

Search Frictions and Efficiency in Decentralized Transport Markets

Giulia Brancaccio*, Myrto Kalouptsi†, Theodore Papageorgiou‡, Nicola Rosaia§¶

April 2022

Abstract

We explore efficiency and optimal policy in decentralized transport markets, such as taxis, trucks, and bulk shipping. We show that in these markets, search frictions distort the transportation network and the dynamic allocation of carriers over space. We derive explicit and intuitive conditions for efficiency, and show how they translate into efficient pricing rules, or optimal taxes and subsidies for the planner who cannot set prices directly. The results imply that destination-based pricing is essential to attain efficiency. Then, using data from dry bulk shipping, we demonstrate that search frictions lead to a sizeable social loss and substantial misallocation of ships over space. Optimal policy can eliminate about half of the welfare loss. Can a centralizing platform, often arising as a market-based solution to search frictions, do better? Interestingly, the answer is no; although the platform eradicates frictions, it exerts market power thus eroding the welfare gains. Finally, we use two recent interventions in the industry (China's Belt and Road Initiative and the environmental initiative IMO 2020) to demonstrate that taking into account the efficiency properties of transport markets is germane for any proposed policy.

*New York University, Stern School of Business, Kaufman Management Center 44 West 4th Street New York NY 10012, USA, giulia.brancaccio@nyu.edu

†Harvard University, Department of Economics, Littauer Center 124, Cambridge MA 02138, USA, myrto@g.harvard.edu

‡Boston College, Department of Economics, Maloney Hall 343, Chestnut Hill, MA, 02467, USA, theodore.papageorgiou@bc.edu

§Princeton University, Department of Economics, Julis Romo Rabinowitz Building, Princeton, NJ 08544, USA, nr1485@princeton.edu

¶We are thankful to Costas Arkolakis, Nick Buchholz, Pablo Fajgelbaum, Manolis Galenianos, Allen Head, Réka Juhász, Marc Melitz, Robin Lee, Ariel Pakes, and Tobias Salz. We gratefully acknowledge the support of the National Science Foundation (under grant SES-1847555), the Social Sciences and Humanities Research Council (under grant #430-2018-00540), as well as the Griswold Center for Economic Policy Studies at Princeton University.

1 Introduction

The transportation sector is vital for economic growth and social development. Generating about 1.2 trillion USD annually and employing 12 million workers in the US alone, the sector is responsible for the movement of both people and goods, making today’s interconnected world possible.¹ A substantial fraction of transport markets are decentralized; lacking a central market place, they feature a fragmented market structure, with a large number of agents meeting and trading in ad hoc fashion. This is the case for taxis, trucking, oceanic bulk shipping (dry bulk carriers and oil/chemical tankers), as well as bike/motorcycle taxis and autorickshaws.^{2,3} These markets can suffer from search frictions, i.e. impediments to trade arising from the meeting process. Are transport agents thus suboptimally allocated over space, distorting transportation flows? And if so, which policies are best at restoring the optimality of the transport network? Can a centralizing platform perform better? In recent years, these questions have captured the interest of industry participants and policymakers alike.

In this paper, we provide an in-depth theoretical and empirical analysis of efficiency in decentralized transport markets. We identify the externalities that stem from search frictions and distort transportation flows. We characterize analytically the conditions for (constrained) efficiency and derive both the set of efficient pricing rules, as well as the optimal taxes and subsidies. Next, using data from the dry bulk

¹U.S. Department of Transportation, Bureau of Transportation Statistics, Transportation Economic Trends and Occupation Employment Statistics of the Bureau of Labor Statistics.

²The transportation sector includes several different segments which can be split into two categories: those that operate on fixed itineraries, and those that operate on flexible routes, on a point-to-point basis. Containerships, airplanes, buses and trains primarily belong to the first group; while taxis, trucks, gas and oil tankers, and dry bulk ships belong to the second, and are the focus of this paper.

³The global taxi industry generated about \$108 billion in revenue in 2016 and \$25 billion in the US alone. In 2019, trucking generated \$800 billion in revenue and moved \$12 billion tons of cargo, representing 80.4% of total freight by value and 72.5% of total freight by weight in the US (Yang, 2022). Finally, seaborne trade accounts for about 70% of trade in terms of value (total of \$18.9 trillion), and 80% in terms of volume (total of 11.08 billion tons). In 2019, bulk shipping (including tankers) accounted for 85% of total seaborne trade in tons and 75% of the total fleet tonnage (it is not straightforward to obtain information on the share of world trade value carried by bulkers; however, mining, agricultural products, fuels, chemicals and iron/steel jointly account for about 52% of total trade value (WTO, 2015). In developing countries, an informal transport sector based on motorcycle and bike taxis (e.g. Africa), autorickshaws (e.g. India), and ojekes (e.g. Indonesia) are often the primary means of transportation in urban centers (Keniston 2011; Dupas et al. 2020).

These markets are decentralized. In many countries, the taxi industry consists of a large number of independent drivers who search for customers on the streets. For instance, in New York City at least 40% of taxis must be operated by individual owners by mandate (while for the remaining size is capped), and dispatchers are not permitted (Frechette et al., 2019). Trucking is also “extremely fragmented”; in the US there are 3.68 million trucks and the top ten carriers only account for around 5% of total revenue. Out of over 892,000 carriers, 97% operate 20 or fewer trucks, and 91% operate less than 6 trucks (source: coyote.com). We discuss shipping in detail in Section 3.1. Finally, bike taxis and autorickshaws are decentralized, and pricing can be either partially regulated by local authorities/local cooperatives, or determined by bilateral bargaining between the driver and the passenger.

shipping industry, we show that search frictions lead to substantial welfare loss and misallocation of ships over space. Recently, centralizing platforms have emerged as a market-based solution to search frictions in many settings, such as taxis. However, we demonstrate that in the case of shipping, such a platform would only eliminate about half of the welfare loss generated by search frictions. Indeed, although platforms eradicate frictions, they exert market power, thus eroding the welfare gains. Importantly, a policymaker who cannot affect the search environment, can attain approximately the same welfare gain as the platform by levying the optimal taxes.

We base our analysis on a model for decentralized transport markets that is widely used in the literature; the model originates in Lagos (2000) and has been used in Lagos (2003), Castillo (2020), Rosaia (2020a) and Buchholz (2021) among others, for taxis, in Yang (2022) for trucking, and in Brancaccio et al. (2020a) for oceanic bulk shipping (henceforth BKP). There is a network of locations at different distances to each other. In every location, available carriers (e.g. ships, taxis) and customers (e.g. exporters, passengers) meet randomly. Carriers that find a customer transport them to their desired destination for a price, and restart there. Carriers that do not find a customer decide whether to wait at their current location or travel empty elsewhere to search. Customers that find a carrier, obtain a valuation from arriving at their destination, while customers that do not, wait another period. Finally, every period, potential customers decide whether to start searching for a carrier, as well as their destination, thus replenishing the customer pool seeking transportation.

Studying efficiency in this setup is not straightforward due to the dynamic nature of decision-making and the spatial network; yet, we are able to characterize efficiency analytically.

First, as is well-known in other settings, search frictions generate “thick market and congestion” externalities: when choosing whether to search, carriers take into account the match surplus they create for themselves, but do not internalize the surplus they create for customers (thickness), nor the negative congestion they generate by making it harder for other carriers to find a match. The same effects hold for customers.

Thick market and congestion externalities are internalized in equilibrium if and only if the positive effect of match creation exactly offsets the negative congestion effect. This amounts to the so-called “Hosios (1990) conditions”: these conditions, which are known to characterize efficiency in search models

with homogeneous agents, require that in each location the surplus share of each side is equal to the corresponding elasticity of the matching function.

Second, search frictions generate “composition externalities.” These stem from customer heterogeneity: customers decide not only whether to search for transportation, but also their destination. Customers heading to different destinations generate different surplus for the carrier they match with. For instance, a customer headed to a remote destination, where it might be difficult to meet another customer, generates a lower surplus for the carrier compared to one going to a central location, all else equal. If customers fail to internalize this effect, the composition of realized trips will be distorted, and carriers will be misallocated over space, leading to suboptimal transport networks.

We show that customers internalize composition externalities if and only if carriers are “insulated” against composition effects, i.e. carriers must receive the same surplus regardless of the customer they match with. This condition for efficiency replicates the no-arbitrage condition obtained in a frictionless world, where competition among carriers ensures that prices coincide with the opportunity cost of each trip, until in equilibrium carriers are indifferent among different customers.

The two efficiency conditions combined characterize the efficient pricing rule, which can be employed by a central price-setting authority, as in the case of taxicabs where prices are regulated. In such environments, our results clarify that the role of optimal destination-specific pricing is to correct composition externalities. Instead, when prices are bilaterally negotiated, composition externalities justify the use of destination-specific taxes and subsidies, as negotiated prices typically fail to fully internalize destination effects.

We exploit the theoretical results to study the impact of search frictions in the dry bulk shipping industry, where a large number of small homogeneous ship(owner)s, often termed “ocean taxis”, meet exporters on a trip by trip basis, through a disperse network of brokers, to arrange the international transportation of raw materials. A number of features of dry bulk shipping, such as information frictions and port infrastructure, can hinder the matching of ships and exporters. Our analysis leverages a rich data set of detailed vessel movements tracking ship exact locations in six-minute intervals, as well as shipping prices.

We begin by documenting the presence of search frictions. To do so, we propose a novel test to argue

that these frictions lead to unrealized potential trade. The test is based on weather shocks at sea that exogenously shift ship arrivals at port: in a frictionless world, in regions with more ships than exporters, the change in the number of ships should not affect matching, since the short side of the market is always matched. Here, instead, matches are indeed affected by weather shocks. We further corroborate the presence of frictions by documenting substantial price dispersion and wasteful ship movements.

We then estimate the model. The unknown parameters of interest include (i) the matching functions; (ii) ship costs; (iii) exporter costs; and (iv) bargaining coefficients. We obtain the first two sets of parameters following BKP, and we bring in additional trade data to obtain the rest.

We first test whether the observed equilibrium is efficient by checking whether our conditions for efficiency hold in the data. Perhaps not surprisingly, we find that neither condition is satisfied, suggesting that the market does not operate efficiently.

