

OPTIMIZING EMISSIONS CONTROL POLICIES
WITH LINEAR PROGRAMMING

BY

CAMERON DAVIES

A thesis submitted in conformity with
the requirements for the degree of
Doctor of Philosophy
Department of Mathematics
University of Toronto

© Copyright by Cameron Davies 2026

ABSTRACT

Optimizing Emissions Control Policies with Linear Programming

Cameron Davies

Doctor of Philosophy

Department of Mathematics

University of Toronto

2026

The climate crisis has created a moral and strategic imperative for countries to rapidly reduce their greenhouse gas emissions, especially within the industrial sector. To maintain public and industrial support, policies must account for the heterogeneity of individual industries.

This thesis uses optimal transportation and linear programming techniques to inform the design of optimal climate policies. Within a specific industry, we extend a model of Newell and Stavins to a linear program which is reminiscent of the Kantorovich problem in optimal transport. Interpreting the primal problem as command-and-control policies, and the dual as market-based policies, we show that strong duality fails for certain emissions targets, meaning that market-based policies may fail to meet emissions targets, even in contexts where these are attainable through a command-and-control approach. We also develop several viable solutions to this shortfall, such as implementing a progressive emissions tax, or carefully designing targets for market-based policies.

ACKNOWLEDGEMENTS

I'd like to begin by thanking my doctoral supervisor, Professor Robert McCann, who has provided me with constant support and invaluable guidance throughout my degree, especially while working on this thesis. His deep and piercing knowledge on the subject of optimal transportation facilitated much of this work. I have learned many things from Professor McCann over the years – how to write and present to a mathematical audience, how to think about and structure complex problems, and that a quick jog is often the best way to get unstuck. Professor McCann has also introduced me to intriguing projects in numerous disciplines, and has provided patient guidance as I tried my hand at each. For this, I am forever grateful.

I also owe a debt of gratitude to Professor Justin Kakeu of the University of Prince Edward Island. Professor Kakeu introduced me to the problem that would later form the bulk of this thesis, and has never hesitated to share his deep knowledge of the economics literature with me. He has also taught me much about how to interpret mathematical results through the lens of economics, providing me with an additional perspective on this research program. Professor Kakeu has been a constant source of intellectual inspiration, and I am glad to have benefitted from his insights.

I am deeply grateful to my supervisory committee members, Professor Almut Burchard and Professor Mary Pugh for their constant help and support throughout my PhD. Their feedback during committee meetings has been immensely valuable for focusing, refining, and scoping this project. I am also grateful to Professor Burchard for introducing me to Professor McCann a number of years ago, and I am grateful to Professor Pugh for showing me the power of numerical analysis, which has helped to furnish and bring life to many results in this thesis.

I would also like to thank my other defense committee members, Professor Luis Seco and Professor Jeffrey Sun, who have provided insightful feedback on this work during past presentations by myself and Professor Kakeu. I am thankful for anyone who has asked questions, or given me feedback on this work, as this has helped to hone my focus and deepen my analysis.

I also wish to extend my deep gratitude to the late Professor Joe Repka and his family. Professor Repka provided me with my first exposure to university-level mathematics, and with my first research opportunity

shortly thereafter. His kindness, generosity, and patience helped me greatly during my Bachelor's degree, and, in both research and teaching, he set an example I have strived to learn from.

I'd also like to thank the staff of the math department, especially Jemima Merisca and Coco Cong, for keeping me aware of, and helping me fulfill the administrative requirements of this degree.

I also want to thank all of my department friends and officemates, who have always been willing to talk with me about my work, and who have been a delight to spend time with during my degree! Thanks Adriano, Alex, Alice, Arthur, Charlie, Fardin, Jonathan, Mikale, Nick, and Vivian!

I also want to acknowledge my friends and family outside of mathematics, who have supported me throughout this degree. I am especially grateful for my partner, Kevin, who always made sure I was well taken care of, and who picked up the slack at home during the busiest portions of this degree.

Finally, I would like to acknowledge the financial support from NSERC, in the form of CGS-M and CGS-D scholarships. These scholarships greatly improved my quality of life, and provided me with financial freedom, while working on this thesis.

*To my caring partner
Kevin Astilla*

CONTENTS

Acknowledgements	iii
Table of Contents	vi
1 INTRODUCTION TO OPTIMAL TRANSPORTATION AND CLIMATE POLICY	1
1.1 Greenhouse Gases and the Role of Regulators	2
1.2 Past Work on Emissions Reduction Policy	4
1.3 Background on Optimal Transportation	7
1.4 Pollution Control for Homogeneous Pollutants	12
1.5 Enforcing Duality through Variance Bounds	18
1.6 Enforcing Duality with Compact Domains	19
1.7 Pollution Control with Heterogeneous Pollutants	20
2 THE DUALITY GAP: FAILURES OF SIMPLISTIC MARKET BASED POLICIES	22
2.1 Summary of Results	22
2.2 Primal and Dual Formulations on General Domains	26
2.3 Solving the Dual Problem on $[0, \infty)^4$	30
2.4 Existence of a Duality Gap on $[0, \infty)^4$	32
3 RESOLVING THE DUALITY GAP WITH VARIANCE BOUNDS: EFFECTIVE MARKET-BASED POLICY THROUGH PROGRESSIVE TAXATION	44
3.1 Summary of Results	44
3.2 Primal and Dual Problems for Variance Bounds	46
3.3 Duality	48
3.4 Primal Attainment and Dual Non-Attainment	54
3.5 Analysis of Variance Bounds	56
3.6 Numerical Methods for the Dual Problem	65
3.7 Other Variance Bound Problems	66
4 RESOLVING THE DUALITY GAP THROUGH COMPACTNESS: EFFECTIVE HYBRID POLICY	69
4.1 Summary of Results	69
4.2 Strong Duality on Compact Domains	69
4.3 Primal and Dual Attainment	73
4.4 Specialization to Squares	79
4.5 Numerical Methods	83

5	MORE THAN CARBON DIOXIDE: A FRAMEWORK FOR GEO- GRAPHICALLY HETEROGENEOUS POLLUTANTS	86
5.1	Optimal Transport Models of Heterogeneous Damages . . .	86
5.2	Duality in the Compact Case	94

PUBLICATIONS

I expect to publish this work as joint publications with Professor Justin Kakeu (University of Prince Edward Island) and Professor Robert McCann (University of Toronto).

INTRODUCTION TO OPTIMAL TRANSPORTATION AND CLIMATE POLICY

The key purpose of this work is to develop mathematical models for low-cost emissions regulation policies, and provide an in-depth study of these models. In doing so, we will draw insights about cases where different policy families may or may not succeed, and we also shed light on which policies are more likely to be successful within a given family. Before doing so, we use this chapter to develop both the economic motivation for the problems we consider, as well as the mathematical background necessary to study them in depth.

We will set the stage for this our work in Section 1.1 by providing some general background on pollutants, most notably greenhouse gases (GHGs) such as carbon dioxide (CO_2), and the effects that these have on the natural environment. In particular, we briefly review of the scientific, economic, and political factors which underscore the importance of pollution reduction for policymakers, and which thereby motivate this current body of work.

In Section 1.2, we review past work on emissions reduction policy. We take particular care to distinguish and explain the development of the two main families of emissions reduction policies: market-based policies, and command-and-control policies. We will also discuss which types of problem are better suited to one policy class over the other, and highlight a broad political shift towards market-based policy over the past half-century.

Next, in Section 1.3, we develop the mathematical background for our work, with a particular focus on duality theory from optimal transportation and infinite-dimensional linear programming. This section also provides the necessary background to contextualize and understand the mathematical novelty of our work.

Section 1.4 builds on the work of Newell and Stavins from [79] to develop a mathematical model for pollution control policies in the case where the pollutant is homogeneously mixed into the atmosphere (for example, CO_2). This problem is geographically homogeneous, in the sense

that policymakers are only concerned with the quantity of emissions, rather than the locations of their sources. This model will be the focus of Chapters 2, 3, and 4, and we provide a high-level overview of our results from those three chapters in Sections 1.4, 1.5, and 1.6, respectively.

Finally, in Section 1.7, we briefly discuss how our models can be adapted to pollutants with heterogeneous geographic effects, such as nitrous oxides or waterborne pollutants, leaving the bulk of this discussion to Chapter 5.

1.1 GREENHOUSE GASES AND THE ROLE OF REGULATORS

Throughout history, economic activity has often generated byproducts, or pollutants, which cause undesirable effects. In the present day, GHGs are the most concerning class of pollutants, and these are primarily produced by fossil fuel consumption, industrial processes, land use practices, forestry practices, and agriculture [84, Figure 2.3]. We refer the reader to the Intergovernmental Panel on Climate Change’s AR6 Synthesis Report [84] for a comprehensive overview of the causes, effects, and projected future effects of climate change, but briefly summarize its findings below as a way of motivating the present work.

To date, greenhouse gas emissions have raised the global average temperature by 1.1°C above the 1850-1900 average, raised sea levels by 0.20 m, and increased the likelihood of extreme weather events such as heatwaves, droughts, flooding, and wildfires. The climate change driven by GHG emissions is currently causing detrimental effects on both oceanic and terrestrial ecosystems, and has exacerbated issues of water scarcity, food supply, disease, and displacement, especially in the developing world [84, Section 2].

The need for action to reduce GHG emissions is further underscored by the projected effects of continued emissions – which are expected to exacerbate the previously mentioned effects, overwhelm natural carbon sinks, acidify and deoxygenate the oceans, and intensify tropical cyclones. From a human perspective, these emissions are projected to cause coastal flooding, decrease food production in some regions, further disrupt ecosystems, and destabilize societies. Some of the most unnerving risks from climate change arise from the potential of driving natural systems to what is known as ‘tipping points,’ where existing negative feedback loops turn positive, leading to irreversible change. For example, climate change could precipitate irreversible biodiversity loss, melting of the Greenland and West Antarctic ice sheets, and sudden disruptions to ocean currents that moderate the world’s weather [84, Section 3].

Fortunately, even though many of these detrimental effects are already baked in, many others can be prevented or mitigated through emissions

reduction efforts. There is a general consensus that the detrimental effects of climate change increase monotonically with the amount of emissions; an immediate corollary is that any reduction in emissions will bring benefits. In particular, emissions reductions will limit, but not eliminate, damages to both the human and natural worlds, and can even lead to positive additional effects, such as improvements in air quality and health [84, Section 4].

Many of the effects of climate change, such as biodiversity loss, are difficult to quantify economically. However, even when looking through a purely economic lens, we find that climate change causes stark negative effects. Depending on the model consulted, a one degree increase in global average temperature is expected to reduce the global gross domestic production (GDP) by anywhere from 1% to over 20% [12, 22, 34, 75, 77, 81]. With the world's estimated GDP estimated at \$111.25 T USD in 2024 [8], this would amount to roughly \$1 to \$22 T USD of annual losses if GDP were to remain constant, and these losses can grow proportionally to global GDP. This underscores that, even from a purely quantitative and economic perspective, climate change is a significant problem. Moreover, it shows that even costly interventions can pay off in the end, so long as they are effective.

As such, there is a clear impetus for climate action on a global level. However, while no individual or firm wishes for the environmental destabilization and disaster which results from unmitigated climate change, certain actors face incentives to continue emitting pollutants such as GHGs. More specifically, it is often the case that the benefits of emissions-generating economic activity accrue directly to the polluter, whereas the harms caused are dispersed, either throughout the polluter's local environment, or globally, as in the case of greenhouse gas emissions. This is a classical case of an economic externality, where a firm's internal accounting does not address the full economic impact of its actions [82, 101, 106]. Under the assumption that a firm's goal is to maximize its own profit, this means that firms are incentivized to emit pollution which causes significant harms (which they only suffer from a small portion of) in order to create a relatively small benefit (which they receive most, if not all, of) [88, Chapter 9]. The presence of such externalities, which are not readily accounted for by market forces, was described by Stern as "the greatest example of market failure we have ever seen" [101, Page 1], and highlights the need for regulatory action to avoid the ensuing harms [90].

From an economic standpoint, reducing GHG emissions can be described as a coordinated action problem [36, 101], which is susceptible to the pitfall of 'free riding,' wherein countries and firms opt not to reduce their emissions, even if reductions lead to preferable global outcomes [101,

Chapter VI]. While collective action at the international level is already a challenging task, requiring the use of international treaties and other mechanisms [101, Chapter VI], it is not the primary focus of this thesis. Rather, we focus on this sort of collective action problem within individual countries, and consider the particular case of industrial regulation, where a regulator strives to meet an emissions target within a prescribed industry. This focus on industrial GHG regulation is especially relevant to Canadian policymakers, given that that country’s federal government has stated that its industrial carbon pricing system is, “expected to deliver more emission reductions than any other policy” [57, Section 1.3].

Within an individual country, cost is an especially salient factor informing policy design – low-cost climate policy allows a government to reallocate money to its other priorities, and stakeholders such as firms and citizens are more likely to support a policy they perceive as cost-effective. Of course, part of this scrutiny arises because the costs of emissions reduction can be substantial – in Canada, for example, it is estimated that firms spent over \$10 billion CAD on environmental protection activities in 2021 alone [26]. These costs can arise from a number of different sources, as firms may need to invest in more efficient equipment, modify supply chains, or even reorganize their business structures in response to emissions reduction policies [44, 97]. In any case, the presence of a substantial mitigation cost is a relevant factor when designing emissions reduction policies, leading policymakers to emphasize the cost, and cost-effectiveness, of their policies. For example, the 2025 Canadian budget asserts that Canada’s climate policies consist of, “measures that will result in the greatest emissions reductions and competitiveness benefits at the lowest cost” [57, Section 1.3]. As such, in light of this policy landscape, it is important for mathematical models to quantify the cost and effects of potential policies.

In any case, with greenhouse gases, as with other pollutants, regulators and other stakeholders find themselves in a complex situation where they must balance global targets with any number of competing interests, which has led to a myriad of emissions control policies being implemented over the years. In the next section, we review past work on industrial emission regulation policy, to contextualize our work.

1.2 PAST WORK ON EMISSIONS REDUCTION POLICY

The purpose of the present work is to introduce the mathematical theory of optimal transport as a tool to shed light on the relative costs of various pollution control policies. However, before we do so, it is useful to outline the current state of regulatory affairs. There are two broad approaches

to emissions regulation which are in use today – command-and-control policies and market-based policies [89]. In a nutshell, command-and-control policies strive to meet a jurisdiction-wide emissions target by setting emissions targets for each firm and providing primarily legal incentives for firms to take action to reduce their emissions [100]. For example, a regulator may assign each firm a pollution emissions cap, require firms to use cleaner technologies, implement a uniform emissions rate standard, or require all polluters to reduce their emissions by a certain percentage, in each case implementing legal or financial penalties for violation [79, 100]. On the other hand, market-based policies aim to meet a jurisdiction-wide emissions target by providing financial incentives to firms. For example, emissions taxes and subsidies, as well as cap-and-trade systems, are considered to be market-based [100].

A. L. Nichols indicates that, as of the 1980s, command-and-control policies were the preferred method of emissions control, and suggests a number of reasons why this was the case – for example, politicians of the time tended not to have the economic training needed to understand market-based policies, elements of the public viewed market-based policies as a ‘permit to pollute,’ and the implementation of such policies would require both data and moral judgments to determine the price of pollution [80, Chapter 1]. However, even at this time, market-based policies tended to be preferred by economists [80, Chapter 1] and, over recent decades, market-based policies have won over policymakers [45, 79, 89, 100].

The key benefit cited by proponents of market-based policies is their increased economic efficiency. In economics, a pollution abatement policy is considered to be economically efficient, if, for each polluter, their marginal abatement cost corresponds to the marginal damages caused by the pollution [80, Chapter 1][1, Section 4.1.1]. In other words, an optimal state is one wherein the cost for each producer to reduce its emissions by one unit below the current level of emissions is equal to the damages caused by one unit of emissions. More concretely, producers that can inexpensively reduce emissions will do so in the face of a financial penalty, whereas those who cannot will simply pay the penalty [80, 89, 100]. As a result, market-based policies can effectively target ‘low-hanging fruit’ and achieve the same goals as command-and-control policies, with less cost to society [80, Chapter 2]. However, it should also be noted that even proponents of market-based policies recognize that there are circumstances where market-based policies are not called for [80, 100]. Nichols, for example, states that, if emissions are to be almost totally eliminated, then almost all control options, from the cheap ones to the most expensive ones, will have to be exercised, reducing the advantage of market-based policies [80, Chapter 2].

An important factor affecting the relative effectiveness of command-and-control and market-based policies is abatement cost heterogeneity, which measures how the cost to reduce emissions from baseline levels varies across firms [79, 80, 97]. The current work is based on the model introduced by Newell and Stavins in [79], which will be introduced in Section 1.4. In this model, the level of cost savings from implementing market-based policies instead of command-and-control policies is a linear function of the level of heterogeneity within an industry [79, Equation (13)]. We also note that, for many pollutants, the amount of damage caused by emitting a unit of pollution is heterogeneous, depending on where the unit was emitted. For example, Nichols asserts that the damages caused by benzene emissions from maleic anhydride plants can vary by a factor of 50, due to factors such as proximity to populated places [80, Chapter 1]. This heterogeneity is discussed in detail in the work of Fowlie and Muller [45], and we address it to some level in Chapter 5. However, our primary focus is in building and studying an undifferentiated model in Chapters 2, 3, and 4 where the damage from one unit of emissions is treated as uniform across polluters, and the only heterogeneity is in the cost of abatement and the quantity of emissions.

It is also important to mention the extensive and rich study of market-based policies through economic equilibrium models. This literature can be distinguished into two veins: partial equilibrium models, where a single commodity market is studied in isolation, and general equilibrium models, where multiple commodities are modelled simultaneously. Partial equilibrium models, which have been studied in, for example [9, 31, 42, 43] are primarily useful in cases where regulation of the studied commodity is unlikely to have overspill effects on other commodities.

General equilibrium models, on the other hand, attempt to model all relevant aspects of the economy, such as related or interchangeable commodities, in order to account for externalities. This theory originated with the work of Walras in the 1870s [105], and its hallmark results include the first and second fundamental theorems of welfare economics: these address cases where socially optimal outcomes can be achieved through decentralized processes. These theorems were developed by a series of economists, most notably Pareto [86] and Lange [67], during the first half of the twentieth century, and proven independently by Arrow [6] and Debreu [32]. The first theorem states that, under suitable assumptions, economic equilibrium (which is attained through market forces) leads to a Pareto optimal outcome (which encapsulates a notion of social optimality, where no agent's position can be improved without weakening that of another). Conversely, the second theorem states that any Pareto optimal outcome can be realized as some competitive equilibrium. Together, these

theorems formalize the notion that socially optimal outcomes can arise from decentralized decision-making, which is a central question about the models of this thesis. For a more comprehensive overview of welfare economics, we refer the reader to [58, Chapter 16].

We now review the history of general equilibrium models of emissions control; a more complete overview and literature review of general and partial equilibrium theory may be found in [63]. General equilibrium, in its modern form, originated in joint work by Arrow and Debreu [7], which Sandmo [95] later applied to optimal taxation theory in order to address a literature gap on the treatment of externalities. This work was adapted to the case of emissions control in the 1990s by a number of authors. In particular, Bovenberg and de Mooij used a general equilibrium framework to study the distortionary effects of environmental taxes in [18], and Bovenberg and van der Ploeg explored the relationship between environmental taxes, employment, and public and private consumption in [19]. Moreover, Nordhaus and Yang built their celebrated RICE model, which models the effects of climate change interventions by disaggregating the world into regions, on general equilibrium theory [83]. General equilibrium theory has not lost steam in the twenty-first century – for instance, Goulder and Williams III highlighted general equilibrium interactions between commodity taxes and the labour market [56], Goulder and Schein studied the relationship between emissions taxes, cap-and-trade systems, and hybrid systems in [55], and Golosov, Hassler, Krusell, and Tsyvinski used a stochastic general equilibrium model to estimate the optimal emissions tax rate in [54]. Anderson and Duanmu continued the comparative study of emissions taxes and cap-and-trade systems in [5]. In particular, this work introduced quota equilibrium and emissions tax equilibrium models, proved a version of the first fundamental theorem of welfare economics for the former, and proved that these notions of equilibrium are not fully equivalent due to multiplicity of equilibria. Ultimately, [5] implies that emissions taxes may be less effective than cap-and-trade systems at meeting a regulator’s target.

In the following section, we complement the economic context for the present work with its mathematical context.

1.3 BACKGROUND ON OPTIMAL TRANSPORTATION

Given that our model is based on optimal transportation, we devote this section to reviewing the foundations and landmark results of that theory. This serves the dual purposes of familiarizing the reader with the mathematical objects which we will work with throughout this thesis, while also providing a point of comparison which will prove especially

interesting when discussing duality results which we will encounter in the remainder of this work. We note to the reader that we only develop this theory insofar as we need it – for more complete references, we refer the reader to [46, 69, 96, 103, 104].

To begin, let X, Y be Polish spaces, and let $c : X \times Y \rightarrow [0, +\infty]$ be lower semicontinuous. Moreover, let $\mathcal{P}(X)$ be the space of probability measures on X (and likewise for $\mathcal{P}(Y)$ and $\mathcal{P}(X \times Y)$). In order to access the optimal transportation problem, we define both pushforward measures and the support of a measure:

Definition 1.1 (Pushforward Measures). Let $\mu \in \mathcal{P}(X)$, and let $T : X \rightarrow Y$ be measurable. Then we can define the pushforward measure $T_{\#}\mu \in \mathcal{P}(Y)$ by the property $(T_{\#}\mu)(A) = \mu(T^{-1}(A))$ for every measurable $A \subseteq Y$.

Definition 1.2 (Support of a Measure; Definition 1.14 of [96]). Let μ be a probability measure on X . We define the support of μ , $\text{spt}(\mu)$, by

$$\text{spt}(\mu) = \bigcap \{A \subseteq X : A \text{ is closed and } \mu(X \setminus A) = 0\}.$$

Let $\mu \in \mathcal{P}(X)$ and $\nu \in \mathcal{P}(Y)$. The Kantorovich Problem in optimal transportation, as originally formulated by Kantorovich in [64], is

$$\inf_{\gamma \in \mathcal{P}(X \times Y)} \left\{ \iint c \, d\gamma \mid \pi_{\#}^X \gamma = \mu, \pi_{\#}^Y \gamma = \nu \right\}, \quad (1.1)$$

where $\pi^X : X \times Y \rightarrow X$ and $\pi^Y : X \times Y \rightarrow Y$ are projection maps defined by $\pi^X(x, y) = x$ and $\pi^Y(x, y) = y$, respectively. This problem has the following folklore interpretation, which is detailed in [103] and helps to explain the terminology ‘optimal transport’: imagine that $X = Y$ represents the land area of a country, $\mu \in \mathcal{P}(X)$ represents how iron mines are distributed across that country, and $\nu \in \mathcal{P}(X) = \mathcal{P}(Y)$ represents the distribution of iron smelters across the country. Moreover, let $c(x, y)$ be the cost to transport one unit of iron ore from a mine located at position x to a smelter located at y . In this case, $\gamma \in \mathcal{P}(X \times Y)$ is known as a ‘transport plan’ and determines how much ore each mine should supply to each factory. Likewise, the quantity $\iint c \, d\gamma$ is the cumulative cost of implementing the transport plan γ and, as such, (1.1) seeks to determine the lowest-cost, or optimal, scheme for transporting ore to smelters. This can be thought of as a centralized problem – one firm or government owns all mines and smelters, and seeks the lowest cost means of transporting iron ore to smelters.

We may ask a number of questions about (1.1). First, is the infimum attained or, equivalently, does there exist an optimal transportation plan? Fortunately, under our standing assumptions, an optimal transport plan

exists [103, Theorem 1.3]. Moreover, we can ask whether the optimal transport plan is unique – this does not hold for general costs [24], but does hold for specific costs that satisfy what is known as a ‘twist condition’ [69, Theorem 2.9]. We may also derive structural results on optimal plans – for example, [69, Theorem 2.9], [96, Theorem 1.43].

Many of these answers are framed in terms of the Monge problem, which is both a historical antecedent of the Kantorovich problem introduced by Gaspard Monge in 1781 [74] and is considered to be both the seminal problem in optimal transportation an interesting problem in its own right. As such, we provide a brief digression introducing the Monge problem and discussing the differences between it and (1.1):

Remark 1.3 (Monge and Kantorovich Problems). We introduce the following Monge problem:

$$\inf \left\{ \int_X c(x, T(x)) d\mu(x) \mid T : \text{spt } \mu \rightarrow Y, T_{\#}\mu = \nu \right\}. \quad (1.2)$$

In short, this can be thought of as a variant of (1.1) where we optimize over maps T instead of transportation plans γ . While (1.2) is still a topic of active research, Kantorovich problems, which are analogous to (1.1), play a dominant role in modern optimal transportation on account of their appealing mathematical properties (for example, linearity in γ) – in fact, many results on (1.2) are deduced through study of (1.1). We refer the reader to [69, Section 1] and [104, Chapter 3] for more complete discussions of the history of optimal transportation especially as it relates to the relationship between (1.1) and (1.2). For now, we remark that the infimum in (1.1) can be no greater than the infimum in (1.2). This is because, given a transportation map T which is admissible in (1.2), the probability measure $\gamma = (\text{id}, T)_{\#}\mu$ is a transportation plan (i.e. it is admissible in (1.1)), and

$$\iint_{[0, \infty)^4} c(x, y) d((\text{id}, T)_{\#}\mu)(x, y) = \int_{[0, \infty)^2} c(x, T(x)) d\mu(x).$$

One of the most important tools for studying optimal transportation problems is Kantorovich-Koopmans duality, which was independently discovered by Kantorovich in [64] and Koopmans in [66]. This result can be derived as a special case of Fenchel-Rockafellar duality for infinite-dimensional linear programming as stated in, for example, [20, Theorem 1.12] or [46, Theorem 6.2.2]. To start, we formally define a dual linear program for (1.1) by

$$\sup_{\substack{\varphi \in C_b(X) \\ \psi \in C_b(Y)}} \left\{ \int \varphi d\mu + \int \psi d\nu \mid \varphi(x) + \psi(y) \leq c(x, y) \text{ on } X \times Y \right\}. \quad (1.3)$$

In this case, C_b refers to the space of continuous bounded functions, and φ and ψ are Lagrange multipliers for the constraints $\pi_{\#}^X \gamma = \mu$ and $\pi_{\#}^Y \gamma = \nu$, respectively. Remarkably, the primal and dual problems have the same value, under some very mild assumptions, and the following is often considered the most important theorem of optimal transport:

Theorem 1.4 (Kantorovich-Koopmans Duality; Theorem 1.42 of [96]). *Let X, Y be Polish spaces, let $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$, and let $c : X \times Y \rightarrow [0, +\infty]$ be lower semicontinuous. Then (1.1)=(1.3).*

This leads to a natural question – do maximizers of (1.3) actually exist? In general, (1.3) lacks some of the compactness properties enjoyed by (1.1), and this adds corresponding difficulty to proving attainment of the dual supremum. Nevertheless, it is possible to derive dual attainment under a wide range of hypotheses: [96, Theorem 1.39] applies if the cost c is uniformly continuous and bounded, and [104, Theorem 5.10(iii)] applies if the infimum (1.3) is finite and $c \leq c_X \oplus c_Y$ for some $(c_X, c_Y) \in L^1(\mu) \times L^1(\nu)$.

Kantorovich-Koopmans duality, as described in Theorem 1.4, becomes especially useful in the case of attainment, as in this case we may use optimal solutions (φ, ψ) to (1.3) to derive powerful structural results on, or even characterize, solutions of (1.1) – some powerful results for quadratic costs are summarized in [103, Theorem 2.12]. Moreover, (1.3) is often a much more tractable optimization problem than (1.1), as it involves optimizing over $C_b(X) \times C_b(Y)$ rather than the measure space $\mathcal{P}(X \times Y)$.

Kantorovich-Koopmans duality is also very powerful from an economic perspective. In the previously mentioned context of coal mines and smelters, Villani shares the following interpretation of the dual problem (1.3), which he credits to Caffarelli [103, Section 1.1.3]. In this case, rather than the central planner transporting ore from mines to smelters at cost $c(x, y)$, contractors enter the market, setting a price $\varphi(x)$ to pick up one ton of ore from a mine located at position x and a separate price $\psi(y)$ to drop a ton of ore at a smelter at position y . The condition $c(x, y) \geq \varphi(x) + \psi(y)$ makes this an appealing bargain to the central planner – it ensures that no additional costs are incurred as a result of outsourcing shipping to contractors. Conversely, Kantorovich-Koopmans duality (Theorem 1.4) ensures that this scheme is also a good deal for the contractors, as they may raise prices to a level where the central planner spends the same (or nearly the same) amount it would have spent shipping the material itself.

In other words, Theorem 1.4 states that it is equally cost-efficient for the owner of the mine and smelters to transport ore itself in a centralized system, or using contractors in a decentralized system of prices and independent agents. As a result of this decentralization, Kantorovich-

Koopmans duality can be interpreted as a welfare theorem of economics [10, 49, 53, 98], analogous to the welfare theorems discussed in the previous section. Note, however, that the notion of social optimality differs from Pareto optimality, which is typically used in the general equilibrium models discussed in Section 1.2. In particular, the central planner seeks an outcome which minimizes the aggregate cost $\iint c d\gamma$, providing a more utilitarian approach to social optimality than Pareto optimality, which only requires that no firm's position can be improved without worsening the position of another.

Duality results, and their absence, form much of our focus the remainder in this thesis. In particular, as we shall see in Chapter 2, the model we introduce in 1.4 differs from the Kantorovich problem (1.1) in that duality only holds under some fairly restrictive assumptions, rather than the broad assumptions of Theorem 1.4. We propose two strategies to solve this issue in Chapters 3 and 4, by tweaking our model to ensure a broader application of duality. Finally, in Chapter 5, we will adapt our model to account for the geographical heterogeneity which we will discuss in Section 1.7.

We conclude this section by reviewing some of the applications of optimal transport, with a focus on economic applications. Within economics, the results and techniques of optimal transport have been applied to matching problems in marriage [10, 33, 52], labour [13, 14, 68, 71], and hedonic markets [29, 98], principal agent or monopolist problems [41, 70, 72, 91], international trade [4, 37, 99], discrete-choice models [15, 28, 30, 52], derivative pricing [50, 60], econometrics [48], auction theory [65], and many other areas. We also refer the reader to [38, 47, 49, 51] for additional surveys of the applications of optimal transportation to economics.

Within the particular context of environmental economics, optimal transport has been applied to model environmental damages caused by road networks [40], study how to encourage consumers to opt for less carbon-intensive modes of transportation [62], use sensitivity analyses to draw new insights from existing energy system optimization models [17, 16], and even model transactions within renewable energy markets [59]. To our knowledge, the present work is the first to apply optimal transportation techniques to study emissions regulation policies, and in particular the efficiency of carbon markets. The closest point of comparison is [3, Chapter 4], which introduced a bilevel linear programming problem to optimize emissions tax levels.

1.4 POLLUTION CONTROL FOR HOMOGENEOUS POLLUTANTS

As alluded to earlier, the starting point of our mathematical study of emissions control policies is the cost function $c : (0, \infty)^2 \times (0, \infty)^2 \rightarrow [0, +\infty)$ defined by

$$c(x_1, x_2, y_1, y_2) = \frac{y_2}{2x_2} \left(x_1 - \frac{y_1}{y_2} \right)^2. \quad (1.4)$$

This cost function first appeared in the work of Newell and Stavins in [79, Section 2.1], where those authors studied heterogeneous pollution abatement costs among firms, with the goal of developing environmental policies which account for this heterogeneity. The key issue here is that, given an environmental target within a particular industry, the cost to attain that target will likely vary among companies. Newell and Stavins identify a number of factors which can affect the cost for a given firm to meet a specified environmental target, including its factories' locations, sizes, ages, and production technologies, as well as local regulations [79, Section 1]. In the context of Equation (1.4), the vector (x_1, x_2, y_1, y_2) represents the state of a single firm, both before and after regulation. In particular, x_1 represents the firm's baseline emissions intensity, or emissions per unit of production, x_2 is measures how easily the firm can change its pre-regulation emissions output, y_1 represents the emissions produced by a given firm post-regulation, and y_2 represents the firm's post-regulation production [79, Section 2.1]. Implicit in this choice of cost function are a number of simplifying assumptions, which we discuss in the following remark:

Remark 1.5 (Derivation and Assumptions for the Newell-Stavins Cost). The discerning reader will notice that the cost c defined in Equation (1.4) only depends on the firm's pre-regulation characteristics (x_1, x_2) – in particular, the structure of the cost function is not firm-dependent. While it is possible for firms to have entirely different cost structures, Newell and Stavins [79] reduce to the formula (1.4) by using a second-order Taylor expansion, an assumption that firms' initial cost is zero, demand levels remain unchanged, and finally an assumption that, in the absence of regulation, firms will have minimized their cost.

It is also worth mentioning how one might empirically determine both the emissions intensity x_1 and the ease of change parameter x_2 for real-life firms.

Remark 1.6 (Obtaining Pre-Regulation Data). The emissions intensity quantity x_1 may be determined from firms' current production and emissions levels. These can be obtained by imposing reporting standards on firms; a

more sophisticated analysis may be facilitated by adapting the methods of [61] from countries to firms. The quantity x_2 is significantly more delicate, as it depends on firms' production methods, locations, sizes, and ages [79]. One could attempt to write x_2 as an explicit function of these (and possibly other) factors, which each may be determined through mandatory reporting, and subsequently use this to generate a database of (x_1, x_2) -values for firms. Alternatively, one could attempt to estimate x_2 -values exogenously, perhaps by measuring the industrial costs of past emissions reduction efforts. In the original work, Newell and Stavins are satisfied with estimating the average value and variance for x_2 , as they do while studying nitrogen oxide emissions by electric utilities in [79, Table 1].

Now that we have provided background on the cost function c from Equation (1.4), we observe that limiting this cost function only on the domain $(x, y) \in (0, \infty)^2 \times (0, \infty)^2$ is restrictive – for example, it does not allow for us to shut down firms by setting their post-regulation emissions and production to 0. As such, we extend c by lower semicontinuity: $x \in [0, \infty)^2$ and $y \in [0, \infty)^2$:

Remark 1.7 (Extended Cost Function). Define a lower semicontinuous function $c : [0, \infty)^2 \times [0, \infty)^2 \rightarrow [0, +\infty]$ by

$$c(x_1, x_2, y_1, y_2) = \begin{cases} \frac{y_2}{2x_2} \left(x_1 - \frac{y_1}{y_2}\right)^2 & \text{if } x_2, y_2 > 0 \\ 0 & \text{if } y_1 = y_2 x_1 \\ +\infty & \text{else} \end{cases} \quad (1.5)$$

In effect, this cost function states that entirely shutting down a producer to meet emissions targets has zero cost if both the output and the emissions of the company become zero, and infinite cost if the company's output becomes zero but its emissions remain positive. In much of the sequel, we effectively ignore the case where $x_2 = 0$ by noting that, under the model proposed by Newell and Stavins in [79], $x_2 = 0$ corresponds to the case of a firm which finds it infinitely hard to adjust its emissions intensity, even by a minuscule amount.

Let $X \subseteq [0, \infty)^2$ and $Y \subseteq [0, \infty)^2$ be domains of interest. Using this notation, we model the heterogeneous state of a given industry, both before and after the implementation of a given regulatory framework, by the probability measure $\gamma \in \mathcal{P}(X \times Y)$. Since the variables $(x_1, x_2) \in X$ reflect the pre-regulation state of the industry, we impose the requirement that $\pi_{\#}^X \gamma = \mu \in \mathcal{P}_c(X)$, where π^X is the projection map from $X \times Y$ to X and μ is a fixed, compactly supported probability measure representing the industry's status quo. On the other hand, the restrictions on the post-regulation industry, which is described by the probability measure $\pi_{\#}^Y \gamma$

(where π^Y is the projection map from $X \times Y$ to Y) are much less stringent. More precisely, it can be assumed that the regulator wishes to enforce an average emissions standard m_1 , as well as an average production level m_2 , where the former is designed to meet the regulator's emissions goals, and where the latter is designed to ensure that the industry cannot respond to the emissions standards by scaling down production to an unacceptable level. As such, these emissions and production targets can be accounted for by the constraint $\iint_{X \times Y} y \, d\gamma = m$.

