lecture Note Ser.*, pages 216–265. Cambridge Univ. Press, Cambridge, 2014. MR3328997
- [7] Mathias Beiglböck, Marcel Nutz, and Nizar Touzi. Complete duality for martingale optimal transport on the line. *Ann. Probab.*, 45(5):3038–3074, 2017. MR3706738
- [8] Patrizia Berti and Pietro Rigo. Sufficient conditions for the existence of disintegrations. *J. Theoret. Probab.*, 12(1):75–86, 1999. MR1674968
- [9] Jean Bertoin. *Lévy Processes*, volume 121 of *Cambridge Tracts in Mathematics*. Cambridge University Press, Cambridge, 1996. MR1406564

- [10] David Blackwell and Persi Diaconis. A non-measurable tail set. In *Statistics, Probability and Game Theory*, volume 30 of *IMS Lecture Notes Monogr. Ser.*, pages 1–5. Inst. Math. Statist., Hayward, CA, 1996. MR1481768
- [11] David Blackwell and Lester E. Dubins. On existence and non-existence of proper, regular, conditional distributions. *Ann. Probability*, 3(5):741–752, 1975. MR0400320
- [12] T. S. Blyth. *Lattices and Ordered Algebraic Structures*. Universitext. Springer-Verlag London, Ltd., London, 2005. MR2126425
- [13] V. I. Bogachev. *Measure Theory. Vols. I, II*. Springer-Verlag, Berlin, 2007. MR2267655
- [14] Alexander I. Bufetov. Rigidity of determinantal point processes with the Airy, the Bessel and the gamma kernel. *Bull. Math. Sci.*, 6(1):163–172, 2016. MR3472852
- [15] Alexander I. Bufetov. Quasi-symmetries of determinantal point processes. *Ann. Probab.*, 46(2):956–1003, 2018. MR3773378
- [16] Alexander I. Bufetov, Andrey V. Dymov, and Hirofumi Osada. The logarithmic derivative for point processes with equivalent Palm measures. *J. Math. Soc. Japan*, 71(2):451–469, 2019. MR3943446
- [17] Alexander I. Bufetov and Grigori Olshanski. A hierarchy of Palm measures for determinantal point processes with gamma kernels. *Studia Math.*, 267(2):121–160, 2022. MR4474774
- [18] Alexander I. Bufetov and Yanqi Qiu. Determinantal point processes associated with Hilbert spaces of holomorphic functions. *Comm. Math. Phys.*, 351(1):1–44, 2017. MR3613499
- [19] Alexander I. Bufetov, Yanqi Qiu, and Alexander Shamov. Kernels of conditional determinantal measures and the Lyons-Peres completeness conjecture. *J. Eur. Math. Soc. (JEMS)*, 23(5):1477–1519, 2021. MR4244512
- [20] Alexander Igorevich Bufetov and Tomoyuki Shirai. Quasi-symmetries and rigidity for determinantal point processes associated with de Branges spaces. *Proc. Japan Acad. Ser. A Math. Sci.*, 93(1):1–5, 2017. MR3590124
- [21] Duncan Dauvergne, Sourav Sarkar, and Bálint Virág. Three-halves variation of geodesics in the directed landscape. *Ann. Probab.*, 50(5):1947–1985, 2022. MR4474505
- [22] Lester E. Dubins and David Heath. With respect to tail sigma fields, standard measures possess measurable disintegrations. *Proc. Amer. Math. Soc.*, 88(3):416–418, 1983. MR0699405
- [23] Frank den Hollander. Probability theory: The coupling method, December 2012.
- [24] H.-J. Engelbert, M. A. Urusov, and M. Walther. A canonical setting and separating times for continuous local martingales. *Stochastic Process. Appl.*, 119(4):1039–1054, 2009. MR2508564
- [25] Philip A. Ernst, Wilfrid S. Kendall, Gareth O. Roberts, and Jeffrey S. Rosenthal. MEXIT: maximal un-coupling times for stochastic processes. *Stochastic Process. Appl.*, 129(2):355–380, 2019. MR3907003
- [26] Jacob Feldman and Calvin C. Moore. Ergodic equivalence relations, cohomology, and von Neumann algebras. I. *Trans. Amer. Math. Soc.*, 234(2):289–324, 1977. MR0578656
- [27] H. Föllmer. Tail structure of Markov chains on infinite product spaces. *Z. Wahrsch. Verw. Gebiete*, 50(3):273–285, 1979. MR0554547
- [28] Joshua Frisch, Alexander Kechris, Forte Shinko, and Zoltán Vidnyánszky. Realizations of countable Borel equivalence relations, 2021.
- [29] Hans-Otto Georgii. Orbit coupling. *Ann. Inst. H. Poincaré Probab. Statist.*, 33(2):253–268, 1997. MR1443958
- [30] Promit Ghosal, Marcel Nutz, and Espen Bernton. Stability of entropic optimal transport and Schrödinger bridges. *J. Funct. Anal.*, 283(9):Paper No. 109622, 22, 2022. MR4460341
- [31] Subhroshekhar Ghosh. Palm measures and rigidity phenomena in point processes. *Electron. Commun. Probab.*, 21:Paper No. 85, 14, 2016. MR3592207
- [32] Subhroshekhar Ghosh and Manjunath Krishnapur. Rigidity hierarchy in random point fields: random polynomials and determinantal processes. *Comm. Math. Phys.*, 388(3):1205–1234, 2021. MR4340929