What are the welfare losses from search frictions and the gains from optimal policy? Can a platform do better? To answer these questions, we begin by comparing the observed equilibrium to the frictionless equilibrium (i.e. the competitive benchmark, or first-best). We find that frictions have a substantial impact on trade and welfare: the first-best welfare is 18% higher compared to the baseline equilibrium, with substantially higher trade volume and value.

Recently, platforms have emerged as a market-based solution to search frictions. By centralizing transactions, they facilitate meetings and can improve market efficiency. Ideally, centralization would coincide with the frictionless equilibrium. However, a centralizing platform does not act like a benevolent planner: it chooses prices to maximize profits, thus substituting one friction (search frictions) for another (market power).

To assess the trade-off between search frictions and market power, we consider the optimal pricing problem of a monopolist platform that owns all ships. This is a complex problem, as it entails choosing prices for all possible origin-destination pairs in order to maximize the monopolist profits, taking into account the dynamic paths of ships. To solve this optimization problem, we build an algorithm relying on Rosaia (2020a). As expected, the increase in welfare is substantially lower than what is achieved in the frictionless equilibrium (around 9%). Importantly, the welfare gains are appropriated to a large extent by the centralizing monopolist in the form of profits.

Strikingly, we find that a policymaker who is not able to directly affect the meeting process, can achieve the same increase in welfare, by levying the optimal taxes and subsidies. In contrast to the centralizing monopolist, now the welfare gains are redistributed to the market agents.

In our final empirical exercise we use two examples to demonstrate that taking into account the efficiency properties of decentralized transport markets is germane for any proposed intervention. First, we consider China’s Belt and Road Initiative (BRI), an initiative that, among other things, provided substantial funding to port infrastructure in Asia, Europe and West Africa. We demonstrate that the BRI, by reducing port costs, acted as a subsidy that encouraged exporter entry, thus indirectly correcting for the thick market and congestion externalities. Indeed, at the observed equilibrium, the entry of an additional exporter has a substantial positive externality on matching rates, but prices remain too high thus leading to too few exporters. As the BRI partially corrects for this, it generates an additional increase in welfare, beyond the direct effect of reducing port costs.

Next, we consider an impactful environmental policy by the International Maritime Organization (“IMO 2020”) that enforced new emissions standards designed to curb ship pollution. Compliance, however, has been lacking in West African ports. We explore how conformity of these ports, currently addressed through international cooperation, would interact with search externalities. A stronger implementation of IMO 2020 by West African ports would increase the ship cost of traveling towards these ports. This, in turn, would act as a tax towards these ports, which offer low continuation value to ships. As such, this tax indirectly alleviates the composition externalities.

Related Literature

This paper broadly relates to several strands of literature: transportation and trade; search and matching; and industry dynamics.

First, our paper contributes to a large, diverse, and rapidly growing literature on transportation. The study of optimal transport networks dates back to Koopmans (1949), while Fajgelbaum and Schaal (2020) and Allen and Arkolakis (2021) have recently revisited this question within a general equilibrium spatial trade framework. Different from our focus, they consider optimal routing and optimal investment in infrastructure.

Within the context of urban transportation, a series of recent empirical papers have studied issues related to efficiency. Frechette et al. (2019) and Buchholz (2021) both study search frictions and regulation frictions in NYC taxicabs. In particular, Buchholz (2021) numerically implements tariff pricing changes in order to explore whether welfare improvements can be achieved. Frechette et al. (2019) investigate the welfare impact of changes in the number of active medallions, as well as the introduction of an “Uber-like” platform (modeled as dispatcher who eradicates frictions taking the regulated prices as given). Shapiro (2018) and Liu et al. (2019) explore the welfare improvements from different centralizing formats; while several papers study issues related to platform pricing (e.g. Bian (2019), Castillo (2020), Ma et al. (2022), as well as Rosaia (2020a) on whose methodology we rely heavily to perform the platform analysis).⁴

We contribute to this literature by fully characterizing the sources of inefficiencies and the optimal policy in the transport setup. The theoretical results convey important policy messages; e.g. efficient prices are a function of destinations, not distance. Moreover, they provide the necessary machinery that allows us to compute the efficient spatial allocation, unlike the previous literature, and therefore quantify the loss from frictions in our empirical application. Finally, this machinery enables us to consider the trade-off between search frictions and market power, which is a first-order issue in understanding the impact of platforms, and yet remains largely unaddressed by the literature.

Moreover, since our empirical application involves international oceanic transportation, we relate to a literature studying transportation in the context of international trade; e.g. Hummels and Skiba (2004), Cosar and Demir (2018), Holmes and Singer (2018), Asturias (2020), Brooks et al. (2021), Lee et al. (2020), Wong (2020) and Ducruet et al. (2022).⁵ In our prior work, BKP, we explore the role of the transportation sector in world trade and spell out the impact of endogenous trade costs. Although we rely on the model setup and empirical strategy employed there, our focus here is different, as this paper provides a formal treatment of efficiency in decentralized transport markets and an empirical analysis of search frictions, welfare loss and the impact of centralization in shipping as a case study.

Second, our paper relates to the literature on efficiency of search models in the spirit of Diamond

⁴In addition, Ghili and Kumar (2021) investigate demand and supply imbalances in ride-sharing platforms; Ostrovsky and Schwarz (2018) focus on carpooling and self-driving cars; Kreindler (2020) studies optimal congestion pricing; Cao et al. (2021) explore competition in bike-sharing platforms.

⁵We also relate to a literature in international trade studying the role of frictions, such as Antras and Costinot (2011), Krolkowski and McCallum (2021), and Eaton et al. (2022) who consider search frictions between importers and exporters and Allen (2014) who investigates information frictions.

(1982), Mortensen (1982) and Pissarides (1985). Hosios (1990) considers efficiency in markets with random search and Nash bargaining. He shows that these markets are generically inefficient and derives the well-known “Hosios condition” that restores efficiency. As shown by Acemoglu and Shimer (1999) and Acemoglu (2001), when agents are heterogeneous, the Hosios condition does not guarantee efficiency because of composition externalities, and labor market policies, such as unemployment benefits or minimum wages, can potentially improve welfare.⁶ In our context, heterogeneity is inherent in the spatial nature of transportation, as it arises endogenously in carriers’ continuation values due to the dynamics of the transport problem; that is, even if customers’ costs and valuations were homogeneous at every region, matching surpluses would still differ across destinations due to differences in carriers’ continuation values. Our paper extends the existing theoretical literature by: showing formally that, together with the thick market and congestion externalities, these are the only sources of externalities in transportation; deriving a novel efficiency condition for composition externalities in models with random search and match heterogeneity; and providing closed-form expressions for the efficient prices (and taxes/subsidies), which, in our context, illustrate the importance of destination pricing.⁷

Finally, we relate to the literature on industry dynamics (Hopenhayn, 1992, Ericson and Pakes, 1995). Related to this paper, a small literature lying in the intersection of search and industry dynamics, has explored trading frictions in decentralized markets (e.g. Gavazza (2011) and (2016) for real assets and Brancaccio et al. (2020c) for over-the-counter financial markets).

The rest of the paper is structured as follows: Section 2 presents the model and provides the efficiency and optimal policy results. Section 3 describes the dry bulk shipping industry and the data used, and presents evidence for search frictions. Section 4 outlines the estimation of the model. Section 5 presents our welfare analysis. Section 6 concludes.

⁶See also Davis (2001), Shimer and Smith (2001), Decreuse (2008), Atkeson et al. (2015) and Bilal (2022).

⁷Extending the literature, we also show how to restore efficiency in the case of decreasing returns to scale in the matching technology. In addition, compared to Hosios (1990), we consider a more general setup with a network of interconnected markets and without the Nash bargaining assumption on pricing.

2 Model

We present a model of decentralized transport markets that focuses on the interaction between carriers (e.g. ships, taxis, trucks) and customers (e.g. exporters, passengers). As mentioned, this framework has been used fairly widely in the literature; here we follow the presentation of BKP.

2.1 Environment

Time is discrete and the horizon is infinite. There are I locations, $i \in \{1, 2, \dots, I\}$, and two types of agents: customers and carriers. Both are risk neutral and have discount factor β . Variables with superscript s refer to carriers and e to customers, in line with our empirical exercise of ships and exporters.

At the beginning of period t , a carrier is either searching in some region i or traveling, full or empty, from some region i to some region j . Carriers searching at i incur a per-period cost c_i^s . Carriers traveling from i to j incur a per-period traveling cost c_{ij}^s . The duration of a trip between location i and location j is stochastic: a traveling carrier arrives at j in the current period with probability d_{ij} . We assume that $d_{ij} < 1$ if $i \neq j$, so that the average duration of the trip is $1/d_{ij} > 1$.⁸

Customers can only be delivered to their destination by carriers and each carrier can carry at most one customer. Following the search and matching literature, we model the number of matches that take place at period t in region i , $m_i(t)$, using a matching function

$$m_i(t) = m_i(s_i(t), e_i(t)) < \min \{s_i(t), e_i(t)\}$$

where $s_i(t)$ and $e_i(t)$ denote the measure of carriers and customers searching in region i , respectively. $m_i(\cdot)$ is increasing, concave and differentiable in both arguments. Note that this captures the potential for unrealized trade: two agents searching in the same location might fail to meet, due to impediments such as information frictions or physical constraints. As Petrongolo and Pissarides (2001) note, “[...] the matching function [...] enables the modeling of frictions [...] with a minimum of added complexity. Frictions derive from information imperfections about potential trading partners, [...] slow mobility, congestion from large numbers, and other similar factors.”

⁸It is straightforward to have deterministic trip duration instead. Our specification, however, preserves tractability and allows for some variability e.g. due to weather/traffic shocks.

Since search is random, the probability according to which customers in i meet a carrier is $\lambda_i^e(t) = m_i(s_i(t), e_i(t))/e_i(t)$, which is the same for all customers. Similarly, the probability according to which carriers meet a customer is $\lambda_i^s(t) = m_i(s_i(t), e_i(t))/s_i(t)$.

When a carrier and a customer meet, if they both accept to match, the customer pays a price $p_{ij}(t)$ upfront and the carrier begins its trip immediately to j . This price is determined differently in different transport markets: for instance, prices are fixed by regulation in taxicabs, while prices are bilaterally negotiated in bulk shipping; our results nest these different cases.

Carriers that remain unmatched decide whether to stay in their current region or travel empty to a different region and search there. Customers that remain unmatched wait in their current region.