The previously mentioned constraints are reminiscent of the classical optimal transportation problem (1.1). In particular, as in that problem, we aim to optimize $\iint_{X \times Y} c \, d\gamma$ subject to the constraint that $\pi_{\#}^X \gamma = \mu$. However, for this problem, we only require that the Y -marginal of γ have fixed mean $m \in (0, \infty)^2$, leading to the following formalized problem:

$$\mathcal{I}_{X,Y}(\mu, m) := \inf_{\gamma \in \mathcal{P}(X \times Y)} \left\{ \iint_{X \times Y} c \, d\gamma \mid \pi_{\#}^X \gamma = \mu, \iint_{X \times Y} y \, d\gamma = m \right\}. \tag{1.6}$$

We describe (1.6) as a ‘Kantorovich-type’ problem in analogy with the classical optimal transportation problem (1.1). If $X = Y = [0, \infty)^2$, we will often suppress subscripts to write $\mathcal{I}(\mu, m) := \mathcal{I}_{[0, \infty)^2, [0, \infty)^2}(\mu, m)$.

We quickly notice that, as in the case of the classical optimal transportation problem (1.1), our problem (1.6), is an infinite-dimensional linear program in γ . In particular, this allows us to introduce a weak dual linear program as follows:

Proposition 1.8. *Let $X \subseteq [0, \infty)^2, Y \subseteq [0, \infty)^2, \mu \in \mathcal{P}(X)$, and $m \in [0, \infty)^2$. Consider the linear program $\mathcal{I}_{X,Y}(\mu, m)$ as defined in Equation (1.6). The (weak) linear programming dual problem to $\mathcal{I}_{X,Y}(\mu, m)$ is*

$$\mathcal{S}_{X,Y}(\mu, m) := \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2}} \left\{ \int_X \varphi \, d\mu + \lambda \cdot m \mid c(x, y) \geq \varphi(x) + \lambda \cdot y, \forall x, y \right\}. \tag{1.7}$$

In particular, $\mathcal{I}_{X,Y}(\mu, m) \geq \mathcal{S}_{X,Y}(\mu, m)$. As with the primal infimum, we define $\mathcal{S}(\mu, m) := \mathcal{S}_{[0, \infty)^2, [0, \infty)^2}(\mu, m)$.

The relation between these primal and dual problems will be the subject of much of this thesis, so it will be useful to rigorously introduce terminology around duality.

Definition 1.9 (Duality). Let $X, Y \subseteq [0, \infty)^2$, let $\mu \in \mathcal{P}([0, \infty)^2)$, and let $m \in [0, \infty)^2$. We say that **weak duality** holds if

$$\mathcal{I}_{X,Y}(\mu, m) \geq \mathcal{S}_{X,Y}(\mu, m),$$

and **strong duality** holds in the case of equality. If weak duality holds but strong duality fails, we say that there is a **duality gap**. Moreover, we say that **duality holds with attainment** if there exists $\gamma_0 \in \mathcal{P}(X \times Y)$ satisfying the constraints of (1.6) with $\mathcal{I}_{X,Y}(\mu, m) = \int \int_{X \times Y} c \, d\gamma$ and there exist $(\varphi, \lambda) \in C_b(X) \times \mathbb{R}^2$ satisfying the constraints of (1.7) with $\mathcal{S}_{X,Y}(\mu, m) = \int \varphi \, d\mu + \lambda \cdot m$.

We postpone the proof of Proposition 1.8 to Section 2.2. In the meantime, we make the following observation:

Remark 1.10. Proposition 1.8 states that weak duality holds for $\mathcal{I}_{X,Y}(\mu, m)$ and $\mathcal{S}_{X,Y}(\mu, m)$.

We now interpret the meaning of (1.7). First, we interpret $-\lambda_1$ as a tax on each unit of emissions, and $-\lambda_2$ as a subsidy on each unit of production, leading to the heuristic expectation that $\lambda_1 \leq 0$ and $\lambda_2 \geq 0$. As such, we interpret the quantity $c(x, y) - \lambda \cdot y$ as the cost for a firm with pre-regulation characteristics x to adopt post-regulation characteristics y under the policy λ . If a firm with pre-regulation characteristics x is rational, we expect it to (formally) choose post-regulation emissions and production levels y such that

$$y \in \operatorname{argmin}_{y \in Y} \{c(x, y) - \lambda \cdot y\}$$

incurring cost

$$\Phi(x, \lambda) := \inf_{y \in Y} [c(x, y) - \lambda \cdot y]. \tag{1.8}$$

As we will see in Chapter 2, the requirement that $\varphi \in C_b(X)$ means that it is not always possible for firms to make, or even approximate, this kind of optimal choice, leading to a duality gap in the sense of Definition 1.9. We also note the relation between the present work and economic equilibrium theory as discussed in Section 1.2

Remark 1.11. The model formed by (1.6) and (1.7) bear some similarities to partial equilibrium models in that they only concern a limited number of commodities (emissions and production). As such, a policymaker who wishes to use these models to inform policy should be conscious of the potential for negative externalities. Moreover, at present, our model does not allow for comparative modelling of cap-and-trade policies and emissions taxes as in the general equilibrium model [5, 55]. However, the purpose of this model differs from that of general equilibrium theory – while the latter strives to model the economy writ large, ours provides advice to a policymaker, such as a cabinet minister or the head of a regulatory body, seeking to implement a preset target. Especially when viewed in light of this context, our model enjoys a number of significant advantages. First,

the explicit modelling of both emissions and production allows our model to avoid the risks around badly chosen emissions quotas identified in [5, Section 1] – in effect, it prevents an industry from shutting down in response to regulation, which is especially important for critical industries such as agriculture, home heating, electricity, and transportation. Second, our model explicitly accounts for abatement costs through the x_2 parameter, which has been identified as a relevant factor for the design of emissions regulation policies by, for example, [25, 79, 80, 97, 107, 108, 109]. Finally, as we will see in Sections 3.6 and 4.5, our model lends itself well to computation; this is a known challenge in general equilibrium modelling [35].

It is also interesting to define a Monge problem associated to (1.6), which corresponds to the case where firms have production levels pre-determined by their initial characteristics (i.e. the case where it is not possible for a single firm to split its production into multiple production methods).

Remark 1.12 (Monge and Kantorovich Problems). We introduce the following Monge problem:

$$\inf \left\{ \int_X c(x, T(x)) d\mu(x) \mid T : \text{spt } \mu \rightarrow Y, \int_X T(x) d\mu(x) = m \right\}. \quad (1.9)$$

In short, this is an optimization over transportation maps T instead of transportation plans γ , and is analogous to the Monge problem in classical optimal transportation discussed in Remark 1.3. For now, we remark that the infimum in (1.6) can be no greater than the infimum in (1.9) – given a transportation map T which is admissible in (1.9), the probability measure $\gamma = (\text{id}, T)_\# \mu$ is a transportation plan (i.e. it is admissible in (1.6)), and

$$\iint_{[0, \infty)^4} c(x, y) d((\text{id}, T)_\# \mu)(x, y) = \int_{[0, \infty)^2} c(x, T(x)) d\mu(x).$$

As we have previously foreshadowed, the duality between (1.6) and (1.7) is much more intricate than Kantorovich-Koopmans duality from Theorem 1.4, so we elaborate on the various cases of duality in the unbounded case $X = Y = [0, \infty)^2$. Proposition 2.1 shows that, in this case, the dual supremum from (1.7) satisfies $\mathcal{S}(\mu, m) = 0$, and moreover, this supremum is attained regardless of the baseline distribution μ or the mean target m . This implies that the presence of a duality gap hinges entirely upon the value of the primal infimum; Theorem 2.3 and Propositions 2.4, 2.5, 2.6, and 2.7 use properties of μ and m to completely classify cases where duality holds and fails, as well as cases where the primal infimum $\mathcal{I}(\mu, m)$ is 0 attained by some $(\mu, m) \in C_b([0, \infty)^2) \times \mathbb{R}^2$. As we will discuss in Section 2.4 market-based policies cannot be expected to meet the emissions

intensity target m in the absence of additional constraints. As such, we use the following two sections, as well as Chapters 3 and 4, to introduce two possible resolutions to the duality gap. We also notice that this result differs from [5, Theorem 1], which categorically asserts the existence of a quota equilibrium. We attribute this discrepancy to the fact that our model explicitly contains a production quota to prevent firms from shutting down essential production; Anderson and Duanmu address this problem through a notion of constrained Pareto optimality and [5, Theorem 2], which implicitly cautions regulators against setting unduly onerous emissions targets.

We close this section by contextualizing (1.6) and (1.7) within the existing literature. The primal problem (1.6) is not the first to augment or modify the optimal transportation problem with marginal constraints. Perhaps the closest mathematical analogue is the problem termed $\mathcal{1S}$ -MCOT in [23, Appendix A] which, in the context of electric vehicle charging, introduces a variant of the optimal transport problem where one marginal is prescribed and the other is taken to satisfy a finite set of moment constraints. Notably, the authors are able to prove strong duality on compact sets by using the Kullback-Leibler divergence to define a regularized problem. However, to our knowledge, the present work is the first to examine which types of marginal constraints can ensure duality on unbounded sets; in Chapter 2, we show that a mean constraint is insufficient to ensure strong duality, whereas in Chapter 3 we show that an additional variance constraint is. Moreover, we also present an alternative proof of duality for compact domains, with scope limited to our marginal constraints of interest, in Chapter 4. It should also be noted that Carlier, Malamut, and Sylvestre discuss this kind of moment constraint problem in [27], citing duality as a particularly delicate challenge.

Moment constraints in optimal transport style problems have also arisen in a number of other works. For instance, ZaeV augmented (1.1) with additional moment constraints in order to study applications in statistical physics and finance [110]. Likewise, (1.1) has been approximated by analogous problems with prescribed moments but otherwise unrestricted marginals in [2, 76].

Finally, we note that, while duality gaps are well-known in the literature (see, for example, [11]), the duality gaps we find in Propositions 2.6 and 2.7 are of interest due to their provenance from a relatively simple economic model – most known duality gaps occur in the presence of pathological costs.

1.5 ENFORCING DUALITY THROUGH VARIANCE BOUNDS

Chapter 3 provides a satisfying cure for the duality gap we identified in Section 2.4. The key result of this chapter is Theorem 1.14 which, in plain language, asserts that the duality gap may be fixed by implementing a progressive tax on emissions; in mathematical terms, this corresponds to placing a variance bound τ on the candidate probability measures from (1.6).

More precisely, given Polish subsets $X, Y \subseteq [0, \infty)^2$, a pre-regulation industry structure $\mu \in \mathcal{P}(X)$, a post-regulation target mean $m \subseteq (0, \infty)^2$, and $\tau > 0$, we introduce the following variance bound problem:

$$\mathcal{I}_{X,Y}(\mu, m; \tau) := \inf_{\gamma \in \mathcal{P}(X \times Y)} \left\{ \iint c \, d\gamma \mid \pi_{\#}^X \gamma = \mu, \overline{\pi_{\#}^Y \gamma} = m, \iint |y - m|^2 \, d\gamma \leq \tau \right\}. \quad (1.10)$$

As with the primal problem (1.6), this problem corresponds to a central planner aiming to meet the emissions-production target m , albeit with the new constraint that $\iint |y - m|^2 \, d\gamma \leq \tau$. In effect, this constraint moderates firms' responses by penalizing excessive heterogeneity – heuristically speaking, it can be expected to prevent extreme responses such as firm closures and unrealistic increases in production.

As the following proposition shows, (1.10) also has a dual problem:

Proposition 1.13. *Let $X, Y \subseteq (0, \infty)^2$ be Polish, let $\mu \in \mathcal{P}(X)$, $m \in [0, \infty)^2$, and $\tau \in [0, \infty)$. Define $\mathcal{S}_{X,Y}(\mu, m; \tau)$ by*

$$\mathcal{S}_{X,Y}(\mu, m; \tau) := \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2 \\ \eta \leq 0}} \left\{ \int \varphi \, d\mu + \lambda \cdot m + \eta \tau \mid c(x, y) - \varphi(x) - \lambda \cdot y - \eta |y - m|^2 \geq 0 \text{ on } X \times Y \right\}. \quad (1.11)$$

Then $\mathcal{I}_{X,Y}(\mu, m; \tau) \geq \mathcal{S}_{X,Y}(\mu, m; \tau)$, and hence weak duality holds.

As before, we interpret $-\lambda_1$ as an emissions tax, and $-\lambda_2$ as a production subsidy. Moreover, in this case, the cost for a firm with pre-regulation characteristics x to adopt the post-regulation emissions-production combination y is given by

$$c(x, y) - \varphi(x) - \lambda_1 \cdot y_1 - \lambda_2 \cdot y_2 - \eta |y - m|^2.$$

Thus, we interpret $-\eta$ as a tax levied on the squared distance $|y - m|^2$ between the firm's post-regulation characteristics y and the regulator's

average target m . In this sense, as the tax bill scales up more than linearly with the distance from the regulatory objective, the tax $-\eta$ behaves broadly similarly to progressive income taxation, where the marginal tax rate increases with earnings. We also notice that, as the only policy instruments in (1.11) are taxes and subsidies, this dual problem may be interpreted as purely market-based policy.

The key result of Chapter 3, and indeed of this entire work, is the following theorem, which states that $\mathcal{I}(\mu, m; \tau) = \mathcal{S}(\mu, m; \tau)$ under some very mild assumptions:

Theorem 1.14 (Duality in the Presence of Variance Bounds). *Let $X = Y = [0, \infty)^2$, $\mu \in \mathcal{P}_c([0, \infty) \times (0, \infty))$, $m \in (0, \infty)^2$, and $\tau \geq 0$. Then*

$$\mathcal{I}(\mu, m; \tau) = \mathcal{S}(\mu, m; \tau).$$

In effect, this result can be interpreted as a welfare theorem, as it states that a regulator may use either command-and-control policy (as in $\mathcal{I}(\mu, m; \tau)$) or market-based policy (as in $\mathcal{S}(\mu, m; \tau)$) to achieve its aims. As with the Kantorovich-Koopmans duality of Theorem 1.4, the notion of optimality differs from Pareto optimality; in this case, an outcome is optimal if it minimizes the aggregate cost to industry. In policy terms, Theorem 1.14 also provides support to the idea that progressive emissions taxes, such as those proposed in [102], are better at achieving policy goals than flat-rate emissions taxes.

We now provide a brief summary of the material to be covered in Chapter 3, which will be expanded somewhat in Section 3.1. First, Section 3.2 introduces the primal problem (1.10) and dual problem (1.11), and proves the weak duality of Proposition 1.13. Then, Section 3.3 uses Fenchel-Rockafellar duality, applied to a suitable Banach space, to upgrade this to the strong duality of Theorem 1.14. We then turn our attention to questions of attainment in Section 3.4, which shows that the primal infimum (1.10) is attained, but provides an example where the dual supremum (1.11) is not attained. With this restriction in mind, we use Section 3.5 to explore the Euler-Lagrange Karush-Kuhn-Tucker (ELKKT) conditions for (1.10) and (1.11), ultimately providing a nonlinear system (3.12) which could ensure primal and dual attainment, should a solution exist. Next, we briefly discuss numerical methods in Section 3.6, and conclude by providing some alternative reformulations of (1.10) and (1.11) in Section 3.7.

1.6 ENFORCING DUALITY WITH COMPACT DOMAINS

Chapter 4 provides an alternative method for curing the duality gap identified in Chapter 2. Rather than nudging firms towards a target with

a progressive tax η , a regulator may instead place an *a priori* limit on firms' post-regulation emissions and production levels by requiring that $Y \subseteq [0, \infty)^2$ be compact. This approach yields the following duality result, which is the highlight of Chapter 4:

Theorem 1.15. *Let $X \subseteq [0, \infty) \times (0, \infty)$ and $Y \subseteq [0, \infty)^2$ be Polish spaces, and assume moreover that Y is compact. Fix $\mu \in \mathcal{P}_c(X)$ and $m \in [0, \infty)^2$. Then strong duality holds, in the sense that $\mathcal{I}_{X,Y}(\mu, m) = \mathcal{S}_{X,Y}(\mu, m)$.*

Continuing with the interpretation of duality that we developed in Chapters 2 and 3, this theorem suggests that well-designed hybrid policies, which contain elements of both command-and-control and market-based policy, can meet a regulator's objective. In particular, while $\mathcal{S}_{X,Y}(\mu, m)$ is typically thought of as market-based policy, the regulator is also limiting post-regulation outcomes to a predefined compact set Y , which is reminiscent of command-and-control policy.

We build off this duality result in Section 4.3, where we consider questions of primal and dual attainment. Fortunately, under some relatively weak hypotheses, we are able to prove Proposition 4.3, where we show that the primal infimum $\mathcal{I}_{X,Y}(\mu, m)$ is attained by some probability measure $\gamma \in \mathcal{P}(X \times Y)$, and Corollary 4.10, which shows that the dual supremum $\mathcal{S}_{X,Y}(\mu, m)$ is attained by some $(\varphi, \lambda) \in C_b(X) \times \mathbb{R}^2$.

This attainment has positive implications for the development of numerical methods; we set the stage for these in Section 4.4 by studying the specific case $Y = [0, R]^2$, and complete a brief foray into numerical methods in Section 4.5.

1.7 POLLUTION CONTROL WITH HETEROGENEOUS POLLUTANTS

In Chapter 5, we suggest how to extend the study in the previous chapters to cover optimal emissions control policies in the context of heterogeneous damages. As we discussed in Section 1.1, economic activity generates pollutants or by-products, such as GHGs which have negative effects on people not directly affiliated with the polluter. GHGs are often assumed to dissipate uniformly into the atmosphere and cause *homogeneous* damages, in the sense that the negative effects caused by a single unit of pollution are independent of the location where it was emitted, and this simplifying assumption underpins our work in Chapters 2, 3, and 4.

However, many pollutants do not dissipate uniformly into the atmosphere, and instead cause heterogeneous damages, or damages that depend on the location where they are emitted. Fowlie and Muller [45] cite nitrous oxides (NO_x) and sulfur dioxide (SO_2) as examples of pollutants which locally affect air quality. In another example, Nichols [80] highlights

the case of benzene, an airborne carcinogen which is emitted as a byproduct of maleic anhydride production. In this case, people who live close to benzene emitters are far more likely to develop cancers such as leukemia than those who do not, placing a premium on locating benzene emitters far away from centres of population. The heterogeneity may be even more stark for pollutants which are released into the ground or the water, rather than the air. For example, environmental historian Kate Brown has identified significant heterogeneity in the damages caused by plutonium production within the United States and the Soviet Union. Namely, the American Hanford plant released radioactive byproducts into the large and quick-flowing Columbia river, which limited the amount of damage to people downstream, whereas the Soviet Mayak plant released radioactive materials into a sluggish and meandering river, causing serious effects for villagers living downstream [21].

To this end, in Chapter 5, we introduce three distinct models for emission harm reduction efforts for heterogeneously spread pollutants, which model various assumptions about whether firms may change location and, if so, how costly it is to do so. As with our work in earlier chapters, each model has a corresponding primal and dual problem, and we prove duality for the most general model, under suitable continuity and compactness assumptions, in Section 5.2. However, we highlight that Chapter 5 as a whole should be taken as a starting point describing some interesting extensions of the homogeneous model studied in Chapters 2, 3, and 4, rather than a complete theory of these extensions.

THE DUALITY GAP: FAILURES OF SIMPLISTIC MARKET BASED POLICIES

2.1 SUMMARY OF RESULTS

In this section, we study the primal and dual problems (1.6) and (1.7) on the unbounded domains $X = [0, \infty)^2$ and $Y = [0, \infty)^2$, and determine conditions on μ and m which either allow duality to hold or cause it to fail. As a broad generalization, these results arise from independently studying the primal infimum $\mathcal{I}(\mu, m)$ and the dual supremum $\mathcal{S}(\mu, m)$.

Section 2.2 is devoted to the proof of Proposition 1.8, which states that, for general subsets $X, Y \subseteq \mathbb{R}^2$, $\mathcal{I}_{X,Y}(\mu, m) \geq \mathcal{S}_{X,Y}(\mu, m)$, or that weak duality holds in the sense of Definition 1.9.

After this, we specialize to the case $X = Y = [0, \infty)^2$. In Section 2.3, we prove that, for this choice of X and Y , the dual problem is effectively trivial in the sense described in the following proposition:

Proposition 2.1 (Dual Triviality). *Fix $\mu \in \mathcal{P}([0, \infty)^2)$ and $m \in [0, \infty)^2$. Then, for the dual problem (1.7),*

$$\mathcal{S}(\mu, m) = 0,$$

and the supremum is attained by $\varphi \equiv 0$, $\lambda = (0, 0)$.

The preceding result, while perhaps surprising, is largely a consequence of the choice of Y and the fact that $c(x, y) - \lambda \cdot y$ is a positively 1-homogeneous function of y .

As Proposition 2.1 effectively reduces questions of duality and attainment to questions about the value and attainment of the primal infimum (1.6), we switch our focus to studying $\mathcal{I}(\mu, m)$ in Section 2.4, which forms the bulk of this chapter.

To study this primal infimum, the following notation will prove convenient:

Definition 2.2 (Emissions Viability Frontiers (EVFs)). Let $\mu \in \mathcal{P}_c([0, \infty)^2)$. Then, for $i = 1, 2$ we define

$$\underline{x}_i := \inf_{(x_1, x_2) \in \text{spt}(\mu)} x_i \quad \text{and} \quad \bar{x}_i := \sup_{(x_1, x_2) \in \text{spt}(\mu)} x_i.$$

We call \underline{x}_1 the **low emissions viability frontier** (LEVF) and \bar{x}_1 the **high emissions viability frontier** (HEVF).

We interpret the emissions viability frontiers as reflecting technological or economic constraints on that industry’s production methods – if a given level of emissions intensity is realistic in an industry, then there should be some firm which attains or exceeds it. In other words, it may simply be the case that, in order to produce one unit of a given good, all known production methods emit at least \underline{x}_1 units of pollution. From a mathematical standpoint, as we show in the following figure, EVFs define a cone-shaped region of Y corresponding to post-regulation outcomes (y_1, y_2) with emissions intensity $\frac{y_1}{y_2}$ lying between \underline{x}_1 and \bar{x}_1 .




Figure 2.1: The Emissions Viability Frontiers from Definition 2.2. The left image illustrates the definition of \underline{x}_1 and \bar{x}_1 for a given probability measure μ , and the right image presents a cone of viable post-regulation outcomes.

More rigorously, Definition 2.2 provides conditions on μ and m which characterize whether the primal infimum $\mathcal{I}(\mu, m)$ is positive (meaning that duality fails between (1.6) and (1.7)) or zero (meaning that duality holds). It will also help us determine if the primal infimum $\mathcal{I}(\mu, m)$ is attained or not. Our first result in this direction is the following, which characterizes when the primal infimum is both zero and attained:

Theorem 2.3 (Characterizing Primal Attainment in $\mathcal{P}_c([0, \infty)^4)$). *Given $\mu \in \mathcal{P}_c([0, \infty) \times (0, \infty))$ and $m \in (0, \infty)^2$, the infimum $\mathcal{I}(\mu, m) = 0$ and is attained by a measure $\gamma \in \mathcal{P}_c([0, \infty)^4)$ if and only if $\underline{x}_1 \leq m_1/m_2 \leq \bar{x}_1$ and $\min\{\mu(x_1 \geq m_1/m_2), \mu(x_1 \leq m_1/m_2)\} > 0$.*

Furthermore, we may split the cases of attainment in Theorem 2.3 into three cases, depending on whether the primal infimum is attained by a probability measure with compact support in $(0, \infty)^4$ or not, as we do in the following proposition:

Proposition 2.4 (Characterizing Primal Attainment in $\mathcal{P}_c((0, \infty)^4)$). *Assume that $\mu \in \mathcal{P}_c([0, \infty) \times (0, \infty))$ and $m \in (0, \infty)^2$ are such that $\underline{x}_1 \leq m_1/m_2 \leq \bar{x}_1$ and $\min\{\mu(x_1 \geq m_1/m_2), \mu(x_1 \leq m_1/m_2)\} > 0$. Then*

- *if $\underline{x}_1 = \bar{x}_1$ or $x_1 < m_1/m_2 < \bar{x}_1$, there exists a measure $\gamma \in \mathcal{P}_c([0, \infty) \times (0, \infty)^3)$ attaining the infimum of 0 in $\mathcal{I}(\mu, m)$.*
- *if neither of the preceding conditions holds, no measure in $\mathcal{P}_c([0, \infty) \times (0, \infty)^3)$ attains the infimum of 0 in $\mathcal{I}(\mu, m)$.*

In particular, Proposition 2.4 yields three cases. The first, where $\underline{x}_1 = \bar{x}_1$ corresponds to an industry where all production methods have the same emissions intensity \underline{x}_1 and the regulator imposes a target m with the same emissions intensity as the industry's production methods. Thus, in order to meet the target m , the industry needs only to scale production to the level specified by m_2 (and emissions to the corresponding level specified by $m_1 = \underline{x}_1 m_2$).

From a practical point of view, the preceding case may be thought of as a simpler version of the second case, $\underline{x}_1 < m_1/m_2 < \bar{x}_1$. In this case, some firms employ production methods which have lower emissions intensity than the emissions intensity target $\frac{m_1}{m_2}$, whereas others use dirtier methods. As a result, command-and-control methods can attain the environmental target m by using a map T to assign high production levels to clean producers and low, but non-zero, production methods to dirty producers. In other words, the regulator can attain the emissions target without making any firm shut down or change the emissions intensity of its production methods. We leave the mathematical mechanisms for doing so, including the formal definition of T , to the proof of Proposition 2.4 in Section 2.4, but illustrate these results in Figure 2.2.

The final case of interest highlighted by Proposition 2.4 is the case where $\underline{x}_1 = m_1/m_2 < \bar{x}_1$. Here, $\mathcal{I}(\mu, m)$ is attained by a compactly supported probability measure in $\mathcal{P}_c([0, \infty)^4)$ but not by any measure in $\mathcal{P}_c((0, \infty)^4)$. This means that any command and control policy which meets the target m without changing the emissions intensity of any firm must shut some

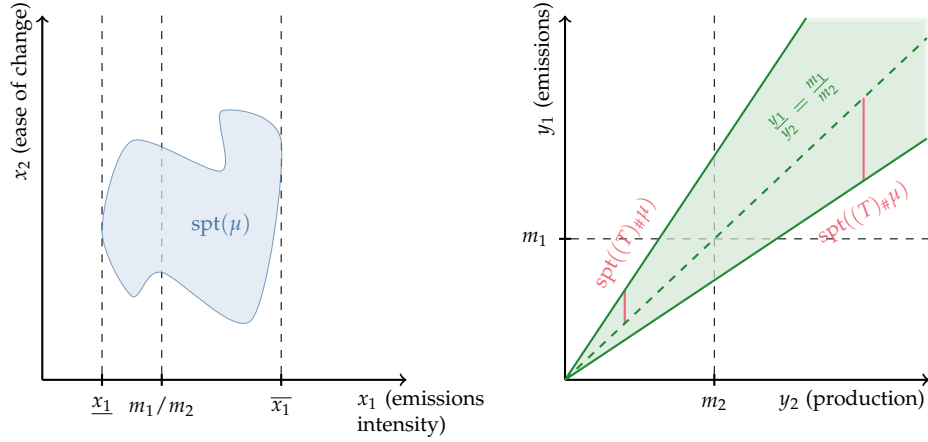


Figure 2.2: The case of Proposition 2.4 where $x_1 < m_1/m_2 < \bar{x}_1$. In this case, the map T defined in Equation (2.4) assigns high production levels to firms with clean production methods and low production levels to firms with dirty production methods.

firms down. In fact, the proof of Proposition 2.4 indicates that the regulator must shut down all firms with suboptimally clean production $x_1 > \underline{x}_1$ and allocate all production to firms with maximally clean production $x_1 = \underline{x}_1$. These results are indicated in Figure 2.3, which follows.

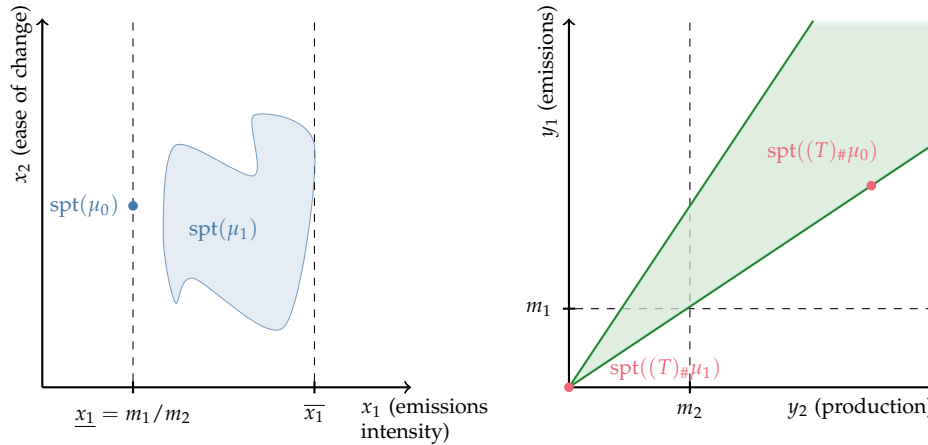


Figure 2.3: The case of Proposition 2.4 where $\underline{x}_1 = m_1/m_2 < \bar{x}_1$, so that the primal infimum is attained in $\mathcal{P}([0, \infty)^4)$ but not $\mathcal{P}((0, \infty)^4)$. For the sake of clarity, we decompose $\mu = \mu_0 + \mu_1$, where $\mu_0 = \mu|_{\{x_1 = \underline{x}_1\}}$ represents firms with maximally clean production methods and $\mu_1 = \mu|_{\{x_1 > \underline{x}_1\}}$ represents those with less clean production. As illustrated in the right figure, Proposition 2.4 indicates that all production should be allocated to μ_0 .

We now describe the cases of non-attainment in Theorem 2.3. The first such case occurs when the target m has emissions intensity m_1/m_2 equal to the lower EVT \underline{x}_1 , but no firms actually reach this target. In this case, the regulator can implement command-and-control policies γ of arbitrarily low but nonzero cost, as the following proposition illustrates:

Proposition 2.5 (Complementary Slackness Failure). *Let $\mu \in \mathcal{P}_c([0, \infty) \times (0, \infty))$ and $m \in (0, \infty)^2$ be such that $\underline{x}_1 \leq m_1/m_2 \leq \bar{x}_1$, but*

$$\min(\mu(x_1 \geq m_1/m_2), \mu(x_1 \leq m_1/m_2)) = 0.$$

Then the $\mathcal{I}(\mu, m) = 0$, but is not attained in $\mathcal{P}_c([0, \infty)^4$).

This means that the regulator can approximate the primal infimum of 0 with a sequence of policies $(\gamma_n)_n$ of decreasing cost as illustrated in Figure 2.4. Each policy γ_n selects a proportion of firms which are sufficiently clean (where the notion of sufficiency becomes more stringent as n increases), allocates all production to these firms, and shuts down all other firms. In practice, this means that, if the regulator is comfortable with their policy having some non-zero cost, it can achieve the target m .

Finally, we characterize the cases in which primal infimum is strictly positive, corresponding to a true duality gap in the sense of Definition 1.9. To start, we quantify the size of the duality gap if all production methods share the same emissions intensity \underline{x}_1 , but the regulatory target m corresponds to a lower emissions intensity $m_1/m_2 < \underline{x}_1$:

Proposition 2.6 (Homogeneous Duality Gap). *If $\mu \in \mathcal{P}_c([0, \infty) \times (0, \infty))$ and $m \in (0, \infty)^2$ are such that $\underline{x}_1 = \bar{x}_1 \neq m_1/m_2$, then*

$$\mathcal{I}(\mu, m) \geq c(\bar{x}_1, \bar{x}_2, m) > 0.$$

The proof of Proposition 2.6 largely boils down to an application of Jensen's inequality, and this case is illustrated in Figure 2.5. We also extend this result to the more general case where $\underline{x}_1 < \bar{x}_1$ and $m_1/m_2 \notin [\underline{x}_1, \bar{x}_1]$:

Proposition 2.7 (Heterogeneous Duality Gap). *If $\mu \in \mathcal{P}_c([0, \infty) \times (0, \infty))$ and $m \in (0, \infty)^2$ are such that $m_1/m_2 \notin [\underline{x}_1, \bar{x}_1]$, then $\mathcal{I}(\mu, m) > 0$.*

Note that, while we do not share a precise lower bound on $\mathcal{I}(\mu, m)$ in Proposition 2.7 at this time, the curious reader will find a (non-sharp) lower bound in Corollary 2.16. We illustrate Proposition 2.7 in Figure 2.6.

2.2 PRIMAL AND DUAL FORMULATIONS ON GENERAL DOMAINS

We will use the following steps, which are analogous to those presented in [69, Section 2.3], to prove Proposition 1.8:




Figure 2.4: Supports of some approximate minimizers γ_n defined in the proof of Proposition 2.5. Each γ_n attains the regulator’s target m and, as $n \rightarrow \infty$, the cost $\int \int c d\gamma_n$ tends to 0.

1. Use Lagrange multipliers to rewrite (1.6) as an unconstrained optimization problem of the form $\inf \sup \mathcal{F}$, for some functional \mathcal{F} .
2. (Formally) exchange the infimum and the supremum.
3. Again use Lagrange multipliers to recover a constrained optimization problem in different variables.

Proof of Proposition 1.8 Fix arbitrary non-empty subsets $X, Y \subseteq [0, \infty)^2$, assume that $\mu \in \mathcal{P}(X)$ and $m \in [0, \infty)^2$, and consider the primal infimum $\mathcal{I}_{X,Y}(\mu, m)$ defined in (1.6).




Figure 2.5: The case of Proposition 2.6 where $\underline{x}_1 = \bar{x}_1 \neq m_1/m_2$. In this case, Jensen’s inequality implies that any γ satisfying the mean constraint $\int \int_{[0,\infty)^4} y d\gamma = m$ has cost bounded below by the positive quantity $c(\bar{x}_1, \bar{x}_2, m)$. In particular, this implies that the primal infimum is at least $c(\bar{x}_1, \bar{x}_2, m) > 0$ and hence, strictly greater than the dual supremum of 0.

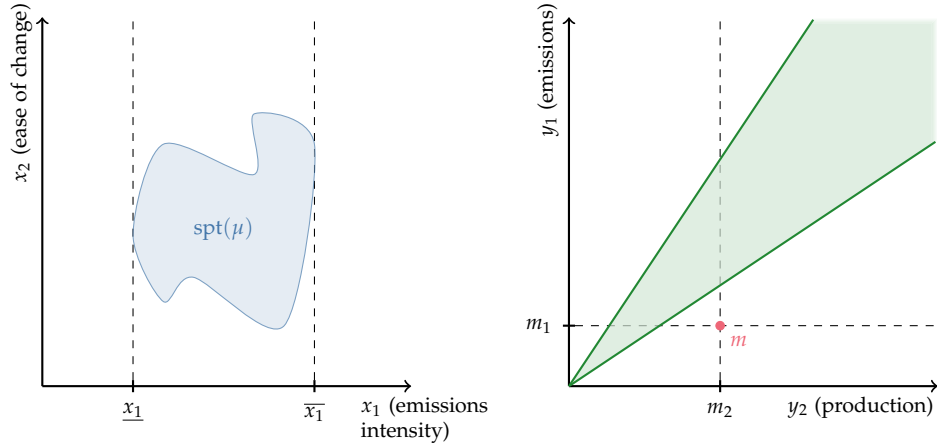


Figure 2.6: An example of the remaining case discussed in Proposition 2.7, where $\underline{x}_1 < \bar{x}_1$ and $m_1/m_2 \notin [\underline{x}_1, \bar{x}_1]$.