- [33] Subhroshekhar Ghosh and Yuval Peres. Rigidity and tolerance in point processes: Gaussian zeros and Ginibre eigenvalues. *Duke Math. J.*, 166(10):1789–1858, 2017. MR3679882
- [34] Sheldon Goldstein. Maximal coupling. *Z. Wahrsch. Verw. Gebiete*, 46(2):193–204, 1978/79. MR0516740
- [35] David Griffeath. A maximal coupling for Markov chains. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, 31:95–106, 1974/75. MR0370771
- [36] Greg Hjorth. Borel equivalence relations. In *Handbook of Set Theory. Vols. 1, 2, 3*, pages 297–332. Springer, Dordrecht, 2010. MR2768683
- [37] J. Hoffmann-Jørgensen. Existence of conditional probabilities. *Math. Scand.*, 28:257–264 (1972), 1971. MR0305445
- [38] Ryuichi Honda and Hirofumi Osada. Infinite-dimensional stochastic differential equations related to Bessel random point fields. *Stochastic Process. Appl.*, 125(10):3801–3822, 2015. MR3373304
- [39] Elton P. Hsu and Karl-Theodor Sturm. Maximal coupling of Euclidean Brownian motions. *Commun. Math. Stat.*, 1(1):93–104, 2013. MR3197874
- [40] Olav Kallenberg. *Foundations of Modern Probability*. Probability and its Applications (New York). Springer-Verlag, New York, second edition, 2002. MR1876169
- [41] Olav Kallenberg. *Random Measures, Theory and Applications*, volume 77 of *Probability Theory and Stochastic Modelling*. Springer, Cham, 2017. MR3642325
- [42] Alexander S. Kechris. *Classical Descriptive Set Theory*, volume 156 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995. MR1321597
- [43] Alexander S. Kechris. The theory of countable borel equivalence relations, May 2019.
- [44] Alexander S. Kechris and Christian Rosendal. Turbulence, amalgamation, and generic automorphisms of homogeneous structures. *Proc. Lond. Math. Soc. (3)*, 94(2):302–350, 2007. MR2308230
- [45] Hans G. Kellerer. Duality theorems for marginal problems. *Z. Wahrsch. Verw. Gebiete*, 67(4):399–432, 1984. MR0761565
- [46] Ali Khezeli. Shift-coupling of random rooted graphs and networks. In *Unimodularity in Randomly Generated Graphs*, volume 719 of *Contemp. Math.*, pages 175–211. Amer. Math. Soc., Providence, RI, 2018. MR3880017
- [47] John M. Lee. *Introduction to Topological Manifolds*, volume 202 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 2011. MR2766102
- [48] Torgny Lindvall. W. Doeblin 1915–1940. *Ann. Probab.*, 19(3):929–934, 1991. MR1112401
- [49] Torgny Lindvall. On Strassen’s theorem on stochastic domination. *Electron. Comm. Probab.*, 4:51–59, 1999. MR1711599
- [50] Torgny Lindvall. *Lectures on the Coupling Method*. Dover Publications, Inc., Mineola, NY, 2002. Corrected reprint of the 1992 original. MR1924231
- [51] Artur O. Lopes and Jairo K. Mengue. Duality theorems in ergodic transport. *J. Stat. Phys.*, 149(5):921–942, 2012. MR2999567
- [52] Russell Lyons. Determinantal probability: basic properties and conjectures. In *Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. IV*, pages 137–161. Kyung Moon Sa, Seoul, 2014. MR3727606
- [53] Russell Lyons. A note on tail trivality for determinantal point processes. *Electron. Commun. Probab.*, 23:Paper No. 72, 3, 2018. MR3866045
- [54] Marcel Nutz and Johannes Wiesel. Entropic optimal transport: convergence of potentials. *Probab. Theory Related Fields*, 184(1-2):401–424, 2022. MR4498514
- [55] Hirofumi Osada. Infinite-dimensional stochastic differential equations related to random matrices. *Probab. Theory Related Fields*, 153(3-4):471–509, 2012. MR2948684
- [56] Hirofumi Osada and Shota Osada. Discrete approximations of determinantal point processes on continuous spaces: tree representations and tail trivality. *J. Stat. Phys.*, 170(2):421–435, 2018. MR3744393