Finally, every period, at each location i , N_i new potential customers decide whether to enter and search for a carrier, subject to an entry cost drawn from $(\kappa_{i1}, \dots, \kappa_{i,J}) \sim F_\kappa$, an absolutely continuous distribution with full support.⁹ We denote by $n_{ij}(t)$ the measure of new entrants, and by $e_{ij}(t)$ the measure of customers in i searching for transportation to j at time t . The total measure of customers searching at i is $e_i(t) = \sum_{j \neq i} e_{ij}(t)$, while $G_{ij}(t)$ denotes the share of demand routed from i to j , i.e. $G_{ij}(t) \equiv e_{ij}(t)/e_i(t)$.

Once they have entered, customers pay a per-period waiting cost c_{ij}^e , and obtain a valuation $w_{ij}(t) = w_{ij}(q(t))$ upon being delivered to destination, where $q(t)$ is the matrix, with typical element $q_{ij}(t)$, denoting the quantity transported (i.e. the measure of accepted matches) from i to j at t . Note that we let customers' valuations for transport be a function of transport flows, allowing the model to capture general equilibrium effects.¹⁰ Formally, the only restriction we impose on $w(\cdot)$ is that there exists a concave welfare function W , such that $w_{ij}(q) = dW(q)/dq_{ij}$. For instance, when $w_{ij}(q) \equiv w_{ij}$ is constant, $W(q) = \sum_{ij} w_{ij} q_{ij}$.

⁹Our results continue to hold under more general specifications for the customers' entry model. Indeed, the proofs in Appendix A only require the existence of a mapping specifying the number of new entrants on each route as a function of the vector of customers' valuations of entering on different routes.

¹⁰For instance, in a general equilibrium trade model, $w(\cdot)$ captures the dependence between commodity prices and trade flows; e.g. when there is an increase in the supply of exports, commodity prices fall and this reduces cargo valuations.

2.2 Behavior and Equilibrium

We begin by describing the optimal behavior of carriers and customers taking the endogenous market variables, $\mathbf{p}, \mathbf{w}, \boldsymbol{\lambda}^s, \boldsymbol{\lambda}^e, \mathbf{G}$, as given.¹¹ We then describe the aggregate transitions and define the market equilibrium. Unless otherwise noted, we assume that infinite sequences are bounded.¹²

Carrier optimality Let $V_{ij}^s(t)$ denote the value of a carrier that begins period t traveling from i to j (empty or loaded), $V_i^s(t)$ the value of a carrier that begins the period searching in location i , and $U_i^s(t)$ the value of a carrier that remained unmatched at i at the end of time t .

A carrier traveling from i to j receives,

$$V_{ij}^s(t) = -c_{ij}^s(t) + d_{ij}\beta V_j^s(t+1) + (1 - d_{ij})\beta V_{ij}^s(t+1) \quad (1)$$

In words, a traveling carrier pays the per period cost c_{ij}^s ; with probability d_{ij} it arrives at destination j , where it begins with value $V_j^s(t)$; with the remaining probability $1 - d_{ij}$, the carrier does not arrive and keeps traveling.

A carrier that starts the period in region i obtains:

$$V_i^s(t) = -c_i^s + \lambda_i^s(t)\sum_j G_{ij}(t) \max\{p_{ij}(t) + V_{ij}^s(t), U_i^s(t)\} + [1 - \lambda_i^s(t)]U_i^s(t) \quad (2)$$

In words, the carrier pays the per period wait cost c_i^s ; with probability $\lambda_i^s(t)G_{ij}(t)$ it meets a customer headed to destination j , in which case it accepts the match if and only if the continuation value from doing so, inclusive of the price $p_{ij}(t)$, is higher than the outside option $U_i^s(t)$ of remaining unmatched. With probability $1 - \lambda_i^s(t)$, the carrier does not meet a customer and receives value $U_i^s(t)$.

If the carrier remains unmatched, it chooses where to search: it can either wait at i or travel empty to another location. The unmatched carrier value function is given by

$$U_i^s(t) = E\max_j[V_{ij}^s(t) + \epsilon_j] \quad (3)$$

¹¹Throughout the Section we use bold to denote sequences. E.g. $\mathbf{p} = (p(t))_{t=0}^\infty$.

¹²I.e. sequence $\mathbf{x} = (x(t))_{t=0}^\infty$ satisfies, $\sup_t |x(t)| < \infty$.

where we set $V_{ii}^s(t) \equiv \beta V_i^s(t+1)$, and the shocks $(\epsilon_1, \dots, \epsilon_J) \sim F_\epsilon$ are drawn from an absolutely continuous distribution with full support.¹³ In words, if the carrier stays in region i , at the beginning of the next period it will be waiting at i ; otherwise, if the carrier chooses another region $j \neq i$, it begins its trip towards j .

We now consider the carriers' optimal behavior with respect to the two decisions they make (whether to accept a match, and where to search if unmatched at the end of the period). Consider first the decision of whether or not to accept a match. If $q_{ij}(t) > 0$, then some match is accepted on ij , hence the price-inclusive continuation value must be weakly higher than the outside option. In contrast, if $q_{ij}(t) < m_i(t)G_{ij}(t)$, then some match is rejected, hence the price-inclusive continuation value must be weakly lower than the outside option. Hence, carriers' optimal match acceptance/rejection policy implies that, for all ij ,

$$q_{ij}(t) > 0 \Rightarrow p_{ij}(t) + V_{ij}^s(t) \geq U_i^s(t), \text{ and } q_{ij}(t) < m_i(t)G_{ij}(t) \Rightarrow p_{ij}(t) + V_{ij}^s(t) \leq U_i^s(t) \quad (4)$$

Second, denote by $b_{ij}(t)$ the measure of carriers who choose to relocate empty from i to j (and by $b_{ii}(t)$ the measure that decides to remain in i). The optimal choice probabilities are given by

$$b_{ij}(t)/\sum_k b_{ik}(t) = \Pr[V_{ij}^s(t) + \epsilon_j \geq V_{ik}^s(t) + \epsilon_k, \forall k] \quad (5)$$

For instance, in our empirical application we assume that ϵ is i.i.d. according to a type-I extreme value distribution with scale parameter σ , yielding the logit expression, $b_{ij}(t)/\sum_k b_{ik}(t) = \exp(V_{ij}^s(t)/\sigma)/\sum_k \exp(V_{ik}(t)/\sigma)$.

Customer optimality We now turn to customers' optimal behavior. We begin with existing customers and then consider customer entry.

Let $V_{ij}^e(t)$ denote the value of a customer that begins period t searching at i towards j :

$$V_{ij}^e(t) = -c_{ij}^e + \lambda_i^e(t) \max\{w_{ij}(t) - p_{ij}(t), \beta V_{ij}^e(t+1)\} + [1 - \lambda_i^e(t)]\beta V_{ij}^e(t+1) \quad (6)$$

¹³Absolute continuity and full support are inconsequential assumptions, but they allow us to simplify the exposition by neglecting tie-breaking in agents' optimization. In particular, our results continue to hold in a model without shocks.

In words, the customer pays the cost c_{ij}^e while waiting; then with probability $\lambda_i^e(t)$ it meets a carrier, accepting to match if and only if the net valuation, $w_{ij}(t) - p_{ij}(t)$, is higher than the outside option of staying unmatched. With the remaining probability it remains unmatched and must search another period.

Similarly to carriers, if $q_{ij}(t) > 0$ then some match is accepted on ij , hence the value of accepting the match must be weakly higher than the outside option. On the other hand, if $q_{ij}(t) < m_i(t)G_{ij}(t)$, then some match is rejected, hence the value of accepting the match must be weakly lower than the outside option. Hence we must have

$$q_{ij}(t) > 0 \Rightarrow w_{ij}(t) - p_{ij}(t) \geq \beta V_{ij}^e(t+1), \text{ and } q_{ij}(t) < m_i(t)G_{ij}(t) \Rightarrow w_{ij}(t) - p_{ij}(t) \leq \beta V_{ij}^e(t+1) \quad (7)$$

Finally, each potential entrant chooses whether and where to enter subject to the random entry costs κ_{ij} , so the number of entrants on ij is given by

$$n_{ij}(t) = N_i \Pr[V_{ij}^e(t) - \kappa_{ij} \geq V_{ik}^e(t) - \kappa_{ik}, \forall k] \quad (8)$$

where choosing i from i is interpreted as the decision not to enter and its mean payoff is normalized to zero. For example, in our empirical application we assume a logit entry model, so that each potential entrant solves, $\max\{\epsilon_{ii}, \max_{j \neq i}[V_{ij}^e(t) - k_{ij} + \epsilon_{ij}]\}$, where ϵ is i.i.d. according to a type-I extreme value distribution, and k_{ij} denotes the average entry cost on ij . In this case, the entry curve satisfies the logit expression, $n_{ij}(t) = N_i \exp[V_{ij}^e(t) - k_{ij}] / \{1 + \sum_{k \neq i} \exp[V_{ik}^e(t) - k_{ik}]\}$.

Feasible dynamic allocations A dynamic allocation for the transportation economy consists of a sequence

$$\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b} = (s(t), e(t), G(t), q(t), n(t), b(t))_{t=0}^{\infty}$$

where, $s(t) = [s_i(t) : i \in I]$ denotes the measure of carriers waiting in each region, $e(t) = [e_i(t) : i \in I]$ the measure of customers waiting in each region, $G(t) = [G_{ij}(t) : ij \in I^2]$ their destination shares, $q(t) = [q_{ij}(t) : ij \in I^2]$ the measure of new matches formed on each route, $n(t) = [n_{ij}(t) : ij \in I^2]$ the measure of new entrants on each route, and $b(t) = [b_{ij}(t) : ij \in I^2]$ the measure of carriers departing

empty on each route.

Let $\tilde{s}(t) = [\tilde{s}_{ij}(t) : ij \in I^2]$ denote the measure of carriers traveling from i to j (or waiting at i if $j = i$) at the beginning of time t . First, feasibility requires that for all ij ,¹⁴

$$s_i(t) = \sum_j d_{ji} \tilde{s}_{ji}(t) \quad (9)$$

$$\sum_j b_{ij}(t) = s_i(t) - \sum_j q_{ij}(t) \quad (10)$$

$$q_{ij}(t) \leq m_i(s_i(t), e_i(t)) G_{ij}(t) \quad (11)$$

In words, $\tilde{s}_{ji}(t)$ carriers traveling from j towards i reach their destination with probability d_{ji} and start searching there; out of the $s_i(t) - \sum_j q_{ij}(t)$ that remained unmatched, $\sum_{j \neq i} b_{ij}(t)$ re-locate, and $b_{ii}(t)$ wait at i ; the number of accepted matches cannot exceed the number of meetings.