Note that, as in the treatment of Kantorovich-Koopmans duality in [96, Section 1.2], any non-negative finite signed measure $\gamma \in \mathcal{M}_+(X \times Y)$ satisfies

$$\sup_{\varphi \in C_b(X)} \left[\int_X \varphi d\mu - \iint_{X \times Y} \varphi(x) d\gamma(x, y) \right] = \begin{cases} 0 & \text{if } \pi_{\#}^X \gamma = \mu \\ +\infty & \text{otherwise.} \end{cases} \quad (2.1)$$

Likewise, we introduce Lagrange multipliers for the mean constraint by noticing that:

$$\sup_{\lambda \in \mathbb{R}^2} \lambda \cdot \left[m - \iint_Y y \, d\gamma(x, y) \right] = \begin{cases} 0 & \text{if } \iint_Y y \, d\gamma(x, y) = m \\ +\infty & \text{otherwise.} \end{cases} \quad (2.2)$$

Thus, by adding in the terms defined in (2.1) and (2.2), we find that

$$\begin{aligned} \mathcal{I}_{X,Y}(\mu, m) = & \inf_{\gamma \in \mathcal{M}_+(X \times Y)} \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2}} \left[\iint_{X \times Y} c \, d\gamma + \int_X \varphi \, d\mu \right. \\ & \left. - \iint_{X \times Y} \varphi(x) \, d\gamma(x, y) + \lambda \cdot \left(m - \iint_Y y \, d\gamma(x, y) \right) \right]. \end{aligned}$$

To find the dual problem to the original problem (1.6), recall that, if A and B are sets and $f : A \times B \rightarrow \mathbb{R} \cup \{\pm\infty\}$, then $\inf_A \sup_B f(a, b) \geq \sup_B \inf_A f(a, b)$ (see [92, Lemma 36.1] for a proof). As such, we find that

$$\begin{aligned} \mathcal{I}_{X,Y}(\mu, m) \geq & \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2}} \left[\int_X \varphi \, d\mu + \lambda \cdot m \right. \\ & \left. + \inf_{\gamma \in \mathcal{M}_+(X \times Y)} \iint_{X \times Y} (c(x, y) - \varphi(x) - \lambda \cdot y) \, d\gamma(x, y) \right]. \end{aligned}$$

Now, notice that

$$\begin{aligned} \inf_{\gamma \in \mathcal{M}_+(X \times Y)} \iint_{X \times Y} c(x, y) - \varphi(x) - \lambda \cdot y \, d\gamma(x, y) \\ = \begin{cases} 0 & \text{if } c(x, y) \geq \varphi(x) + \lambda \cdot y \text{ on } X \times Y \\ -\infty & \text{otherwise.} \end{cases} \end{aligned}$$

Thus, we pass to the following dual optimization problem:

$$\begin{aligned} \mathcal{I}_{X,Y}(\mu, m) \geq & \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2}} \left\{ \int_X \varphi \, d\mu + \lambda \cdot m \mid c(x, y) \geq \varphi(x) + \lambda \cdot y \text{ on } X \times Y \right\} \\ = & \mathcal{S}_{X,Y}(\mu, m), \end{aligned}$$

which is the desired conclusion. ■

2.3 SOLVING THE DUAL PROBLEM ON $[0, \infty)^4$

We have now shown that $\mathcal{S}_{X,Y}(\mu, m)$ is indeed the linear programming dual problem to $\mathcal{I}(\mu, m)$, and thus turn our attention to evaluating $\mathcal{S}_{X,Y}(\mu, m)$ in the case $X = Y = [0, \infty)^2$. In this section, we prove Proposition 2.1, which states that $\mathcal{S}(\mu, m) = 0$, independent of the choice of $\mu \in \mathcal{P}([0, \infty)^2)$ and $m \in (0, \infty)^2$.

This proposition is, in large part, a consequence of the function $G_\lambda(x, y) := c(x, y) - \lambda \cdot y$ being 1-homogeneous in y , as will become clear in the following lemma:

Lemma 2.8. *Define the function $G_\lambda : [0, \infty)^2 \times [0, \infty)^2 \rightarrow \mathbb{R} \cup \{+\infty\}$ by $G_\lambda(x, y) = c(x, y) - \lambda \cdot y$, where c is as defined in Equation (1.5), and define $g_\lambda : [0, \infty)^2 \rightarrow \mathbb{R} \cup \{-\infty\}$ by $g_\lambda(x) = \inf_{y \in [0, \infty)^2} G_\lambda(x, y)$. Then,*

$$g_\lambda(x) = \begin{cases} 0 & \text{if } c(x, y) - \lambda \cdot y \geq 0 \text{ for all } y \in [0, \infty)^2 \\ -\infty & \text{otherwise.} \end{cases}$$

Remark 2.9. The function g_λ can be thought of as a \bar{c} -transform of the function $\psi_\lambda(y) = \lambda \cdot y$, in the sense of [96, Definition 1.10]. However, the present work does not make use of this fact, so we avoid introducing this notation.

Proof of Lemma 2.8. The key observation for this proof is that $c(x, y) - \lambda \cdot y$ is a positively 1-homogeneous function of y , in the sense that, for any $(x, y) \in [0, \infty)^2$ and $t \in [0, \infty)$,

$$c(x, ty) - \lambda \cdot (ty) = t[c(x, y) - \lambda \cdot y],$$

with the appropriate adjustments made for the remaining cases of (1.5).

With this in mind, assume that x is such that $c(x, y) - \lambda \cdot y \geq 0$ for all $y \in [0, \infty)^2$, and observe that, by Equation (1.5), $c(x, 0) - \lambda \cdot 0 = 0$ and hence

$$0 \leq \inf_{y \in [0, \infty)^2} [c(x, y) - \lambda \cdot y] \leq c(x, 0) - \lambda \cdot 0 = 0,$$

as desired.

On the other hand, if x admits $y_0 \in [0, \infty)^2$ for which $c(x, y_0) - \lambda \cdot y_0 < 0$, homogeneity implies that, for any $t \in [0, \infty)$,

$$\inf_{y \in [0, \infty)^2} [c(x, y) - \lambda \cdot y] \leq c(x, ty_0) - \lambda \cdot (ty_0) = t[c(x, y_0) - \lambda \cdot y_0],$$

so taking $t \rightarrow +\infty$, we recover the other case of Lemma 2.8. \blacksquare

We now prove the equivalence of a number of conditions which will be useful for the proof of Proposition 2.1.

Lemma 2.10. Fix $\lambda \in \mathbb{R}^2$. The following are equivalent:

1. $c(x, y) - \lambda \cdot y \geq 0$ on $[0, \infty)^4$.
2. $g_\lambda \equiv 0$.
3. There exists $\varphi \in C_b([0, \infty)^2)$ such that $c(x, y) \geq \varphi(x) + \lambda \cdot y$ on $[0, \infty)^4$.
4. $\lambda_1 \leq 0$ and $\lambda_2 \leq 0$.

Proof. (1) \implies (2) This follows directly from Lemma 2.8.

(2) \implies (3): Take $\varphi \equiv 0$. Then

$$c(x, y) - \varphi(x) - \lambda \cdot y = c(x, y) - \lambda \cdot y \geq g_\lambda(x) = 0$$

for any $(x, y) \in [0, \infty)^4$

(3) \implies (4): We prove the contrapositive. As such, assume without loss of generality that $\lambda_1 > 0$. Fix $x_1 \in [0, \infty)$ such that $x_1 > -\frac{\lambda_2}{\lambda_1}$, to ensure that $\lambda_1 x_1 + \lambda_2 > 0$. Now define a sequence $\{y(k)\}_{k \in \mathbb{N}}$ by $y(k) = (x_1 k, k)$. By Equation (1.5), we compute

$$c(x, y(k)) - \lambda \cdot y(k) = -(\lambda_1 x_1 + \lambda_2)k.$$

Hence, since $\lambda_1 x_1 + \lambda_2 < 0$, we deduce that any φ satisfying the inequality of (3) necessarily satisfies $\varphi(x) \leq -(\lambda_1 x_1 + \lambda_2)k$ for any $k \in \mathbb{N}$, contradicting the assumption that $\varphi \in C_b([0, \infty)^2)$.

(4) \implies (1): Notice that $c(x, y) \geq 0$ for any $(x, y) \in [0, \infty)^4$. Moreover, since $\lambda_1 \leq 0$, $\lambda_2 \leq 0$, $y_1 \geq 0$, and $y_2 \geq 0$ by assumption, we deduce that

$$c(x, y) - \lambda \cdot y \geq 0$$

for any $(x, y) \in [0, \infty)^4$, as desired. \blacksquare

Finally, we arrive at the main objective of this section, the proof of Proposition 2.1.

Proof of Proposition 2.1. First, observe that, since $c \geq 0$, $\lambda = (0, 0)$ and $\varphi \equiv 0$ satisfy the constraints of (1.7) on $[0, \infty)^4$. Moreover, this pair witnesses that

$$\mathcal{S}(\mu, m) \geq \int 0 \, d\mu + \lambda \cdot (0, 0) = 0,$$

so it suffices to prove that $\mathcal{S}(\mu, m) \leq 0$.

To do so, fix (φ, λ) admissible for $\mathcal{S}(\mu, m)$, so that $c(x, y) \geq \varphi(x) + \lambda \cdot y$ on $[0, \infty)^4$. Observe that the (3) \implies (4) implication of Lemma 2.10 implies that $\lambda \in (-\infty, 0]^2$. Rearranging the constraint inequality from (1.7), we find that $\varphi(x) \leq c(x, y) - \lambda \cdot y$ on $[0, \infty)^4$ and hence

$$\varphi(x) \leq \inf_{y \in [0, \infty)^2} [c(x, y) - \lambda \cdot y] = g_\lambda(x).$$

Applying the (3) \implies (2) implication of Lemma 2.10, we find that

$$\varphi(x) \leq g_\lambda(x) = 0$$

for any $x \in [0, \infty)^2$. Hence, we conclude that

$$\int_{[0, \infty)^2} \varphi \, d\mu + \lambda \cdot m \leq \int 0 \, d\mu + (0, 0) \cdot m = 0,$$

where we have used that $\lambda_1 \leq 0$, $\lambda_2 \leq 0$, $m_1 \geq 0$, and $m_2 \geq 0$. \blacksquare

2.4 EXISTENCE OF A DUALITY GAP ON $[0, \infty)^4$

We now turn our attention to examining the relationship between $\mathcal{I}(\mu, m)$ as defined in (1.6) and the dual supremum $\mathcal{S}(\mu, m)$ defined in (1.7). In Section 2.2, we proved that $\mathcal{I}(\mu, m) \geq \mathcal{S}(\mu, m)$, and in Section 2.3 we proved Proposition 2.1, which states that $\mathcal{S}(\mu, m) = 0$ and moreover the supremum is attained by $(\varphi, \lambda) = (0, 0) \in C_b(\mathbb{R}^2) \times \mathbb{R}^2$. As such, the relationship between $\mathcal{I}(\mu, m)$ and $\mathcal{S}(\mu, m)$ is determined entirely by the primal problem $\mathcal{I}(\mu, m)$. More specifically, any of the three following possibilities, based on Definition 1.9, can occur:

- **Strong duality holds with primal attainment:** $\mathcal{I}(\mu, m) = 0$ and the infimum is attained by $\gamma \in \mathcal{P}_c([0, \infty)^4)$. In this case, a regulator may scale up the output of low-emissions producers, and scale down the output of high-emissions producers to meet the emissions and production standards m without incurring any additional cost. Moreover, the regulator may use either command-and-control or market-based policies to do so.
- **Strong duality holds without primal attainment:** $\mathcal{I}(\mu, m) = 0$, but the infimum is not attained on $\mathcal{P}_c([0, \infty)^4)$. In this case, command-and-control policies can be designed to have arbitrarily low cost.
- **Existence of a duality gap:** $\mathcal{I}(\mu, m) > 0$. In this case, meeting the emissions target will require command-and-control policies, and all command-and-control policies incur or exceed some baseline cost.

Perhaps surprisingly, any of these three possibilities can occur, depending on the measure μ and target m . In what follows, we simplify the analysis by focusing on the case where $\mu \in \mathcal{P}_c([0, \infty) \times (0, \infty))$, corresponding to the assumption that $x_2 > \underline{x}_2 > 0$ for all $x \in \text{spt}(\mu)$. In plain language, this assumption implies that all firms have some minimal capacity $C > 0$ to change their emissions intensity. Even if a firm's production methods are rigid, one can imagine a firm doing so by, for example, investing in

carbon offsets, or even purchasing a houseplant to remove a trace amount of carbon dioxide out of the atmosphere.

With this convention standing, we briefly recall the results which we will prove in this section. In Theorem 2.3, we use μ and m to completely characterize when duality holds with primal attainment. Then, in Proposition 2.4, we split the case of strong duality with primal attainment into two completely characterized cases – one where there exists γ attaining the infimum in $\mathcal{P}_c([0, \infty) \times (0, \infty)^3)$, and one where any admissible γ has support intersecting $\{y_1 = 0\} \cup \{y_2 = 0\}$. Then, in Proposition 2.5, we characterize when strong duality holds without attainment, i.e. when $\mathcal{I}(\mu, m) = 0$, but is not attained on $\mathcal{P}_c([0, \infty)^4)$. Finally, we provide sufficient conditions for the third possibility (a true duality gap) in Proposition 2.6 and Proposition 2.7.

We begin by proving Theorem 2.3:

Proof of Theorem 2.3. Assume first that $\underline{x}_1 = m_1/m_2 = \bar{x}_1$. In this case, we verify that the product measure $\gamma = \mu \otimes \delta_m$ attains the infimum. It is immediately clear that both constraints of (1.6) are satisfied, as $\pi_{\#}^X \gamma = \mu$ and $\iint_{X \times Y} y \, d\gamma(x, y) = m$. Moreover, as all $(x, y) \in \text{spt}(\gamma)$ satisfy $y_1 - y_2 x_1 = m_1 - \frac{m_1}{m_2} m_2 = 0$, the definition of c in Equation (1.5) implies that

$$\iint_{(0, \infty)^4} c \, d\gamma = 0,$$

so that the primal infimum is attained.

Otherwise, we assume without loss of generality that $\underline{x}_1 \leq m_1/m_2 < \bar{x}_1$. Define functions $F : \mathbb{R} \rightarrow [0, 1]$ and $H : \mathbb{R} \cup \{+\infty\} \rightarrow \mathbb{R}$ by

$$F(t) := \mu(x_1 \leq t)$$

and

$$H(t) := \begin{cases} \int_{x_1 \leq t} x_1 \, d\mu(x) & t \in \mathbb{R} \\ \int_X x_1 \, d\mu(x) & t = +\infty \end{cases}.$$

Notice that F is the cumulative distribution function of the X_1 -marginal $\mu_1 := \pi_{\#}^{X_1} \mu$ of μ . Likewise, H can be interpreted as an antiderivative of $x_1 d\mu_1(x_1)$. Using this notation, we define constants $r_+ = r_+(\mu, m)$ and $r_- = r_-(\mu, m)$ by

$$\begin{pmatrix} r_+ \\ r_- \end{pmatrix} := \begin{pmatrix} H(+\infty) - H\left(\frac{m_1}{m_2}\right) & H\left(\frac{m_1}{m_2}\right) \\ 1 - F\left(\frac{m_1}{m_2}\right) & F\left(\frac{m_1}{m_2}\right) \end{pmatrix}^{-1} \begin{pmatrix} m_1 \\ m_2 \end{pmatrix} \quad (2.3)$$

In effect, this operator takes a regulatory target (m_1, m_2) and returns new production levels (r_+, r_-) , where r_+ corresponds to a uniform high level

of production for firms which exceed the emissions intensity target $\frac{m_1}{m_2}$, and r_- corresponds to a uniform low level of production for firms which only meet or underperform this target.

We note that the assumption that $m_1/m_2 < \bar{x}_1$ ensures that both r_+ and r_- are well-defined – more precisely, it allows us to conclude that

$$\begin{aligned} F\left(\frac{m_1}{m_2}\right) \left(H(+\infty) - H\left(\frac{m_1}{m_2}\right) \right) &> \frac{m_1}{m_2} F\left(\frac{m_1}{m_2}\right) \left(1 - F\left(\frac{m_1}{m_2}\right) \right) \\ &\geq \left(1 - F\left(\frac{m_1}{m_2}\right) \right) H\left(\frac{m_1}{m_2}\right), \end{aligned}$$

so that the matrix appearing in Equation (2.3) is indeed invertible. Now, use these constants to define a transport map $T : [0, \infty)^2 \rightarrow [0, \infty)^2$ by

$$T(x_1, x_2) = \begin{cases} (r_+x_1, r_+) & \text{for } x_1 > m_1/m_2 \\ (r_-x_1, r_-) & \text{for } x_1 \leq m_1/m_2. \end{cases} \quad (2.4)$$

We will show that $\gamma_0 = (id, T)_\# \mu$ attains the primal infimum of 0. First, by construction, it is clear that $\pi_\#^X \gamma = \mu$. Moreover, we find that

$$\begin{aligned} \iint_{[0, \infty)^4} c \, d\gamma_0 &= \int_{[0, \infty)^2} c(x, T(x)) \, d\mu(x) \\ &= \int_{x_1 > m_1/m_2} c(x, (r_+x_1, r_+)) \, d\mu(x) \\ &\quad + \int_{x_1 \leq m_1/m_2} c(x, (r_-x_1, r_-)) \, d\mu(x) \\ &= 0, \end{aligned}$$

where we have used the definition of c in Equation (1.5) to conclude that

$$c(x, (r_+x_1, r_+)) = c(x, (r_-x_1, r_-)) = 0.$$

Finally, we check the mean constraint on the y -marginal by computing $\iint_{[0, \infty)^2} y \, d\gamma(x, y) = \int_{[0, \infty)^2} T \, d\mu$. By definition,

$$\begin{aligned} \int_{[0, \infty)^2} T \, d\mu &= r_+ \int_{x_1 > m_1/m_2} \begin{pmatrix} x_1 \\ 1 \end{pmatrix} \, d\mu(x) + r_- \int_{x_1 \leq m_1/m_2} \begin{pmatrix} x_1 \\ 1 \end{pmatrix} \, d\mu(x) \\ &= \begin{pmatrix} H(+\infty) - H\left(\frac{m_1}{m_2}\right) & H\left(\frac{m_1}{m_2}\right) \\ 1 - F\left(\frac{m_1}{m_2}\right) & F\left(\frac{m_1}{m_2}\right) \end{pmatrix} \begin{pmatrix} r_+ \\ r_- \end{pmatrix} \\ &= \begin{pmatrix} m_1 \\ m_2 \end{pmatrix}, \end{aligned}$$

where the second and third steps follow from the definitions of F and H and that of (r_+, r_-) , respectively. Thus, in this case, there is no duality gap and the primal infimum is attained.

Next, we show that, in all other cases, the primal infimum is not attained. As such, assume without loss of generality that μ satisfies $\mu(x_1 \leq m_1/m_2) = 0$ (so that no status quo producers meet the prescribed emissions target). Moreover, assume by way of contradiction that $\gamma \in \mathcal{P}([0, \infty)^4)$ satisfies the constraints in (1.6) with $\int \int_{[0, \infty)^4} c \, d\gamma = 0$. In particular, the marginal constraint implies that $\gamma(x_1 > m_1/m_2) = 1$, the mean constraint implies that $\int \int y \, d\gamma(x, y) = m$, and the assumption that $\int \int c \, d\gamma = 0$ implies that $\gamma(y_1 = y_2 x_1) = 1$. Combining these assumptions, we deduce that

$$m_1 = \int \int y_1 \, d\gamma(x, y) = \int \int y_2 x_1 \, d\gamma(x, y) > \frac{m_1}{m_2} \int \int y_2 \, d\gamma(x, y) = m_1, \quad (2.5)$$

furnishing us with our desired contradiction. \blacksquare

Remark 2.11. Note that if, instead of assuming that $\underline{x}_1 \leq m_1/m_2 < \bar{x}_1$, we had assumed that $\underline{x}_1 < m_1/m_2 \leq \bar{x}_1$, then we would have to swap all strict inequalities with non-strict inequalities, and vice versa, in the preceding definitions of r_+ and r_- . Of course, this case is less economically realistic – it reflects an emissions-intensity target which is equal to or higher than the emissions intensity of any producer before regulation, and there would be little reason for government to impose such a target. Nevertheless, such a model may be useful if a regulator wishes to increase production of a useful byproduct – for example, one could imagine a government providing incentives for fishers to use nets which are more effective at removing garbage from the ocean, making the garbage an effective bycatch.

Remark 2.12 (Economic Implications of Theorem 2.3). The first case, where $\underline{x}_1 = m_1/m_2 = \bar{x}_1$, can be interpreted as the case where the entire industry's emissions profile is homogeneous – each producer has the same emissions intensity, and moreover this shared emissions intensity is precisely in line with the regulated emissions intensity standards. In this case, the transportation plan corresponds to scaling production to meet the desired balance of output and emissions.

The second case, where $\underline{x}_1 \leq m_1/m_2 \leq \bar{x}_1$, corresponds to a case where the desired emissions intensity target can be met or exceeded by existing technology. In this case, the transportation plan T defined in Equation (2.4) increases the relative proportion of production from existing low-intensity emitters. In other words, this transportation plan reallocates high production levels to firms which outperform the regulator's emis-

sions intensity target. We also note that T can be defined regardless of the statistical properties of the measure μ , so long as the assumption that $\min(\mu(x_1 \geq m_1/m_2), \mu(x_1 \leq m_1/m_2)) > 0$ is met. More precisely, that assumption ensures the existence of producers who outperform the emissions intensity target and, since we are considering measures $\gamma \in \mathcal{P}_c([0, \infty)^4)$, we may scale up these producers' production levels without bound, in order to meet the regulator's target m .

Finally, if $\mu(x_1 \leq m_1/m_2) = 0$, this corresponds to the case where the emissions intensity target is not being met by any current producers in the industry. In this case, it is not possible for a simplistic market-based policy like the one framed in $\mathcal{S}(\mu, m)$ to attain the regulator's target. In this chapter, the only solution to this issue is for the regulator to institute command-and-control policy, represented by $\mathcal{I}(\mu, m)$. However, as we will see in Chapter 3, the regulator's target can also be achieved through sophisticated market-based policy involving progressive emissions taxation – in light of this, the main issue with $\mathcal{S}(\mu, m)$ is the flat-rate tax on emissions represented by λ_1 . Likewise, the results of Chapter 4 indicate that regulators may also resolve this issue with hybrid policy, introducing hard caps on the amount each firm may emit/produce, but otherwise allowing them to work within a market-based system. Finally, an innovation-based strategy could also expand the range of economically viable production methods, which would have the effect of switching an industry to one of the earlier cases.

From here, we move directly into the proof of Proposition 2.4:

Proof of Proposition 2.4. If $\underline{x}_1 = \bar{x}_1$, then we note that $\gamma = \mu \otimes \delta_m$ as defined in the proof of Theorem 2.3 belongs to $\mathcal{P}_c([0, \infty) \times (0, \infty)^3)$. Likewise, if $\underline{x}_1 < m_1/m_2 < \bar{x}_1$, then both r_+ and r_- as defined in (2.3) are positive, so the corresponding measure γ is supported in $\mathcal{P}_c([0, \infty) \times (0, \infty)^3)$.

Thus, we need only to show that if neither of the preceding conditions hold (i.e. if $\underline{x}_1 < m_1/m_2 = \bar{x}_1$ or $\underline{x}_1 = m_1/m_2 < \bar{x}_1$), the primal infimum cannot be attained in $\mathcal{P}_c([0, \infty) \times (0, \infty)^3)$. As such, assume without loss of generality that $\underline{x}_1 = m_1/m_2$ so that μ satisfies

$$\mu(x_1 \leq m_1/m_2) = \mu(x_1 = m_1/m_2) > 0$$

and $\mu(x_1 > m_1/m_2) > 0$. Moreover, assume by way of contradiction that $\gamma \in \mathcal{P}_c([0, \infty) \times (0, \infty)^3)$ satisfies the constraints in (1.6) and $\iint_{[0, \infty)^4} c \, d\gamma = 0$. These constraints imply that, by using reasoning analogous to that of (2.5),

$$m_1 = \frac{m_1}{m_2} \iint_{x_1=m_1/m_2} y_2 \, d\gamma(x, y) + \iint_{x_1>m_1/m_2} x_1 y_2 \, d\gamma(x, y) > m_1,$$

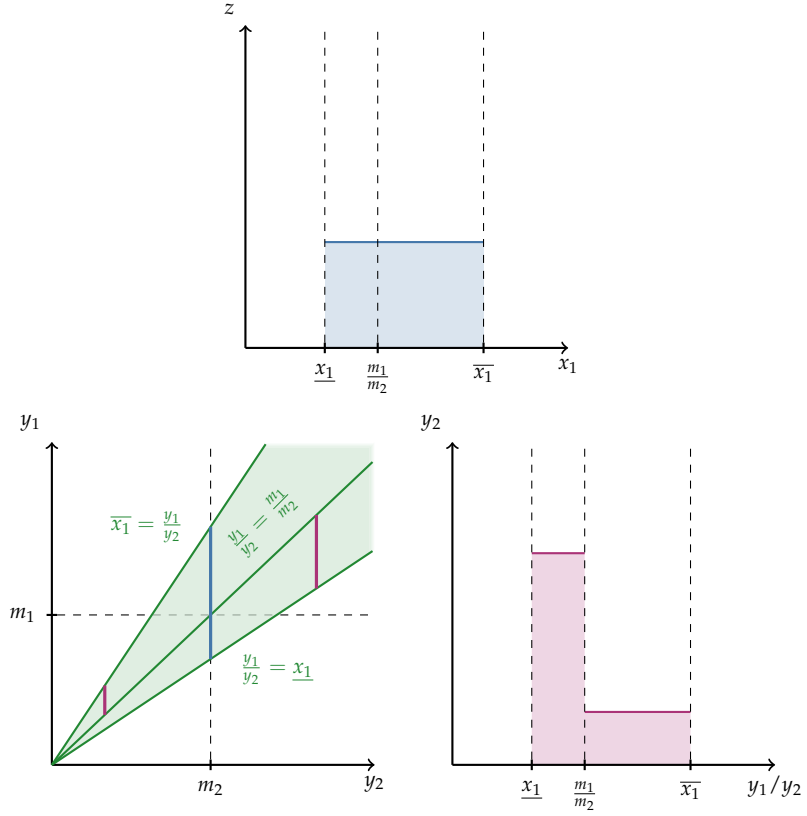


Figure 2.7: The results of applying the map T to adjust an industry's emissions intensity profile. The first subfigure depicts the density of $\pi_{\#}^{X_1} \mu$ for a sample measure μ , representing the industry's pre-regulation emissions intensity profile. Here, z references the share of production at each emissions intensity level. For the purpose of illustration, production is uniformly distributed over the range $[x_1, \bar{x}_1]$ of emissions intensities. The second subfigure depicts the EVF for μ in (y_1, y_2) -coordinates. Here, the blue line segment represents the status quo emissions intensity profile of the industry and the purple line segments represent the support of $T_{\#} \mu$. The map T prescribes new emissions and productions levels to firms which, when averaged across the industry, simultaneously meet the emissions target m_1 and the production target m_2 . Finally, if we define $q(y_1, y_2) = \frac{y_1}{y_2}$, the third graph represents the density of $q_{\#} T_{\#} \mu$, which profiles post-regulation production in the industry, segmented by emissions intensity. In particular, the share of production undertaken using methods which meet and exceed the target has increased.

establishing the desired contradiction. ■

Remark 2.13 (Economic Interpretation of Proposition 2.4). The existence of a transportation plan $\gamma \in \mathcal{P}_c([0, \infty) \times (0, \infty)^3)$ as in this proposition

means that it is possible to attain the given target without phasing out any existing production methods – high-emissions methods can simply be scaled down to produce a small nonzero fraction of the total output. If, on the other hand, the only transportation plans attaining an infimum of 0 in equation (1.6) lie in $\mathcal{P}_c([0, \infty)^4) \setminus \mathcal{P}_c([0, \infty) \times (0, \infty)^3)$, then no producers exceed the emissions intensity target prescribed by m . In this case, Proposition 2.4 implies that, to meet the target m , all production methods which fail to meet the emissions-intensity target must m_1/m_2 be phased out entirely. In essence, this is because there are no economically viable methods which outperform the emissions intensity targets (and which can be used to compensate for poor performers elsewhere in the industry). Hence, allowing even a small proportion of producers to fall short of the prescribed emissions intensity targets would mean that the whole industry, on average, fails to meet these targets.

Next, we move into the proof of Proposition 2.5.

Proof of Proposition 2.5. By Theorem 2.3, the infimum is either nonzero or non-attained, so it will suffice to define a sequence $\{\gamma_n\}_n$ of measures satisfying the constraints of (1.6), but for which $\iint c \, d\gamma_n \rightarrow 0$. Without loss of generality, assume that $m_1/m_2 = \underline{x}_1$, and hence that $\mu(x_1 \leq m_1/m_2) = 0$. For each $n \in \mathbb{N}$, define the set $S_n = \{x \in [0, \infty)^2 \mid x_1 \leq \underline{x}_1 + \frac{1}{n}\}$, and let $S_n^c = \{x \in [0, \infty)^2 \mid x_1 > \underline{x}_1 + \frac{1}{n}\}$, be its set complement with respect to $[0, \infty)^2$. Define a measure

$$\gamma_n := \mu|_{S_n^c} \otimes \delta_0 + \mu|_{S_n} \otimes \delta_{m/\mu(S_n)}.$$

We now verify that each γ_n indeed satisfies both marginal constraints. Using the Riesz Representation Theorem and Definition 1.1 of push-forward measures, we check that $\pi_{\#}^X \gamma_n = \mu$ by checking that, for any $f \in C_b([0, \infty)^2)$, $\int_{[0, \infty)^2} f \, d(\pi_{\#}^X \gamma_n) = \int_{[0, \infty)^2} f \, d\mu$ as follows:

$$\begin{aligned} \int_{[0, \infty)^2} f \, d\pi_{\#}^X \gamma_n &= \iint_{[0, \infty)^4} f(x) \, d(\mu|_{S_n^c} \otimes \delta_0)(x, y) \\ &\quad + \iint_{[0, \infty)^4} f(x) \, d(\mu|_{S_n} \otimes \delta_{m/\mu(S_n)})(x, y) \\ &= \int_{[0, \infty)^2} f(x) \, d\mu|_{S_n^c}(x) + \int_{[0, \infty)^2} f(x) \, d\mu|_{S_n}(x) \\ &= \int_{[0, \infty)^2} f(x) \, d\mu(x), \end{aligned}$$

as desired. Likewise, we may check the mean constraint on the second marginal by noting that

$$\begin{aligned} \int_{[0, \infty)^2} y d(\pi_{\#}^Y \gamma_n) &= \iint_{[0, \infty)^4} y d\mu|_{S_n^c}(x) d\delta_0(y) \\ &\quad + \iint_{[0, \infty)^4} y d\mu|_{S_n}(x) d\delta_{m/\mu(S_n)}(y) \\ &= \mu(S_n^c) \int_{[0, \infty)^2} y d\delta_0(y) + \mu(S_n) \int_{[0, \infty)^2} y d\delta_{m/\mu(S_n)}(y) \\ &= m. \end{aligned}$$

Moreover, Equation (1.5) implies that $c(x, 0) = 0$, so we deduce that

$$\begin{aligned} \iint c d\gamma_n &= \iint_{[0, \infty)^4} c d(\mu|_{S_n} \otimes \delta_{m/\mu(S_n)}) \\ &= \int_{S_n} c(x, m/\mu(S_n)) d\mu(x) \\ &= \frac{m_2}{2\mu(S_n)} \int_{S_n} \frac{1}{x_2} \left(x_1 - \frac{m_1}{m_2}\right)^2 d\mu(x) \\ &= \frac{m_2}{2\mu(S_n)} \int_{S_n} \frac{1}{x_2} (x_1 - \underline{x}_1)^2 d\mu(x) \\ &\leq \frac{m_2}{2n^2\mu(S_n)} \int_{S_n} \frac{1}{x_2} d\mu(x), \end{aligned}$$

where we have used the fact that $\underline{x}_1 \leq x_1$ by definition and $x_1 - \underline{x}_1 \leq \frac{1}{n}$, by the domain of integration. Finally, recalling Definition 2.2 of \underline{x}_2 , and noticing that $\underline{x}_2 > 0$ by the assumption that $\mu \in \mathcal{P}_c([0, \infty) \times (0, \infty))$, we extend our previous estimates to conclude that

$$\iint c d\gamma_n \leq \frac{m_2}{2n^2\mu(S_n)} \int_{S_n} \frac{1}{\underline{x}_2} d\mu(x) = \frac{m_2}{2n^2\underline{x}_2},$$

which tends to 0 as $n \rightarrow \infty$. As such, the infimum of 0 is approached in this case. \blacksquare

Next, we prove Proposition 2.6.

Proof of Proposition 2.6. Assume that $\underline{x}_1 = \overline{x}_1 \neq m_1/m_2$. Moreover, fix γ satisfying the constraints of $\mathcal{I}(\mu, m)$ with $\gamma(\{x_2, y_2 > 0\} \cup \{y_1 = y_2 x_1\}) = 1$, so that $c < +\infty$ on $\text{spt}(\gamma)$. This assumption is satisfied by, for example, $\gamma = \mu \otimes \delta_m$, and saves us from dealing with the case $c = +\infty$ in the following calculations. Applying the constraint that $\iint y d\gamma = m$, the

definition of \underline{x}_2 from Definition 2.2, and the assumption that $x_1 \equiv \underline{x}_1$ on $\text{spt}(\gamma)$, we deduce that

$$\begin{aligned} \iint c \, d\gamma &= \iint_{([0, \infty) \times (0, \infty))^2} \frac{y_2}{2x_2} \left(x_1 - \frac{y_1}{y_2}\right)^2 d\gamma(x, y) \\ &\geq \frac{1}{2\bar{x}_2} \iint_{([0, \infty) \times (0, \infty))^2} y_2 \left(\bar{x}_1^2 - 2\frac{y_1}{y_2}\bar{x}_1 + \frac{y_1^2}{y_2^2}\right) d\gamma(x, y) \\ &= \frac{\bar{x}_1^2 m_2}{2\bar{x}_2} - \frac{\bar{x}_1 m_1}{\bar{x}_2} + \frac{1}{2\bar{x}_2} \int_{[0, \infty)^2} \frac{y_1^2}{y_2} d(\pi_{\#}^Y \gamma)(y). \end{aligned}$$

Now, notice that the map $(y_1, y_2) \mapsto \frac{y_1^2}{y_2}$ is a convex function from \mathbb{R}^2 to \mathbb{R} . As such, we apply Jensen's inequality to the preceding line to deduce that

$$\begin{aligned} \iint c \, d\gamma &\geq \frac{\bar{x}_1^2 m_2}{2\bar{x}_2} - \frac{\bar{x}_1 m_1}{\bar{x}_2} + \frac{1}{2\bar{x}_2} \frac{m_1^2}{m_2} \\ &= c(\bar{x}_1, \bar{x}_2, m), \end{aligned}$$

as desired. \blacksquare

Remark 2.14 (Interpretation of Proposition 2.6). In Proposition 2.6, the assumption that $\underline{x}_1 = \bar{x}_1$ reflects an assumption that the industry in question has homogeneous emissions intensity x_1 . Moreover, the assumption $m_1/m_2 \neq \underline{x}_1$ means that the emissions target is not in line with the industry-standard production method. As such, our discovery of a duality gap in this case belies the need for command-and-control policies or more refined market-based methods, such as the ones we will discuss in Chapters 3 and 4.