- [57] Hirofumi Osada and Tomoyuki Shirai. Absolute continuity and singularity of Palm measures of the Ginibre point process. *Probab. Theory Related Fields*, 165(3-4):725–770, 2016. MR3520017
- [58] J. W. Pitman. On coupling of Markov chains. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete*, 35(4):315–322, 1976. MR0415775
- [59] Luca Pratelli and Pietro Rigo. A strong version of the Skorohod representation theorem. *J. Theoret. Probab.*, 36(1):372–389, 2023. MR4585120
- [60] Chris Preston. Some notes on standard borel and related spaces, March 2007. MR2572291
- [61] Svetlozar T. Rachev and Ludger Rüschendorf. *Mass Transportation Problems. Vol. I. Probability and its Applications* (New York). Springer-Verlag, New York, 1998. Theory. MR1619171
- [62] D. Ramachandran and L. Rüschendorf. A general duality theorem for marginal problems. *Probab. Theory Related Fields*, 101(3):311–319, 1995. MR1324088
- [63] Daniel Revuz and Marc Yor. *Continuous Martingales and Brownian Motion*, volume 293 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, third edition, 1999. MR1725357
- [64] Pietro Rigo. A note on duality theorems in mass transportation. *J. Theoret. Probab.*, 33(4):2337–2350, 2020. MR4166202
- [65] Forte Shinko. Equidecomposition in cardinal algebras. *Fund. Math.*, 253(2):197–204, 2021. MR4205971
- [66] Tomoyuki Shirai and Yoichiro Takahashi. Random point fields associated with certain Fredholm determinants. I. Fermion, Poisson and boson point processes. *J. Funct. Anal.*, 205(2):414–463, 2003. MR2018415
- [67] S. M. Srivastava. *A Course on Borel Sets*, volume 180 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1998. MR1619545
- [68] V. Strassen. The existence of probability measures with given marginals. *Ann. Math. Statist.*, 36:423–439, 1965. MR0177430
- [69] Yoichiro Takahashi. Absolute continuity of Poisson random fields. *Publ. Res. Inst. Math. Sci.*, 26(4):629–647, 1990. MR1081507
- [70] Hermann Thorisson. On maximal and distributional coupling. *Ann. Probab.*, 14(3):873–876, 1986. MR0841589
- [71] Hermann Thorisson. Transforming random elements and shifting random fields. *Ann. Probab.*, 24(4):2057–2064, 1996. MR1415240
- [72] M. A. Urusov and A. S. Chernyĭ. Separating times for measures on spaces with filtration. *Teor. Veroyatnost. i Primenen.*, 48(2):416–427, 2003. MR2015464
- [73] Cédric Villani. *Topics in Optimal Transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2003. MR1964483
- [74] Cédric Villani. *Optimal Transport*, volume 338 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2009. Old and new. MR2459454
- [75] Florian Völlering. On maximal agreement couplings, 2016.
- [76] D. A. Zaev. On the Monge-Kantorovich problem with additional linear constraints. *Mat. Zametki*, 98(5):664–683, 2015. MR3438523

Acknowledgments. We thank Steve Evans for many useful conversations, and we thank Jonathan Niles-Weed and Marcel Nutz for their insights about the relationship between strong duality and Kantorovich duality. We also thank the anonymous reviewer for many suggestions that improved the quality of this paper.

Electronic Journal of Probability

Electronic Communications in Probability

Advantages of publishing in EJP-ECP

- Very high standards
- Free for authors, free for readers
- Quick publication (no backlog)
- Secure publication (LOCKSS¹)
- Easy interface (EJMS²)

Economical model of EJP-ECP

- Non profit, sponsored by IMS³, BS⁴, ProjectEuclid⁵
- Purely electronic

Help keep the journal free and vigorous

- Donate to the IMS open access fund⁶ (click here to donate!)
- Submit your best articles to EJP-ECP
- Choose EJP-ECP over for-profit journals

¹LOCKSS: Lots of Copies Keep Stuff Safe <http://www.lockss.org/>

²EJMS: Electronic Journal Management System: <https://vtex.lt/services/ejms-peer-review/>

³IMS: Institute of Mathematical Statistics <http://www.imstat.org/>

⁴BS: Bernoulli Society <http://www.bernoulli-society.org/>

⁵Project Euclid: <https://projecteuclid.org/>

⁶IMS Open Access Fund: <https://imstat.org/shop/donation/>