In addition, at the end of time t , $q_{ij}(t) + b_{ij}(t)$ carriers start traveling from i to j , joining the $\tilde{s}_{ij}(t)(1 - d_{ij})$ traveling carriers which have not reached their destination yet; $e_i(t)G_{ij}(t) - q_{ij}(t)$ unmatched customers must search another period, joined by $n_{ij}(t + 1)$ new entrants at the beginning of $t + 1$. Hence,

$$\tilde{s}_{ij}(t + 1) = \tilde{s}_{ij}(t)(1 - d_{ij}) + q_{ij}(t) + b_{ij}(t), \text{ for all } ij \quad (12)$$

$$e_i(t + 1)G_{ij}(t + 1) = e_i(t)G_{ij}(t) - q_{ij}(t) + n_{ij}(t + 1), \text{ for all } ij \quad (13)$$

Equations (9)-(13) characterize the entire set of dynamic allocations that can arise by aggregating agents' individual behavior, taking as given initial conditions $\tilde{s}_{ij}(0), \tilde{e}_{ij}(0) > 0$ for traveling carriers and existing customers on each route ij at time 0.

Definition 1. A dynamic allocation $\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}$ is feasible if it satisfies constraints (9)-(13). The set of feasible allocations is denoted by \mathcal{A} .¹⁴

Equilibrium In equilibrium, customers and carriers respond optimally to their expectations, and the latter are consistent with the true realizations of the endogenous market variables.

¹⁴Note that we can omit the sequence $\tilde{\mathbf{s}} = (\tilde{\mathbf{s}}(t))_{t \geq 1}$ from the description of dynamic allocations, since $\tilde{\mathbf{s}}$ is pinned down by $\tilde{\mathbf{s}}(0), \mathbf{q}, \mathbf{b}$ from (12): $\tilde{s}_{ij}(t) = \tilde{s}_{ij}(0)(1 - d_{ij})^t + \sum_{\tau=0}^{t-1} [q_{ij}(\tau) + b_{ij}(\tau)](1 - d_{ij})^{t-\tau-1}$, for all $t \geq 1$. Note also that we do not consider carrier entry; the constraint on the fleet size is consistent with most applications of interest, and can be due to either regulatory constraints (e.g. fixed number of medallions) or time to build.

Definition 2. $s, e, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}, \mathbf{p}$ is an equilibrium if: (i) (Feasibility) $s, e, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b} \in \mathcal{A}$; (ii) (Optimality) Conditions (1)-(8) are satisfied under $\mathbf{p}, \mathbf{w}, \mathbf{G}, \boldsymbol{\lambda}^s, \boldsymbol{\lambda}^e$, where $w_{ij}(t) = w_{ij}(q(t))$, $\lambda_i^s(t) = m_i(s_i(t), e_i(t))/s_i(t)$, and $\lambda_i^e(t) = m_i(s_i(t), e_i(t))/e_i(t)$, for all ijt .

$s, e, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}$ is an equilibrium allocation if there exists a sequence \mathbf{p} of prices such that $s, e, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}, \mathbf{p}$ is an equilibrium.

Definition 2 characterizes the entire set of equilibria that can arise under different assumptions on the pricing mechanism. For instance, both exogenous prices in the case of taxicabs, as well as bilateral bargaining in the case of bulk shipping, select a subset of the equilibria satisfying Definition 2.

2.3 Efficiency

In this Section we present our efficiency results and discuss the role of optimal pricing and taxation. First, we identify the sources of inefficiency, by comparing the social planner's problem to the market equilibrium allocation, and we state the conditions for efficiency. We then derive taxes and subsidies restoring efficiency when prices are Nash bargained. Finally, we derive the efficient pricing rules.

Social planner The present discounted social welfare of a dynamic allocation $s, e, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}$ is given by

$$u(s, e, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}) \equiv \sum_{t=0}^{\infty} \beta^t \{ W(q(t)) - K(n(t)) + \mathcal{E}(b(t)) - \sum_i s_i(t) c_i^s - \sum_{ij} e_i(t) G_{ij}(t) c_{ij}^e - \sum_{ij} [q_{ij}(t) + b_{ij}(t)] c_{ij}^s / [1 - \beta(1 - d_{ij})] \} \quad (14)$$

where $W(q(t))$ denotes the sum of customers' valuations at $q(t)$ - that is, the gross total customer welfare, before accounting for search and entry costs; $K(n(t))$ denotes the sum of customers' entry costs; and $\mathcal{E}(b(t))$ denotes carriers' total expected utility from the i.i.d. shocks when their relocation choices are described by the vector $b(t)$. The functions $K(\cdot)$ and $\mathcal{E}(\cdot)$ are formally defined and discussed in Appendix A.1.2.

In words, at time t , $n_{ij}(t)$ new customers enter on ij , generating total entry costs $K(n(t))$; $s_i(t)$ searching carriers pay the per-period cost c_i^s ; $e_i(t)G_{ij}(t) = e_{ij}(t)$ searching customers pay the per-period cost c_{ij}^e . Then matching takes place, and $q_{ij}(t)$ customers depart on each route ij , generating social value

$W(q(t))$. At the end of the period, unmatched carriers face the relocation choice, generating a total expected utility $\mathcal{E}(b(t))$ from the idiosyncratic shocks. Finally, $q_{ij}(t) + b_{ij}(t)$ carriers start traveling, and they pay traveling cost c_{ij}^s until reaching their destination.¹⁵

The social planner, who is subject to the same search frictions as the market, maximizes the discounted social welfare over the set of all feasible allocations. That is, the (constrained) efficient allocations solve $\sup_{\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b} \in \mathcal{A}} u(\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b})$. It is useful to formulate the planner's problem in two distinct sub-problems as follows. First, define the discounted social welfare corresponding to the triplet $\mathbf{s}, \mathbf{e}, \mathbf{G}$ by

$$v(\mathbf{s}, \mathbf{e}, \mathbf{G}) \equiv \sup_{\mathbf{q}, \mathbf{n}, \mathbf{b}} u(\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}) \text{ s.t. } \mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b} \in \mathcal{A} \quad (15)$$

Problem (15) solves for the carriers' optimal relocation decisions, \mathbf{b} , the decision of whether to accept a match or not, \mathbf{q} , and customer entry, \mathbf{n} , while taking as given the number of searching carriers (\mathbf{s}) and customers (\mathbf{e}), as well as the customers' destination shares (\mathbf{G}). We can then write the social planner's problem as follows:

$$\sup_{\mathbf{s}, \mathbf{e}, \mathbf{G}} v(\mathbf{s}, \mathbf{e}, \mathbf{G}) \quad (16)$$

Intuitively, first, the planner chooses the allocation of searching agents $\mathbf{s}, \mathbf{e}, \mathbf{G}$. Second, conditional on this choice, the planner faces the logistics problem of choosing carriers' movements, (\mathbf{q}, \mathbf{b}) , optimally so that they are available in different regions as described by \mathbf{s} , and consistent with demand for transport as described by \mathbf{e}, \mathbf{G} .¹⁶ In other words, the planner solves for the optimal \mathbf{q}, \mathbf{b} that are consistent with $\mathbf{s}, \mathbf{e}, \mathbf{G}$.

This formulation allows us to show, in Proposition 1 below, that there are no externalities involved with Problem 15. That is, conditional on $\mathbf{s}, \mathbf{e}, \mathbf{G}$, the market solves for the efficient carrier movements and match acceptance/rejections.

Proposition 1. *Let $\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}$ be an equilibrium allocation. Then $\mathbf{q}, \mathbf{n}, \mathbf{b}$ solves Problem 15.*

Proof. See Appendix A. □

¹⁵Hence the expected discounted sum of future traveling costs is given by $\sum_{\tau=0}^{\infty} (1 - d_{ij})^{\tau} \beta^{\tau} c_{ij}^s = c_{ij}^s / [1 - \beta(1 - d_{ij})]$.

¹⁶Note that \mathbf{n} is added as a control variable for notational convenience, even though it is pinned down by $\mathbf{e}, \mathbf{G}, \mathbf{q}$ from equation (13).

Proposition 1 implies that the only sources of inefficiency stem from $(\mathbf{s}, \mathbf{e}, \mathbf{G})$, i.e. from the allocation of searching agents. Hence, it is useful to “optimize out” $(\mathbf{q}, \mathbf{n}, \mathbf{b})$ and focus on the planner’s problem (16), which optimizes over $(\mathbf{s}, \mathbf{e}, \mathbf{G})$.

Efficiency Conditions Having pin-pointed the sources of inefficiency, we next discuss why they arise, and derive efficiency conditions. We do so by first discussing intuitively the marginal social values of additional searching agents, and how they compare to the market. Theorem 1 contains the formal result.¹⁷

To facilitate this comparison, it is useful to rewrite the agents’ value functions in terms of their private surplus from matching. Let $\Delta_{ij}^s(t)$, $\Delta_{ij}^e(t)$ and $\Delta_{ij}(t)$ denote the carrier’s, customer’s and joint surplus respectively, so that:

$$\Delta_{ij}^s(t) = \max\{p_{ij}(t) + V_{ij}^s(t) - U_i^s(t), 0\} \quad (17)$$

$$\Delta_{ij}^e(t) = \max\{w_{ij}(t) - p_{ij}(t) - \beta V_{ij}^e(t+1), 0\} \quad (18)$$

$$\Delta_{ij}(t) = \Delta_{ij}^s(t) + \Delta_{ij}^e(t)$$

Then the private value functions $V_i^s(t)$ and $V_i^e(t)$ - equations (2) and (6)- can be written as

$$V_i^s(t) = -c_i^s + \lambda_i^s(t) \sum_j G_{ij}(t) \Delta_{ij}^s(t) + U_i^s(t) \quad (19)$$

$$V_i^e(t) = -c_{ij}^e + \lambda_i^e(t) \Delta_{ij}^e(t) + \beta V_{ij}^e(t+1) \quad (20)$$

so that agents always receive their outside option, and on top of it, with probability $\lambda_i^s(t)$ and $\lambda_i^e(t)$, they also obtain their respective match surplus.