We introduce the following technical linear algebra lemma as a key step to prove Proposition 2.7:

Lemma 2.15. *Let $u, v, w \in [0, \infty) \times (0, \infty)$ and $\kappa > 0$ satisfy the following system:*

$$\begin{cases} w = u + v \\ v_1/v_2 \geq \kappa \\ w_1/w_2 < \kappa. \end{cases}$$

Then $u_2 \geq w_2 - w_1/\kappa > 0$.

Proof of Lemma 2.15. We compute

$$u_2 = w_2 - v_2 \geq w_2 - \frac{v_1}{\kappa} = w_2 - \frac{w_1}{\kappa} + \frac{u_1}{\kappa} > 0,$$

where we have used, in order, that $w = u + v$, $v_1/v_2 \geq \kappa$, $w = u + v$ and finally that $w_1/w_2 < \kappa$. \blacksquare

Proof of Proposition 2.7 Without loss of generality, assume that $m_1/m_2 < \underline{x}_1$ and, as in the proof of 2.6, fix γ admissible for $\mathcal{I}(\mu, m)$ such that $c < +\infty$ on $\text{spt}(\gamma)$. Using the Disintegration Theorem (stated, for example, in [85, Theorem 1.9]), disintegrate $d\gamma(x, y) = d\gamma_x(y)d\mu(x)$, so that

$$\iint_{[0, \infty)^4} c(x, y) d\gamma(x, y) = \int_{[0, \infty)^2} \int_{[0, \infty)^2} c(x, y) d\gamma_x(y) d\mu(x).$$

For each $x \in [0, \infty)^2$, we may use calculations which are similar to those in the proof of Proposition 2.6 to find that

$$\int_{[0, \infty)^2} c(x, y) d\gamma_x(y) \geq c\left(x, \int_{[0, \infty)^2} y d\gamma_x(y)\right) = c(x, m_\gamma(x)),$$

where we introduce the convenient notation $m_\gamma(x) := \int_{[0, \infty)^2} y d\gamma_x(y)$ and $m_{\gamma,i}(x) := \int_{[0, \infty)^2} y_i d\gamma_x(y)$. Using these estimates, along with the fact that $c \geq 0$, we deduce that

$$\begin{aligned} \iint c d\gamma &\geq \int_{[0, \infty)^2} c(x, m_\gamma(x)) d\mu(x) \\ &\geq \int_{m_{\gamma,1}(x)/m_{\gamma,2}(x) \leq \frac{1}{2}\left(\frac{m_1}{m_2} + \underline{x}_1\right)} \frac{m_{\gamma,2}(x)}{2x_2} \left(x_1 - \frac{m_{\gamma,1}(x)}{m_{\gamma,2}(x)}\right)^2 d\mu(x) \\ &\geq \int_{m_{\gamma,1}(x)/m_{\gamma,2}(x) \leq \frac{1}{2}\left(\frac{m_1}{m_2} + \underline{x}_1\right)} \frac{m_{\gamma,2}(x)}{8\bar{x}_2} \left(\underline{x}_1 - \frac{m_1}{m_2}\right)^2 d\mu(x) \\ &= \frac{1}{8\bar{x}_2} \left(\underline{x}_1 - \frac{m_1}{m_2}\right)^2 \int_{m_{\gamma,1}(x)/m_{\gamma,2}(x) \leq \frac{1}{2}\left(\frac{m_1}{m_2} + \underline{x}_1\right)} m_{\gamma,2} d\mu. \end{aligned}$$

Note that for the last inequality we have used that $0 < x_2 \leq \bar{x}_2$, $\frac{m_{\gamma,2}(x)}{m_{\gamma,1}(x)} \leq \frac{1}{2}\left(\frac{m_1}{m_2} + \underline{x}_1\right)$, and $\frac{m_1}{m_2} < \underline{x}_1 < x_1$ on $\text{spt}(\mu)$, which together imply that

$$x_1 - \frac{m_{\gamma,2}(x)}{m_{\gamma,1}(x)} \geq \frac{1}{2} \left(\underline{x}_1 - \frac{m_1}{m_2}\right) > 0.$$

Thus, all that remains is to bound $\int_{m_{\gamma,1}(x)/m_{\gamma,2}(x) \leq \frac{1}{2}\left(\frac{m_1}{m_2} + \underline{x}_1\right)} m_{\gamma,2} d\mu$ from below by a positive quantity which is independent of γ . To this end, we use

the assumption that $\iint y \, d\gamma = m$, the definition of m_γ , the Disintegration Theorem, and a change of variables to deduce that:

$$\begin{aligned} m &= \int_{[0, \infty)^2} m_\gamma(x) \, d\mu(x) \\ &= \int_{m_{\gamma,1}(x)/m_{\gamma,2}(x) \leq \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right)} m_\gamma \, d\mu + \int_{m_{\gamma,1}(x)/m_{\gamma,2}(x) > \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right)} m_\gamma \, d\mu \\ &= \int_{y_1/y_2 \leq \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right)} y \, d((m_\gamma)_\# \mu)(y) + \int_{y_1/y_2 > \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right)} y \, d((m_\gamma)_\# \mu)(y). \end{aligned}$$

Now, notice that both $\left\{ y \in [0, \infty)^2 \mid y_1/y_2 \leq \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right) \right\}$ and $\left\{ y \in [0, \infty)^2 \mid y_1/y_2 > \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right) \right\}$ are convex sets whose respective closures contain 0. Therefore, we deduce that

$$\int_{m_{\gamma,1}(x)/m_{\gamma,2}(x) \leq \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right)} m_\gamma \, d\mu \in \left\{ y \in [0, \infty)^2 \mid y_1/y_2 \leq \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right) \right\}$$

and

$$\int_{m_{\gamma,1}(x)/m_{\gamma,2}(x) > \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right)} m_\gamma \, d\mu \in \left\{ y \in [0, \infty)^2 \mid y_1/y_2 \geq \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right) \right\},$$

with these two vectors summing to m . To conclude, we apply Lemma 2.15 with $\kappa = \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right)$, $u = \int_{m_{\gamma,1}(x)/m_{\gamma,2}(x) \leq \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right)} m_\gamma \, d\mu$, $v = \int_{m_{\gamma,1}(x)/m_{\gamma,2}(x) > \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right)} m_\gamma \, d\mu$, and $w = m$ to deduce the estimate

$$\int_{m_{\gamma,1}(x)/m_{\gamma,2}(x) \leq \frac{1}{2} \left(\frac{m_1}{m_2} + \underline{x}_1 \right)} m_{\gamma,2}(x) \, d\mu(x) \geq m_2 - \frac{2m_1 m_2}{m_1 + m_2 \underline{x}_1} > 0.$$

Thus, combining this with our previous estimates, we conclude that

$$\iint c \, d\gamma \geq \frac{m_2}{8\underline{x}_2} \left(\underline{x}_1 - \frac{m_1}{m_2} \right)^2 \left(1 - \frac{2m_1}{m_1 + m_2 \underline{x}_1} \right) > 0,$$

which provides the desired γ -independent lower bound. \blacksquare

Corollary 2.16. *If $\mu \in \mathcal{P}_c([0, \infty) \times (0, \infty))$ and $m \in (0, \infty)^2$ are such that $m_1/m_2 < \underline{x}_1$, then*

$$\mathcal{I}(\mu, m) \geq \frac{m_2}{8\underline{x}_2(m_1 + m_2 \underline{x}_1)} \left(\underline{x}_1 - \frac{m_1}{m_2} \right)^3.$$

Proof. Rearrange the final estimate in the preceding proof and minimize over γ . ■

RESOLVING THE DUALITY GAP WITH VARIANCE BOUNDS: EFFECTIVE MARKET-BASED POLICY THROUGH PROGRESSIVE TAXATION

3.1 SUMMARY OF RESULTS

In this chapter, we prove that the duality issues which we discovered in Chapter 2 can be resolved by introducing a uniform bound on the variances of the candidate measures γ in the definition of $\mathcal{I}_{X,Y}(\mu, m)$ from Equation (1.6), yielding the new primal problem $\mathcal{I}_{X,Y}(\mu, m; \tau)$ from Equation (1.10) and the new dual problem $\mathcal{S}_{X,Y}(\mu, m; \tau)$ from Equation (1.11).

We begin this chapter in Section 3.2 by introducing the variance bound problem (1.10). As we noted in the introduction, this problem corresponds to a central planner, who is seeking a transport plan $\gamma \in \mathcal{P}(X \times Y)$ which will transform an industry with pre-regulation characteristics described by μ to have mean emissions m_1 and mean production m_2 , all while ensuring that the post-regulation industry satisfies $\iint |y - m|^2 d\gamma \leq \tau$. In this section, we also prove Proposition 1.13, which establishes (1.11) as a weak dual problem to (1.10), and briefly examine the properties of (1.11) relative to the previously established dual problem (1.7).

The proof of strong duality is much more technical than that of weak duality, and we use Section 3.3 to develop the necessary machinery for a proof of Theorem 1.14. In particular, we heavily use Friecke's formulation of Fenchel-Rockafellar duality from [46, Theorem 6.2.2], and apply it to a suitably defined Banach space. Mathematically, this section shows that a variance bound is sufficient to resolve the duality issues we observed in Chapter 2. In the context of economics, Theorem 1.14 asserts that, provided that there is a penalty on heterogeneity, a regulator may achieve its aims using either command-and-control or market-based policy.

After this proof, we briefly use Section 3.4 to prove the following two propositions, which show that $\mathcal{I}(\mu, m; \tau)$ is attained, but that $\mathcal{S}(\mu, m; \tau)$ is not attained in general. In particular, primal attainment is as follows:

Proposition 3.1 (Primal Attainment). *Let $X = Y = [0, \infty)^2$, let $\mu \in \mathcal{P}_c([0, \infty) \times (0, \infty))$, let $m \in (0, \infty)^2$, and let $\tau \geq 0$. Then $\mathcal{I}(\mu, m; \tau)$ is attained, in the sense that there exists a measure $\gamma \in \mathcal{P}(X \times Y)$ satisfying the constraints in the definition of $\mathcal{I}(\mu, m; \tau)$, and for which $\iint c \, d\gamma = \mathcal{I}(\mu, m; \tau)$.*

However, the following proposition shows that we cannot rely on the existence of dual maximizers:

Proposition 3.2 (Dual Non-Attainment). *Let $x_0 \in (0, \infty)^2$, let $\mu = \delta_{x_0}$, and let $\tau > 0$. Choose $m \in (0, \infty)^2$ such that $c(x_0, m) > 0$. Then the dual supremum $\mathcal{S}(\mu, m; \tau)$ is not attained.*

From here, we move into Section 3.5, where we explore Euler-Lagrange-Karush-Kuhn-Tucker (ELKKT) conditions that ensure that (1.10) and (1.11) are attained, reducing the question of attainment to a system of nonlinear equations. Our first key result here is Proposition 3.13, which provides complementary slackness conditions for a quartet $(\gamma, \varphi, \lambda, \eta) \in \mathcal{P}(X \times Y) \times C_b(X) \times \mathbb{R}^2 \times (-\infty, 0]$ to attain both the primal infimum (1.10) and the dual supremum (1.11). The next significant result is Proposition 3.19 which, in the $\eta < 0$ case, uses the conditions of Proposition 3.13 to explicitly identify a unique optimal choice $Y = Y(x; \lambda, \eta)$ of post-regulation emissions and production for each firm, given a tax-subsidy structure $(\lambda, \eta) \in \mathbb{R}^2 \times (-\infty, 0]$. Here, $-\lambda_1$ is a flat-rate emissions tax, λ_2 is a flat-rate production subsidy, and $-\eta$ is a non-flat tax, corresponding to the marginal shadow cost of complying with the variance bound τ . Next, we use Corollary 3.20 to explicitly identify optimal choices of γ and φ for given $(\lambda, \eta) \in \mathbb{R}^2 \times (-\infty, 0]$. Finally, we conclude this section with Corollary 3.21, which provides an explicit system of three nonlinear equations in three unknowns which, if solved, would guarantee that the conditions of Proposition 3.13 are met.

Next, in Section 3.6, we briefly introduce some numerical methods for (1.10) and (1.11), which are based on Proposition 3.13. These numerical methods centre around finding the optimal choice of $(\lambda, \eta) = (\lambda(\mu, m, \tau), \eta(\mu, m, \tau))$ given the Newell-Stavins cost c , the pre-regulation industry structure μ , the regulator's mean target m , and the regulator's variance bound τ . In particular, we provide examples indicating that, in the event of primal and dual attainment, the Nelder-Mead algorithm from [78] can be used to construct approximate optimizers, whereas this algorithm fails if dual attainment does.

Finally, we conclude with Section 3.7, which discusses two alternative methods for defining variance bounds and the associated primal and dual problems.

3.2 PRIMAL AND DUAL PROBLEMS FOR VARIANCE BOUNDS

We use this section to establish formulas and relationships between (1.10) and (1.11). When doing so, it will be helpful to establish notation for the mean vector and covariance matrix of a probability measure in $\mathcal{P}([0, \infty)^d)$:

Definition 3.3. Given a probability measure $\nu \in \mathcal{P}([0, \infty)^d)$, its mean is

$$\bar{\nu} := \int y d\nu(y) \in [0, \infty)^d$$

and its covariance matrix $\text{cov}(\nu) \in \mathbb{R}^d \times \mathbb{R}^d$ is given by

$$\text{cov}(\nu) := \left(\int_{[0, \infty)^d} (y_i - \bar{\nu}_i)(y_j - \bar{\nu}_j) d\nu(y) \right)_{i,j=1}^d.$$

Recall that $\text{cov}(\nu)$ is the closest multivariate analogue of the variance of a probability measure in $\mathcal{P}(\mathbb{R})$, as it tracks all information about the second moments of ν . While we will explore more ways to implement multivariate variance bounds in Section 3.7, for now, we will use the following simple notion:

Definition 3.4 (Variance Bounds on $\mathcal{P}([0, \infty)^d)$). Given $\tau \in [0, \infty)$, we say that $\nu \in \mathcal{P}([0, \infty)^d)$ **satisfies a (trace) variance bound with parameter τ** if:

$$\text{tr}(\text{cov}(\nu)) = \int_{[0, \infty)^d} |y - \bar{\nu}|^2 d\nu(y) \leq \tau.$$

In light of this definition, we can think of $\mathcal{I}_{X,Y}(\mu, m; \tau)$ as a version of $\mathcal{I}_{X,Y}(\mu, m)$ where we introduce the additional constraint that $\pi_{\#}^Y \gamma$ satisfies a variance bound with parameter τ . From this, it is immediately clear that $\mathcal{I}_{X,Y}(\mu, m; \tau) \geq \mathcal{I}_{X,Y}(\mu, m)$.

We now prove Proposition 1.13, which establishes weak duality between the primal problem (1.10) and the dual problem (1.11) using an identical strategy to the proof of Proposition 1.8.

Proof of Proposition 1.13. Assume that $m \in \text{conv}(Y)$, as otherwise $\mathcal{I}_{X,Y}(\mu, m; \tau) = \mathcal{S}_{X,Y}(\mu, m; \tau) = +\infty$. As before, introduce the terms defined in (2.1) and (2.2) to account for the prescribed first marginal and the

mean constraint on the second marginal, respectively. To account for the variance bound $\iint |y - m|^2 d\gamma(x, y) \leq \tau$, we introduce the term

$$\sup_{\eta \leq 0} \left\{ \eta \left(\tau - \iint |y - m|^2 d\gamma(x, y) \right) \right\}.$$

This yields an unconstrained inf-sup optimization problem. Using [92, Lemma 36.1] to interchange the infimum and the supremum, and grouping the terms that depend on γ , we recover a dual saddle problem of the form

$$\begin{aligned} \mathcal{I}_{X,Y}(\mu, m; \tau) \geq & \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2 \\ \eta \leq 0}} \left[\int \varphi d\mu + \lambda \cdot m + \eta \tau \right. \\ & \left. + \inf_{\gamma \in \mathcal{M}_+(X \times Y)} \left[\iint c(x, y) - \varphi(x) - \lambda \cdot y - \eta |y - m|^2 d\gamma(x, y) \right] \right]. \end{aligned}$$

Thus, we conclude by computing

$$\begin{aligned} & \inf_{\gamma \in \mathcal{M}_+(X \times Y)} \left[\iint c(x, y) - \varphi(x) - \lambda \cdot y - \eta |y - m|^2 d\gamma(x, y) \right] \\ & = \begin{cases} 0, & \text{if } c(x, y) \geq \varphi(x) + \lambda \cdot y + \eta |y - m|^2 \text{ on } X \times Y \\ +\infty, & \text{else,} \end{cases} \end{aligned}$$

yielding the desired result. \blacksquare

We immediately note that the variance bound resolves the main technical issue with the original dual problem (1.7) that we exploited in Section 2.3. Recall that, in that problem, the constraint takes the form

$$c(x, y) - \lambda \cdot y \geq \varphi(x) \text{ for } (x, y) \in X \times Y.$$

As we saw in Proposition 2.1 and Lemma 2.10, the 1-homogeneity of $y \mapsto c(x, y) - \lambda \cdot y$ severely restricts the choices of φ that are admissible in (1.7) – in particular, the best possible choices of φ and λ are $\varphi \equiv 0$ and $\lambda = (0, 0)$, leading to triviality of the dual problem.

For Problem 1.11, the constraint takes the form

$$c(x, y) - \lambda \cdot y - \eta |y - m|^2 \geq \varphi(x),$$

where $\eta \leq 0$. Here, a strictly negative choice of η means that the term $-\eta |y - m|^2$ dominates $-\lambda \cdot y$ as $|y| \rightarrow \infty$, greatly expanding the range of possible choices of φ ; the space of admissible functions is no longer as sensitive to the limiting behaviour of $y \mapsto c(x, y) - \lambda \cdot y$.

3.3 DUALITY

We now prove that duality holds in the presence of variance bounds. Our proof relies heavily on Friescheke's formulation of Fenchel-Rockafellar duality in [46, Theorem 6.2.2], which uses the following definitions:

Definition 3.5 (Infimal Convolution). Let V be a normed vector space, and let $\mathcal{F}, \mathcal{G} : V \rightarrow \mathbb{R} \cup \{+\infty\}$. Then we define the infimal convolution $\mathcal{F} \square \mathcal{G} : V \rightarrow \mathbb{R} \cup \{\pm\infty\}$ by

$$\mathcal{F} \square \mathcal{G}(v) = \inf_{v \in V} \{\mathcal{F}(u - v) + \mathcal{G}(v)\}.$$

Definition 3.6 (Fenchel-Legendre Transforms). Let V be a normed vector space and V^* be its topological dual. If $\mathcal{F} : V \rightarrow \mathbb{R} \cup \{+\infty\}$ is convex, define its Fenchel-Legendre transform $\mathcal{F}^* : V^* \rightarrow \mathbb{R} \cup \{\pm\infty\}$ by

$$\mathcal{F}^*(\ell) = \sup_{u \in V} [\langle \ell, u \rangle - \mathcal{F}(u)],$$

where $\langle \cdot, \cdot \rangle$ is the duality pairing.

For convenience, we also recall Friescheke's version of Fenchel-Rockafellar duality:

Theorem 3.7 (Fenchel-Rockafellar Duality without Continuity; Theorem 6.2.2 of [46]). *Let V be a normed vector space, and let V^* be its topological dual. Assume that:*

1. $\mathcal{F}, \mathcal{G} : V \rightarrow \mathbb{R} \cup \{+\infty\}$ are convex and proper.
2. The infimal convolution $\mathcal{H} := \mathcal{F} \square \mathcal{G}_- : V \rightarrow \mathbb{R} \cup \{\pm\infty\}$ does not take on the value $-\infty$, where \mathcal{G}_- is the reflection defined by $\mathcal{G}_-(u) = \mathcal{G}(-u)$.
3. \mathcal{H} is lower semicontinuous.

Then

$$\inf_{u \in V} (\mathcal{F}(u) + \mathcal{G}(u)) = \mathcal{H}(0) = \mathcal{H}^{**}(0) = \sup_{\ell \in V^*} (-\mathcal{F}^*(\ell) - \mathcal{G}^*(-\ell)).$$

In Friescheke's treatment of the classical optimal transportation problem [46, Theorem 3.2.3], he chooses $V = C_0(X) \times C_0(Y)$ (where $C_0(Z)$ refers to the space of continuous functions on Z vanishing at ∞), and chooses \mathcal{F} and \mathcal{G} so that $\inf_{u \in V} (\mathcal{F}(u) + \mathcal{G}(u))$ is a proxy for the dual problem in optimal transportation. However, our dual problem involves optimizing over functions $u \in \mathcal{C}(X \times Y)$ of the form

$$u(x, y) = \varphi(x) + \lambda \cdot y + \eta |y - m|^2,$$

so we use the following, tailored, topological vector space in the sequel:

Definition 3.8 (A Weighted Banach Space). Given $m \in [0, \infty)^2$, define $w_m : X \times Y \rightarrow \mathbb{R}$ by $w_m(x, y) = 1 + |y - m|^2$, and define the weighted Banach space $(C_{b,m}(X \times Y), \|\cdot\|_{\infty,m})$ by

$$C_{b,m}(X \times Y) = \left\{ u \in C(X \times Y) \mid \frac{u}{w_m} \in C_b(X \times Y) \right\}$$

and $\|u\|_{\infty,m} = \left\| \frac{u}{w_m} \right\|_{\infty}$.

We observe that $(C_{b,m}(X \times Y), \|\cdot\|_{\infty,m})$ is indeed a Banach space, as it may be checked that $T_m : (C_{b,m}(X \times Y), \|\cdot\|_{\infty,m}) \rightarrow (C_b(X \times Y), \|\cdot\|_{\infty})$ given by $T_m(u) = \frac{u}{w_m}$ is an isometric isomorphism.

Next, we compute the dual space $(C_{b,m}(X \times Y), \|\cdot\|_{\infty,m})^*$ in terms of that of $(C_b(X \times Y), \|\cdot\|_{\infty})$. In what follows, we use $\langle \cdot, \cdot \rangle$ for the duality pairing for $(C_b(X \times Y), \|\cdot\|_{\infty})$ and $\langle \cdot, \cdot \rangle_m$ for the duality pairing for $(C_{b,m}(X \times Y), \|\cdot\|_{\infty,m})$:

Lemma 3.9. *The adjoint mapping*

$$T_m^* : (C_b(X \times Y), \|\cdot\|_{\infty})^* \rightarrow (C_{b,m}(X \times Y), \|\cdot\|_{\infty,m})^*$$

defined in [94, Theorem 4.10] is an isometric isomorphism. As a consequence, for

$$\ell_m = T_m^*(\ell) \in (C_{b,m}(X \times Y), \|\cdot\|_{\infty,m})^*,$$

the action of ℓ_m on $u \in (C_{b,m}(X \times Y), \|\cdot\|_{\infty,m})$ is given by

$$\langle \ell_m, u \rangle_m = \left\langle \ell, \frac{u}{w_m} \right\rangle.$$

Proof of Lemma. This follows directly from the fact that T_m is an isometric isomorphism from $(C_{b,m}(X \times Y), \|\cdot\|_{\infty,m})$ to $(C_b(X \times Y), \|\cdot\|_{\infty})$ (see [73, Theorem 1.10.12]), and the definition of the adjoint in [94, Theorem 4.10].

■

At last, we prove duality in the sense of Theorem 1.14, under the mild additional assumption that μ have compact support in $[0, \infty) \times (0, \infty)$. For simplicity of notation, we will suppress the topologies throughout the following proof (for example, $C_{b,m}(X \times Y)^*$ will be understood to mean $(C_{b,m}(X \times Y), \|\cdot\|_{\infty,m})^*$ with the dual norm).

Proof of Theorem 1.14. Define $\mathcal{F}, \mathcal{G} : C_{b,m}(X \times Y) \rightarrow \mathbb{R} \cup \{+\infty\}$ by

$$\mathcal{F}(u) = \begin{cases} -(\int_X \varphi d\mu + \lambda \cdot m + \eta\tau) & \text{if } u(x, y) = -(\varphi(x) + \lambda \cdot y \\ & + \eta|y - m|^2) \text{ on } X \times Y \text{ where} \\ & \varphi \in C_b(X), \lambda \in \mathbb{R}^2, \text{ and } \eta \leq 0 \\ +\infty & \text{else,} \end{cases} \quad (3.1)$$

and

$$\mathcal{G}(u) = \begin{cases} 0 & \text{if } u \geq -c \text{ on } X \times Y \\ +\infty & \text{else,} \end{cases}.$$

Our proof proceeds in four steps. First, we check that \mathcal{F} is well-defined (it is clear that \mathcal{G} is). Second, we verify that the hypotheses of Theorem 3.7 indeed apply to \mathcal{F} and \mathcal{G} . Third, we directly apply Theorem 3.7 to get equality between $\mathcal{S}(\mu, m; \tau)$ and a suitable minimization problem on $(C_{b,m}(X \times Y), \|\cdot\|_{\infty, m})^*$. Finally, we use techniques broadly similar to [103, Lemma 1.25] to conclude that this minimization problem is equivalent to $\mathcal{I}(\mu, m; \tau)$.

We begin by checking that \mathcal{F} is well-defined. To this end, we assume that

$$\varphi(x) + \lambda \cdot y + \eta|y - m|^2 = \tilde{\varphi}(x) + \tilde{\lambda} \cdot y + \tilde{\eta}|y - m|^2 \quad (3.2)$$

on $X \times Y$, and show that $\int_X \varphi d\mu + \lambda \cdot m + \eta\tau = \int_X \tilde{\varphi} d\mu + \tilde{\lambda} \cdot m + \tilde{\eta}\tau$. First, rearrange Equation (3.2) so that all terms involving x are on one side, and all terms involving y are on the other, to deduce that there exists $s \in \mathbb{R}$ such that $\varphi(x) \equiv \tilde{\varphi}(x) + s$ and $\lambda \cdot y + \eta|y - m|^2 \equiv \tilde{\lambda} \cdot y + \tilde{\eta}|y - m|^2 - s$. Observing that the latter equation equates two quadratic polynomials in $y = (y_1, y_2)$, we deduce that $(\lambda, \eta, s) = (\tilde{\lambda}, \tilde{\eta}, 0)$. Finally, since $s = 0$, we deduce that $\varphi = \tilde{\varphi}$.

We now enter the second stage of our proof, which involves checking the hypotheses of Theorem 3.7. To begin, we note that, by construction, \mathcal{F} and \mathcal{G} are both linear, and hence convex, on their domains. Moreover, by taking $(\varphi, \lambda, \eta) = (0, 0, 0)$, we see that $\mathcal{F}(0) = \mathcal{G}(0) = 0$. Hence, both \mathcal{F} and \mathcal{G} are proper.

Next, we check that the infimal convolution $\mathcal{H} = \mathcal{F} \square \mathcal{G}_-$ does not take on the value $-\infty$. By definition, for $u \in C_{b,m}(X \times Y)$,

$$\mathcal{H}(u) = \mathcal{F} \square \mathcal{G}_-(u) = \inf_{v \in C_{b,m}(X \times Y)} [\mathcal{F}(u - v) + \mathcal{G}(-v)].$$

Using the definitions of \mathcal{F} and \mathcal{G} , we minimize over the set of v for which $u - v$ is in the effective domain of \mathcal{F} and $-v$ is in the effective domain of \mathcal{G} . That is, we may assume that

$$v(x, y) = u(x, y) + \varphi(x) + \lambda \cdot y + \eta|y - m|^2$$

for some $(\varphi, \lambda, \eta) \in C_b(X) \times \mathbb{R}^2 \times (-\infty, 0]$ and that $v(x, y) \leq c(x, y)$. Combining these conditions, and factoring out the negative sign to replace the infimum with a supremum, we find that

$$\begin{aligned} \mathcal{H}(u) = - \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2, \eta \geq 0}} \left\{ \int \varphi \, d\mu + \lambda \cdot m + \eta\tau \mid \varphi(x) + \lambda \cdot y + \eta|y - m|^2 \right. \\ \left. \leq c(x, y) - u(x, y) \text{ on } X \times Y \right\}. \end{aligned} \quad (3.3)$$

Now, choose $\gamma_0 \in \mathcal{P}_c(X \times Y)$ satisfying $\pi_{\#}^X \gamma_0 = \mu$, $\iint y \, d\gamma_0 = m$, and $\iint |y - m|^2 \, d\gamma_0 = \tau$ (for example, one can take $\gamma_0 = \mu \otimes \nu$, where $\nu = \frac{\tau}{|m|^2 + \tau} \delta_0 + \frac{|m|^2}{|m|^2 + \tau} \delta_{\frac{|m|^2}{|m|}}$). This allows us to rewrite and estimate

$$\begin{aligned} \mathcal{H}(u) = - \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2, \eta \geq 0}} \left\{ \iint \varphi(x) + \lambda \cdot y + \eta|y - m|^2 \, d\gamma_0(x, y) \mid \varphi(x) \right. \\ \left. + \lambda \cdot y + \eta|y - m|^2 \leq c(x, y) - u(x, y) \text{ on } X \times Y \right\} \\ \geq - \iint (c - u) \, d\gamma_0 \\ > -\infty, \end{aligned}$$

where the penultimate inequality arises from plugging in the constraint, and the last inequality follows from continuity of c and u and the compact support of γ_0 .

Finally, we check the last hypothesis of Theorem 3.7, lower semicontinuity of \mathcal{H} on $C_{b,m}(X \times Y)$. To do so, we note that, as in [46, Section 6.2.4], it suffices to check continuity of \mathcal{H} by verifying that, if $u, v \in C_{b,m}(X \times Y)$ satisfy $\|u - v\|_{\infty, m} \leq \varepsilon$, then $\mathcal{H}(v) \geq \mathcal{H}(u) - (1 + \tau)\varepsilon$. Take $(\varphi, \lambda, \eta) \in C_b(X) \times \mathbb{R}^2 \times (-\infty, 0]$ admissible for $H(v)$ as in (3.3) so that

$$c(x, y) - \varphi(x) - \lambda \cdot y - \eta|y - m|^2 \geq v(x, y) \geq u(x, y) - \varepsilon(1 + |y - m|^2),$$

where the last inequality comes from the assumption that $\|u - v\|_{\infty, m} \leq \varepsilon$. Moving $\varepsilon(1 + |y - m|^2)$ to the left side, we see that the trio $(\varphi - \varepsilon, \lambda, \eta - \varepsilon)$ is admissible for $\mathcal{H}(u)$. Thus, by definition,

$$\begin{aligned} - \int \varphi \, d\mu - \lambda \cdot m - \eta\tau + (1 + \tau)\varepsilon &= - \int (\varphi - \varepsilon) \, d\mu - \lambda \cdot m - (\eta - \varepsilon)\tau \\ &\geq \mathcal{H}(u). \end{aligned}$$

By rearranging and minimizing over (φ, λ, η) , we recover the desired conclusion.

Next, we move onto the third part of the proof, which is extracting the raw conclusion from Theorem 3.7. First, we notice that, by the definition of \mathcal{F} and \mathcal{G} ,

$$\inf_{u \in C_{b,m}(X \times Y)} (\mathcal{F}(u) + \mathcal{G}(u)) = -\mathcal{S}(\mu, m; \tau).$$

Thus, we proceed to compute Fenchel-Legendre transforms \mathcal{F}^* and \mathcal{G}^* , as well as $\sup_{\ell_m \in C_{b,m}(X \times Y)^*} (-\mathcal{F}^*(\ell_m) - \mathcal{G}^*(-\ell_m))$. Under our previous notation, Lemma 3.9 allows us to write each $\ell_m \in C_{b,m}^*$ in the form $\ell_m = T_m^*(\ell)$ for some $\ell \in C_b^*$. Moreover, by the definition of Fenchel-Legendre transforms and Lemma 3.9,

$$\begin{aligned} \mathcal{F}^*(\ell_m) &= \sup_{u \in C_{b,m}} [\langle \ell_m, u \rangle_m - \mathcal{F}(u)] \\ &= \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2, \eta \geq 0}} \left\{ \langle \ell_m, -(\varphi(x) + \lambda \cdot y + \eta|y - m|^2) \rangle_m \right. \\ &\quad \left. + \int \varphi \, d\mu + \lambda \cdot m + \eta\tau \right\} \\ &= \begin{cases} 0 & \text{if } \langle \ell_m, \varphi(x) \rangle_m = \int \varphi \, d\mu \, \forall \varphi \in C_b(X), \\ & \langle \ell_m, y \rangle_m = m, \text{ and } \langle \ell_m, |y - m|^2 \rangle_m \leq \tau. \\ +\infty & \text{else} \end{cases} \end{aligned}$$

Note that, in the first line, we applied the conditions for \mathcal{F} to be finite from Equation (3.1), and in the final line we scaled the dual variables to recover conditions on ℓ_m .

Likewise, we compute

$$\begin{aligned} \mathcal{G}^*(-\ell_m) &= \sup_{u \in C_{b,m}} [\langle -\ell_m, u \rangle_m - G(u)] \\ &= \sup_{-c \leq u \in C_{b,m}} \langle \ell_m, -u \rangle_m \\ &= \sup_{c \geq u \in C_{b,m}} \langle \ell_m, u \rangle_m. \end{aligned}$$

We observe that $\mathcal{G}^*(-\ell_m) < +\infty$ only if ℓ_m is non-negative in the sense that if $u \geq 0$, then $\langle \ell_m, u \rangle_m \geq 0$, in which case we write $\ell_m \geq 0$. Otherwise, there exists some $u_0 \in C_{b,m}$ satisfying $u_0 \leq 0 \leq c$ but $\langle \ell_m, u_0 \rangle_m > 0$, and we may take the supremum to $+\infty$ by rescaling.

