

I am an applied economist working at the intersection of industrial organization, trade, and environmental economics. My research combines theory with empirics, focusing on market environments where equilibrium mechanisms interact with the efficiency and distributional impacts of public policy. I have taken this approach to study environmental policy in markets subject to pre-existing distortions, zoning and housing regulation in urban markets, and industrial policy in natural resource sectors. A summary of my research in each of these areas is presented below.

I. ENVIRONMENTAL POLICY IN MARKETS SUBJECT TO PRE-EXISTING DISTORTIONS

In markets where the only imperfection is an environmental externality, a Pigouvian tax equal to marginal damages attains the first-best allocation. However, environmental externalities typically co-exist with other distortions in real world markets. In the presence of imperfect competition, a Pigouvian tax needs to be trimmed down because market power is already depressing quantities part of the way towards the first-best (Buchanan, 1969). In the presence of imperfect information between regulators and polluters, price-based incentives may be suboptimal altogether and dominated by quantity-based mandates (Weitzman, 1974).

The first chapter of my PhD dissertation, “Efficiency and Redistribution in Environmental Policy: An Equilibrium Analysis of Agricultural Supply Chains” (accepted at the *Journal of Political Economy*), tackles the question of how to design environmental policy in an industry where Pigouvian taxation is not feasible at the externality’s source, and where pre-existing distortions show up in the form of market power. The empirical context is the South American agricultural sector, a global commodity powerhouse with a major contribution to greenhouse gas emissions, primarily due to deforestation. This setting has two key features that shape the feasibility of environmental policy as well as its abatement performance.

First, agricultural supply chains are funnel-shaped. The upstream stage consists of millions of atomistic farmers facing a concentrated sector of downstream agribusiness firms. Second, carbon footprints of agricultural commodities are mostly driven by the deforestation and production decisions of farmers. Hence, the externality’s source is upstream and varies geographically with the carbon intensity of land. These features suggest that while regulatory compliance is more feasibly ensured at the funnel’s bottleneck for a few large firms, downstream interventions can be poorly targeted because the incentives they ultimately need to correct are those of upstream farmers.

To evaluate the trade-off between policy feasibility and targeting, I develop an equilibrium model of agricultural commodity trade where upstream atomistic farmers face downstream monopolistic firms, thus giving the supply chain its funnel-shape. I estimate the model using granular South American data, where the key empirical object is the supply elasticity of farmers and its heterogeneity across space. The main empirical finding is that supply is less elastic for farmers in

geographically remote areas, resulting in wider markdowns on the farm-gate prices they receive. These remote areas also happen to be where the carbon density of land is highest, resulting in a positive spatial correlation between market power and the externality's intensity.

Motivated by recent policy discussions on carbon border adjustments, I use the model to evaluate how a carbon tax on downstream agribusiness firms is transmitted to upstream farmers. Market power results in incomplete pass-through of the tax to farmers, and least complete of all to regions where markdowns are widest, which is also where carbon-intensities are highest. As a result, the Pigouvian signal, as measured by the pass-through rate, is eroded most in locations where the social cost is highest. Hence, the pre-existing distortion—imperfect competition—worsens targeting for the policy tool that aims to correct the primary distortion, the environmental externality.

To the extent that market power unambiguously depresses output in a laissez-faire equilibrium, there is a classic intuition that “the monopolist is the conservationist's friend” (Solow, 1974). Hence, my results introduce nuance to this view by showing that the presence of market power is more ambiguous for the *transmission* of environmental policy. Moreover, the sufficient statistic that resolves this ambiguity is the correlation between the two distortions—the degree of market power and the environmental externality's intensity. In my empirical context market power worsens targeting because the correlation is positive. In other settings the correlation could be negative, in which case market power improves targeting by differentially raising Pigouvian pass-through where social cost is highest. Understanding how specific industries vary in the primitives shaping such correlations can help reveal general insights about the efficiency of corrective policies.

II. SPATIAL INEQUALITY IN URBAN MARKETS

An important driver of socioeconomic inequality is the extent to which different types of households sort across different locations. Across cities, differences in labor market characteristics, such as sectoral composition and skill-bias, determine inequality in income. Within cities, differences in neighborhood characteristics, such as local public goods and consumption amenities, determine inequality in access to non-tradable consumption. Moreover, local labor and consumption markets are endogenously shaped by the residential choices of households, resulting in a feedback loop between spatial sorting and local amenities (Diamond, 2016).

The second chapter of my PhD dissertation, “*Location Sorting and Endogenous Amenities: Evidence from Amsterdam*” (forthcoming at *Econometrica*), co-authored with Milena Almagro, asks how residential sorting across neighborhoods is reinforced by the endogenous response of consumption amenities to local socioeconomic composition. To answer this question, we estimate a dynamic spatial equilibrium model of residential choice for the city of Amsterdam. In the model, consumption amenities in each location are the equilibrium outcome of a market for non-tradables, where supply caters to the heterogeneous tastes of residents. Hence, neighborhoods are akin to horizontally differentiated products whose multi-dimensional characteristics—the amenities—are endogenously shaped by the agents choosing to live in them.

To estimate our structural model, we use detailed micro-data from the Netherlands, which allows us to track the universe of Amsterdam’s residents over time and the evolution of a rich set of neighborhood amenities. Moreover, we construct shifters of neighborhood characteristics by leveraging spatial variation in the entry of short-term rental platforms such as Airbnb. Our empirical results indicate significant heterogeneity across households in their preferences for different types of amenities, as well as in the supply response of amenities to demographic composition.