While agents focus on their *private* surplus from matching, the social planner also takes into account how agents’ search behavior affects total welfare. Indeed, intuitively, the social value of an additional searching carrier in location i satisfies

$$V_i^{*s}(t) = -c_i^s + \underbrace{\lambda_i^s(t) \sum_j G_{ij}(t) \Delta_{ij}^*(t)}_{\text{average match creation}} + s_i(t) \underbrace{\frac{d\lambda_i^s(t)}{ds_i(t)} \sum_j G_{ij}(t) \Delta_{ij}^*(t)}_{\text{congestion effect}} + U_i^{*s}(t) \quad (21)$$

¹⁷We are grateful to a referee for their helpful suggestions in the exposition of this section.

In words, the additional carrier entails a waiting cost, c_i^s ; then at rate $\lambda_i^s(t)$ the carrier is matched, generating expected social surplus $\Sigma_j G_{ij}(t) \Delta_{ij}^*(t)$. In addition, the carrier changes the matching probability $\lambda_i^s(t)$ for all other $s_i(t)$ searching carriers by $d\lambda_i^s(t)/ds_i(t)$, so the total expected impact of this congestion is $s_i(t) d\lambda_i^s(t)/ds_i(t) \Sigma_j G_{ij}(t) \Delta_{ij}^*(t)$. Finally, $U_i^{*s}(t)$ is the social value of an unmatched carrier in i . As we show in Appendix A, these social values can be derived directly from the planner's first order conditions.¹⁸

Efficiency requires that the agents' private values are equal to the planner's (e.g. $V_i^{*s}(t) = V_i^s(t)$, etc.). If this is the case, agents internalize their marginal social impact on the equilibrium allocation. Comparing the social value of an additional searching carrier in (21) to the private one in (19), we notice two differences: first, the carrier takes into account the positive match surplus they create for themselves, $\lambda_i^s(t) \Sigma_j G_{ij} \Delta_{ij}^s(t)$, but not the surplus generated for the customer, $\lambda_i^s(t) \Sigma_j G_{ij}(t) \Delta_{ij}^e(t)$. Second, the carrier ignores the negative congestion effect on other matches, $d\lambda_i^s(t)/ds_i(t) s_i(t) \Sigma_j G_{ij}(t) \Delta_{ij}(t)$. These are the “thick market and congestion” externalities, which are well-known in the search literature.

In order for the social values to equal the private values, it must be that these two effects exactly offset each other. That is, $\lambda_i^s(t) \Sigma_j G_{ij}(t) \Delta_{ij}^e(t) = -s_i(t) d\lambda_i^s(t)/ds_i(t) \Sigma_j G_{ij}(t) \Delta_{ij}(t)$, which can be rewritten as,

$$\frac{\Sigma_j G_{ij}(t) \Delta_{ij}^s(t)}{\Sigma_j G_{ij}(t) \Delta_{ij}(t)} = \eta_i^s(t) \quad (22)$$

where $\eta_i^s(t) = s_i(t) (dm_i(t)/ds_i(t))/m_i(t)$ denotes the elasticity of the matching function with respect to searching carriers at it . This condition, which requires that the carrier's surplus share equals the elasticity of the matching function with respect to carriers, is a restatement of the well-known Hosios (1990) condition. As shown formally below, when this condition holds, thick market and congestion externalities are internalized.

¹⁸Formally, $V_i^{*s}(t)$, $U_i^{*s}(t)$, $\Delta_{ij}^*(t)$ and $V_{ij}^{*e}(t)$ correspond to the shadow values - that is, (minus) the Lagrange multipliers - associated with the planner's constraints (9)-(11) and (13), respectively. With this notation, the first order conditions of the planner's problem with respect to $s_i(t)$ and $e_{ij}(t) = e_i(t) G_{ij}(t)$ yield recursive expressions for the shadow values of searching agents:

$$\begin{aligned} V_i^{*s}(t) &= -c_i^s + dm_i(t)/ds_i(t) \Sigma_j G_{ij}(t) \Delta_{ij}^*(t) + U_i^{*s}(t) \\ V_{ij}^{*e}(t) &= -c_{ij}^e + dm_i(t)/de_i(t) \Sigma_k G_{ik}(t) \Delta_{ik}^*(t) + \beta V_{ij}^{*e}(t+1) + \lambda_i^e(t) [\Delta_{ij}^*(t) - \Sigma_k G_{ik}(t) \Delta_{ik}^*(t)] \end{aligned}$$

Substituting $dm_i(t)/ds_i(t) = \lambda_i^s(t) + s_i(t) d\lambda_i^s(t)/ds_i(t)$ and $dm_i(t)/de_i(t) = \lambda_i^e(t) + e_i(t) d\lambda_i^e(t)/de_i(t)$ yields equations (21) above, as well as (23) below. Efficiency requires that $(V_i^{*s}(t), U_i^{*s}(t), V_{ij}^{*e}(t), \Delta_{ij}^*(t)) = (V_i^s(t), U_i^s(t), V_{ij}^e(t), \Delta_{ij}(t))$.

Next, the social value of an additional searching customer in location i heading to destination j satisfies

$$V_{ij}^{*e}(t) = -c_{ij}^e + \underbrace{\lambda_i^e(t)\Delta_{ij}^*(t)}_{\text{match creation}} + e_i(t) \underbrace{\frac{d\lambda_i^e(t)}{de_i(t)}\sum_k G_{ik}(t)\Delta_{ik}^*(t)}_{\text{congestion effect}} + \beta V_{ij}^{*e}(t+1) \quad (23)$$

Note that customers decide not only whether to enter or not, but also their destination. For this reason, differently from the case of the carriers - equation (21) - the customers' match creation effect is destination-specific. This creates a "composition effect": an additional customer affects not only the total number of matches at origin, but also the destination shares faced by carriers. To see this, equation (23) can be re-written as

$$V_{ij}^{*e}(t) = -c_{ij}^e + \underbrace{\lambda_i^e(t)\sum_k G_{ik}(t)\Delta_{ik}^*(t)}_{\text{average match creation}} + e_i(t) \underbrace{\frac{d\lambda_i^e(t)}{de_i(t)}\sum_k G_{ik}(t)\Delta_{ik}^*(t)}_{\text{congestion effect}} + \underbrace{\lambda_i^e(t)[\Delta_{ij}^*(t) - \sum_k G_{ik}(t)\Delta_{ik}^*(t)]}_{\text{composition effect}} + \beta V_{ij}^{*e}(t+1)$$

which is identical to equation (21) except for an additional component stemming from customer heterogeneity, which captures how an additional customer affects the composition of matches. As opposed to the planner, customers do not fully internalize their composition effect, $\Delta_{ij}^*(t) - \sum_k G_{ik}(t)\Delta_{ik}^*(t)$, as they ignore how they affect carriers' average surplus. These externalities are akin to the composition externalities sometimes found in the labor literature. It is worth emphasizing that in this context, the match surplus includes the entire dynamic path of carriers (who after taking a customer to their destination, restart their search there), and as such captures the endogenous structure of the transport network.

In order for the social value (23) to equal the private one (20), it must be that

$$\lambda_i^e(t)\Delta_{ij}^s(t) = -e_i(t)\frac{d\lambda_i^e(t)}{de_i(t)}\sum_k G_{ik}(t)\Delta_{ik}(t) \quad (24)$$

There are now as many such conditions, as there are customer destinations. Compared to the equivalent condition for the carrier, the positive match surplus for the carrier that each additional customer generates, $\Delta_{ij}^s(t)$, is now destination-specific. At the optimum, condition (24) dictates that the positive match creation effect must exactly offset the negative congestion effect for every destination, j . Since the

congestion effect is the same for all j , at the optimum it must be that,

$$\Delta_{ij}^s(t) = \Delta_{ik}^s(t), \text{ for all } j, k \quad (25)$$

i.e. carriers must be indifferent across different types of customers. Intuitively, this guarantees that customers internalize their composition effect on $G(t)$, since differences in their surplus fully reflect the differences in social surplus across destinations, i.e.

$$\Delta_{ij}^e(t) - \Delta_{ik}^e(t) = \Delta_{ij}(t) - \Delta_{ik}(t), \text{ for all } j, k \quad (26)$$

As we show formally below, this new condition guarantees that customers internalize the composition externalities.

Finally, averaging condition (24) by customer type j leads to

$$\frac{\sum_j G_{ij} \Delta_{ij}^e(t)}{\sum_j G_{ij} \Delta_{ij}(t)} = \eta_i^e(t) \quad (27)$$

which is similar to (22) from the customers' perspective, and guarantees that customers also internalize the thick market and congestion externalities.

We now state the formal result, which shows that the planner's first order conditions with respect to (s, e, G) lead to the above conditions ensuring that the market equilibrium internalizes all externalities.

Definition 3. At an allocation $\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}$, we say that: (i) Carriers internalize thick market and congestion externalities if \mathbf{s} maximizes $v(\cdot, \mathbf{e}, \mathbf{G})$; (ii) Customers internalize thick market and congestion externalities if \mathbf{e} maximizes $v(\mathbf{s}, \cdot, \mathbf{G})$; and (iii) Customers internalize composition externalities if \mathbf{G} maximizes $v(\mathbf{s}, \mathbf{e}, \cdot)$.

Theorem 1. *At an equilibrium allocation $\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}$: (i) Carriers internalize thick market and congestion externalities if and only if condition (22) holds; (ii) Customers internalize thick market and congestion externalities if and only if condition (27) holds; (iii) Customers internalize composition externalities if and only if condition (25) holds. Hence, $\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}$ is efficient only if conditions (22), (27)*

and (25) hold jointly.¹⁹

Proof. See Appendix A. □

Conditions (i) and (ii) require that the share of the surplus appropriated by agents on each side of the market be equal to the elasticity of matches with respect to agents searching on the same side, recasting the familiar Hosios (1990) conditions in our setup. They are always satisfied under the surplus sharing rule $\Delta_{ij}^s(t) = \eta_i^s(t)\Delta_{ij}(t)$ and $\Delta_{ij}^e(t) = \eta_i^e(t)\Delta_{ij}(t)$, which corresponds to Nash bargaining when carriers' and customers' bargaining weights equal $\eta_i^s(t)$ and $\eta_i^e(t)$.²⁰

Condition (iii) is novel. Note that a similar condition emerges in a frictionless environment, where competition among carriers ensures that prices coincide with their opportunity cost; in that case, as in condition (iii), carriers are indifferent among serving different types of customers in equilibrium. Also similarly, in a world with search frictions, but where carriers can direct their search to a specific customer type (i.e. a model of directed search which is efficient, see Moen 1997), a no-arbitrage condition makes carriers indifferent between searching across destinations in equilibrium.²¹

Unlike a typical random search and matching setup with homogeneous agents, a simple constant surplus sharing rule is no longer sufficient to guarantee efficiency, since it penalizes high value matches too much (respectively, low value matches too little). As we discuss in the next section, this justifies the use of subsidies for high value matches (respectively, taxes for low value matches).