Thus, in summary,

$$\begin{aligned} \mathcal{S}(\mu, m; \tau) &= \inf_{\ell_m \in C_{b,m}^*} (\mathcal{F}^*(\ell_m) + \mathcal{G}^*(-\ell_m)) \tag{3.4} \\ &= \inf_{\ell_m \in C_{b,m}^*} \left\{ \sup_{c \geq u \in C_{b,m}} \langle \ell_m, u \rangle_m \mid \ell_m \geq 0, \langle \ell_m, \varphi(x) \rangle_m = \int \varphi d\mu \right. \\ &\quad \left. \forall \varphi \in C_b(X), \langle \ell_m, y \rangle_m = m, \text{ and} \right. \\ &\quad \left. \langle \ell_m, |y - m|^2 \rangle_m \leq \tau \right\}. \end{aligned}$$

We now check that the desired duality result holds by showing that, for any ℓ_m admissible in (3.4), there exists some measure γ_m which is also admissible in (3.4) and which satisfies

$$\sup_{c \geq u \in C_{b,m}} \langle \ell_m, u \rangle_m \geq \sup_{c \geq u \in C_{b,m}} \langle \gamma_m, u \rangle = \iint c d\gamma_m.$$

Writing $\ell_m = \tau_m^*(\ell)$ for $\ell \in C_b^*$ we recall that, if $u \in C_{b,m}$, then $\langle \ell_m, u \rangle_m = \langle \ell, \frac{u}{w_m} \rangle$, where $\frac{u}{w_m} \in C_b$. By [103, Lemma 1.24], we may decompose $\ell = \gamma + R$, where $\gamma \in \mathcal{M}(X \times Y)$ is a non-negative measure and R is a non-negative continuous linear functional supported at infinity, in the sense that $\langle R, v \rangle = 0$ for $v \in C_0(X \times Y)$. Thus, we define our candidate γ_m by $\iint u d\gamma_m = \iint \frac{u}{w_m} d\gamma$. By non-negativity of R , we have that

$$\langle \ell_m, u \rangle_m = \langle \ell, u \rangle = \left\langle \gamma, \frac{u}{w_m} \right\rangle + \left\langle R, \frac{u}{w_m} \right\rangle \geq \langle \gamma_m, u \rangle_m,$$

and hence, optimizing over $u \leq c$, we have that

$$\sup_{c \geq u \in C_{b,m}} \langle \ell_m, u \rangle_m \geq \sup_{c \geq u \in C_{b,m}} \langle \gamma_m, u \rangle_m = \iint c d\gamma.$$

Thus, we need only to check admissibility of γ_m . Without loss of generality, we assume that X is compact, as μ is compactly supported. First, for the X -marginal constraint, we observe that, for any $\varphi \in C_b(X)$,

$$\int \varphi d\mu = \langle \ell_m, \varphi \rangle_m = \iint \varphi d\gamma_m + \langle R, \frac{\varphi}{1 + |y - m|^2} \rangle.$$

Since X is compact, $\frac{\varphi}{1 + |y - m|^2} \in C_0(X \times Y)$, and hence the term involving R vanishes. Likewise, for the mean constraint, we find that

$$m = \iint y d\gamma_m + \langle R, \frac{y}{1 + |y - m|^2} \rangle,$$

and the term involving R vanishes as $\frac{y}{1 + |y - m|^2} \in C_b(X \times Y)$. Finally, γ_m also satisfies the variance constraint, as R and γ_m are non-negative, and hence

$$\begin{aligned} \tau &\geq \langle \ell_m, |y - m|^2 \rangle_m \\ &= \iint |y - m|^2 d\gamma_m + \langle R, \frac{|y - m|^2}{1 + |y - m|^2} \rangle \\ &\geq \iint |y - m|^2 d\gamma_m. \end{aligned}$$

From this, we conclude that

$$\begin{aligned} \mathcal{S}(\mu, m; \tau) &= \inf_{\gamma \in \mathcal{M}_+(X \times Y)} \left\{ \iint c d\gamma \mid \pi_{\#}^X \gamma = \mu, \right. \\ &\quad \left. \iint y d\gamma = m, \iint |y - m|^2 d\gamma \leq \tau \right\} \\ &= \mathcal{I}(\mu, m; \tau), \end{aligned}$$

where in the last line we have used the condition that $\pi_{\#}^X \gamma = \mu$ to deduce that γ is a probability measure. ■

3.4 PRIMAL ATTAINMENT AND DUAL NON-ATTAINMENT

Now that we have established the duality relation $\mathcal{I}(\mu, m; \tau) = \mathcal{S}(\mu, m; \tau)$, it is natural to ask whether $\mathcal{I}(\mu, m; \tau)$ and $\mathcal{S}(\mu, m; \tau)$ are attained. We start by proving Proposition 3.1 by showing attainment of $\mathcal{I}(\mu, m; \tau)$ using the direct method of the calculus of variations:

Proof of Proposition 3.1. We use the direct method of the calculus of variations, where we endow the space $\mathcal{M}_+(X \times Y)$ of non-negative finite signed measures with the weak topology in duality with $C_b(X \times Y)$. Since

c is non-negative and lower semicontinuous by construction in Equation (1.5), [96, Lemma 1.6] implies that the map $\gamma \mapsto \iint c \, d\gamma$ is weakly lower semicontinuous.

Now, let $(\gamma_n)_n$ be a minimizing sequence for $\mathcal{I}(\mu, m; \tau)$. We check that this is a tight sequence in the sense that, for any $\varepsilon > 0$, there exists a compact $K_\varepsilon \subseteq X \times Y$ such that, for any n , $\gamma_n(K_\varepsilon^c) \leq \varepsilon$. By Prokhorov's Theorem, as stated in [96, Box 1.4], there exists a compact $X_\varepsilon \subseteq X$ such that $\mu(X_\varepsilon^c) \leq \frac{\varepsilon}{2}$. Hence, by the constraint $\pi_{\#}^X \gamma = \mu$, we have that $\gamma_n(X_\varepsilon^c \times Y) \leq \frac{\varepsilon}{2}$ for any n . Likewise, we define $R := \sqrt{2\tau/\varepsilon}$ and $Y_\varepsilon := Y \cap \overline{B_R(m)}$, and note that the constraint $\iint |y - m|^2 \, d\gamma \leq \tau$ implies that, for any n ,

$$\tau \geq \iint_{X \times Y_\varepsilon^c} |y - m|^2 \, d\gamma_n \geq R^2 \gamma_n(X \times Y_\varepsilon^c) = \frac{2\tau}{\varepsilon} \gamma_n(X \times Y_\varepsilon^c),$$

so that $\gamma_n(X \times Y_\varepsilon^c) \leq \frac{\varepsilon}{2}$ after rearranging. Thus, taking the compact set $K_\varepsilon = X_\varepsilon \times Y_\varepsilon$, we find that

$$\gamma_n(K_\varepsilon^c) \leq \gamma_n(X_\varepsilon^c \times Y) + \gamma_n(X \times Y_\varepsilon^c) \leq \varepsilon$$

for any n . Combining this with the continuity of $\gamma \mapsto \iint c \, d\gamma$, we conclude that γ_n converges subsequentially to some $\gamma_\infty \in \mathcal{P}(X \times Y)$ attaining the minimum for $\mathcal{I}(\mu, m; \tau)$. \blacksquare

However, Proposition 3.2 provides a counterexample to dual attainment in the case $\mu = \delta_{x_0}$, as we will see in the following proof:

Proof of Proposition 3.2. For the primal problem (1.10), Jensen's inequality and convexity of $c(x_0, \cdot)$ imply that, for any candidate measure γ ,

$$\iint c(x_0, y) \, d\gamma(x_0, y) \geq c\left(x_0, \iint y \, d\gamma(x, y)\right) = c(x_0, m)$$

and hence $\mathcal{I}(\mu, m; \tau) \geq c(x_0, m) > 0$. By taking $\gamma = \delta_{(x_0, m)}$, we recover equality.

For the dual problem, we substitute $(x, y) = (x_0, m)$ into the constraint equation from (1.11) to deduce that $\varphi(x_0) \leq c(x_0, m) - \lambda \cdot m$. Hence, since $\mu = \delta_{x_0}$,

$$\int \varphi \, d\mu + \lambda \cdot m + \eta\tau \leq c(x_0, m) + \eta\tau \leq c(x_0, m) = \mathcal{I}(\mu, m; \tau).$$

Moreover, equality requires that both $\varphi(x_0) = c(x_0, m) - \lambda \cdot m$ and $\eta\tau = 0$. Using that $\tau > 0$ and $\eta \leq 0$, we deduce that dual attainment may hold only if $\eta = 0$. However, in this case, Lemma 2.10 implies that $c(x_0, y) - \lambda \cdot y$ is bounded below (and hence an admissible φ exists) if and only if $\lambda_1, \lambda_2 \leq 0$.

Moreover, as Lemma 2.10 states that, in this case, $\varphi(x_0) \leq \inf_y [c(x_0, y) - \lambda \cdot y] = 0$. Combining these observations, we deduce that

$$\varphi(x_0) + \lambda \cdot m + \eta\tau \leq 0 < \mathcal{I}(\mu, m; \tau),$$

as desired. ■

Remark 3.10 (Interpretation of Proposition 3.2). While, at first, there seems to be some tension between Theorem 1.14 and Proposition 3.2, we reconcile this apparent discrepancy as follows. In particular, we interpret Proposition 3.2 as stating that the approximate optimizers of $\mathcal{S}(\mu, m; \tau)$ (which are necessary for the duality statement in Theorem 1.14) must universally satisfy $\eta < 0$. In other words, this witnesses that the regularization term $-\eta|y - m|^2$ in $\mathcal{S}(\mu, m; \tau)$ is indeed vital for duality.

3.5 ANALYSIS OF THE VARIANCE BOUND PROBLEM WITH ELKKT CONDITIONS

We now define Euler-Lagrange Karush-Kuhn-Tucker conditions for trace variance bounds, in order to partially characterize optimizers for $\mathcal{I}(\mu, m; \tau)$ and $\mathcal{S}(\mu, m; \tau)$, if they exist. To begin, it will be useful to express both the primal problem (1.10) and the dual problem (1.11) as unconstrained optimization problems. To do so, define $H : \mathcal{M}_+(X \times Y) \times (C_b(X) \times \mathbb{R}^2 \times (-\infty, 0]) \rightarrow \mathbb{R}$ by

$$\begin{aligned} H(\gamma; \varphi, \lambda, \eta) &= \iint c \, d\gamma + \int \varphi \, d\mu - \iint \varphi(x) \, d\gamma \\ &\quad + \lambda \cdot \left(m - \iint y \, d\gamma \right) + \eta \left(\tau - \iint |y - m|^2 \, d\gamma \right). \end{aligned}$$

Under this notation, we observe that

$$\mathcal{I}_{X,Y}(\mu, m; \tau) = \inf_{\gamma \in \mathcal{M}_+(X \times Y)} \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2 \\ \eta \in (-\infty, 0]}} H(\gamma; \varphi, \lambda, \eta) \quad (3.5)$$

and

$$\mathcal{S}_{X,Y}(\mu, m; \tau) = \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2 \\ \eta \in (-\infty, 0]}} \inf_{\gamma \in \mathcal{M}_+(X \times Y)} H(\gamma; \varphi, \lambda, \eta). \quad (3.6)$$

This leads us to the following proposition about the primal problem:

Proposition 3.11 (ELKKT Conditions for the Primal Trace Problem). *Given $\gamma \in \mathcal{P}(X \times Y)$, assume that $(\varphi_0, \lambda_0, \eta_0) \in C_b(X) \times \mathbb{R}^2 \times (-\infty, 0]$ attains the inner supremum in (3.5), i.e.*

$$H(\gamma; \varphi_0, \lambda_0, \eta_0) = \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2 \\ \eta \in (-\infty, 0]}} H(\gamma; \varphi, \lambda, \eta). \quad (3.7)$$

Then, $\pi_{\#}^X \gamma = \mu$, $\iint y \, d\gamma = m$, and $\iint |y - m|^2 \, d\gamma \leq \tau$. Furthermore, if $\eta_0 < 0$, then $\iint |y - m|^2 \, d\gamma = \tau$.

Proof. If $\eta_0 < 0$ then, given $(\varphi_1, \lambda_1, \eta_1) \in C_b(X) \times \mathbb{R}^2 \times (0, \infty)$, we deduce that $(\varphi_0 + t\varphi_1, \lambda_0 + t\lambda_1, \eta_0 + t\eta_1) \in C_b(X) \times \mathbb{R}^2 \times (-\infty, 0]$ for small t . Since $(\varphi_0, \lambda_0, \eta_0)$ is assumed to satisfy (3.7), it follows that $\left. \frac{d}{dt} \right|_{t=0} H(\gamma; \varphi_0 + t\varphi_1, \lambda_0 + t\lambda_1, \eta_0 + t\eta_1) = 0$. Expanding and computing, we deduce that

$$0 = \int \varphi_1 \, d\mu - \iint \varphi_1 \, d\gamma + \lambda_1 \cdot \left(m - \iint y \, d\gamma \right) + \eta_1 \left(\tau - \iint |y - m|^2 \, d\gamma \right).$$

Since $(\varphi_1, \lambda_1, \eta_1) \in C_b(X) \times \mathbb{R}^2 \times \mathbb{R}$ was arbitrary, we use the Riesz Representation Theorem and linear algebra to conclude that $\pi_{\#}^X \gamma = \mu$, $\iint y \, d\gamma = m$, and $\iint |y - m|^2 \, d\gamma = \tau$, as claimed.

On the other hand, if $\eta_0 = 0$, then $\eta_0 + t\eta_1 \leq 0$ if and only if $t\eta_1 \leq 0$. Thus, we fix $\eta_1 \leq 0$ and instead use one-sided derivative condition

$$0 \geq \lim_{t \rightarrow 0^+} \frac{H(\gamma; \varphi_0 + t\varphi_1, \lambda_0 + t\lambda_1, \eta_0 + t\eta_1) - H(\gamma; \varphi_0, \lambda_0, \eta_0)}{t}.$$

By considering $\pm\varphi_1$ and $\pm\lambda_1$, we deduce that $\pi_{\#}^X \gamma = \mu$ and $\iint y \, d\gamma = m$. However, for $\eta_1 \leq 0$, the one-sided condition only implies that $0 \geq \eta_1 \left(\tau - \iint |y - m|^2 \, d\gamma \right)$, or rather, since η_1 is arbitrary, that $\iint |y - m|^2 \, d\gamma \leq \tau$, as claimed. \blacksquare

Proposition 3.12 (ELKKT Conditions for the Dual Trace Problem). *Given $(\varphi, \lambda, \eta) \in C_b(X) \times \mathbb{R}^2 \times (-\infty, 0]$, assume that $\gamma_0 \in \mathcal{P}(X \times Y)$ attains the inner infimum in (3.6), i.e.*

$$H(\gamma_0; \varphi, \lambda, \eta) = \inf_{\gamma \in \mathcal{M}_+(X \times Y)} H(\gamma; \varphi, \lambda, \eta).$$

Then, $c(x, y) - \varphi(x) - \lambda \cdot y - \eta|y - m|^2 \geq 0$, with equality on $\text{spt}(\gamma_0)$.

Proof. First fix $\gamma_1 \in \mathcal{M}_+(X \times Y)$, and notice that $\gamma_0 + t\gamma_1 \in \mathcal{M}_+(X \times Y)$ for any $t \geq 0$. Thus, by the condition that γ_0 is a minimum,

$$\begin{aligned} 0 &\leq \lim_{t \rightarrow 0^+} \frac{H(\gamma_0 + t\gamma_1; \varphi, \lambda, \eta) - H(\gamma_0; \varphi, \lambda, \eta)}{t} \\ &= \iint c(x, y) - \varphi(x) - \lambda \cdot y - \eta|y - m|^2 d\gamma_1(x, y). \end{aligned}$$

Since γ_1 is arbitrary, we deduce that $c(x, y) - \varphi(x) - \lambda \cdot y - \eta|y - m|^2 \geq 0$ on $X \times Y$.

To get equality on $\text{spt}(\gamma_0)$, we fix $(x_0, y_0) \in \text{spt}(\gamma_0)$ and $r > 0$, and consider the renormalized restriction $\gamma_1 := \frac{1}{\gamma_0(B_r(x_0, y_0))} \gamma_0|_{B_r(x_0, y_0)}$ of γ_0 to $B_r(x_0, y_0)$. Since $\gamma_0 + t\gamma_1 \in \mathcal{M}_+(X \times Y)$ for small t , we conclude that

$$\begin{aligned} 0 &= \left. \frac{d}{dt} \right|_{t=0} H(\gamma_0 + t\gamma_1; \varphi, \lambda, \eta) \\ &= \iint c(x, y) - \varphi(x) - \lambda \cdot y - \eta|y - m|^2 d\gamma_1(x, y) \\ &= \frac{1}{\gamma(B_r(x_0, y_0))} \iint_{B_r(x_0, y_0)} c(x, y) - \varphi(x) - \lambda \cdot y - \eta|y - m|^2 d\gamma_0(x, y). \end{aligned}$$

Thus, by (for example) Theorem 1 in [39, Section 1.7], we conclude that $c(x, y) - \varphi(x) - \lambda \cdot y - \eta|y - m|^2 = 0$ for γ -a.e. (x, y) . ■

Fortunately, the conditions we have derived are both necessary and sufficient for optimality of γ .

Proposition 3.13 (Sufficient Conditions for a Saddle Point). *Assume that $\gamma_0 \in \mathcal{M}_+(X \times Y)$ and $(\varphi_0, \lambda_0, \eta_0) \in C_b(X) \times \mathbb{R}^2 \times (-\infty, 0]$ satisfy the primal constraints $\pi_{\#}^X \gamma_0 = \mu$, $\iint y d\gamma_0(x, y) = m$, and $\iint |y - m|^2 d\gamma_0(x, y) \leq \tau$ from (1.10). If, moreover, both of the following hold:*

- $c(x, y) - \varphi_0(x) - \lambda_0 \cdot y - \eta_0|y - m|^2 \geq 0$, with equality γ_0 -a.e.
- Either $\eta_0 = 0$ or $\iint |y - m|^2 d\gamma_0(x, y) = \tau$,

then

$$\sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2 \\ \eta \in (-\infty, 0]}} H(\gamma_0; \varphi, \lambda, \eta) = H(\gamma_0; \varphi_0, \lambda_0, \eta_0) = \inf_{\gamma \in \mathcal{M}_+(X \times Y)} H(\gamma; \varphi_0, \lambda_0, \eta_0). \quad (3.8)$$

As a consequence,

$$\mathcal{S}_{X,Y}(\mu, m; \tau) = H(\gamma_0; \varphi_0, \lambda_0, \eta_0) = \mathcal{I}_{X,Y}(\mu, m; \tau).$$

In other words, γ_0 and $(\varphi_0, \lambda_0, \eta_0)$ optimize the primal and dual problems, respectively.

Proof. We begin by checking the first identity in (3.8). As such, we fix $(\varphi_1, \lambda_1, \eta_1) \in C_b(X) \times \mathbb{R}^2 \times (-\infty, 0]$ and show that $H(\gamma_0; \varphi_1, \lambda_1, \eta_1) \leq H(\gamma_0; \varphi_0, \lambda_0, \eta_0)$. After expanding and rearranging, we estimate

$$\begin{aligned} H(\gamma_0; \varphi_0, \lambda_0, \eta_0) &= H(\gamma_0; \varphi_1, \lambda_1, \eta_1) + \iint c \, d\gamma_0 - \iint c \, d\gamma_0 \\ &\quad + \int (\varphi_0 - \varphi_1) \, d\mu - \iint (\varphi_0 - \varphi_1) \, d\gamma_0 \\ &\quad + (\lambda_0 - \lambda_1) \cdot \left(m - \iint y \, d\gamma_0 \right) \\ &\quad + (\eta_0 - \eta_1) \left(\tau - \iint |y - m|^2 \, d\gamma_0 \right) \\ &\geq H(\lambda_0; \varphi_1, \lambda_1, \eta_1). \end{aligned}$$

Here, the given assumptions ensure that all terms except $H(\gamma_0; \varphi_1, \lambda_1, \eta_1)$ and $(\eta_0 - \eta_1) (\tau - \iint |y - m|^2 \, d\gamma_0)$ vanish, and the latter is non-negative since $\eta_1 \leq 0$, $\iint |y - m|^2 \, d\gamma_0(x, y) \leq \tau$, and either $\eta_0 = 0$ or $\iint |y - m|^2 \, d\gamma_0(x, y) = \tau$.

Next, we check the second identity in (3.8). To do so, fix $\gamma_1 \in \mathcal{M}_+(X \times Y)$ and compute

$$\begin{aligned} H(\gamma_1; \varphi_0, \lambda_0, \eta_0) &= H(\gamma_0, \varphi_0, \lambda_0, \eta_0) \\ &\quad + \iint c(x, y) - \varphi_0(x) - \lambda \cdot y - \eta_0 |y - m|^2 \, d\gamma_1 \\ &\quad - \iint c(x, y) - \varphi_0(x) - \lambda \cdot y - \eta_0 |y - m|^2 \, d\gamma_0 \\ &\geq H(\gamma_0, \varphi_0, \lambda_0, \eta_0), \end{aligned}$$

where the inequality follows from the assumption that $c(x, y) - \varphi_0(x) - \lambda \cdot y - \eta_0 |y - m|^2 \geq 0$, with equality γ_0 -a.e.

Finally, we show that $(\gamma_0, \varphi_0, \lambda_0, \eta_0)$ optimizes both the primal and dual problems. To see why, we apply Proposition 1.13 as well as Equations (3.5), (3.6), and (3.8) to find that

$$\begin{aligned} \mathcal{S}_{X,Y}(\mu, m; \tau) &\leq \mathcal{I}_{X,Y}(\mu, m; \tau) \\ &\leq \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2 \\ \eta \in (-\infty, 0]}} H(\gamma_0; \varphi, \lambda, \eta) \\ &= H(\gamma_0; \varphi_0, \lambda_0, \eta_0) \\ &= \inf_{\gamma \in \mathcal{M}_+(X \times Y)} H(\gamma; \varphi_0, \lambda_0, \eta_0) \\ &\leq \mathcal{S}_{X,Y}(\mu, m; \tau), \end{aligned}$$

and hence equality holds throughout. \blacksquare

Remark 3.14 (Relation to Strong Duality). Recall that Theorem 1.14 only applies to $X = Y = [0, \infty)^2$, but Proposition 3.13 applies to general $X, Y \subseteq [0, \infty)^2$. Hence, we may only use the weak duality of Proposition 1.13 in the preceding proof. Nevertheless, this proof suggests that a careful examination of the ELKKT conditions might provide an alternative proof of duality in some contexts; Proposition 3.2 indicates that any such result would not apply in full generality.

Remark 3.15 (Focus on the Case $\eta_0 \neq 0$). If $(\gamma_0, \varphi_0, \lambda_0, 0)$ satisfies the conditions in Proposition 3.13, then, by tweaking the proof of that proposition, we find that $(\gamma_0, \varphi_0, \lambda_0)$ witnesses strong duality between (1.6) and (1.7), in the sense that $\mathcal{I}_{X,Y}(\mu, m) = H(\gamma_0, \varphi_0, \lambda_0, 0) = \mathcal{S}_{X,Y}(\mu, m)$. As such, when there is a duality gap (as is the case if of Proposition 2.6 or Proposition 2.7 applies), quartets of the form $(\gamma_0, \varphi_0, \lambda_0, 0)$ cannot satisfy the conditions of Proposition 3.13. Moreover, in this case, Proposition 2.1 implies that $\varphi_0 \equiv 0$ and $\lambda_0 = (0, 0)$, so it is unlikely that useful information can be recovered from the dual problem. In light of this, we restrict some of the following results to the case $\eta_0 < 0$.

Remark 3.16 (Motivating Cost is not Twisted). At this stage, it is natural to use a twist condition on c to write the unknown second marginal $\pi_{\#}^Y \gamma$ in terms of the known marginal μ and the optimal dual variables $(\varphi_0, \lambda_0, \eta_0)$ (see [69, Section 2.5] for a full discussion of the theory). However, even on $(0, \infty)^2 \times (0, \infty)^2$, we compute

$$\nabla_x c(x, y) = \frac{y_2}{x_2} \left(x_1 - \frac{y_1}{y_2} \right) \left[-\frac{1}{2x_2} \left(x_1 - \frac{y_1}{y_2} \right) \right],$$

which implies that, for any $t > 0$, $\nabla_x c(x, (tx_1, t)) = (0, 0)$. In other words, for fixed $x \in (0, \infty)^2$, the map $y \mapsto \nabla_x c(x, y)$ fails to be injective, meaning that we must take more care to derive structural results on γ .

At this point, we fix domains $X = Y = [0, \infty)^2$, assume the existence of a quartet $(\gamma_0, \varphi_0, \lambda_0, \eta_0)$ satisfying the conditions of Proposition 3.13, and study its properties. We first use the condition

$$c(x, y) - \varphi_0(x) - \lambda_0 \cdot y - \eta_0 |y - m|^2 \geq 0, \quad (3.9)$$

with equality on $\text{spt}(\gamma_0)$, to characterize γ_0 and φ_0 in terms of λ_0 and η_0 . Once we've done so, we indicate how the mean and variance constraints might be used to recover the values (λ_0, η_0) , again assuming the primal and dual attainment of Proposition 3.13.

We begin by defining a function corresponding to (3.9), identifying some of its useful properties, and proving that it has a unique minimum:

Lemma 3.17 (A Useful Function). *Given $(\varphi, \lambda, \eta) \in C_b([0, \infty)^2) \times \mathbb{R}^2 \times (-\infty, 0)$, define the function*

$$g_x(y) := c(x, y) - \varphi(x) - \lambda \cdot y - \eta|y - m|^2.$$

This function is proper, strictly convex, lower semicontinuous, and coercive on $[0, \infty)^2$ in the sense that $\lim_{\|y\| \rightarrow \infty} g_x(y) = +\infty$. As such, for any $x \in [0, \infty)^2$, $g_x(y)$ attains a unique minimum on $[0, \infty)^2$.

Proof. First note that, by Equation (1.5), $c(x, (y_2x_1, y_2)) < +\infty$, and hence g_x is proper. Likewise, strict convexity follows from observing that $c(x, y)$, $\lambda \cdot y$, and $-\eta|y - m|^2$ are all convex, with the latter being strictly convex for $\eta < 0$. Moreover, g_x is lower semicontinuous as the sum of lower semicontinuous functions. Finally, coercivity comes from the $-\eta|y - m|^2$ term and the assumption that $\eta < 0$. Hence, by a variant of the Extreme Value Theorem (e.g. [93, Theorem 1.9]), g_x attains a minimum on $[0, \infty)^2$; uniqueness follows from strict convexity. ■

We will later explicitly identify where the minimum of g_x is attained, and the following lemma summarizes some properties of roots of cubic equations which are useful for this task.

Lemma 3.18 (Properties of Cubic Roots). *Let u_0 be the largest root of the cubic $u^3 + pu + q = 0$. Then $u_0 = u_0(p, q)$ is continuous for $q \leq 0$, and C^1 for $q < 0$. Moreover, if $q = 0$, then $u_0(p, 0) = \sqrt{\max(0, -p)}$.*

Proof. For fixed (p, q) , $f(u, p, q) := u^3 + pu + q$ is convex for $u \in [0, \infty)$. If we also assume that $q < 0$, then $f(0) = q < 0$ and $\lim_{u \rightarrow \infty} f(u) = +\infty$. Thus, in this range, $f(\cdot, p, q)$ has a unique non-negative root u_0 . Likewise, if $q = 0$, we factor the cubic $u^3 + pu$ to find that $u_0 = \sqrt{\max(0, -p)}$.

We now check the claimed continuity and differentiability properties of our function u_0 , mainly using the Implicit Function Theorem. Observe that, if $\frac{\partial}{\partial u} f(u, p, q) = 3u^2 + p = 0$, then $f(u, p, q) = -2u^3 + q$. Thus, the Implicit Function Theorem allows us to write u as a C^1 function of (p, q) unless

$$3u^2 + p = 0 \quad \text{and} \quad -2u^3 + q = 0.$$

In particular, the second equation has no solutions on $[0, \infty)$ for $q < 0$, so the Implicit Function Theorem applies in this case. If $q = 0$, the second equation implies that $u = 0$, and the first equation implies that $p = 0$. In other words, the only solution to this system for $u \geq 0$, $p \in \mathbb{R}$, and $q \leq 0$ is $(0, 0, 0)$, and hence u_0 is C^1 for $q > 0$, and continuous away from $(p, q) = (0, 0)$.

Thus, all that remains is checking continuity of u_0 at $(p, q) = (0, 0)$. For this, fix $\delta \in (0, \frac{1}{4})$ and assume that $(p, q) \in (-\delta, \delta) \times (-\delta, 0]$. Then, at $u = (2\delta)^{1/3}$,

$$\begin{aligned} u^3 + pu + q &= 2\delta + 2^{1/3}\delta^{1/3}p + q \\ &> 2\delta - 2^{1/3}\delta^{4/3} - \delta \\ &= \delta(1 - (2\delta)^{1/3}) \\ &> \delta(1 - 2^{-1/3}) \\ &> 0. \end{aligned}$$

In other words, $u_0 \in [0, (2\delta)^{1/3})$, proving continuity as $\delta \rightarrow 0$. \blacksquare

We now explicitly characterize the minimizer of g_x , limiting to the case $x_2 > 0$ to reflect our standing assumption that $\mu \in \mathcal{P}_c([0, \infty) \times (0, \infty))$ (although it is possible to use similar steps to derive a formula if $x_2 = 0$).

Proposition 3.19 (Explicit Minimizers of g_x). *Let $(\varphi, \lambda, \eta) \in C_b([0, \infty)^2) \times \mathbb{R}^2 \times (-\infty, 0)$ and define the notation*

$$\begin{aligned} A &= A(x; \eta) = -4x_2\eta > 0 \\ B_1 &= B_1(x; \lambda, \eta) = 2x_1 + 2x_2\lambda_1 - 4x_2\eta m_1 \\ B_2 &= B_2(x; \lambda, \eta) = x_1^2 - 2x_2\lambda_2 + 4x_2\eta m_2 \end{aligned}$$

Let $u_0 = u_0(x; \lambda, \eta)$ be the largest real root of $u^3 + (2 - B_2)u - B_1 = 0$. Then the function $g_x(y) = c(x, y) - \varphi(x) - \lambda \cdot y - \eta|y - m|^2$ attains a minimum at

$$Y(x; \lambda, \eta) := \begin{cases} \frac{u_0^2 - B_2}{A}(u_0, 1) & \text{if } B_1 > 0, \text{ and } B_2 < \frac{B_1^2}{4} \\ (0, -\frac{B_2}{A}) & \text{if } B_1 \leq 0 \text{ and } B_2 < 0 \\ (0, 0) & \text{if } B_2 \geq 0 \text{ and } B_1 \leq 2\sqrt{B_2} \end{cases}.$$

Proof. The first case corresponds to when g_x has a critical point on $(0, \infty)^2$. Such a point y must solve $\nabla g_x(y) = 0$, and hence the following system:

$$\begin{cases} -\frac{x_1}{x_2} + \frac{1}{x_2} \frac{y_1}{y_2} - \lambda_1 - 2\eta y_1 + 2\eta m_1 &= 0 \\ \frac{x_1^2}{2x_2} - \frac{1}{2x_2} \frac{y_1^2}{y_2^2} - \lambda_2 - 2\eta y_2 + 2\eta m_2 &= 0 \end{cases}.$$

Changing variables to $u = \frac{y_1}{y_2}$, rewriting coefficients in terms of A , B_1 , and B_2 , and rearranging, we rewrite this system as:

$$\begin{cases} y_2 &= \frac{B_1}{A} \frac{1}{u} - \frac{2}{A} \\ y_2 &= \frac{1}{A} u^2 - \frac{B_2}{A} \end{cases}. \quad (3.10)$$

By the assumption that $B_1 > 0$, the first equation in (3.10) defines y_2 as a strictly decreasing reciprocal function on $u > 0$, and the second defines y_2 as an increasing quadratic function. Thus, the condition $B_2 < \frac{B_1^2}{4}$ ensures that the unique solution (u, y_2) to (3.10) satisfies $y_2 > 0$. In particular, either $B_2 < 0$ and $\frac{1}{A}u^2$ is positive everywhere, or $B_2 \geq 0$ and the positive root, $u = \sqrt{B_2}$, of the quadratic equation is less than the root $u = B_1$ of the reciprocal equation, so that the graphs intersect in the region $y > 0$.

In this case, we solve this system by setting the two equations of (3.10) equal and multiplying through by Au to get the depressed cubic

$$u^3 + (2 - B_2)u - B_1 = 0. \quad (3.11)$$

Since $B_1 > 0$ by assumption, Lemma 3.18 establishes a unique positive solution to this cubic, so that u_0 is well-defined. Moreover, by recalling that $u = \frac{y_1}{y_2}$ and applying the second equation of (3.10), we deduce that $(y_1, y_2) = \frac{u_0^2 - B_2}{A}(u_0, 1) \in (0, \infty)^2$ is the unique solution to $\nabla g_x(y) = 0$ on $(0, \infty)^2$.

We next use Equation (3.10) to check that $\nabla g_x = 0$ has no solutions on $(0, \infty)^2$ unless $B_1 > 0$ and $B_2 < \frac{B_1^2}{4}$. If $B_1 > 0$ and $B_2 \geq \frac{B_1^2}{4}$, then notice that $\frac{B_1}{A} \frac{1}{u} - \frac{2}{A} > 0$ only for $u \in (0, \frac{B_1}{2})$, but $\frac{1}{A}u^2 - \frac{B_2}{A} > 0$ only for $u \in (\sqrt{B_2}, \infty)$. The condition $B_2 \geq \frac{B_1^2}{4}$ implies that there is no intersection between these intervals, and hence no solution to $\nabla g_x = 0$ on $(0, \infty)^2$. Similarly, if $B_1 \leq 0$, then Equation (3.11) implies that

$$u(u^2 - B_2 + 2) = B_1 \leq 0,$$

and hence either $u \leq 0$ or $Ay_2 = u^2 - B_2 \leq 0$, again implying that $\nabla g_x \neq 0$ on $(0, \infty)^2$.

Having ruled out solutions on $(0, \infty)^2$, we search for a minimum with $y_1 = 0$ and $y_2 > 0$. Under these conditions,

$$g_x(0, y_2) = \frac{x_1^2}{2x_2}y_2 - \varphi(x) - \lambda_2 y_2 - \eta(y_2 - m_2)^2,$$

where we have used that $c(x, (0, y_2)) = \frac{x_1^2}{2x_2}y_2$. Differentiating, we find that

$$\frac{d}{dy_2}g_x(0, y_2) = \frac{x_1^2}{2x_2} - \lambda_2 - 2\eta(y_2 - m_2),$$

which is equal to 0 if and only if $y_2 = -\frac{B_2}{A}$. Of course, this is positive if and only if $B_2 < 0$, and hence we have our minimum in this case.

Because of the definition of c in (1.5), only other possible minimum of g_x could occur is at the origin, which is the minimum in all other cases. ■

We immediately observe the following structural result.

Corollary 3.20 (Optimal γ_0 and φ_0). *If $X = Y = [0, \infty)^2$ and $(\gamma_0, \varphi_0, \lambda_0, \eta_0)$ satisfy the conditions of Proposition 3.13 and $\eta_0 < 0$, then*

$$\gamma_0 = (\text{id}, Y(\cdot; \lambda_0, \eta_0))_{\#}\mu$$

and

$$\varphi_0(x) = c(x, Y(x; \lambda_0, \eta_0)) - \lambda_0 \cdot Y(x; \lambda_0, \eta_0) - \eta_0 |Y(x; \lambda_0, \eta_0) - m|^2$$

on $\text{spt}(\mu)$.

Proof. This follows from construction of Y in Proposition 3.19 and the condition $c(x, y) - \lambda_0 \cdot y - \eta_0 |y - m|^2 \geq 0$ with equality γ_0 -a.e. from Proposition 3.13. ■

For the last result in this section, we derive a system of equations which, if solved, would be enough to prove existence of an optimal quartet:

Corollary 3.21 (System of Equations Guaranteeing Attainment). *Fix $m \in [0, \infty)^2$ and $\tau \geq 0$. Let $(\lambda_0, \eta_0) \in \mathbb{R}^2 \times (-\infty, 0)$ satisfy the following system of equations:*

$$\begin{cases} m &= \int Y(\cdot; \lambda_0, \eta_0) d\mu \\ \tau &= \int |Y(\cdot; \lambda_0, \eta_0) - m|^2 d\mu \end{cases} \quad (3.12)$$

Then $(\gamma_0, \varphi_0, \lambda_0, \eta_0)$ are admissible in Proposition 3.13, where γ_0 is as defined in Corollary 3.20, and where φ_0 is taken to be a continuous bounded extension of the function defined in Corollary 3.20, so long as such an extension exists.

Proof. The given system ensures that $\gamma_0 = (\text{id}, Y(\cdot; \lambda_0, \eta_0))_{\#}\mu$ from Corollary 3.20 satisfies the marginal constraints in Proposition 3.13. The construction of φ_0 and γ_0 from Corollary 3.20 ensures that the remaining constraint of Proposition 3.13 is satisfied. ■

Remark 3.22 (Application of Corollary 3.21). By looking at components, we see that (3.12) is a system of three non-linear equations in three unknowns, and hence there is hope to solve it using, for example, a fixed point theorem. However, since we learned in Section 3.4 that there are examples of (μ, m) for which the dual supremum is not attained, we anticipate that the existence of a solution to (3.12) will depend subtly on the relationship between μ and m .