Consistent with previous work in local public finance, we show that allowing for endogenous amenities reinforces spatial sorting. However, the implications for welfare inequality are ambiguous when amenities are multi-dimensional and rely on the patterns of preference heterogeneity. Namely, endogenous amenities may increase inequality between demographic groups if their preferences are closely aligned, but may decrease it if substantially misaligned. The reason is that if preferences are misaligned, high and low income groups sort into different locations. As they sort, amenities respond and neighborhoods become even more horizontally differentiated, so the two groups avoid competing for housing in the same location. Hence, low income groups obtain their preferred amenities without having the high income groups bid up their rents. To summarize, when preferences are heterogeneous and endogenous amenities are multi-dimensional, there are two mechanisms reducing the welfare gap across groups: tailored amenities due to neighborhood differentiation and lower rental prices due to housing market segmentation.

We conclude by using our framework to evaluate how the entry of short-term rentals for tourists affects the welfare of local residents in Amsterdam. We decompose welfare impacts into i) the direct effect on rent due to the reduction in housing supply available to locals, and ii) the indirect effects from the endogenous change in amenities induced by the rise in tourist presence. The key insight behind our results is that while all local residents lose from higher rents, their losses may be compensated or amplified depending on how they value the amenities that tourists bring along. Hence, the correlation between income and alignment with tourist preferences determines whether the entry of short-term rentals is regressive or progressive. Specifically, if the lowest-income groups value the amenities that tourists bring, welfare impacts can be progressive. These results underscore the importance of taking into account heterogeneity in demand and supply primitives for distributional analysis, not only quantitatively but also qualitatively.

III. INDUSTRIAL POLICY AND NATURAL RESOURCE EXTRACTION

The state often plays an important role in natural resource extraction. Natural resources can be an important source of rents, especially for developing countries where alternative streams of fiscal revenue are limited by informality and small tax bases. Moreover, governments often use natural resources as a political tool, both to redistribute rents internally across domestic actors, but also externally as geopolitical leverage over other sovereign states.

The third chapter of my PhD dissertation, “[The Political Economy of Migration Restrictions under Apartheid](#)”, co-authored with Leonard Le Roux, asks how natural resource sectors shape

domestic institutions in the context of Apartheid-era South Africa. South African economic growth during the late 19th century was driven by natural resource extraction, most notably through gold exports. As the first globalization came to an abrupt end with World War I and commodity markets collapsed, many commodity-export dependent countries responded by diversifying away from primary sectors into manufacturing. Like many other developing countries during the post-war period, South Africa adopted import substitution as part of a broader industrial policy initiative to boost its manufacturing sector. While favoring urban manufacturing, these policies came at the expense of rural export-oriented sectors—mining and agriculture—by raising the cost of intermediate inputs that either had to be imported at higher tariff rates or purchased domestically from inefficient local manufacturers.

Distributional tensions across sectors are common to industrial policy experiments everywhere, almost by definition. We argue that the South African context is unique in how these sectoral tensions were resolved. Starting in 1952, the Apartheid government began to restrict internal migration, with much of the black population forced to live in rural homelands while cities were reserved mostly for whites. A primary objective was to divert captive labor to the mining and agricultural sectors. Thus, migration restrictions on the disenfranchised partially compensated the losers from import substitution, white farm owners and miners. While mining was an important source of fiscal revenue, agricultural districts mattered because they accounted for most of the regime’s political support. Hence, this setting shows how distortions to the domestic labor market can complement distortions to the international terms of trade for the purpose of political gain. Migration policy and industrial policy can be complementary, especially for non-democratic governments that rely on the support of a few key industries to remain in power, and where restricting internal migration for a subset of the population is available as a policy instrument.

For the first stage of this project, we have assembled internal migration micro-data from historic South African population censuses. We use this data to show dramatic changes in migration rates before and after restrictions were removed, and primarily to estimate migration elasticities. For the second stage, we have complemented our migration dataset with data from the labor demand side, which we have constructed from mine-level archival records and previously un-digitized agricultural censuses. Finally, for the third stage, which is currently in progress, we are building a multi-sector migration model to study how migration restrictions are endogenously chosen by the regime to complement its broader industrial policy goals.

One research direction that branches off from the previous project concerns the political economy of natural resource extraction. As already mentioned, natural resources can shape domestic institutions, but may also have international consequences to the extent they are used by governments for geopolitical objectives. In ongoing work with Jonathan Elliott and Allan Hsiao, “*Geopolitics, Critical Minerals, and the Energy Transition*”, we zoom in on the upstream stages of the global electric vehicle (EV) supply chain i.e., the extraction of critical mineral commodities. While existing work on the EV industry has focused on downstream stage issues, such as network

effects due to complementarities between vehicles and charging stations, or learning-by-doing in the manufacturing stage, our focus is on the geopolitical aspects of the upstream commodity stage.

References

Buchanan, J.M. (1969). External diseconomies, corrective taxes, and market structure. *American Economic Review* 59(1):174–177.

Diamond, R. (2016). The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000. *American Economic Review* 106(3):479–524.

Solow, R.M. (1974). The Economics of Resources or the Resources of Economics. *American Economic Review* 64(2):1–14.

Weitzman, M.L. (1974). Prices vs. quantities. *The Review of Economic Studies* 41(4):477–491.