This logic holds in random search models with heterogeneity, and it is not specific to the transport problem studied here. It is important to note, however, that match surplus heterogeneity is inherent in the spatial nature of transportation, as it arises endogenously in carriers' continuation values. That is, even if customers' costs and valuations were homogeneous at every region, matching surpluses would still differ across destinations due to differences in carriers' continuation values.

Indeed, the importance of pricing destination effects has been recognized in the transportation liter-

¹⁹Formally, conditions (22),(27) and (25) are necessary when $v(s, e, \mathbf{G})$ is differentiable in s , e and \mathbf{G} . As we show in the proof in Appendix A, this is the case almost everywhere.

²⁰In particular note that, in the absence of taxes, a necessary condition for thick market and congestion externalities to be internalized by both carriers and customers is that the matching functions exhibit constant returns to scale, i.e. $\eta_i^s(t) + \eta_i^e(t) = 1$.

²¹In this case, there is a separate market for each customer destination; carriers enter different markets, until in equilibrium they are (ex ante) indifferent across different choices. Markets of more desirable destinations entail longer waiting times for carriers and vice versa.

ature.²² Hence, Theorem 1 draws an important connection between the latter and the search literature. In environments with regulated-prices, it clarifies that the task of optimal destination-specific pricing is to correct composition externalities in customers' decisions. In decentralized environments, composition externalities justify the use of destination-specific taxes and subsidies, as negotiated prices typically fail to fully internalize destination effects. The remainder of this Section formalizes this intuition by deriving explicit formulae for the optimal prices, and optimal taxation in the case of bargained prices.²³

Optimal Taxes under Nash bargaining We first consider the case of a central authority that cannot directly control prices, but can impose taxes and subsidies. We derive the optimal instruments when market prices are Nash bargained, a price mechanism which is commonly employed to capture bilateral negotiations. In that case, prices are determined by the usual surplus sharing conditions $\Delta_{ij}^s(t) = \gamma_i[\Delta_{ij}^s(t) + \Delta_{ij}^c(t)]$ and $\Delta_{ij}^c(t) = (1 - \gamma_i)[\Delta_{ij}^s(t) + \Delta_{ij}^c(t)]$, where γ_i denotes carriers' bargaining coefficient at i . We first provide an intuitive derivation, and we conclude by stating the formal result in Corollary 1 below.

Suppose that, for every ijt , the planner can impose a tax, $\tau_{ij}^q(t)$, on every match on ij , a per-period tax, $\tau_i^s(t)$, on carriers searching at i , and a per-period tax, $\tau_i^c(t)$, on customers searching at i . Notice that, regardless of who pays $\tau_{ij}^q(t)$, we can think of $\gamma_i\tau_{ij}^q(t)$ and $(1 - \gamma_i)\tau_{ij}^q(t)$ as the incidence of this tax on carriers and customers, respectively, since Nash bargaining implies that agents split their joint surplus according to γ_i . Moreover, the pre-tax social match surplus includes the planner's tax revenues and so is now given by $\Delta_{ij}(t) = \Delta_{ij}^s(t) + \Delta_{ij}^c(t) + \tau_{ij}^q(t)$.

First, Theorem 1 indicates that, to correct the composition externalities, destination-specific taxes must be such that customers' surplus shares fully reflect the differences in social surplus across destinations, i.e. $\Delta_{ik}^c(t) - \Delta_{ij}^c(t) = \Delta_{ik}(t) - \Delta_{ij}(t)$. Rearranging this yields $(1 - \gamma_i)[\tau_{ij}^q(t) - \tau_{ik}^q(t)] = \gamma_i[\Delta_{ik}(t) - \Delta_{ij}(t)]$, so that the customer is taxed more for destinations with low (gross) carrier surplus,

²²For instance, this idea can be found in Koopmans (1949), while BKP observed the presence of destination effects in oceanic shipping prices.

²³Our baseline model can be generalized along several dimensions, including general types of heterogeneity, one-to-many matching and multiple stops (results available upon request). In this general model, we allow both carriers and customers to have a general dynamic type; while here we only allow this type to include current location. This can include many different attributes, such as the carrier's residual capacity, the list of its existing contracts with customers, or the customer's preference for priority. We can show that an analogous theorem holds, as the same externalities arise (with the notable difference that composition externalities now arise from both customers and carriers, since both can be heterogeneous).

$\gamma_i \Delta_{ij}(t)$.²⁴

This system has one degree of freedom, as it pins down only the differences between destination-specific taxes. Hence τ^q can be chosen so that the average match tax $\sum_j G_{ij}(t) \tau_{ij}^q(t) = 0$; this yields,

$$\tau_{ij}^q(t) = \frac{\gamma_i}{1 - \gamma_i} [\sum_k G_{ik}(t) \Delta_{ik}(t) - \Delta_{ij}(t)] \quad (28)$$

In words, composition externalities can be corrected under a balanced budget, by means of simple transfers from low to high value customers, penalizing the former proportionally to their externality on the average match surplus. Note that naturally, if customers were homogeneous, $\tau^q = 0$.

Second, rewriting the conditions of Theorem 1 to explicitly account for Nash bargaining, we find that thick market and congestion externalities are internalized when,

$$\tau_i^s(t) / \lambda_i^s(t) = [\gamma_i - \eta_i^s(t)] \sum_j G_{ij}(t) \Delta_{ij}(t) \quad (29)$$

$$\text{and } \tau_i^e(t) / \lambda_i^e(t) = [1 - \gamma_i - \eta_i^e(t)] \sum_j G_{ij}(t) \Delta_{ij}(t) \quad (30)$$

The left hand side of (29) is the search tax per match levied on carriers. To restore efficiency, search taxes must fill the gap between carriers' pre-tax average surplus shares, $\gamma_i \sum_j G_{ij}(t) \Delta_{ij}(t)$, and their marginal social impact through match creation, $\eta_i^s(t) \sum_j G_{ij}(t) \Delta_{ij}(t)$. The intuition for (30) is similar. Consistently with the discussion following Theorem 1, when $\tau^s, \tau^e = 0$, the conditions required to internalize the thick market and congestion externalities boil down to $\gamma = \eta^s(t)$ and $1 - \gamma = \eta^e(t)$, which are the well-known Hosios conditions.

Corollary 1. *If prices are Nash bargained, at an equilibrium with taxes τ : (i) If τ^s satisfies condition (29), carriers internalize thick market and congestion externalities; (ii) If τ^e satisfies condition (30), customers internalize thick market and congestion externalities; (iii) If τ^q satisfies condition (28), customers internalize composition externalities.*²⁵

²⁴Note that, consistently with Theorem 1, when $\tau^q = 0$, the condition required to internalize the composition externalities boils down to the carriers' surplus $\gamma_i \Delta_{ij}(t)$ being the same across all destinations (Condition 25 above).

²⁵As noted in the discussion introducing condition (28), the set of efficient tax systems has one degree of freedom at every region, since condition (25) only pins down the differences between destination-specific taxes. Hence conditions (28) - (30) describe the unique efficient tax system satisfying $\sum_j G_{ij}(t) \tau_{ij}^q(t) = 0$, for all it . The full set of efficient tax systems is characterized in the proof in Appendix A.

Proof. See Appendix A. □

Before concluding this Section, note that adding conditions (29) and (30) yields the following expression for total tax revenues

$$e_i(t)\tau_i^e(t) + s_i(t)\tau_i^s(t) = m_i(t)[1 - \eta_i^s(t) - \eta_i^e(t)] \sum_j G_{ij}(t)\Delta_{ij}(t)$$

Hence, in the case where the matching function exhibits constant returns to scale ($\eta_i^s(t) + \eta_i^e(t) = 1$), thick market and congestion externalities can be restored under a balanced budget, by means of simple transfers between carriers and customers. If the matching function has decreasing returns to scale, then tax revenues are positive at that location, capturing the fact that matches become harder to form as the market size grows. On the contrary, matches are subsidized when the matching functions have increasing returns.

Price-setting authority One implication of the results in the previous Section is that, when the matching functions exhibit constant returns to scale, efficiency can be restored by means of re-distributive transfers. This suggests that, in this case, efficiency can be restored through prices. We thus consider the case where a central authority can mandate prices, as is typical for taxis.

Under constant returns to scale, Theorem 1 yields a closed-form expression for the optimal prices:

$$p_{ij}(t) = \eta_i^s(t) \sum_k G_{ik}(t)\Delta_{ik}(t) + U_i^s(t) - V_{ij}^s(t) \quad (31)$$

This expression ensures that prices reflect differences in carrier continuation values across destinations, i.e. $p_{ij}(t) - p_{ik}(t) = V_{ik}^s(t) - V_{ij}^s(t)$, for all jk . Equivalently, this guarantees that prices are higher for customers with low value destinations.²⁶

²⁶One can see this by writing the optimal prices as

$$p_{ij}(t) = w_{ij}(t) - \beta V_{ij}^e(t+1) - \eta_i^e(t) \sum_k G_{ik}(t)\Delta_{ik}(t) + [\sum_k G_{ik}(t)\Delta_{ik}(t) - \Delta_{ij}(t)]$$

Note that, under constant returns, this is equivalent to (31). This illustrates that customers' charges must reflect shipping values net of their outside options; in addition, customers must be compensated for their marginal welfare impact through match creation, and penalized proportionally to their negative effect on the composition of matches.

To gain more intuition, we can show that, when prices are set according to (31), we have

$$\eta_i^e(t)\Delta_{ij}^s(t) = \eta_i^s(t)\{\Delta_{ij}^e(t) - [\Delta_{ij}(t) - \sum_k G_{ik}(t)\Delta_{ik}(t)]\}$$

where the terms $\Delta_{ij}^s(t)$ and $\Delta_{ij}^e(t)$ depend on the price $p_{ij}(t)$. This relationship is reminiscent of a surplus sharing condition under Nash bargaining, where however we have (i) replaced the bargaining coefficients with the respective matching function elasticities (this amounts to satisfying the Hosios condition under Nash bargaining); and (ii) adjusted customers' outside option by the deviation of route ij 's total surplus from the average across destinations. Similar to the discussion of optimal taxes, adjusting customers' outside option ensures that they fully internalize the social value of their destination in their decision-making: if the customer has chosen a destination with higher social surplus, he should enjoy a higher outside option (and thus a lower price), and vice versa.