3.6 NUMERICAL METHODS FOR THE DUAL PROBLEM

We dedicate this brief section to discuss some numerical insights on dual attainment, that we develop in the Trace Variance Bounds.ipynb Jupyter notebook at <https://github.com/Camdav/Thesis-Numerics>. In brief, these numerical methods use the conclusions of Proposition 3.19 and in particular the formula for φ_0 in Corollary 3.20 to express the dual objective function $\int \varphi d\mu + \lambda \cdot y + \eta|y - m|^2$ as a trivariate function of (λ, η) . Then, we use the SciPy implementation of the Nelder-Mead optimization algorithm [78] to solve this system. In particular, we describe one case of success and one case of failure in the following remarks:

Remark 3.23 (Successful Numerical Optimization). Let $m = (1, 2)$ and $\tau = 1.5$, and let $\mu = \frac{1}{4}(\delta_{(1,1)} + \delta_{(2,1)} + \delta_{(1,2)} + \delta_{(2,2)})$. Then our numerical experiments indicate that $\lambda_0 \approx (-0.4728, 0.5477)$ and $\eta_0 \approx -0.09$ optimize the dual problem, and we may compute φ_0 on $\text{spt}(\mu)$ according to Corollary 3.20. Piecing this all together, we find that

$$\mathcal{S}(\mu, m; \tau) \approx \int \varphi_0 d\mu + \lambda_0 \cdot m + \eta_0 \cdot \tau \approx 0.324028.$$

Likewise, we find that γ_0 as defined in Corollary 3.20 satisfies

$$\mathcal{I}(\mu, m; \tau) \approx \iint c d\gamma_0 \approx 0.324037.$$

This represents an approximately 0.003% discrepancy, which we attribute to early termination of the algorithm.

Remark 3.24 (Unsuccessful Numerical Optimization). Let $m = (1, 2)$ and $\tau = 1.5$, and let $\mu = \delta_{(2,4)}$. Then the Nelder-Mead optimization procedure [78] fails to terminate, and gives final values $\lambda_0 \approx (-0.375, 0.4688)$ and $\eta_0 \approx -1.74 \times 10^{-16} \approx 0$. Moreover, these parameters yield

$$\int \varphi_0 d\mu + \lambda_0 \cdot m + \eta_0 \cdot \tau \approx -0.9375 < 0.5625 = \iint c d\gamma_0,$$

meaning that the output primal and dual variables do not even come close to witnessing duality. The small value of η_0 should be noted – one possible interpretation is that the optimization procedure fails to account for the discontinuity of the objective function

$$F(\lambda, \eta) = \int \varphi_0(\cdot; \lambda, \eta) d\mu + \lambda \cdot y + \eta\tau$$

at $\eta = 0$ and attempts to reach a non-existent maximum on that set. One can imagine that this failure can be partially attributed to the fact that μ

only has one point in its support, as the system (3.12) might be easier to solve for a measure with multiple points in its support.

3.7 OTHER VARIANCE BOUND PROBLEMS

As we alluded to in Section 3.2, there are multiple possible ways to generalize the notion of variance bounds to multidimensional spaces, due to the fact that probability measures on $\mathcal{P}(\mathbb{R}^d)$ have covariance matrices, rather than scalar variances. For the purposes of simplifying the analysis in this chapter, we have limited our focus to the trace bound case from Definition 3.4. Nevertheless, the following methods for bounding variances may also prove useful, so we provide an abbreviated and informal discussion of their theory:

Definition 3.25 (Alternative Variance Bounds on $\mathcal{P}([0, \infty)^d)$). We start by defining the following alternative notions of variance bounds on $\mathcal{P}([0, \infty)^d)$:

- **Componentwise Bound:** Given $\sigma \in [0, \infty)^d$, we say $\nu \in \mathcal{P}([0, \infty)^d)$ satisfies a componentwise variance bound with parameters σ if

$$\int_{[0, \infty)^d} (y_i - \bar{v}_i)^2 d\nu(y) \leq \sigma_i \text{ for all } i \in \{1, \dots, d\}.$$

- **Matrix Bound:** Given a symmetric positive-semidefinite matrix $\Sigma \in \mathbb{R}^{d \times d}$, we say that ν satisfies a matrix variance bound with matrix Σ if

$$\text{cov}(\nu) \leq \Sigma,$$

in the sense that $\Sigma - \text{cov}(\nu)$ is a symmetric positive-semidefinite matrix.

These two bounds are related to each other, as well as the trace variance bound from Definition 3.4, in the following ways, which we summarize without proof:

Lemma 3.26 (Relations Between Bounds). *The following statements hold:*

- If ν satisfies a componentwise bound with parameters σ then it satisfies a trace bound with parameter $\sum_i \sigma_i$.
- Conversely, if ν satisfies a trace bound with parameter τ , it also satisfies a componentwise bound with parameters (τ, \dots, τ) .
- If ν satisfies a matrix bound with matrix Σ , then it also satisfies a trace bound with parameter $\sum_i \Sigma_{ii}$.

We now introduce the primal and dual problems for componentwise and matrix bounds in the $d = 2$ case. As before, fix two Polish subspaces $X, Y \subseteq \mathbb{R}^2$, $\mu \in \mathcal{P}(X)$ and $m \in \text{conv}(Y)$. We first introduce the componentwise bound problem, which uses the parameter $\sigma \in [0, \infty)^2$:

$$\mathcal{I}_{X,Y}(\mu, m; \sigma) := \inf_{\gamma \in \mathcal{P}(X \times Y)} \left\{ \iint c \, d\gamma \mid \pi_{\#}^X \gamma = \mu, \overline{\pi_{\#}^Y \gamma} = m, \right. \\ \left. \iint (y_i - m_i)^2 d\gamma(x, y) \leq \sigma_i \text{ for } i = 1, 2 \right\}.$$

It can be shown that this primal problem has weak dual given by:

$$\mathcal{S}_{X,Y}(\mu, m; \sigma) = \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2 \\ \kappa \in (-\infty, 0]^2}} \left\{ \int \varphi \, d\mu + \lambda \cdot m + \kappa \cdot \sigma \mid \right. \\ \left. c(x, y) - \varphi(x) - \lambda \cdot y - \kappa_1 (y_1 - m_1)^2 - \kappa_2 (y_2 - m_2)^2 \geq 0 \text{ on } X \times Y \right\}. \quad (3.13)$$

Likewise, we introduce the matrix bound problem, which takes a positive semidefinite matrix Σ as a parameter but adds infinitely many constraints:

$$\mathcal{I}_{X,Y}(\mu, m; \Sigma) := \inf_{\gamma \in \mathcal{P}(X \times Y)} \left\{ \iint c \, d\gamma \mid \pi_{\#}^X \gamma = \mu, \overline{\pi_{\#}^Y \gamma} = m, \right. \\ \left. z^t (\Sigma - \text{cov}(\pi_{\#}^Y \gamma)) z \geq 0 \, \forall z \in \mathbb{R}^2 \right\}.$$

This yields the dual problem

$$\mathcal{S}_{X,Y}(\mu, m; \Sigma) := \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2 \\ 0 \geq \psi \in C_b(\mathbb{R}^2)}} \left\{ \int \varphi \, d\mu + \lambda \cdot m + \int_{\mathbb{R}^2} \psi(z) z^T M z \, dz \mid \right. \\ \left. c(x, y) - \varphi(x) - \lambda \cdot y - \int_{\mathbb{R}^2} \psi(z) (y \cdot z)^2 \, dz \geq 0 \text{ on } X \times Y \right\}. \quad (3.14)$$

We conclude this section, and this chapter, by remarking on why we expect these problems to be viable alternatives to trace variance bounds:

Remark 3.27. As with the trace bound dual problem (1.11), componentwise and matrix bounds cure the main technical issue with the original unbounded problem (1.7) that we explored in Proposition 2.1 and Lemma 2.10.

In particular, Problem (3.13) introduces a constraint of the form

$$c(x, y) - \lambda \cdot y - \kappa_1(y_1 - m_1)^2 - \kappa_2(y_2 - m_2)^2 \geq \varphi(x),$$

where $\kappa_1 \leq 0$ and $\kappa_2 \leq 0$. If κ_1 and κ_2 are both nonzero, the leading order terms grow quadratically in y_1 and y_2 , respectively. This implies that φ may satisfy the preceding constraint even if $c(x, y) - \lambda \cdot y$ is negative for some y , as the terms of the form $-\kappa_i(y_i - m_i)^2$ dominate as $|y| \rightarrow \infty$.

Similarly, for Problem (3.14), the relevant constraint is

$$c(x, y) - \lambda \cdot y - \int_{\mathbb{R}^2} \psi(z)(y \cdot z)^2 dz \geq \varphi(x),$$

where $0 \geq \psi \in C_b(\mathbb{R}^2)$. If ψ is continuous and not identically zero, then $\int_{\mathbb{R}^2} \psi(z)(y \cdot z)^2 dz < 0$ for any $y \in [0, \infty)^2$. This is because continuity of ψ implies that $\text{spt}(\psi)$ contains some open set where $(y \cdot z) > 0$. In other words, there exists an open set of \mathbb{R}^2 where $\psi(z)(y \cdot z)^2 < 0$. Thus, since the integrand is always non-positive, we deduce the desired inequality. Since, moreover, the term $\int_{\mathbb{R}^2} \psi(z)(y \cdot z)^2 dz$ is 2-homogeneous in y , it dominates as $|y| \rightarrow \infty$.

RESOLVING THE DUALITY GAP THROUGH COMPACTNESS: EFFECTIVE HYBRID POLICY

4.1 SUMMARY OF RESULTS

In this section, we study how restricting the domain Y of possible post-regulation outcomes to be compact affects the primal problem $\mathcal{I}_{X,Y}(\mu, m)$ and the dual problem $\mathcal{S}_{X,Y}(\mu, m)$. The cornerstone of this chapter is Theorem 1.15, which asserts that this restriction is also enough to guarantee duality. In other words, the policies $\mathcal{S}_{X,Y}(\mu, m)$, which allow market-based mechanisms to operate under the *a priori* restriction of post-regulation outcomes to a compact Y , can achieve the same regulatory success as pure command-and-control policy. To this end, Section 4.2 is dedicated to the proof of Theorem 1.15, using methods which are broadly similar to those we employed in the proof of Theorem 1.14 in Section 3.3.

Section 4.3, and in particular Proposition 4.3, introduces a number of relatively weak standing assumptions which ensure the attainment of the primal infimum $\mathcal{I}_{X,Y}(\mu, m)$. Under these assumptions, Lemma 4.4 identifies the optimal dual potential $\varphi^\lambda \in C_b(X)$ for a fixed Lagrange multiplier λ . This allows us to reduce the dual problem $\mathcal{S}_{X,Y}(\mu, m)$ to a bivariate unconstrained optimization problem in λ in Corollary 4.6, and ultimately prove attainment of the dual supremum $\mathcal{S}_{X,Y}(\mu, m)$ in Corollary 4.10.

Next, Section 4.4 specializes further to the case where $Y = [0, R]^2$, providing an explicit computation of φ^λ in Proposition 4.13. This step is important for computing $\mathcal{S}_{X,Y}(\mu, m)$ using the numerical methods which we will briefly explore in Section 4.5.

4.2 STRONG DUALITY ON COMPACT DOMAINS

We consider the version of problems (1.6) and (1.7) where $X \subseteq [0, \infty) \times (0, \infty)$ is a Polish space and $Y \subseteq [0, \infty)^2$ is compact. To prove duality in

this case, we use an approach similar to our proof of 1.14, in that we apply Frieesecke's version of Fenchel-Rockafellar duality (stated in Theorem 3.7) to recover strong duality:

Proof of Theorem 1.15. We first address the case where m does not lie in the convex hull of Y , i.e. $m \notin \text{conv}(Y)$. In this case, no probability measure $\gamma \in \mathcal{P}(X \times Y)$ is admissible for $\mathcal{I}_{X,Y}(\mu, m)$, and hence $\mathcal{I}_{X,Y}(\mu, m) = +\infty$ by convention. Likewise, since $\text{conv}(Y)$ is compact, as the convex hull of a compact set, [20, Theorem 1.7] guarantees the existence of a hyperplane strictly separating $\text{conv}(Y)$ and $\{m\}$, in the sense that there exists $\alpha \in \mathbb{R}$ and $\lambda^1 \in \mathbb{R}^2$ with $\lambda^1 \cdot m > -\alpha$ but $\lambda^1 \cdot y < -\alpha$ for all $y \in Y$. We observe that, for any $t > 0$, $(\varphi^t, \lambda^t) := (t\alpha, t\lambda^1)$ satisfies the constraints of (1.7), as

$$c(x, y) \geq 0 = t\alpha - t\alpha > \varphi^t(x) + \lambda^t \cdot y \text{ on } X \times Y.$$

However, we compute

$$\int \varphi^t d\mu + \lambda^t \cdot m = t(\alpha + \lambda^1 \cdot m),$$

which tends to $+\infty$ as $t \rightarrow +\infty$ by the assumption that $\lambda^1 \cdot m > -\alpha$. Hence, in this case, $\mathcal{S}_{X,Y}(\mu, m) = +\infty = \mathcal{I}_{X,Y}(\mu, m)$.

Thus, we focus on the case where $m \in \text{conv}(Y)$. We begin by defining functions to which we can apply Theorem 3.7. In particular, define $\mathcal{F}, \mathcal{G} : C_b(X \times Y) \rightarrow \mathbb{R} \cup \{+\infty\}$ by

$$\mathcal{F}(u) = \begin{cases} -(\int_X \varphi d\mu + \lambda \cdot m) & \text{if } u(x, y) = -(\varphi(x) + \lambda \cdot y) \text{ on } X \times Y \\ & \text{where } \varphi \in C_b(X) \text{ and } \lambda \in \mathbb{R}^2 \\ +\infty & \text{else,} \end{cases} \quad (4.1)$$

and

$$\mathcal{G}(u) = \begin{cases} 0 & \text{if } u \geq -c \text{ on } X \times Y \\ +\infty & \text{else,} \end{cases}.$$

We immediately observe that \mathcal{G} is well defined, but it is worth quickly noting that \mathcal{F} is as well. Let $\nu \in \mathcal{P}(Y)$ have mean $\bar{\nu} = m$. Then, if $\varphi(x) + \lambda \cdot y = \tilde{\varphi}(x) + \tilde{\lambda} \cdot y$ on $X \times Y$,

$$\begin{aligned} \int \varphi d\mu + \lambda \cdot m &= \iint_{X \times Y} \varphi(x) + \lambda \cdot y d(\mu \otimes \nu)(x, y) \\ &= \iint_{X \times Y} \tilde{\varphi}(x) + \tilde{\lambda} \cdot y d(\mu \otimes \nu)(x, y) \\ &= \int \tilde{\varphi} d\mu + \tilde{\lambda} \cdot m, \end{aligned}$$

as required.

Next, we verify the hypotheses of Theorem 3.7. To start, we observe that, like their analogues defined in the proof of 1.14, \mathcal{F} and \mathcal{G} are both linear, and hence convex, and that $\mathcal{F}(0) = \mathcal{G}(0) = 0$, so that both are proper.

Now, we check that $\mathcal{H} := \mathcal{F} \square \mathcal{G}_-$ does not take on the value $-\infty$. As in the proof of Theorem 1.14, we deduce that, for $u \in C_b(X \times Y)$,

$$\begin{aligned} \mathcal{H}(u) &= \inf_{v \in C_b(X \times Y)} [\mathcal{F}(u - v) + \mathcal{G}(-v)] \\ &= - \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2}} \left\{ \int \varphi \, d\mu + \lambda \cdot m \mid \varphi(x) + \lambda \cdot y \leq c(x, y) - u(x, y) \right. \\ &\qquad \qquad \qquad \left. \text{on } X \times Y \right\}, \end{aligned} \tag{4.2}$$

where we have used the definitions of \mathcal{F} and \mathcal{G} to parametrize admissible functions $v \in C_b(X \times Y)$ by $(\varphi, \lambda) \in C_b(X) \times \mathbb{R}^2$, and where we have factored out a negative sign to change the infimum to a supremum. By fixing $\gamma_0 = \mu \times \nu$, where $\bar{\nu} = m$, we deduce that

$$\begin{aligned} \mathcal{H}(u) &= - \sup_{\substack{\varphi \in C_b(X) \\ \lambda \in \mathbb{R}^2}} \left\{ \iint \varphi(x) + \lambda \cdot y \, d\gamma(x, y) \mid \varphi(x) + \lambda \cdot y \right. \\ &\qquad \qquad \qquad \left. \leq c(x, y) - u(x, y) \text{ on } X \times Y \right\} \\ &\geq - \iint (c - u) \, d\gamma_0 \\ &> -\infty, \end{aligned}$$

where the last estimate arises from the fact that c and u are continuous, combined with the assumption that $\text{spt}(\mu)$ and Y are compact.

Finally, we check the lower semicontinuity hypothesis of Theorem 3.7. As before, it suffices to check that if $u, v \in C_b(X \times Y)$ satisfy $\|u - v\|_\infty \leq \varepsilon$, then $\mathcal{H}(v) \geq \mathcal{H}(u) - \varepsilon$. Thus, take $(\varphi, \lambda) \in C_b(X) \times \mathbb{R}^2$ admissible for the reformulation of $\mathcal{H}(v)$ in (4.2), and observe that

$$c(x, y) - \varphi(x) - \lambda \cdot y \geq v(x, y) \geq u(x, y) - \varepsilon$$

on $X \times Y$. In particular, this implies that $(\varphi - \varepsilon, \lambda)$ is admissible for $\mathcal{H}(u)$ and hence

$$- \int \varphi \, d\mu - \lambda \cdot m \geq \mathcal{H}(u) - \varepsilon.$$

Minimizing over the set of (φ, λ) which are admissible for $\mathcal{H}(v)$, we recover that $\mathcal{H}(v) \geq \mathcal{H}(u) - \varepsilon$, as desired.

Having checked its hypotheses, we now unpack the conclusion from Theorem 1.14. We immediately observe that \mathcal{F} and \mathcal{G} are defined so that $\inf_{u \in C_b(X \times Y)} (\mathcal{F}(u) + \mathcal{G}(u)) = -\mathcal{S}_{X,Y}(\mu, m)$.

To recover the primal problem, we may assume without loss of generality that X is compact, owing to the assumption that $\mu \in \mathcal{P}_c(X)$. This, combined with the fact that Y is compact, implies that $X \times Y$ is compact, and hence $C_b(X \times Y) = C_0(X \times Y)$. Hence, by the Riesz Representation Theorem, $(C_b(X \times Y), \|\cdot\|_\infty)$ has topological dual space $(\mathcal{M}(X \times Y), \|\cdot\|_{TV})$. This allows us to compute, for $\gamma \in \mathcal{M}(X \times Y)$,

$$\mathcal{F}^*(\gamma) = \begin{cases} 0 & \text{if } \iint \varphi(x) d\gamma(x, y) = \int \varphi d\mu \forall \varphi \in C_b(X) \\ & \text{and } \iint y d\gamma(x, y) = m \end{cases}$$

and

$$\mathcal{G}^*(-\gamma) = \sup_{c \geq u \in C_b(X \times Y)} \iint u d\gamma,$$

which is equal to $+\infty$ unless $\gamma \in \mathcal{M}_+(X \times Y)$ and $\iint c d\gamma$ otherwise (as we may approximate the lower semicontinuous function c from below with functions in C_b). Thus,

$$\begin{aligned} \mathcal{S}_{X,Y}(\mu, m) &= - \sup_{\gamma \in \mathcal{M}(X \times Y)} (-\mathcal{F}^*(\gamma) - \mathcal{G}^*(-\lambda)) \\ &= \inf_{\gamma \in \mathcal{M}(X \times Y)} (\mathcal{F}^*(\gamma) + \mathcal{G}^*(-\lambda)) \\ &= \inf_{\gamma \in \mathcal{M}_+(X \times Y)} \left\{ \iint c d\gamma \mid \pi_{\#}^X \gamma = \mu \text{ and } \iint y d\gamma(x, y) = m \right\} \\ &= \mathcal{I}_{X,Y}(\mu, m), \end{aligned}$$

where in the last line we have noted that the constraint $\pi_{\#}^X \gamma = \mu$ ensures that γ is a probability measure. ■

We conclude this section by making a few notes about the implications of our results.

Remark 4.1 (Differences between the proofs of Theorem 1.14 and Theorem 1.15). The proofs of Theorem 1.14 and Theorem 1.15 are broadly similar, and differ only in a few key respects. First, the unbounded domains in Theorem 1.14 make it relatively easy to conclude that \mathcal{F} is well-defined by looking at polynomial coefficients, whereas in Theorem 1.15 we need to explicitly introduce a probability measure on Y with mean m . On the other hand, the compact domains in 1.15 completely remove the technical issues

with dual spaces that we encountered in Theorem 1.14, and allow us to identify the topological dual space of $C_b(X \times Y)$ as a space of measures.

Remark 4.2 (The $x_2 = 0$ Case). Throughout this work, we typically avoid the $x_2 = 0$ case, as it implies that it is impossible for a firm to change the emissions intensity of its production processes; in other words, emissions must scale proportionally to production. However, it is possible to identify some niche cases corresponding to $x_2 = 0$. For example, in the context of the metalworking industry, some small-scale producers continue to use traditional blacksmithing methods, despite these methods being superseded by modern methods for most commercial applications. Policymakers may deem the total loss of millennia-old blacksmithing traditions to be unacceptable, and hence may wish to build a framework which allows it to continue at a small scale. Mathematically, we anticipate that the duality result in Theorem 1.15 extends to the $x_2 = 0$ case, but have omitted this to simplify exposition.

4.3 PRIMAL AND DUAL ATTAINMENT

In this section, we show that both the primal infimum $\mathcal{I}_{X,Y}(\mu, m)$ and the dual supremum $\mathcal{S}_{X,Y}(\mu, m)$ are attained. We start with the former, as this is the more standard result.

Proposition 4.3 (Primal Attainment for Compact Domains). *Assume that (i) $X \subseteq [0, \infty) \times (0, \infty)$ and $Y \subseteq [0, \infty)^2$ are Polish, (ii) Y is compact and contains some point (y_1, y_2) with $y_2 > 0$, (iii) $\mu \in \mathcal{P}_c(X)$, and (iv) $m \in \text{int}(\text{conv}(Y))$. Then $\mathcal{I}_{X,Y}(\mu, m)$ is attained, in the sense that there exists $\gamma_* \in \mathcal{P}(X \times Y)$ satisfying $\int \int_{X \times Y} c \, d\gamma_* = \mathcal{I}_{X,Y}(\mu, m)$, as well as the constraints $\pi_{\#}^X \gamma_* = \mu$ and $\int \int_{X \times Y} y \, d\gamma_*(x, y) = m$.*

Proof. We need only to show that, with respect to the weak-* topology on $\mathcal{M}(X \times Y)$ (in duality with $C_b(X \times Y)$), the map $\gamma \mapsto \int \int c \, d\gamma$ is lower semicontinuous and the space of admissible measures is compact, as this allows us to apply Weierstrass' Theorem from [96, Box 1.1]. Lower semicontinuity is an immediate consequence of [96, Lemma 1.6]. Likewise, compactness can be checked by applying the Banach-Alaoglu Theorem (see [96, Box 1.2]) and verifying that the conditions $\pi_{\#}^X \gamma = \mu$ and $\int \int_{X \times Y} y \, d\gamma(x, y) = m$ are closed. ■

Dual attainment is somewhat more technical. To begin, we define a function h_x^λ on Y by

$$h_x^\lambda(y) := c(x, y) - \lambda \cdot y. \quad (4.3)$$

We immediately note that the condition $c(x, y) - \lambda \cdot y \geq \varphi(x)$ suggests that, for a given λ , we may maximize the objective functional $\int \varphi \, d\mu + \lambda \cdot m$ by choosing

$$\varphi^\lambda(x) := \inf_Y h_x^\lambda = \inf_{y \in Y} [c(x, y) - \lambda \cdot y]. \quad (4.4)$$

Formally speaking, this choice reduces $\mathcal{S}_{X, Y}(\mu, m)$ to a bivariate optimization problem in λ , but we will of course need to verify that φ^λ is actually continuous:

Lemma 4.4. *Let $Y \subseteq [0, \infty)^2$ be compact, assume that there exists $(y_1, y_2) \in Y$ with $y_2 > 0$, and let $\lambda \in \mathbb{R}^2$. Then $\operatorname{argmin}_Y h_x^\lambda$ is nonempty and compact and φ^λ is continuous on $[0, \infty) \times (0, \infty)$.*

Proof. We proceed by applying [93, Theorem 1.17] on parametric minimizers to a suitable function. To this end, define a proper, lower semicontinuous function $g : \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R} \cup \{+\infty\}$ by

$$g(x, y) = \begin{cases} h_x^\lambda(y) & \text{for } (x, y) \in ([0, \infty) \times (0, \infty)) \times Y \\ +\infty & \text{else.} \end{cases}$$

Note that g is proper as a result of the definition of c in (1.5) and the assumption that Y contains some point with $y_2 > 0$.

We next check that $g(x, y)$ is level bounded in y locally uniformly in x , in the sense of [93, Definition 1.16]. To this end, observe that, for any $x \in \mathbb{R}^2$ and $\alpha \in \mathbb{R}$, the construction of g ensures that

$$\{y \in \mathbb{R}^2 \mid g(x, y) \leq \alpha\} \subseteq Y,$$

which is bounded, as desired.

Thus, [93, Theorem 1.17] implies that $p(x) := \inf_{y \in \mathbb{R}^2} g(x, y)$ is a proper and lower semicontinuous function on \mathbb{R}^2 and moreover, for any x with $p(x) < +\infty$, $\operatorname{argmin}_{y \in \mathbb{R}^2} g(x, y)$ is nonempty and compact. We observe that, by the definition of g ,

$$\inf_{y \in \mathbb{R}^2} g(x, y) = \begin{cases} \inf_Y h_x^\lambda & \text{if } x \in [0, \infty) \times (0, \infty) \\ +\infty & \text{else.} \end{cases}.$$

Moreover, by the definition of c in (1.5), and by the assumption that Y contains some point with $y_2 > 0$, we have that $\inf_Y h_x^\lambda = \inf_{y \in Y} [c(x, y) - \lambda \cdot y] < +\infty$ for any $x \in [0, \infty) \times (0, \infty)$. This allows us to recover that, for any $x \in [0, \infty) \times (0, \infty)$, $\operatorname{argmin}_Y h_x^\lambda$ is nonempty and compact.

Finally, for continuity on $[0, \infty) \times (0, \infty)$, we apply [93, Theorem 1.17(c)], which, specialized to our case, states that φ^y is continuous at \bar{x} if there

exists $\bar{y} \in \operatorname{argmin}_{y \in Y} [c(\bar{x}, y) - \lambda \cdot y]$ such that $c(x, \bar{y}) - \lambda \cdot \bar{y}$ is continuous as a function of x . To this end, fix $\bar{x} \in [0, \infty) \times (0, \infty)$ and let $\bar{y} \in \operatorname{argmin}_{y \in Y} [c(\bar{x}, y) - \lambda \cdot y]$. If $\bar{y}_2 \neq 0$, then (1.5) implies that

$$c(x, \bar{y}) - \lambda \cdot \bar{y} = \frac{\bar{y}_2}{2x_2} (x_1 - \frac{\bar{y}_1}{\bar{y}_2}),$$

which is clearly continuous in x , including at x_0 . On the other hand, if $\bar{y}_2 = 0$, then (1.5), combined with the assumption that $\bar{y} \in \operatorname{argmin}_{y \in Y} [c(\bar{x}, y) - \lambda \cdot y]$, allows us to conclude that $\bar{y}_1 = 0$, and hence

$$h_x^\lambda(y) = c(x, \bar{y}) - \lambda \cdot \bar{y} \equiv 0$$

on $[0, \infty) \times (0, \infty)$, again proving the desired continuity, and concluding this proof. ■

This continuity allows us to write $\mathcal{S}_{X,Y}(\mu, m)$ as a bivariate optimization problem; we start by defining its objective function:

Definition 4.5. Given $\lambda \in \mathbb{R}^2$, and (X, Y, μ, m) satisfying the assumptions of Proposition 4.3, define an objective functional $J_{X,Y,\mu,m} : \mathbb{R}^2 \rightarrow \mathbb{R}$ by

$$J_{X,Y,\mu,m}(\lambda) = \lambda \cdot m + \int_X \varphi^\lambda d\mu. \quad (4.5)$$

Note that, in most cases, X, Y, μ and m are clear from context, and in such cases we suppress these subscripts.

The following result states that this bilevel optimization problem yields the same result as the dual supremum from (1.7):

Corollary 4.6. *If the assumptions of Proposition 4.3 hold, then*

$$\mathcal{S}_{X,Y}(\mu, m) = \sup_{\lambda \in \mathbb{R}^2} J(\lambda).$$

Proof. Given φ, λ admissible for $\mathcal{S}_{X,Y}(\mu, m)$, we observe that, by the constraint in the definition of $\mathcal{S}_{X,Y}(\mu, m)$,

$$\varphi(x) \leq \min_Y h_x^\lambda = \varphi^\lambda(x),$$

which implies that $\mathcal{S}_{X,Y}(\mu, m) \leq \sup_{\lambda \in \mathbb{R}^2} J(\lambda)$. Continuity of φ^λ implies that the pair $(\lambda, \varphi^\lambda)$ is admissible for $\mathcal{S}_{X,Y}(\mu, m)$ and hence the desired equality holds. ■

The preceding corollary effectively reduces questions of attainment for $\mathcal{S}_{X,Y}(\mu, m)$ to those of attainment for J . To answer these, we use the following technical lemma:

Lemma 4.7. *Let $Y \subseteq \mathbb{R}^n$ be compact, and let $m \in \mathbb{R}^n$. Then the following are equivalent:*

1. *There exists a finite set $\{y^1, \dots, y^k\} \subseteq Y$ with the property that, for any $\lambda \in \mathbb{R}^n \setminus \{0\}$, there exists some $i \in \{1, \dots, k\}$ such that $\lambda \cdot (y^i - m) < 0$.*
2. *For any $\lambda \in \mathbb{R}^n \setminus \{0\}$, there exists $y \in Y$ with $\lambda \cdot (y - m) < 0$.*
3. *$m \in \text{int}(\text{conv}(Y))$, the open convex hull of Y .*

Proof. We first prove equivalence of statements (1) and (2). The implication (1) \implies (2) is trivial, since we may choose $y \in \{y^1, \dots, y^k\} \subseteq Y$.

On the other hand, for the implication (2) \implies (1), let S^1 denote the unit circle in \mathbb{R}^2 . Given $y \in Y$, define $S_y := \{\lambda \in S^1 : \lambda \cdot (y - m) < 0\}$, and note that the assumption (2) implies that $\{S_y : y \in Y\}$ is an open cover of S^1 . By compactness of S^1 , this reduces to a finite subcover, so that there exists $\{y^1, \dots, y^k\} \subseteq Y$ with $S^1 \subseteq \bigcup_{i=1}^k S_{y^i}$. By decomposing $\lambda \cdot (y^i - m) = |\lambda| \cdot \frac{\lambda}{|\lambda|} \cdot (y^i - m)$, we deduce that (1) also holds.

Next, we prove the implication (2) \implies (3) using the contrapositive. To this end, assume that $m \notin \text{int}(\text{conv}(Y))$. If $\text{int}(\text{conv}(Y)) = \emptyset$, then Y must be contained in some affine hyperplane and hence, in particular, there exists some $\lambda_0 \in \mathbb{R}^n$ and some $\kappa \in \mathbb{R}$ such that $\lambda_0 \cdot y = \kappa$ for all $y \in Y$. Thus, $\lambda_0 \cdot (y - m) = \kappa - \lambda_0 \cdot m$ for any $y \in Y$, and we may assume that this quantity is non-negative by considering $-\lambda_0$ if necessary. Otherwise, we assume that $\text{int}(\text{conv}(Y))$ is non-empty. Combining this with the convexity of $\text{int}(\text{conv}(Y))$, we may apply the special case of the Hahn-Banach Separation Theorem described in [20, Lemma 1.3] to choose $\lambda_0 \in \mathbb{R}^n$ satisfying $\lambda_0 \cdot y \geq \lambda_0 \cdot m$ for all $y \in Y$, thereby proving that (2) does not hold.

Finally, for the implication (3) \implies (2), we again prove the contrapositive. Thus, assume that there exists $\lambda \in \mathbb{R}^n \setminus \{0\}$ such that $\lambda \cdot (y - m) \geq 0$ for all $y \in Y$. Then the closed half-space $H := \{y \in \mathbb{R}^2 : \lambda \cdot (y - m) \geq 0\}$ is such that $Y \subseteq \text{conv}(Y) \subseteq H$ and $m \in \partial H$. However, this implies that $\text{int}(\text{conv}(Y)) \subseteq \text{int}(H)$ and, since $m \notin \text{int}(H)$, we deduce that $m \notin \text{int}(\text{conv}(Y))$, as desired. \blacksquare

Lemma 4.8. *In general, $J : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a concave function. Moreover, if any of the equivalent conditions of Lemma 4.7 hold, then J is coercive, in the sense that $\lim_{|\lambda| \rightarrow \infty} J(\lambda) = -\infty$.*

Proof. For concavity, we observe that the term $\lambda \cdot m$ is linear. Thus, it suffices to show concavity of the second term, $\int \varphi^\lambda d\mu$. To do so, we fix

$\lambda^0, \lambda^1 \in \mathbb{R}^2$, and, for $t \in [0, 1]$, define $\lambda^t = (1 - t)\lambda^0 + t\lambda^1$. This allows us to compute

$$\begin{aligned} \int_X \min_{y \in Y} [c(x, y) - \lambda^t \cdot y] d\mu(x) &\geq \int_X \left[(1 - t) \min_{y \in Y} [c(x, y) - \lambda^0 \cdot y] \right. \\ &\quad \left. + t \min_{y \in Y} [c(x, y) - \lambda^1 \cdot y] \right] d\mu(x) \\ &= (1 - t) \int_X \min_{y \in Y} [c(x, y) - \lambda^0 \cdot y] d\mu(x) \\ &\quad + t \int_X \min_{y \in Y} [c(x, y) - \lambda^1 \cdot y] d\mu(x), \end{aligned}$$

proving concavity.

For coercivity, assume that condition (1) of Lemma 4.7 holds, and let $\{y^1, \dots, y^k\} \subseteq Y$ be a finite set with the property that, for any $\lambda \in \mathbb{R}^2 \setminus \{0\}$, there exists $i \in \{1, \dots, k\}$ with $\lambda \cdot (y^i - m) < 0$. Applying the definition of φ^λ , we deduce that, for any $x \in X$,

$$\varphi^\lambda(x) \leq \min_{i=1, \dots, k} \{c(x, y^i) - \lambda \cdot y^i\}.$$

Substituting this into Definition 4.5 of J and simplifying somewhat, we deduce that, for $\lambda \neq (0, 0)$,

$$\begin{aligned} J(\lambda) &\leq \int_X \min_{i=1, \dots, k} \{c(x, y^i) - \lambda \cdot (y^i - m)\} d\mu(x) \\ &\leq \max_{i=1, \dots, k} \left\{ \int c(x, y^i) \right\} d\mu(x) + |\lambda| \min_{i=1, \dots, k} \frac{\lambda}{|\lambda|} \cdot (y^i - m) \end{aligned}$$

Now, observe that $u \mapsto \min_{i=1, \dots, k} u(y^i - m)$ is continuous on the compact set S^1 , and hence attains a maximum at some $u_0 \in S^1$. Moreover, Condition (1) of Lemma 4.7 implies that $\kappa := \min_{i=1, \dots, k} u_0(y^i - m) < 0$. This, combined with our previous work, allows us to deduce that

$$J(\lambda) \leq \max_{i=1, \dots, k} \left\{ \int c(x, y^i) d\mu(x) \right\} + \kappa |\lambda|.$$

Hence, since the first term in the preceding expression does not depend on λ , J is coercive. \blacksquare

We now conclude with our desired results.

Corollary 4.9. *Under the hypotheses of Lemma 4.8, J attains a maximum at some point $\lambda_0 \in \mathbb{R}^2$.*

Proof. This is a standard result, which follows from using coercivity to show that J must be small outside of a sufficiently large ball, and then applying the Extreme Value Theorem to J on this ball. ■

Corollary 4.10. *Under the assumptions of Proposition 4.3, $\mathcal{S}_{X,Y}(\mu, m)$ is attained by $(\varphi^{\lambda_0}, \lambda_0) \in C_b(X) \times \mathbb{R}^2$.*

Proof. This follows immediately from construction, and the continuity of φ^λ from Lemma 4.4 which, combined, imply that $(\varphi^{\lambda_0}, \lambda_0)$ are admissible for $\mathcal{S}_{X,Y}(\mu, m)$. Applying Corollary 4.6, we deduce that

$$\int \varphi^{\lambda_0} d\mu + \lambda_0 \cdot m = J(\lambda_0) = \sup_{\lambda \in \mathbb{R}^2} J(\lambda) = \mathcal{S}_{X,Y}(\mu, m),$$

as desired. ■

We close this section with some structural results on optimizers. The first one, which follows, is analogous to the ELKKT conditions we derived in Proposition 3.12 and Proposition 3.13.

Corollary 4.11. *Under the assumptions of Proposition 4.3, let γ_0 and (φ_0, λ_0) be optimal for $\mathcal{I}_{X,Y}(\mu, m)$ and $\mathcal{S}_{X,Y}(\mu, m)$, respectively. Then $c(x, y) - \varphi_0(x) - \lambda_0 \cdot y = 0$ for γ_0 -a.e. (x, y) . As such, we may assume without loss of generality that, if $(x_0, y_0) \in \text{spt}(\gamma_0)$, then $y_0 \in \text{argmin}_Y h_{x_0}^{\lambda_0}$.*

Proof. By construction,

$$\iint c d\gamma_0 = \int_X \varphi_0 d\mu + \lambda_0 \cdot m = \iint_{X \times Y} \varphi_0(x) + \lambda_0 \cdot y d\gamma_0 \leq \iint c d\gamma_0,$$

where the first equality holds by Theorem 1.15, and hence equality holds throughout. Since, moreover, $\varphi_0(x) + \lambda_0 \cdot y \leq c(x, y)$ on $X \times Y$, we conclude that $c(x, y) = \varphi_0(x) + \lambda_0 \cdot y$ for γ_0 -a.e. (x, y) , as desired. The final conclusion follows immediately from noticing that $c(x, y) - \varphi(x) - \lambda \cdot y \geq 0$, with equality γ_0 -a.e.; for fixed x_0 , we need $y_0 \in \text{argmin}_Y h_{x_0}^{\lambda_0}$ to ensure the equality $c(x, y) - \varphi(x) - \lambda \cdot y = 0$. ■

We close by noticing that, as a consequence of the 1-homogeneity of h_x^λ in y , this function is minimized on ∂Y .

Lemma 4.12. *Under the assumptions of Proposition 4.3,*

$$\varphi^\lambda(x) = \min_Y h_x^\lambda = \min_{\partial Y} h_x^\lambda,$$

or, in other words, $\partial Y \cap \text{argmin}_Y h_x^\lambda \neq \emptyset$.

Proof. The function φ_λ attains a minimum at some point $y_0 \in Y$ by Lemma 4.4. If $y_0 \in \partial Y$, the conclusion is immediate. Otherwise, if $y_0 \in \text{int}(Y)$, then we consider the function

$$f(t) := c(x_0, ty_0) - \lambda_0 \cdot ty_0 = t(c(x_0, y_0) - \lambda_0 \cdot y_0).$$

If $c(x_0, y_0) - \lambda_0 \cdot y_0 \neq 0$, then $f(t)$ is either increasing or decreasing at $t = 1$, contradicting the assumption that y_0 is a minimum. On the other hand, if $c(x_0, y_0) = \lambda_0 \cdot y_0$, then $f(t) \equiv 0$, which implies that $c(x_0, y) - \lambda \cdot y$ is identically zero on the ray $y = ty_0$. As this ray necessarily intersects ∂Y , we recover the desired conclusion. ■

4.4 SPECIALIZATION TO SQUARES

We now turn our attention to the case where $Y = [0, R]^2$ is a compact square. While this particular choice of domain is arbitrary, such a choice nevertheless must be made in order to have any hope of using numerical methods to study $\mathcal{I}_{X,Y}(\mu, m)$ and $\mathcal{S}_{X,Y}(\mu, m)$.

We begin our study by explicitly computing the minimal values and arguments of the function $h_x^\lambda(y) = c(x, y) - \lambda \cdot y$ for $y \in [0, R]^2$ in the following proposition, the results of which are depicted in Figures 4.1 and 4.2:

Proposition 4.13. *Let $x \in [0, \infty) \times (0, \infty)$, let $Y = [0, R]^2$, and let $\lambda \in \mathbb{R}^2$. Define the notation*

$$\begin{aligned}\Lambda_1 &= \Lambda_1(x; \lambda) = x_1 + x_2 \lambda_1 \\ \Lambda_2 &= \Lambda_2(x; \lambda) = x_1^2 - 2x_2 \lambda_2.\end{aligned}$$

Then

$$\text{argmin}_Y h_x^\lambda = \begin{cases} \{(0, R)\} & \text{if } \Lambda_1 \leq 0 \text{ and } \Lambda_2 < 0 \\ \{t(0, R) \mid t \in [0, 1]\} & \text{if } \Lambda_1 \leq 0 \text{ and } \Lambda_2 = 0 \\ \{(0, 0)\} & \text{if } \Lambda_2 > 0 \text{ and } \Lambda_1 < \sqrt{\Lambda_2} \\ \{t(R\Lambda_1, R) \mid t \in [0, 1]\} & \text{if } 0 \leq \Lambda_1 \leq 1 \text{ and } \Lambda_2 = \Lambda_1^2 \\ \{(R\Lambda_1, R)\} & \text{if } 0 \leq \Lambda_1 \leq 1 \text{ and } \Lambda_2 < \Lambda_1^2 \\ \{(R, R)\} & \text{if } \Lambda_2 \leq 1 \leq \Lambda_1 \\ \left\{t \left(R, \frac{R}{\sqrt{\Lambda_2}}\right) \mid t \in [0, 1]\right\} & \text{if } \Lambda_1 \geq 1 \text{ and } 1 \leq \Lambda_2 = \Lambda_1^2 \\ \{(R, R/\sqrt{\Lambda_2})\} & \text{if } \Lambda_1 \geq 1 \text{ and } 1 \leq \Lambda_2 < \Lambda_1^2. \end{cases} \quad (4.6)$$

As a result, $\varphi^\lambda(x) = \min_{y \in [0, R]^2} h_x^\lambda(y)$ has explicit formula given by:

$$\varphi^\lambda(x) = \begin{cases} \frac{R}{2x_2} \Lambda_2 & \text{if } \Lambda_1 \leq 0 \text{ and } \Lambda_2 \leq 0 \\ 0 & \text{if } \Lambda_2 \geq 0 \text{ and } \Lambda_1 \leq \sqrt{\Lambda_2} \\ \frac{R}{2x_2} (\Lambda_2 - \Lambda_1^2) & \text{if } 0 \leq \Lambda_1 \leq 1 \text{ and } \Lambda_2 \leq \Lambda_1^2 \\ \frac{R}{2x_2} (\Lambda_2 - 2\Lambda_1 + 1) & \text{if } \Lambda_2 \leq 1 \leq \Lambda_1 \\ \frac{R}{x_2} (\sqrt{\Lambda_2} - \Lambda_1) & \text{if } \Lambda_1 \geq 1 \text{ and } 1 \leq \Lambda_2 \leq \Lambda_1^2. \end{cases} \quad (4.7)$$

Proof. In light of Lemma 4.12, we need only to consider the values of $h_x^\lambda = c(x, y) - \lambda \cdot y$ on $\partial Y = \partial[0, R]^2$. In particular, we compute

$$h_x(y) = \begin{cases} y_2 \left(\frac{x_1^2}{2x_2} - \lambda_2 \right) & y \in \{0\} \times [0, R] \\ \frac{R}{2x_2} \left(x_1 - \frac{y_1}{R} \right)^2 - \lambda_1 y_1 - \lambda_2 R & y \in [0, R] \times \{R\} \\ \frac{y_2}{2x_2} \left(x_1 - \frac{R}{y_2} \right)^2 - \lambda_1 R - \lambda_2 y_2 & y \in \{R\} \times (0, R] \\ +\infty & y \in (0, R] \times \{0\}. \end{cases} \quad (4.8)$$

The first case is a linear function of y_2 , and can be expressed in the form $\frac{y_2}{2x_2} \Lambda_2$. Hence, it is minimized at $y_2 = 0$ if $\Lambda_2 \geq 0$ or $y_2 = R$ if $\Lambda_2 \leq 0$. The second case is a strictly convex quadratic function of y_1 with vertex at $y_1 = R(x_1 + x_2 \lambda_1) = R\Lambda_1$. Hence, the minimum of h_x^λ for $y \in [0, R] \times \{R\}$ occurs at $(0, R)$ if $\Lambda_1 \leq 0$, $(R\Lambda_1, R)$ if $0 \leq \Lambda_1 \leq 1$, and (R, R) if $\Lambda_1 \geq 1$. The third case is a strictly convex function of y_2 , which tends to $+\infty$ as $y_2 \rightarrow 0^+$. We may check that this function has a unique critical point on $(0, \infty)$ at $y_2 = R/\sqrt{x_1^2 - 2\lambda_2 x_2} = R/\sqrt{\Lambda_2}$ if $\Lambda_2 > 0$, and no critical points on $(0, \infty)$ if $\Lambda_2 \leq 0$. Hence, the minimum occurs at $y_2 = R$ if $\Lambda_2 \leq 1$ (as in this case, $R/\sqrt{\Lambda_2} \geq R$) and at $y_2 = R/\sqrt{\Lambda_2}$ if $\Lambda_2 \geq 1$.

Thus, we have up to five candidate points for the minimum to occur; $(0, 0)$, $(0, R)$, $(R\Lambda_1, R)$, (R, R) , and $(R, R/\sqrt{\Lambda_2})$. Explicitly computing and comparing the values of h_x^λ at these points, taking care to exclude $(R\Lambda_1, R)$ if $\Lambda_1 \notin [0, 1]$ and $(R, R/\sqrt{\Lambda_2})$ if $\Lambda_2 < 1$, we arrive at the desired formulas. Notice also that, if the minimum is attained at both $(0, 0)$ and at some other point on the boundary, then the entire line segment between these two points also minimizes h_x^λ , as a consequence of the 1-homogeneity of φ^λ . In all other cases, 1-homogeneity rules out interior minima. ■

Remark 4.14 (Interpretation of Proposition 4.13). Within the electricity industry, Proposition 4.13 can be interpreted as prescribing different outcomes to power plants depending on their pre-regulation characteristics. In particular, firms with $Y(x; \lambda, \eta) = 0$ are made obsolete by the policy (λ, η) and totally shutter as a result. Firms with $Y_1(x; \lambda, \eta) = 0$ will switch to totally clean production as a response to regulation. Other firms, with

$Y_1(x; \lambda, \eta) \neq 0$, will likely need to make more incremental modifications to their production processes (e.g. by retooling their production processes, or by capturing their emissions before they are released into the atmosphere).

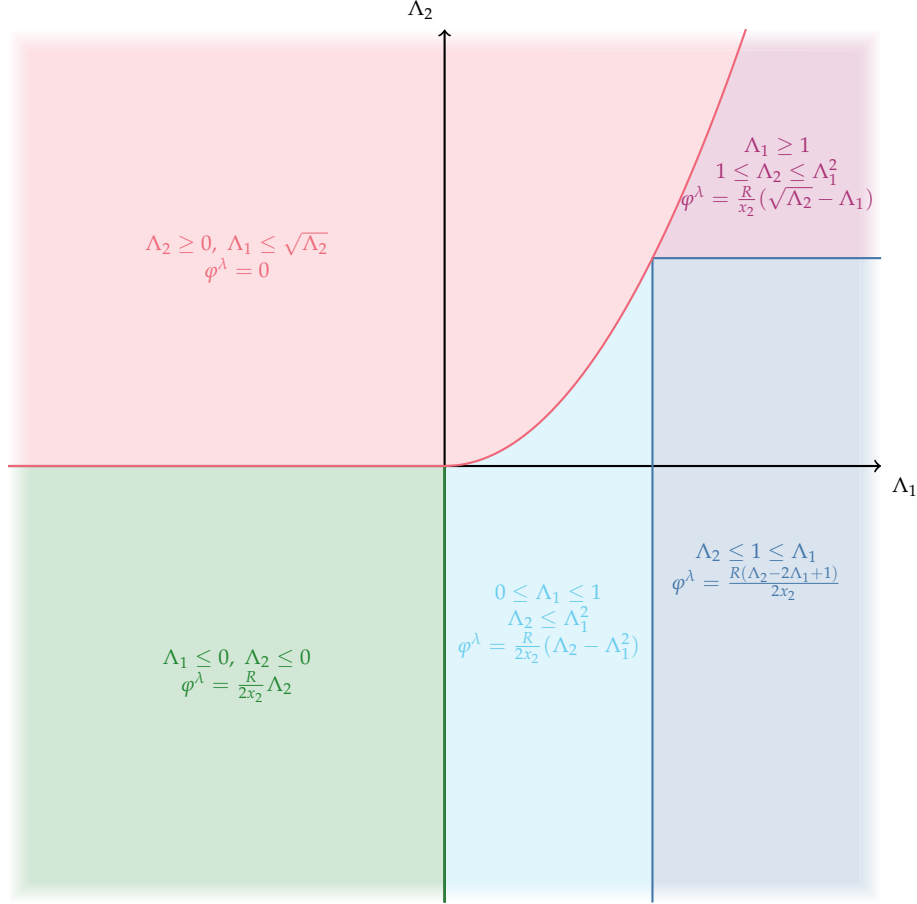


Figure 4.1: A figure demonstrating the various cases of Proposition 4.13, and the corresponding optimal choice of φ^λ from Equation (4.4). In other words, φ^λ is the largest continuous bounded function satisfying $c(x, y) + \lambda \cdot y \geq \varphi(x)$.

Remark 4.15. The classification of minimizers of h_x^λ in Proposition 4.13 should remind the reader of Proposition 3.19, where we classified the minimal values of the function $g_x(y) = c(x, y) - \lambda \cdot y - \eta|y - m|^2$. In particular, $\Lambda_1(x; \lambda) = B_1(x; \lambda, 0)/2$ and $\Lambda_2(x; \lambda) = B_2(x; \lambda, 0)$, where B_1 and B_2 are defined in the statement of Proposition 3.19. One key difference, however, is that, in the absence of a variance bound, the dual problem $\mathcal{S}_{X,Y}(\mu, m)$ lacks the strict convexity of $\mathcal{S}(\mu, m; \tau)$, and hence minimizers are not always unique.




Figure 4.2: The partition of X induced by Proposition 4.13 with $\lambda = (0.5, -1)$. This figure illustrates that all five of the regions depicted in Figure 4.1 can appear for a fixed λ (although, for other choices of λ , only some regions may appear).

Remark 4.16 (Primal Minimizers). Given $\lambda \in \mathbb{R}^2$ for which the pair $(\varphi^\lambda, \lambda)$ is optimal in the dual problem $\mathcal{S}_{X,Y}(\mu, m)$, Proposition 4.13 sheds light on how to select γ optimal for the primal problem $\mathcal{I}_{X,Y}(\mu, m)$. In particular, given a firm with pre-regulation characteristics x , Proposition 4.13 implies that this firm should be assigned post-regulation characteristics $y \in \operatorname{argmin}_Y h_x^\lambda$. The fact that $\operatorname{argmin}_Y h_x^\lambda$ is not always a singleton implies that, at least for some (μ, m) , $\mathcal{I}_{X,Y}(\mu, m)$ admits minimizers that cannot be written as transportation maps. More particularly, if $\operatorname{argmin}_Y h_x^\lambda$ is a

line segment, rather than a singleton, then the mass at x_1 should be distributed across this line segment in a way that meets the mean constraint $\iint y d\gamma(x, y) = m$. Through the lens of economics, this implies that firms may use mixed strategies to respond to the regulator's actions. For example, in the case where $\operatorname{argmin}_Y h_x^\lambda = \{t(R\Lambda_1, R) \mid t \in [0, 1]\}$, this could correspond to requiring some threshold firms to maximize production, while requiring others to decrease production or shutter. Alternatively, this could correspond to scaling all firms' production by a specific fraction.

4.5 NUMERICAL METHODS

In this section, we describe some results obtained from numerically implementing Proposition 4.13 as in the `Compact Domains.ipynb` Jupyter notebook at <https://github.com/Camdav/Thesis-Numerics>. For some background, these methods use Equation (4.6) to compute $\operatorname{argmin}_{y \in Y} [c(x, y) - \lambda \cdot y]$ and Equation (4.8) to compute $\varphi^\lambda(x) = \min_{y \in Y} [c(x, y) - \lambda \cdot y]$. With these quantities computed, the dual problem is identified with an unconstrained two-variable optimization problem in λ as in Corollary 4.6. This optimization problem is then solved using the Nelder-Mead method [78] to obtain an optimal λ , from which we recover the optimal $\varphi = \varphi^\lambda$, and hence the value of the dual supremum $\mathcal{S}_{X,Y}(\mu, m)$. This optimal choice of λ also allows us to use Equation (4.6) to identify measures which should, in theory, minimize the primal problem $\mathcal{I}_{X,Y}(\mu, m)$. However, as we will see in Remark 4.18, the fact that the argmin in Equation (4.6) may contain multiple points, as well as the discontinuity of this argmin in x , leads to some numerical instability. Our experience suggests that this instability is especially serious for discrete measures with a small number of points in their support, and therefore partially solve these issues by replacing point masses with clouds of nearby point masses in Remark 4.19.

Remark 4.17 (A Case of Successful Optimization). Let $Y = [0, 3]^2$, let $\mu = \frac{1}{4}(\delta_{(1,1)} + \delta_{(2,1)} + \delta_{(1,2)} + \delta_{(2,2)})$, and let $m = (1.3, 1.5)$. Then the objective function J from Corollary 4.6 is maximized at $\lambda \approx (-0.089, 0.146)$, and we can recover the optimal dual potential φ^λ from Equation (4.8). Moreover, in this case, the primal infimum and dual supremum are both approximately 0.0089, and the candidate γ_0 resulting from our numerical methods satisfies the mean constraint $\iint y d\gamma_0 = m = (1.3, 1.5)$.

Remark 4.18 (Primal Optimization Failure). Again let $Y = [0, 3]^2$ and $\mu = \frac{1}{4}(\delta_{(1,1)} + \delta_{(2,1)} + \delta_{(1,2)} + \delta_{(2,2)})$, but instead let $m = (1, 2)$. In this case, the objective function J from Corollary 4.6 is maximized at $\lambda \approx (-0.462, 0.710)$ — as expected, the regulator must scale up the emissions tax to meet the decreased emissions target, and increase the production

subsidy to meet the increased production quota. As before, we recover φ^λ from Equation (4.8), which allows us to compute that $\mathcal{S}_{X,Y}(\mu, m) \approx 0.3462$. However, the probability measure γ_0 which our algorithm suggests as a primal optimizer fails the mean constraint, as $\iint y \, d\gamma_0(x, y) \approx (0.46, 1.5)$. Moreover, $\iint c \, d\gamma_0 \approx 0.2397$ which is about 30% smaller than $\mathcal{S}_{X,Y}(\mu, m)$, and which further highlights that γ_0 is inadmissible for the primal problem $\mathcal{I}_{X,Y}(\mu, m)$. Closer inspection reveals that, for the optimal λ ,

$$\Lambda_2((2, 2), \lambda) \approx 1.1597 \approx (1.0769)^2 \approx (\Lambda_1((2, 2), \lambda)).$$

In other words, the point $(2, 2)$ lies close to the boundary of the third case from Equation (4.6), (where $\Lambda_2 > 0$, $\Lambda_1 < \sqrt{\Lambda_2}$, and $\operatorname{argmin}_{y \in Y}[c(x, y) - \lambda \cdot y] = \{(0, 0)\}$), and the last case, (where $\Lambda_1 \geq 1$, $1 \leq \Lambda_2 \leq \Lambda_1^2$, and $\operatorname{argmin}_{y \in Y}[c(x, y) - \lambda \cdot y] = \{(3, 3/\sqrt{\Lambda_2})\}$). We suspect that the point $(2, 2)$ lies precisely on the boundary between these two cases, which corresponds to the second-to-last case of Equation (4.6), where $\operatorname{argmin}_{y \in Y}[c(x, y) - \lambda \cdot y] = \{t(3, 3/\sqrt{\Lambda_2}) : t \in [0, 1]\}$. However, numerical errors mean that it is misidentified as belonging to the third case of Equation (4.6), where $\operatorname{argmin}_{y \in Y}[c(x, y) - \lambda \cdot y] = \{(0, 0)\}$. Given that the source measure μ assigns a quarter of its mass to $(2, 2)$, this causes significant errors in the computation of both $\iint y \, d\gamma_0(x, y)$ and $\iint c \, d\gamma_0$. The second-to-last case of Equation (4.6) suggests that we may modify γ_0 by assigning mass at $(2, 2)$ to $0.72(3, 3/\sqrt{\Lambda_2})$, rather than $(0, 0)$. In this case, we manually check that the mean constraint is satisfied and $\iint c \, d\gamma_0 \approx 0.3768 > 0.3462$. We suspect that this small discrepancy implies that the true optimizer γ_0 must contain multiple points of $\{(2, 2)\} \times \{t(3, 3/\sqrt{\Lambda_2}) : t \in [0, 1]\}$ in its support.

Remark 4.19 (Better Numerics With Point Clouds). Given that the numerical issues we encountered in Remark 4.18 arise when (Λ_1, Λ_2) are on or near thresholds between cases of Equation (4.6), it is natural to ask what happens if we replace μ with a measure with more points in its support. Heuristically, a smaller fraction of the mass should wind up on or near thresholds, leading to more accurate numerical methods. To demonstrate this, we take $Y = [0, 3]^2$ and $m = (1, 2)$ as in Remark 4.18, but split each point mass of μ into a large number (in our case 1024) of point masses randomly distributed over a small square centred at the original point (we choose its side length to be 0.002). In this case, our numerical methods suggest yield a candidate primal optimizer with mean $(0.998, 1.998) \approx (1, 2)$, and with $\iint c \, d\gamma_0 \approx 0.3457$. Given that our methods also yield that $\mathcal{S}_{X,Y}(\mu, m) \approx 0.3461$, we find that $\iint c \, d\gamma_0$ is around 0.12% smaller than what $\mathcal{I}_{X,Y}(\mu, m)$ should be. Thus, as this error is around 250 times smaller than the error from Remark 4.18, we conclude that the issues identified in

Remark 4.18 are partially, but not fully, resolved by splitting point masses into point clouds. It is reasonable to expect that the results would improve further if we had split each point mass of μ into an even larger number of point masses.

MORE THAN CARBON DIOXIDE: A FRAMEWORK FOR GEOGRAPHICALLY HETEROGENEOUS POLLUTANTS

5.1 OPTIMAL TRANSPORT MODELS OF HETEROGENEOUS DAMAGES

Our starting point for modelling heterogeneous damages is the cost function (1.5). In using this cost function, we implicitly assume that the cost for a firm to change its behaviour is independent of the external damages caused.

On the other hand, regulators will wish to limit total damages, rather than just total emissions, so the regulator's target should be affected by heterogeneity. To this end, we model the initial state of an industry by a probability measure $\mu \in \mathcal{P}(X \times G)$, where $X \subseteq [0, \infty)^2$ represents the space of possible (x_1, x_2) -values, which have the same interpretation as in the preceding chapters, and where G refers to the geographic location of emitters, and will typically be a subset of either the sphere S^2 or the plane \mathbb{R}^2 . We interpret the mass fraction prescribed by μ as representing a number of firms, regardless of the firms' emissions and production values. It is particularly instructive to consider the discrete case where $\mu = \sum_{i=1}^n \delta_{(x^i, g^i)}$, where $(x^i, g^i) \in X \times G$ for each i . This represents the case of an industry with a total of n firms, where the i -th firm is located at the position $g^i \in G$, has pre-regulation emissions intensity x_1^i , can change its emissions intensity with ease x_2^i .

We will also introduce a continuous bounded function $p \in C_b(G)$, where, for $z \in G$, $p(z)$ represents the damages caused by emitting one unit of pollutants at position z . As a motivating example, we may think of p as representing a mollified version of the population density over geography G , but, of course, practical applications may need to consider additional sources of damage. For example, prevailing wind patterns can lead to increased exposure for people living downwind of airborne pollution

sources. Waterborne pollution may have impacts on people living far downstream, while sparing people who live even slightly upstream from the polluters. Moreover, risks such as ecological degradation in sensitive ecosystems do impact humans in an easily quantifiable way, but may be of great interest to policymakers. With all of this being said, determining an appropriate choice of p is not the focus of the present work, so we will work with an arbitrary $p \in C_b(G)$ for the duration.

We now introduce three models which quantify the impact that emissions control regulation has on the behaviour of polluting firms. In particular, each model corresponds to one of the following hypotheses about the possibility of relocating firms:

- (H1)** Polluting firms cannot be relocated, and new polluting firms cannot be opened. That is, all post-regulation emissions must occur in locations where there were already emitters pre-regulation. In this case, polluters can only respond to regulation by modifying their production methods or scaling their production levels.
- (H2)** Polluting firms can be freely relocated and new polluting firms can be freely opened at any location in G .
- (H3)** Polluting firms can be relocated at a cost. We can model this by introducing a geographic cost function $\tilde{c} : X \times G \times Y \times G \rightarrow \mathbb{R}$, where $\tilde{c}(x, g, y, g')$ represents the cost of replacing a firm with pre-regulation characteristics x and location g with a firm with post-regulation characteristics y and location g' . For example, if we assume that the costs of changes to the production process are independent from the costs of relocation, we can represent $\tilde{c}(x, g, y, g') = c(x, y) + r(g, g')$, where c is the cost from [79] and r is the cost to relocate from location g to location g' .

We also provide some remarks on these assumptions. It should be relatively clear to the reader that the assumption **(H3)** is the most comprehensive and realistic – with modern technology, it is technically feasible to produce goods almost anywhere, but a reasonable choice of \tilde{c} will reflect that it is costly to produce goods in remote areas. In fact, as we shall see later, assumption **(H1)** is a special case of **(H3)** where we enforce that $\tilde{c}(x, g, y, g') = +\infty$ if $g \neq g'$. otherwise. Likewise, assumption **(H2)** can be interpreted as a special case of **(H3)** when $\tilde{c}(x, g, y, g') = c(x, y)$. On the other hand, the assumption **(H2)** may be overly permissive in many contexts, especially if G represents a large geographical region. For example, if the damage function p represents population density, it is plausible to suspect that a model based on **(H2)** will suggest that a regulator move

firms to extremely remote areas – for example, it could suggest that Denmark move its polluting companies to the Greenlandic Ice Cap, or that China move its polluters to the Gobi Desert. Despite this, **(H2)** could be a useful assumption to make when G represents a small jurisdiction and the cost of relocating firms is small.

While we will later introduce models corresponding to each of these assumptions, our primary focus will be on assumption **(H1)**. This assumption is appealing because it is much simpler to work with than assumption **(H3)** and, as we will see later, it can be formulated as a type of multi-marginal optimal transport problem. However, we should also address the assumption's key drawback, which is that it does not allow for polluting firms to change location, or for new polluting firms to open. These restrictions are less stringent than they seem at first glance for the following reasons. First, it is likely that economic and environmental factors are already partially reflected by the pre-regulation distribution of firms, so it is plausible to suggest that many of these firms are already located in economically viable areas, and that at least some of them are located in areas where they will not cause egregious environmental damage. Second, the mere act of opening a polluting firm in a new location may be unpalatable to regulators – they answer to politicians whose constituents may be upset by policies which could cause them direct harm. As such, this case could correspond to a case where a government places a moratorium on new emissions while allowing existing emitters to continue operations. On the other hand, it is reasonable to assume that there would be less opposition to increasing production and emissions at existing facilities.

We now turn our attention to introducing mathematical models for each of the assumptions **(H1)**, **(H2)**, and **(H3)**.

*A Model for **(H1)***

In this case, we encode the pre-regulation, post-regulation, and geographical information about the industry by $\gamma \in \mathcal{P}(X \times Y \times G)$, and we will seek to optimize $\int \int \int c(x, y) d\gamma(x, y, g)$ subject to suitable constraints on γ . We first impose the constraint that $\pi_{\#}^{(X, G)} \gamma = \mu$, which ensures that γ encapsulates information about the pre-regulation state and geographic distribution of the industry. Next, we seek to impose constraints on the post-regulation state $\pi_{\#}^{(Y, G)} \gamma$ of the industry, which reflect the regulator's production quota, and a cap on aggregate damages. In the case of homogeneous damages, we imposed a regulatory constraint $\int \int_{X \times Y} y d\gamma(x, y) = m$ to prescribe an aggregate post-regulation emissions level of m_1 and an ag-

gregate post-regulation production level of m_2 . We adopt a post-regulation production constraint

$$\iiint_{X \times Y \times G} y_2 d\gamma(x, y, g) = m_2$$

which is a direct analogue of the homogeneous case, reflecting an assumption that the amount of production is independent of the geography of the producer. In contrast, an aggregate post-regulation emissions standard is not relevant to the case of heterogeneous damages, as emissions are not directly proportional to damages. Instead, we set an ‘acceptable damage threshold’ of $D \geq 0$, and bound the aggregate damage by requiring that

$$\iiint_{X \times Y \times G} y_1 p(g) d\gamma(x, y, g) \leq D.$$

This leads us to the following primal minimization problem:

$$\inf_{\gamma \in \mathcal{P}(X \times Y \times G)} \left\{ \iiint c(x, y) d\gamma(x, y, g) \mid \pi_{\#}^{(X, G)} \gamma = \mu, \right. \\ \left. \iiint y_1 p(g) d\gamma(x, y, g) \leq D, \iiint y_2 d\gamma(x, y, g) = m_2 \right\}. \quad (5.1)$$

We notice that this problem is analogous to the primal problems discussed in the earlier chapters of this work, but here we prescribe the (X, G) -marginal of γ instead of its X -marginal, and we replace the constraints on the Y -marginal of γ with constraints on its (Y, G) -marginal. This suggests an interpretation of (5.1) in terms of command-and-control policies, where the regulator prescribes new emissions and production levels to polluting firms. More precisely, the regulator’s problem is that of assigning companies new post-regulation emissions and production levels which minimize the aggregate cost $\iiint c(x, y) d\gamma(x, y, g)$ to companies, while maintaining the production level m_2 and limiting aggregate damages to be at most D .

The primal minimization problem also bears some resemblance to multi-marginal optimal transportation problems of the type discussed, for example, by Pass in [87]. In the classical multi-marginal optimal transport problem, one would prescribe each marginal of γ , and then minimize the cost subject to these marginal constraints. In (5.1), we instead prescribe the (X, G) -marginal of γ and then impose a weaker constraint on the (Y, G) -marginal. This has the effect of prescribing the X -marginal and G -marginal of γ (as well as their joint distribution), while leaving a weaker constraint on the Y -marginal of γ , and in particular, the integral of $(y_1 p(g), y_2)$ with respect to the joint distribution $\pi_{\#}^{(Y, G)} \gamma$.

Again, as in earlier chapters, we define a dual problem and interpret it in terms of market-based policies. To do so, we introduce Lagrange multipliers corresponding to each of the constraints, and notice that, for a non-negative measure $\gamma \in \mathcal{M}_+(X \times Y \times G)$

$$\sup_{\varphi \in C_b(X, G)} \left\{ \iint \varphi \, d\mu - \iiint \varphi \, d\gamma \right\} = \begin{cases} 0 & \text{if } \pi_{\#}^{(X, G)} \gamma = \mu \\ +\infty & \text{else} \end{cases},$$

$$\sup_{\lambda_1 \leq 0} \left\{ \lambda_1 \left(D - \iiint y_1 p(g) \, d\gamma \right) \right\} = \begin{cases} 0 & \text{if } \iiint y_1 p(g) \, d\gamma \leq D \\ +\infty & \text{else} \end{cases},$$

and

$$\sup_{\lambda_2 \in \mathbb{R}} \left\{ \lambda_2 \left(m_2 - \iiint y_2 \, d\gamma \right) \right\} = \begin{cases} 0 & \text{if } \iiint y_2 \, d\gamma = m_2 \\ +\infty & \text{else} \end{cases}.$$

Implementing each of these constraints, we rewrite (5.1) as a saddle-point inf-sup optimization problem, formally exchange the infimum and the supremum, and finally group terms to arrive at the following dual problem:

$$\sup_{\substack{\varphi \in C_b(X \times G) \\ \lambda_1 \leq 0, \lambda_2 \in \mathbb{R}}} \left\{ \iint \varphi \, d\mu + \lambda_1 D + \lambda_2 m_2 \mid \begin{array}{l} c(x, y) \geq \varphi(x, g) \\ + \lambda_1 y_1 p(g) + \lambda_2 y_2 \text{ on } X \times Y \times G \end{array} \right\}. \quad (5.2)$$

We interpret λ_1 as a tax on damages (rather than just emissions) and λ_2 as a subsidy for production.

In a formal manner, if we define

$$\varphi^\lambda(x, g) := \inf_{y \in Y} \{ c(x, y) - \lambda_1 y_1 p(g) - \lambda_2 y_2 \},$$

we may rewrite (5.2) as the following bi-level optimization problem

$$\sup_{\substack{\lambda_1 \leq 0 \\ \lambda_2 \in \mathbb{R}}} \left\{ \iint \varphi^\lambda(x, g) \, d\mu(x, g) + \lambda_1 D + \lambda_2 m_2 \right\}.$$

This allows us to interpret $y(x, g) \in \operatorname{argmin}_{y \in Y} \{ c(x, y) - \lambda_1 y_1 p(g) - \lambda_2 y_2 \}$ as optimal post-regulation emissions and production levels for a firm at position g with pre-regulation characteristics x , when subject to the combination of taxes and subsidies represented by λ . In turn, φ^λ describes the minimal cost for such a firm to respond to the tax-subsidy combination

λ . As in the geographically homogeneous case, we expect the optimal value of λ , if it exists, to depend on the underlying parameters of the model – in this case, the cost function c , the pre-regulation distribution μ , the acceptable damages D , and the production quota m_2 .

A Model for (H2)

In this case, the pre-regulation locations of firms are entirely irrelevant, as we only aim to cap post-regulation damages, and firms may change their locations for free. As such, we again encode the pre-regulation and post-regulation characteristics and geography of the industry by $\gamma \in \mathcal{P}(X \times Y \times G)$ and seek to optimize $\iiint c(x, y) d\gamma(x, y, g)$. Since the pre-regulation geographic distribution of the industry is irrelevant, we encode the pre-regulation state of the industry by requiring that $\pi_{\#}^X \gamma = \pi_{\#}^X \mu$, where we recall that we originally defined $\mu \in \mathcal{P}(X \times G)$. We could, of course, instead prescribe that $\pi_{\#}^X \gamma = \tilde{\mu}$, for some probability measure $\tilde{\mu} \in \mathcal{P}(X)$ representing the pre-regulation distribution of firm characteristics.