Corollary 2. *When the matching functions exhibit constant returns to scale, all externalities are internalized at equilibrium if and only if prices satisfy condition (31).*

This result provides a simple recipe to assess efficiency of mandated pricing rules.²⁷ For instance, in the case of taxicabs, prices are regulated and are roughly set equal to a tariff plus a fee proportional to travel time and distance. Corollary 2 indicates that this pricing rule is unlikely to be efficient, since the efficient prices should be origin-destination specific.

In the case where the matching functions exhibit increasing or decreasing returns to scale, efficiency requires that there is a wedge between the price a customer pays and the price the carrier receives.²⁸ As in the discussion of optimal taxes above, these wedges capture the social cost (benefit) of an additional searching agent on both sides when there are decreasing (increasing) returns.

²⁷Equation (31) does not provide the prices in closed-form, since the right-hand-side also depends on them through the value functions. Thus, the prices can be obtained as the fixed point of an algorithm which, at every iteration, first computes the equilibrium under the current set of prices, and then updates the latter according to condition (31). There is however a second interpretation of the efficient prices, which follows our proof in Appendix A and recognizes that the value functions can be interpreted as social planner's objects, corresponding to a primal-dual solution to the planner's problem. As such, the prices can be obtained by solving the planner's problem, computing the optimal Lagrange multipliers $\mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta}$, and setting prices according to (31). The efficient taxes derived in previous Section can be also obtained through a similar algorithm.

²⁸In particular, the prices received by carriers are those given by (31), while the prices paid by customers are those given in footnote 26.

3 Empirical application: dry bulk shipping

In this section we describe our empirical application using data from the dry bulk shipping industry. We begin with a brief description of the industry and the available data in Section 3.1. In Section 3.2 we discuss search frictions in this market. In Section 4 we discuss estimation and results.²⁹ Throughout the following sections, unless otherwise noted, a time period is a week and we split ports into 15 geographical regions, depicted in Figure 8 of Appendix D.³⁰

3.1 Industry description and data

Industry description Dry bulk shipping involves vessels designed to carry a homogeneous unpacked dry cargo for individual shippers. It is the prevalent mode of transportation for international trade in commodities, such as iron ore, steel, coal, bauxite, phosphates, but also grain, sugar, chemicals, lumber and wood chips; it accounts for about half of total seaborne trade in tons (UNCTAD) and 45% of the total world fleet, which includes also containerships and oil tankers.³¹

Bulk carriers operate much like taxi cabs: a specific cargo is transported individually by a specific ship, for a trip between a single origin and a single destination. Ships carry the cargo of one customer at a time, who fills up the entire ship. In this paper, we focus on spot contracts and in particular the so-called “trip-charters”, in which the shipowner is paid on a per day rate.³²

The industry is unconcentrated, consisting of a large number of small shipowning firms (Kalouptsidi, 2014): the maximum fleet share is around 4%, while the firm size distribution features a large tail of small shipowners. Dry bulk shipping is a decentralized market: shipowners and exporters find each other using brokers. “The broker’s task is to discover what cargoes or ships are available....[M]ost owners and charterers use one or more brokersoften in tense competition with [each] other” (Stopford, 2009). Finally, there are four size categories of dry bulk carriers: Handysize (10,000–40,000 DWT), Handymax (40,000–60,000 DWT), Panamax (60,000–100,000 DWT) and Capesize (larger than 100,000 DWT). Beyond this,

²⁹The material in Section 3.2 used to be included in a previous working version of BKP (Brancaccio et al. 2018).

³⁰Our results are robust to different specifications of space and time aggregation.

³¹Bulk ships are different from containerships, which carry cargo (mostly manufactured goods) from many different cargo owners in container boxes, along fixed itineraries according to a timetable. It is generally not technologically possible to substitute bulk with container shipping.

³²Trip-charters are the most common type of contract. Long-term contracts (“time-charters”), however, do exist: on average, about 10% of contracts signed are long-term. Interestingly, though, ships in long-term contracts, are often “relet” in a series of spot contracts, suggesting that arbitrage is possible.

shipping services are largely perceived as homogeneous. In his lifetime, a shipowner will contract with hundreds of different exporters, carry a multitude of different products and visit numerous countries. We discuss these issues further at the end of Section 3.2.

Data We combine several data sets. First, we use a data set of shipping contracts, from 2010 to 2016, collected by Clarksons Research. An observation is a transaction between a shipowner and a charterer for a specific trip and reports the vessel, the charterer, the contract signing date, the loading and unloading dates, the price in dollars per day, as well as some information on the origin and destination.

Second, we use AIS (Automatic Identification System) data from exactEarth Ltd for the ships in the Clarksons data set between May 2012 and March 2016. AIS transceivers on the ships automatically broadcast information, such as their position (longitude and latitude), speed, and level of draft (the vertical distance between the waterline and the bottom of the ship’s hull), at regular intervals of at most six minutes. The draft is a crucial variable, as it allows us to determine whether a ship is loaded or not at any point in time.

Finally, we augment the two ship data sets with international trade data from Comtrade on export value and volume by country pair for bulk commodities.

Our final dataset involves 5,398 ships, which corresponds to about half the world bulk fleet. We end up with 12,007 shipping contracts, for which we know the price, as well as the exact origin and destination. The average price is 14,000 dollars per day (or 290,000 dollars for the entire trip), with substantial variation. Trips last on average 2.9 weeks. Contracts are signed close to the loading date, on average 6 days before. We have 393,058 ship-week observations at which the ship decides to either ballast (i.e. travel empty) someplace or stay at its current location. Finally, Clarksons reports the product carried in about 20% of the sample. The main commodity categories are grain (29%), ores (21%), coal (25%), steel (8%) and chemicals/fertilizers (6%).³³

Some salient facts An important feature of this market revealed by the AIS data, is that trade in commodities is greatly imbalanced, so that most countries are either large net importers or large net exporters. For instance, Australia, Brazil and Northwest America, the world’s biggest exporters of

³³We refer the interested reader to BKP for further details on the construction of the final dataset, which requires using the AIS data to construct ship histories of port calls, as well as for detailed summary statistics (see Table 1 in BKP).

commodities, are rich in minerals, grain, coal, etc. At the same time, China and India, the world’s biggest importers, require raw materials to grow further. As a result, commodities flow out of the former, towards the latter.

The trade imbalances have implications for both ship ballast behavior and shipping prices. Indeed, at any point in time, 42% of ships are traveling without cargo. At the same time, prices are largely asymmetric and depend on the destination’s trade imbalance: all else equal, the prospect of having to ballast after offloading is associated with higher shipping prices. This is shown in Table 7, in which we regress the shipping prices from i to j on the probability of having to ballast after j , as well as the average distance traveled ballast after j ; we find that both variables significantly raise the price to export to j . This is intuitive: a shipowner commands a higher price for a trip towards an unattractive destination (i.e. one with a low continuation value). These facts are documented in more detail in BKP.

These patterns have important implications for efficiency. The observed heterogeneity of different regions foreshadows that composition externalities might be a serious concern: for instance, a ship in Australia may be matched to an exporter headed to China, or an exporter headed to North Europe. These destinations have very different implications for the ship’s continuation value.

3.2 Search frictions in dry bulk shipping

A number of features of dry bulk shipping, such as information frictions and port infrastructure, can hinder the matching of ships and exporters. In this section we argue that these frictions indeed lead to unrealized potential trade. Consider a geographical region, such as a country or a set of neighboring countries, where there are s ships available to pick up cargo and e exporters searching for a ship to transport their cargo. We define search frictions by the inequality:

$$m < \min \{s, e\} \tag{32}$$

where m is the number of matched ships and exporters. In other words, under frictions there is potential trade that remains unrealized; in contrast, in a frictionless world, the entire short side of the market gets matched, so that $m = \min \{s, e\}$. When inequality (32) holds, matches are often modeled via a matching function, $m = m(s, e)$, as is done in Section 2 above, and also extensively in the search literature.

We present three facts consistent with frictions, as defined by (32). In particular, we (i) provide a direct test for inequality (32); (ii) we document wastefulness in ship loadings; (iii) we document substantial price dispersion.

First, we provide a simple test that directly verifies the presence of search frictions. If we observed all variables s, e, m , it would be straightforward to test (32); this is essentially what is done in the labor literature, where the co-existence of unemployed workers and vacant firms is interpreted as evidence of frictions. However in our setup, we observe m (i.e. ships leaving loaded) and s , but not e ; we thus need to adopt a different approach.

Assume there are more ships than exporters, i.e. $\min(s, e) = e$. We begin with this assumption, because our sample period is one of low shipping demand and severe ship oversupply due to high ship investment between 2005 and 2008 (see Kalouptsidei, 2014). If there are no search frictions, so that $m = \min(s, e) = e$, small exogenous changes in the number of ships should not affect the number of matches. In contrast, if there are search frictions, an exogenous change in the number of ships changes the number of matches, through the matching function $m = m(s, e)$. We approximate an exogenous change in the number of ships, with unpredictable ocean weather conditions. The intuition is that wind affects the speed at which ships travel and therefore exogenously shifts the supply of ships at port. We therefore explore whether exogenously changing the number of ships in regions with a lot more ships than exporters affects the realized number of matches.³⁴ Since we do not observe exporters directly, to select periods in which there are more ships than exporters, for each region we consider weeks when there are at least twice as many ships waiting in port as matches. Table 1 presents the results: matches are affected by weather conditions in all but one region, consistent with the presence of search frictions. It is worth noting that although search frictions have been prevalent in the recent literature on transport markets, to our knowledge we are the first to propose a formal test for their presence.

Second, we document simultaneous arrivals and departures of empty ships. Indeed, the first two panels of Figure 1 display the weekly number of ships that arrive empty and load, as well as the number of ships that leave empty, in two net exporting countries: Norway and Chile. In Norway, several ships arrive empty and load, but almost no ship departs empty. In Chile, however, the picture is quite different: it

³⁴We collect global data on daily sea weather from the ERA-Interim archive, from the European Centre for Medium-Range Weather Forecasts (CMWF). This allows us to construct weekly data on the wind speed (in each direction) across all oceans (see Table 1 for details on constructing the instrument).