As we did with model (H1), we account for the post-regulation production quota m_2 by requiring $\iiint y_2 d\gamma = m_2$, and we account for the damage threshold by requiring that $\iint y_1 p(g) d(\pi_{\#}^{YG} \gamma)(y, g) \leq D$. These conventions lead us to the minimization problem

$$\inf_{\gamma \in \mathcal{P}(X \times Y \times G)} \left\{ \iiint c(x, y) d\gamma(x, y, g) \mid \pi_{\#}^X \gamma = \pi_{\#}^X \mu, \right. \\ \left. \iint y_1 p(g) d\gamma(x, y, g) \leq D, \iiint y_2 d\gamma(x, y, g) = m_2 \right\}.$$

This is quite similar to the problem we defined in Equation (5.1), and we can use analogous steps to find that it has a formal dual problem given by

$$\sup_{\substack{\varphi \in C_b(X) \\ \lambda_1 \leq 0, \lambda_2 \in \mathbb{R}}} \left\{ \iint \varphi(x) d\mu(x, g) + \lambda_1 D + \lambda_2 m_2 \mid c(x, y) \geq \varphi(x) \right. \\ \left. + \lambda_1 y_1 p(g) + \lambda_2 y_2 \text{ on } X \times Y \times G \right\}. \quad (5.3)$$

In this case, we again interpret λ_1 as a tax on damages, and λ_2 as a subsidy on production. A key distinction here is that, when rewriting (5.3) as a

bilevel optimization problem, we must take into account the fact that firms have free choice of g . As such, we define

$$\varphi^\lambda(x) := \inf_{y \in Y, g \in G} \{c(x, y) - \lambda_1 y_1 p(g) - \lambda_2 y_2\}$$

to get the following bi-level optimization problem.

$$\sup_{\substack{\lambda_1 \leq 0 \\ \lambda_2 \in \mathbb{R}}} \left\{ \iint \varphi^\lambda(x) d\mu(x, g) + \lambda_1 D + \lambda_2 m_2 \right\}.$$

Here, $(y(x), g(x)) \in \operatorname{argmin}_{Y \times G} \{c(x, y) - \lambda_1 y_1 p(g) - \lambda_2 y_2\}$ represents the optimal post-regulation emissions levels, production levels, and location for a firm with pre-regulation characteristics x responding to the market-based policy λ , and φ^λ represents the cost of this optimal response.

A Model for (H3)

Finally, we account for the more general assumption **(H3)**. In this case, both pre-regulation and post-regulation locations of firms are relevant so, for the purpose of avoiding confusion, we introduce the notation $X_G := X \times G$ and $Y_G := Y \times G$. Under this assumption, we can model a command-and-control policy by $\gamma \in \mathcal{P}(X_G \times Y_G)$; for this model we establish a convention of using single integrals to refer to integration over X_G or Y_G and double integrals to integrate over their product. Note that, unlike the models based on assumptions **(H1)** and **(H2)**, we should now consider a cost which depends on the geographic distribution of firms. In other words, where (x, g) represent coordinates of X_G and (y, g') represent coordinates of Y_G , we now wish to optimize a cost functional of the form $\iint \tilde{c}(x, g, y, g') d\gamma(x, g, y, g')$, where $\tilde{c}(x, g, y, g')$ refers the cost to transition a firm with pre-regulation characteristics x and location g to have post-regulation characteristics y at location g' . For example, if we assume that relocation costs are independent from the transition costs c , we could write $\tilde{c}(x, g, y, g') = c(x, y) + r(g, g')$, where r represents the cost of geographic relocation from position g to position g' . Our constraints, on the other hand, share more similarities with those we introduced earlier. In particular, we require that $\pi_{\#}^{X_G} \gamma = \mu$, implement the damage

constraint $\iint y_1 p(g') d\gamma(x, g, y, g') \leq D$, and finally the production quota $\iint y_2 d\gamma(x, g, y, g') = m_2$. This leads us to the primal problem

$$\inf_{\gamma \in \mathcal{P}(X_G \times Y_G)} \left\{ \iint \tilde{c} d\gamma \mid \pi_{\#}^{X_G} \gamma = \mu, \iint y_1 p(g') d\gamma(x, g, y, g') \leq D, \right. \\ \left. \iint y_2 d\gamma(x, g, y, g') = m_2 \right\}, \quad (5.4)$$

which has dual problem

$$\sup_{\substack{\varphi \in C_b(X_G) \\ \lambda_1 \leq 0, \lambda_2 \in \mathbb{R}}} \left\{ \int \varphi d\mu + \lambda_1 D + \lambda_2 m_2 \mid \tilde{c}(x, g, y, g') \geq \varphi(x, g) \right. \\ \left. + \lambda_1 y_1 p(g') + \lambda_2 y_2 \text{ on } X_G \times Y_G \right\}. \quad (5.5)$$

We write this as a bilevel optimization problem by defining $\varphi^\lambda(x, g) := \inf_{(y, g') \in Y_G} \{ \tilde{c}(x, g, y, g') - \lambda_1 y_1 p(g') + \lambda_2 y_2 \}$, which yields

$$\sup_{\substack{\lambda_1 \leq 0 \\ \lambda_2 \in \mathbb{R}}} \left\{ \int \varphi^\lambda d\mu + \lambda_1 D + \lambda_2 m_2 \right\}.$$

In this case,

$$(y(x, g), g'(x, g)) \in \underset{Y_G}{\operatorname{argmin}} \{ \tilde{c}(x, g, y, g') - \lambda_1 y_1 p(g') + \lambda_2 y_2 \}$$

formally represents the optimal post-regulation emissions levels, production levels, and location for a firm with pre-regulation characteristics x and location g . As usual, $\Phi(x, g, \lambda)$ represents the cost of this optimal response. We now notice that each of the other models can be interpreted as special cases of the model **(H3)**.

Remark 5.1 (**(H1)** and **(H2)** as Special Cases of **(H3)**). It is not difficult to check that, if we consider the cost function

$$\tilde{c}(x, g, y, g') = \begin{cases} c(x, y) & \text{if } g = g' \\ +\infty & \text{else.} \end{cases},$$

and replace γ with its (X, Y, G) -marginal $\pi_{\#}^{(X, Y, G)} \gamma$, the primal and dual problems for model **(H3)** reduce to those for model **(H1)** with appropriate constraints. In other words, **(H1)** can be viewed as a special case of **(H3)**, with an infinite cost penalty for relocation.

We can also recover **(H2)** as a special case of **(H3)** by considering a cost of the form $\tilde{c}(x, g, y, g) = c(x, y)$, and replacing γ with its (X, Y_G) -marginal $\pi_{\#}^{X, Y_G} \gamma$. Thus, **(H2)** can be thought of as a special case of **(H3)** with no penalty for relocation.

5.2 DUALITY IN THE COMPACT CASE

We follow a strategy analogous to that of [103, Chapter 1.1] to establish duality. The key element in doing so is the following version of Rockafellar duality, stated in [103, Theorem 1.9]:

Theorem 5.2 (Fenchel-Rockafellar Duality; Theorem 1.9 of [103]). *Let E be a normed vector space, E^* its topological dual space, and Θ, Ξ two convex functions on E with values in $\mathbb{R} \cup \{+\infty\}$. Let Θ^*, Ξ^* be the Legendre-Fenchel transforms of Θ and Ξ , respectively. Assume that there exists $z_0 \in E$ such that $\Theta(z_0) < +\infty$, $\Xi(z_0) < +\infty$, and Φ is continuous at z_0 . Then*

$$\inf_E [\Theta + \Xi] = \max_{z^* \in E^*} [-\Theta^*(-z^*) - \Xi^*(z^*)]. \quad (5.6)$$

We use this theorem to prove the following result, which is analogous to [103, Theorem 1.3]:

Proposition 5.3. *Let $X, Y \subseteq \mathbb{R}^2$ be compact, let G be compact and denote $X_G = X \times G$ and $Y_G = Y \times G$. Moreover, let $D \geq 0$, $m_2 \geq 0$, $\mu \in \mathcal{P}(X_G)$, $0 \leq p \in C_b(g)$, and $0 \leq \tilde{c} \in C(X_G \times Y_G)$. Furthermore, assume that p is not identically zero and that Y is not entirely contained within any line (i.e. that there are three non-collinear points in Y). Then the primal infimum (5.4) and dual supremum (5.5) agree.*

Proof. We apply Theorem 5.2 with $E = (C_b(X_G \times Y_g), \|\cdot\|_{\infty})$, where $\|\cdot\|_{\infty}$ is the supremum norm. In this case, the dual space is the space of regular Radon measures with the total variation norm. In other words, $E^* = (C_b(X_G \times Y_g), \|\cdot\|_{\infty}) = (\mathcal{M}(X \times Y), \|\cdot\|_{TV})$.

We apply the theorem with the functions $\Theta, \Xi : C_b(X_G \times Y_g) \rightarrow \mathbb{R}$ defined by

$$\Theta(u) = \begin{cases} 0, & \text{if } u \geq -\tilde{c} \text{ on } X_G \times Y_G \\ +\infty, & \text{else} \end{cases}$$

and

$$\Xi(u) = \begin{cases} \int \varphi d\mu + \lambda_1 D + \lambda_2 m_2, & \text{if } u(x, g, y, g') = \varphi(x, g) + \lambda_1 y_1 p(g') \\ & + \lambda_2 y_2 \text{ for some } \varphi \in C_b(X_G), \\ & \lambda_1 \geq 0, \text{ and } \lambda_2 \in \mathbb{R} \\ +\infty & \text{else.} \end{cases}.$$

Under the assumption that p is not identically zero, and that Y is not contained within a line, Ξ is well-defined. To see why, assume

$$\varphi(x, g) + \lambda_1 y_1 p(g') + \lambda_2 y_2 = \tilde{\varphi}(x, g) + \tilde{\lambda}_1 y_1 p(g') + \tilde{\lambda}_2 y_2.$$

By bringing all terms depending on (x, g) to one side and all terms depending on (y, g') to the other, we deduce that there exists $s \in \mathbb{R}$ such that $\varphi = \tilde{\varphi} + s$ and

$$\lambda_1 y_1 p(g') + \lambda_2 y_2 = \tilde{\lambda}_1 y_1 p(g') + \tilde{\lambda}_2 y_2 + s$$

for all $(y, g') \in Y \times G$. In particular, the vector $(y_1 p(g'), y_2)$ must lie in the space

$$\ell := \{v \in \mathbb{R}^2 : (\lambda_1 - \tilde{\lambda}_1, \lambda_2 - \tilde{\lambda}_2) \cdot v = s\},$$

which is a line unless $\lambda_1 = \tilde{\lambda}_1$ and $\lambda_2 = \tilde{\lambda}_2$. Taking g' such that $p(g') > 0$, and using the assumption that Y is not contained within any line, we deduce that $\lambda_1 - \tilde{\lambda}_1 = \lambda_2 - \tilde{\lambda}_2 = s = 0$. Thus, $\varphi = \tilde{\varphi}$, $\lambda_1 = \tilde{\lambda}_1$, and $\lambda_2 = \tilde{\lambda}_2$, proving that the function is well-defined.

Additionally, note that if $u_0 \equiv 1$, then $\Theta \equiv 0$ on some neighbourhood of u_0 , and hence is continuous at u_0 . Likewise, by taking $\varphi \equiv 1$ and $\lambda_1 = \lambda_2 = 0$, we deduce that $\Xi(u_0) = 1 < +\infty$, which proves that Θ and Ξ satisfy the hypotheses of Theorem 5.2.

We now compute the infimum on the left hand side of (5.6). The definitions of Θ and Ξ allow us to write it as

$$\inf_{\substack{\varphi \in C_b(X_G) \\ \lambda_1 \geq 0, \lambda_2 \in \mathbb{R}}} \left\{ \int \varphi d\mu + \lambda_1 D + \lambda_2 m_2 \mid -\tilde{c}(x, g, y, g') \leq \varphi(x, g) \right. \\ \left. + \lambda_1 y_1 p(g') + \lambda_2 y_2 \right\}.$$

Switching the signs of φ , λ_1 , and λ_2 , we find that this is equal to

$$- \sup_{\substack{\varphi \in C_b(X_G) \\ \lambda_1 \leq 0, \lambda_2 \in \mathbb{R}}} \left\{ \int \varphi d\mu + \lambda_1 D + \lambda_2 m_2 \mid \tilde{c}(x, g, y, g') \geq \varphi(x, g) \right. \\ \left. + \lambda_1 y_1 p(g') + \lambda_2 y_2 \text{ on } X_G \times Y_G \right\},$$

or rather the negative of the dual supremum.

To compute the right hand side of (5.6), we need to calculate Θ^* and Ξ^* . By an identical calculation to the one Villani uses in the proof of [103, Theorem 1.3], we deduce that, for $\gamma \in \mathcal{M}(X_G \times Y_G)$,

$$\Theta^*(-\gamma) = \begin{cases} \iint \tilde{c} \, d\gamma & \text{if } \gamma \in \mathcal{M}_+(X \times Y) \text{ is non-negative} \\ +\infty & \text{else.} \end{cases}$$

We can also compute Ξ^* . Applying the definitions of Ξ and its Fenchel-Legendre transform,

$$\begin{aligned} \Xi^*(\gamma) &= \sup_{u \in C_b(X_G \times Y_G)} \left[\iint u \, d\gamma - \Xi(u) \right] \\ &= \sup_{\substack{\varphi \in C_b(X_G) \\ \lambda_1 \geq 0, \lambda_2 \in \mathbb{R}}} \left\{ \iint [\varphi(x, g) + \lambda_1 y_1 p(g') + \lambda_2 y_2] \, d\gamma \right. \\ &\quad \left. - \int \varphi \, d\mu - \lambda_1 D - \lambda_2 m_2 \right\} \\ &= \begin{cases} 0 & \text{if } \pi_{\#}^{X_G} \gamma = \mu, \iint y_1 p(g') \, d\gamma \leq D, \text{ and } \iint y_2 \, d\gamma = m_2 \\ +\infty & \text{else.} \end{cases} \end{aligned}$$

Hence,

$$\begin{aligned} &\max_{\gamma \in \mathcal{M}(X_G \times Y_G)} [-\Theta^*(-\gamma) - \Xi^*(\gamma)] \\ &= \max_{\gamma \in \mathcal{M}(X_G) \times Y_G} \left\{ -\iint \tilde{c} \, d\gamma \mid \pi_{\#}^{X_G} \gamma = \mu, \iint y_1 p(g') \, d\gamma \leq D, \right. \\ &\quad \left. \iint y_2 \, d\gamma = m_2 \right\} \\ &= -\min \left\{ \iint \tilde{c} \, d\gamma \mid \pi_{\#}^{X_G} \gamma = \mu, \iint y_1 p(g') \, d\gamma \leq D, \right. \\ &\quad \left. \iint y_2 \, d\gamma = m_2 \right\}' \end{aligned}$$

which is the negative of the primal infimum, concluding our proof of duality. \blacksquare

We conclude this chapter with some comments on the application and limitations of our present approach:

Remark 5.4 (Domain Restrictions of Proposition 5.3). For the purpose of simplifying the exposition, Proposition 5.3 currently only applies to continuous costs \tilde{c} . However, the Newell-Stavins cost c may not be continuous if $x_2 = 0$ or $y_2 = 0$, and is extended by lower semicontinuity to these regions in (1.5). Thus, with the current result, if the definition of \tilde{c} involves the

Newell-Stavins cost c in a non-trivial way, we would expect to have to take $X, Y \subseteq [0, \infty) \times (0, \infty)$. We suspect that Proposition 5.3 will generalize to the case where $Y \subseteq [0, \infty)^2$ is compact; variance bounds of the kind implemented in Chapter 3 are also likely to ensure an analogous duality result.

Remark 5.5 (Continuity and Semicontinuity). Even on a domain where c is continuous, the function \tilde{c} which we would use to model the case **(H1)** where firms are not allowed to relocate is not continuous, and hence Proposition 5.3 does not directly apply. However, it should easily extend to lower semicontinuous \tilde{c} by approximating from below by bounded Lipschitz functions, as in [96, Box 1.5].

BIBLIOGRAPHY

- [1] United States Environmental Protection Agency, *Guidelines for preparing economic analyses*, Tech. report, United States Environmental Protection Agency, 2016.
- [2] Aurélien Alfonsi, Rafaël Coyaud, Virginie Ehrlacher, and Damiano Lombardi, *Approximation of optimal transport problems with marginal moments constraints*, *Math. Comp.* **90** (2021), no. 328, 689–737. MR 4194160
- [3] Hossa Almutairi, *Carbon tax based on the emission factor*, Ph.D. thesis, University of Waterloo, 2013.
- [4] James E. Anderson and Eric van Wincoop, *Gravity with gravitas: A solution to the border puzzle*, *The American Economic Review* **93** (2003), no. 1, 170–192.
- [5] Robert M. Anderson and Haosui Duanmu, *Cap-and-trade and carbon tax meet Arrow-Debreu*, *Econometrica* **93** (2025), no. 2, 357–393. MR 4889419
- [6] Kenneth J. Arrow, *An extension of the basic theorems of classical welfare economics*, *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, 1950, Univ. California Press, Berkeley-Los Angeles, Calif., 1951, pp. 507–532. MR 44815
- [7] Kenneth J. Arrow and Gerard Debreu, *Existence of an equilibrium for a competitive economy*, *Econometrica* **22** (1954), 265–290. MR 77069
- [8] World Bank, *GDP (current USD)*, 2024.
- [9] A. H. Barnett, *The Pigouvian tax rule under monopoly*, *The American Economic Review* **70** (1980), no. 5, 1037–1041.
- [10] Gary S. Becker, *A theory of marriage: Part i*, *Journal of Political Economy* **81** (1973), no. 4, 813–846.
- [11] Mathias Beiglböck, Christian Léonard, and Walter Schachermayer, *A general duality theorem for the Monge-Kantorovich transport problem*, *Studia Math.* **209** (2012), no. 2, 151–167. MR 2943840

- [12] Adrien Bilal and Diego R. Känzig, *The macroeconomic impact of climate change: Global vs. local temperature*, Working Paper 32450, National Bureau of Economic Research, July 2024.
- [13] Job Boerma, Aleh Tsyvinski, and Alexander P. Zemin, *Bunching and taxing multidimensional skills*, Working Paper 30015, National Bureau of Economic Research, May 2022.
- [14] ———, *Sorting with teams*, *Journal of Political Economy* **133** (2025), no. 2, 421–454.
- [15] Odran Bonnet, Alfred Galichon, Yu-Wei Hsieh, Keith O’Hara, and Matt Shum, *Yogurts choose consumers? Estimation of random-utility models via two-sided matching*, *Rev. Econ. Stud.* **89** (2022), no. 6, 3085–3114. MR 4506554
- [16] Emanuele Borgonovo, Valeria Di Cosmo, Daniele Mosso, Matteo Nicoli, Elmar Plischke, Laura Savoldi, and Anderson Rodrigo de Queiroz, *A framework for global sensitivity analysis in long-term energy systems planning using optimal transport*, *Energy* **338** (2025), 138788.
- [17] Emanuele Borgonovo, Alessio Figalli, Elmar Plischke, and Giuseppe Savaré, *Global sensitivity analysis via optimal transport*, *Management Science* **71** (2024), no. 5, 3809–3828.
- [18] A. Lans Bovenberg and Ruud A. de Mooij, *Environmental levies and distortionary taxation*, *The American Economic Review* **84** (1994), no. 4, 1085–1089.
- [19] A.L. Bovenberg and F. van der Ploeg, *Environmental policy, public finance and the labour market in a second-best world*, *Journal of Public Economics* **55** (1994), no. 3, 349–390.
- [20] Haim Brezis, *Functional analysis, Sobolev spaces and partial differential equations*, Universitext, Springer, New York, 2011. MR 2759829
- [21] Kate Brown, *Plutopia*, Oxford University Press, Oxford, 2013.
- [22] Marshall Burke, Solomon M. Hsiang, and Edward Miguel, *Global non-linear effect of temperature on economic production*, *Nature* **527** (2015), 235–239.
- [23] Ana Bušić, Thomas Le Corre, and Sean Meyn, *Moment constrained optimal transport for control applications*, Preprint at arXiv: 2208.01958, 2025.

- [24] Luis A. Caffarelli, Mikhail Feldman, and Robert J. McCann, *Constructing optimal maps for Monge's transport problem as a limit of strictly convex costs*, J. Amer. Math. Soc. **15** (2002), no. 1, 1–26. MR 1862796
- [25] Emanuele Campiglio, Simon Dietz, and Frank Venmans, *Optimal climate policy as if the transition matters*, Working Paper 387, Grantham Research Institute on Climate Change and the Environment, December 2022.
- [26] Statistics Canada, *Table 38-10-0131-01 capital and operating expenditures on environmental activities by geographic region (x 1,000,000)*.
- [27] Guillaume Carlier, Hugo Malamut, and Maxime Sylvestre, *Weak optimal transport with moment constraints: constraint qualification, dual attainment and entropic regularization*, Preprint at arXiv: 2511.16211, 2026.
- [28] Victor Chernozhukov, Alfred Galichon, Marc Henry, and Brendan Pass, *Identification of hedonic equilibrium and nonseparable simultaneous equations*, Journal of Political Economy **129** (2021), no. 3, 842–870.
- [29] Pierre-André Chiappori, Robert J. McCann, and Lars P. Nesheim, *Hedonic price equilibria, stable matching, and optimal transport: equivalence, topology, and uniqueness*, Econom. Theory **42** (2010), no. 2, 317–354. MR 2564439
- [30] Khai Xiang Chiong, Alfred Galichon, and Matt Shum, *Duality in dynamic discrete-choice models*, Quant. Econ. **7** (2016), no. 1, 83–115. MR 3491285
- [31] Pavel Ciaian and Johan F.M. Swinnen, *Land market imperfections and agricultural policy impacts in the new EU member states: A partial equilibrium analysis*, American Journal of Agricultural Economics **88** (2006), no. 4, 799–815.
- [32] Gerard Debreu, *The coefficient of resource utilization*, Econometrica **19** (1951), no. 3, 273–292.
- [33] Colin Decker, Elliott H. Lieb, Robert J. McCann, and Benjamin K. Stephens, *Unique equilibria and substitution effects in a stochastic model of the marriage market*, J. Econom. Theory **148** (2013), no. 2, 778–792. MR 3041038
- [34] Melissa Dell, Benjamin F. Jones, and Benjamin A. Olken, *Temperature shocks and economic growth: Evidence from the last half century*, American Economic Journal **4** (2012), no. 3, 66–95.

- [35] Xiaotie Deng, Christos Papadimitriou, and Shmuel Safra, *On the complexity of equilibria*, Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing (New York, NY, USA), STOC '02, Association for Computing Machinery, 2002, p. 67–71.
- [36] Andreas Duit, *Patterns of environmental collective action: Some cross-national findings*, Political Studies **59** (2010), no. 4, 900–920.
- [37] Arnaud Dupuy, Alfred Galichon, and Yifei Sun, *Estimating matching affinity matrices under low-rank constraints*, Inf. Inference **8** (2019), no. 4, 677–689. MR 4045480
- [38] Ivar Ekeland, *Notes on optimal transportation*, Econom. Theory **42** (2010), no. 2, 437–459. MR 2564444
- [39] Lawrence C. Evans and Ronald F. Gariepy, *Measure theory and fine properties of functions*, Studies in Advanced Mathematics, CRC Press, Boca Raton, FL, 1992. MR 1158660
- [40] Pablo D. Fajgelbaum and Edouard Schaal, *Optimal transport networks in spatial equilibrium*, Econometrica **88** (2020), no. 4, 1411–1452. MR 4131226
- [41] Alessio Figalli, Young-Heon Kim, and Robert J. McCann, *When is multidimensional screening a convex program?*, J. Econom. Theory **146** (2011), no. 2, 454–478. MR 2888826
- [42] Carolyn Fischer, *Combining rate-based and cap-and-trade emissions policies*, Climate Policy **3** (2003), S89–S103, Special Supplement on Defining and Trading Emission Targets.
- [43] Christian Flachsland, Robert Marschinski, and Ottmar Edenhofer, *To link or not to link: benefits and disadvantages of linking cap-and-trade systems*, Climate Policy **9** (2009), no. 4, 358–372.
- [44] Organisation for Economic Co-operation and Development, *Assessing the economic impacts of environmental policies: Evidence from a decade of OECD research*, Tech. report, OECD Publishing, 2021.
- [45] Meredith Fowlie and Nicholas Muller, *Market-based emissions regulation when damages vary across sources: What are the gains from differentiation?*, Journal of the Association of Environmental and Resource Economists **6** (2019), no. 3, 593–632.
- [46] Gero Friesecke, *Optimal transport—a comprehensive introduction to modeling, analysis, simulation, applications*, Society for Industrial and Applied Mathematics, Philadelphia, PA, [2025] ©2025. MR 4904240

- [47] Alfred Galichon, *Optimal transport methods in economics*, Princeton University Press, Princeton, NJ, 2016. MR 3586373
- [48] ———, *A survey of some recent applications of optimal transport methods to econometrics*, *Econom. J.* **20** (2017), no. 2, C1–C11. MR 3685644
- [49] ———, *Advances in economics and econometrics: Twelfth world congress*, ch. The Unreasonable Effectiveness of Optimal Transport in Economics, pp. 90–114, Cambridge University Press, 2025.
- [50] Alfred Galichon, Pierre Henry-Labordère., and Nizar Touzi, *A stochastic control approach to no-arbitrage bounds given marginals, with an application to lookback options*, *Ann. Appl. Probab.* **24** (2014), no. 1, 312–336. MR 3161649
- [51] Alfred Galichon and Robert McCann, *Special Issue: Optimal transportation, equilibrium, and applications to economics [Editorial]*, *Econom. Theory* **67** (2019), no. 2, 345–347. MR 3936050
- [52] Alfred Galichon and Bernard Salanié, *Cupid’s invisible hand: social surplus and identification in matching models*, *Rev. Econ. Stud.* **89** (2022), no. 5, 2600–2629. MR 4478278
- [53] Roy Gardner, *L. V. Kantorovich: The price implications of optimal planning*, *Journal of Economic Literature* **28** (1990), no. 2, 638–648.
- [54] Mikhail Golosov, John Hassler, Per Krusell, and Aleh Tsyvinski, *Optimal taxes on fossil fuel in general equilibrium*, *Econometrica* **82** (2014), no. 1, 41–88.
- [55] Lawrence H. Goulder and Andrew R. Schein, *Carbon taxes versus cap and trade: a critical review*, *Climate Change Economics* **04** (2013), no. 03, 1350010.
- [56] Lawrence H. Goulder and Roberton C. Williams III, *The substantial bias from ignoring general equilibrium effects in estimating excess burden, and a practical solution*, *Journal of Political Economy* **111** (2003), no. 4, 898–927.
- [57] Government of Canada, *Canada strong budget 2025*, <https://budget.canada.ca/2025/report-rapport/toc-tdm-en.html>, 2025, Federal Budget.
- [58] Jerry R. Green, Andreu Mas-Colell, and Michael D. Whinston, *Microeconomic theory*, Oxford University Press, New York, 1995.

- [59] Sreenivasulu Gumpu, Balakrishna Pamulaparthi, and N C Sahoo, *An optimal transport theory based approach for efficient dispatch of transactions in energy markets*, *Energy Conversion and Economics* **4** (2023), no. 3, 213–231.
- [60] Pierre Henry-Labordère, *Model-free hedging*, Chapman & Hall/CRC Financial Mathematics Series, CRC Press, Boca Raton, FL, 2017, A martingale optimal transport viewpoint. MR 3699668
- [61] Ulrich Hounyo, Johnson Kakeu, and Li Lu, *Heterogeneity in carbon intensity patterns: A subsampling approach*, *Energy Economics* **138** (2024), 107819.
- [62] Abdullahi Adinoyi Ibrahim, Daniela Leite, and Caterina De Bacco, *Sustainable optimal transport in multilayer networks*, *Phys. Rev. E* **105** (2022), no. 6, Paper No. 064302, 8. MR 4455341
- [63] Justin Johnson Kakeu, *Valuing exhaustible resource ownership: General equilibrium assets-markets versus partial equilibrium*, *Theoretical Economics Letters* **8** (2018), no. 5, 844–853.
- [64] Leonid Kantorovich, *On the translocation of masses*, *C. R. (Doklady) Acad. Sci. URSS (N.S.)* **37** (1942), 199–201. MR 9619
- [65] Alexander V. Kolesnikov, Fedor Sandomirskiy, Aleh Tsyvinski, and Alexander P. Zimin, *Beckmann’s approach to multi-item multi-bidder auctions*, Preprint at arXiv: 2203.06837, 2022.
- [66] Tjalling C. Koopmans, *Optimum utilization of the transportation system*, *Econometrica* **17** (1949), 136–146.
- [67] Oscar Lange, *The foundations of welfare economics*, *Econometrica* **10** (1942), no. 3/4, 215–228.
- [68] Robert J. McCann, *Academic wages, singularities, phase transitions and pyramid schemes*, *Proceedings of the International Congress of Mathematicians—Seoul 2014*. Vol. III, Kyung Moon Sa, Seoul, 2014, pp. 835–848. MR 3729054
- [69] Robert J. McCann and Nestor Guillen, *Five lectures on optimal transportation: geometry, regularity and applications*, *Analysis and geometry of metric measure spaces*, CRM Proc. Lecture Notes, vol. 56, Amer. Math. Soc., Providence, RI, 2013, pp. 145–180. MR 3060502
- [70] Robert J. McCann, Cale Rankin, and Kelvin Shuangjian Zhang, *$C^{1,1}$ regularity for principal-agent problems*, *Adv. Math.* **478** (2025), Paper No. 110396, 22. MR 4924768

- [71] Robert J. McCann and Maxim Trokhimtchouk, *Optimal partition of a large labor force into working pairs*, *Econom. Theory* **42** (2010), no. 2, 375–395. MR 2564441
- [72] Robert J. McCann and Kelvin Shuangjian Zhang, *On concavity of the monopolist's problem facing consumers with nonlinear price preferences*, *Comm. Pure Appl. Math.* **72** (2019), no. 7, 1386–1423. MR 3957395
- [73] Robert E. Megginson, *An introduction to Banach space theory*, Graduate Texts in Mathematics, vol. 183, Springer-Verlag, New York, 1998. MR 1650235
- [74] Gaspard Monge, *Mémoire sur la théorie des déblais et des remblais*, *Histoire de l'Académie Royale des Sciences de Paris, avec les Mémoires de Mathématique et de Physique pour la même année (1781)*, 666–704.
- [75] Frances C. Moore and Delavane B. Diaz, *Temperature impacts on economic growth warrant stringent mitigation policy*, *Nature Climate Change* **5** (2015), 127–131.
- [76] Olga Mula and Anthony Nouy, *Moment-SoS methods for optimal transport problems*, *Numer. Math.* **156** (2024), no. 4, 1541–1578. MR 4777943
- [77] Ishan B. Nath, Valerie A. Ramey, and Peter J. Klenow, *How much will global warming cool global growth?*, Working Paper 32761, National Bureau of Economic Research, July 2024.
- [78] J. A. Nelder and R. Mead, *A simplex method for function minimization*, *Comput. J.* **7** (1965), no. 4, 308–313. MR 3363409
- [79] Richard G. Newell and Robert S. Stavins, *Cost heterogeneity and the potential savings from market-based policies*, *Journal of Regulatory Economics* **23** (2003), 43–59.
- [80] Albert L. Nichols, *Targeting economic incentives for environmental protection*, MIT Press, Cambridge, 1984.
- [81] William Nordhaus, *Rolling the 'dice': an optimal transition path for controlling greenhouse gases*, *Resource and Energy Economics* **15** (1993), no. 1, 27–50.
- [82] ———, *Climate change: The ultimate challenge for economics*, *The American Economic Review* **109** (2019), no. 6, 1991–2014.
- [83] William D. Nordhaus and Zili Yang, *A regional dynamic general-equilibrium model of alternative climate-change strategies*, *The American Economic Review* **86** (1996), no. 4, 741–765.

- [84] Intergovernmental Panel on Climate Change, *Climate change 2023: Synthesis report. contribution of working groups i, ii and iii to the sixth assessment report of the intergovernmental panel on climate change*, Tech. report, IPCC, 2023.
- [85] Dmitry Panchenko, *Lecture notes on probability theory*, Self-Published, Amazon, 2019.
- [86] Vilfredo Pareto, *Manuel d'économie politique*, vol. 38, Giard & Brière, 1909.
- [87] Brendan Pass, *Multi-marginal optimal transport: theory and applications*, ESAIM Math. Model. Numer. Anal. **49** (2015), no. 6, 1771–1790. MR 3423275
- [88] Charles Perrings and Ann Kinzig, *Conservation: economics, science, and policy*, Oxford University Press, Oxford, 2021.
- [89] Paul R. Portney, *Market-based approaches to environmental policy: A "refresher" course*, Resources Magazine (2020), 44–47.
- [90] Joseph Pryor, Paolo Agnolucci, Mariza Montes de Oca Leon, Carolyn Fischer, and Dirk Heine, *Carbon pricing around the world*, ch. 5, International Monetary Fund, 2023.
- [91] Jean-Charles Rochet and Philippe Choné, *Ironing, sweeping, and multidimensional screening*, Econometrica **66** (1998), no. 4, 783–826.
- [92] R. Tyrrell Rockafellar, *Convex analysis*, Princeton Landmarks in Mathematics, Princeton University Press, Princeton, NJ, 1997, Reprint of the 1970 original, Princeton Paperbacks. MR 1451876
- [93] R. Tyrrell Rockafellar and Roger J.-B. Wets, *Variational analysis*, Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 317, Springer-Verlag, Berlin, 1998. MR 1491362
- [94] Walter Rudin, *Functional analysis*, second ed., International Series in Pure and Applied Mathematics, McGraw-Hill, Inc., New York, 1991. MR 1157815
- [95] Agnar Sandmo, *Optimal taxation in the presence of externalities*, The Swedish Journal of Economics **77** (1975), no. 1, 86–98.
- [96] Filippo Santambrogio, *Optimal transport for applied mathematicians*, Progress in Nonlinear Differential Equations and their Applications, vol. 87, Birkhäuser/Springer, Cham, 2015, Calculus of variations, PDEs, and modeling. MR 3409718

- [97] Ronald J. Shadbegian and Wayne B. Gray, *Pollution abatement expenditures and plant-level productivity: A production function approach*, *Ecological Economics* **54** (2005), no. 2, 196–208, Technological Change and the Environment.
- [98] L. S. Shapley and M. Shubik, *The assignment game. I. The core*, *Internat. J. Game Theory* **1** (1972), no. 2, 111–130. MR 311290
- [99] J.M.C. Santos Silva and Silvana Tenreyro, *The log of gravity*, *The Review of Economics and Statistics* **88** (2006), no. 4, 641–658.
- [100] Robert S. Stavins and Bradley Whitehead, *Market-based environmental policies*, ch. 7, Yale University Press, 1997.
- [101] Nicholas Stern, *The economics of climate change: the Stern review*, Cambridge University Press, Cambridge, UK, 2007.
- [102] Meng Sun, Yingdong Xu, Xiao Yu, and Yanzhe Zhang, *The green-innovation-inducing effect of a unit progressive carbon tax*, *Sustainability* **13** (2021), no. 21, 11708.
- [103] Cédric Villani, *Topics in optimal transportation*, Graduate Studies in Mathematics, vol. 58, American Mathematical Society, Providence, RI, 2003. MR 1964483
- [104] ———, *Optimal transport*, Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 338, Springer-Verlag, Berlin, 2009, Old and new. MR 2459454
- [105] Leon Walras, *Elements of pure economics: Or the theory of social wealth*, Routledge, [1874] 2013.
- [106] Martin L. Weitzman, *On a world climate assembly and the social cost of carbon*, *Economica* **84** (2017), no. 336, 559–586.
- [107] Robertson C. Williams III, *Setting the initial time-profile of climate policy: The economics of environmental policy phase-ins*, *The Design and Implementation of US Climate Policy*, University of Chicago Press, 09 2012.
- [108] A.P Xepapadeas, *Environmental policy, adjustment costs, and behavior of the firm*, *Journal of Environmental Economics and Management* **23** (1992), no. 3, 258–275.
- [109] Yun-Ting Yeh, *Adjustment frictions and the cost of environmental regulatory uncertainty*, Job market paper, 2026.

- [110] Danila A. Zaev, *On the Monge-Kantorovich problem with additional linear constraints*, *Mat. Zametki* **98** (2015), no. 5, 664–683. MR 3438523

COLOPHON

This thesis was typeset using the typographical look-and-feel `classicthesis` developed by André Miede and Ivo Pletikosić.

The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*".

Figures were created using the `tikz` package and Desmos graphing software.