	N	Joint Significance	$\frac{s}{m}$
North America West Coast	193	0	2.706
North America East Coast	200	0	3.172
Central America	199	0.001	3.451
South America West Coast	198	0	2.913
South America East Coast	200	0	4.083
West Africa	200	0.001	5.862
Mediterranean	200	0	4.244
North Europe	200	0	3.577
South Africa	200	0	2.862
Middle East	200	0	3.86
India	200	0.34	8.58
South East Asia	200	0	3.334
China	200	0.038	6.194
Australia	187	0	2.457
Japan-Korea	200	0	5.311

Table 1: Test for search frictions. Regressions of the number of matches in each region on the unpredictable component of weather conditions in the surrounding seas. For each region we use weeks in which there are at least twice as many ships as matches. The first column reports the number of observations; the second column joint significance; and the third column the average ratio of ships over matches in each region during these weeks. To proxy for the unpredictable component of weather, we partition the globe into cells of $9^\circ \times 9^\circ$, and for each cell we collect data on the speed of the horizontal (E/W) and vertical (N/S) component of wind, as well as wave period and height. To control for seasonality, we residualize the weather measurements for each cell on a quarter fixed effect. The potential regressors include one and two weeks lagged values of all the weather measurements for cells in the sea. Finally, we follow Belloni et al. (2012) to select the relevant instruments in each region.

frequently happens that an empty ship arrives and picks up cargo, while at the same time another ship departs empty. This is suggestive of wastefulness in Chile: although there was both an empty ship in Chile, as well as a cargo awaiting transport, a different ship came empty from elsewhere to pick up this cargo.

This pattern is observed in many countries. Indeed, the third panel of Figure 1 depicts the histogram of the bi-weekly ratios of outgoing empty ships over incoming empty and loading ships for net exporting countries. In the absence of frictions, one would expect this ratio to be close to zero. However, we see that the average ratio is well above zero. Moreover, this pattern is quite robust in a number of dimensions.³⁵

³⁵This figure is robust to alternative market definitions, time periods and when constructed separately for each vessel size. Capesize vessels exhibit somewhat larger mass towards zero, consistent with the somewhat higher concentration of ships and charterers, as well as the ships' ability to approach fewer ports. The figure is also similar if done by port rather than country. To control for repairs we remove stops longer than 6 weeks. Finally, we only consider as "ships arriving empty" the ships arriving empty and sailing full towards another region, and we consider as "ships leaving empty" ships sailing empty toward

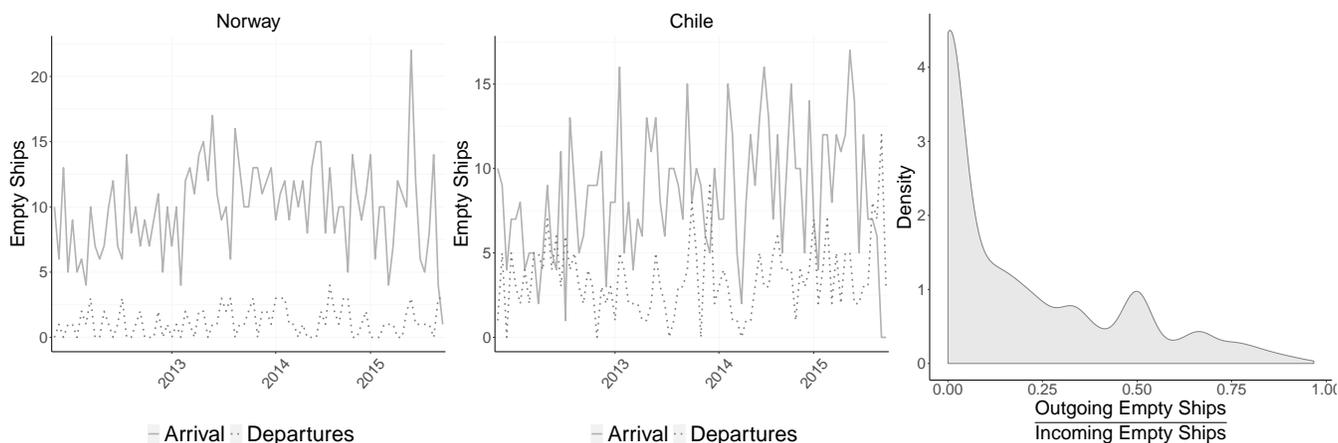


Figure 1: Simultaneous arrivals and departures of empty ships: The first two panels depict the flow of ships arriving empty and loading, and ships leaving empty in two-week intervals in Norway and Chile. The last panel shows the histogram of the ratio of outgoing empty over incoming empty and loading ships across all net exporting countries.

Third, again inspired by the labor literature, we investigate dispersion in prices. In markets with no frictions, the law of one price holds, so that there is a single price for the same service. This does not hold, for instance, in labor markets, where there is large wage dispersion among workers who are observationally identical.³⁶ Table 7 in Appendix D shows that there is substantial price dispersion in shipping contracts. More specifically, at best we can account for about 70% of price variation, controlling for ship size, as well as quarter, origin and destination fixed effects. Moreover, the coefficient of variation of prices within a given quarter, origin and destination triplet is about 30% (23%) on average (median). In the most popular trip, from Australia to China, the *weekly* coefficient of variation is on average 34% and ranges from 15% to 65% across weeks. In addition, the type of product carried affects the price paid and overall more valuable goods lead to higher contracted prices, as shown in the same table. In the absence of frictions, if there are more ships than exporters, as is the case during our sample period, we would expect prices to be bid down to the ships' opportunity cost.³⁷ In contrast, in markets with frictions and bilateral bargaining, since ships now have market power, the price depends on the exporter's valuation and higher value exporters pay more.

a different country; so movements to nearby ports are excluded. This definition also implies that refueling cannot explain the fact either- though there are very small differences in fuel prices across space anyway (less than 10%).

³⁶This fact has generated an influential literature on frictional wage inequality. See for instance Mortensen (2003) and references therein.

³⁷In a frictionless market with more ships than exporters and homogeneous ships, in equilibrium the price from an origin to a destination would be such that ships are indifferent between transporting the cargo and staying unmatched.

As in labor markets, a multitude of factors can lead to frictions (i.e. unrealized matches) in shipping. First, the decentralized and unconcentrated nature of the market and the mere existence of brokers, suggests that information frictions are present. The meeting process involves a disperse network of brokers; oftentimes more than two brokers intervene to close a deal, suggesting that the ship’s and the exporter’s brokers do not always find each other, and that an “intermediate broker” was necessary to bring the two together (Panayides, 2016). In interviews, brokers claimed to receive 5,000-7,000 emails per day; sorting through these emails is reminiscent of an unemployed worker sorting through hundreds of vacancy postings.³⁸ Port infrastructure, congestion and regulations may also hinder matching.³⁹

Are bulk ships similar to taxis? It may seem surprising at first that we make a parallel between the bulk shipping industry and the taxi industry. Yet, it is remarkable how similar the operation of ships is to that of taxicabs (and trucks). In fact, shipowners themselves call bulk ships “the ocean taxis.” Much like taxis, bulk ships “do not operate on a fixed schedule but merely trade in all parts of the world in search of suitable cargoes” (Plakantonaki, 2010, which is the internal manual of a large shipping firm). Moreover, bulk ships transport one cargo at a time and are homogeneous; they largely operate on a voyage basis; they sign contracts very close to loading; and they operate in a decentralized environment without dispatchers. Let us discuss two assumptions that are particularly important: ship homogeneity and random search.

Ships do not specialize neither geographically, nor in terms of products: during the period of our data, ships deliver cargo to 13 out of 15 regions on average and carry at least 2 of the 3 main products (coal, ore and grain). Moreover, neither shipowner characteristics, nor shipowner fixed effects have any explanatory power in price regressions, as shown in Table 5 in Appendix B.1. Furthermore, we fail to

³⁸In response to trading frictions, almost all decentralized markets have intermediaries. Indeed, since Demsetz (1968), search frictions have been used to explain the existence of intermediaries. In addition, a large literature has shown that intermediaries are not able to entirely wipe out frictions (e.g. for theory see Duffie et al. (2005) and for empirics see Gavazza (2016)). As such, a long literature has modeled a multitude of markets (e.g. real estate, OTC financial markets, marriage markets) as having search frictions despite the presence of intermediaries (Albrecht et al., 2016, Allen et al., 2018, Brancaccio et al., 2020c, Farboodi et al., 2021). None of this is entirely surprising; the mere fact that centralized platforms emerge in decentralized markets with intermediaries suggests that intermediaries may not be able to eradicate frictions.

³⁹That said, it is worth pointing out that the matching function, should be interpreted as a “black box” capturing in a reduced form way the real complexities of ship-exporter matching; as such, if institutions respond to the environment (e.g. improvements in communications or port infrastructure) the matching function may change. We thus consider our results to be short- to medium-run.

reject that ships' ballast destination decisions are the same across ships, from a given origin.⁴⁰

Next, we argue that random search is a reasonable approximation in shipping. Meetings occur through an unconcentrated network of brokers. The contracting process begins when a shipowner's broker "is contacted by a competitive broker by email or phone or yahoo/MSN messenger for a specific cargo" (Plakantonaki, 2010). In Appendix B.2, we copy an extract from Plakantonaki (2010) with an example of a representative email notifying brokers of an available cargo. Roughly, the email has the following format: "Please indicate or offer a vessel to be chartered for CORUS to load a full cargo of bulk coal, size 50,000 +/- 5%; load in Poland and discharge in Immingham; vessel must appear between 4/23 and 4/30". As mentioned above, the broker receives thousands such emails a day, suggesting that agents cast a wide net when searching rather than focusing on specific markets. Continuing from Plakantonaki (2010), the broker "considers whether he has a ship available, *usually in the vicinity*, which can safely arrive at the load port during the requested laycan", where the laycan is the requested arrival window. In our data, contracts are signed on average only 6 days prior to loading (compared to an average trip length of 3 weeks), and ships are in the vicinity of the loading area when they sign (they are in the loading region on average 12 days earlier); this suggests that ships do not plan ahead sequences of multiple trips.⁴¹

Finally, we examine the random search assumption more formally. Directed search would imply that carriers can meet only customers heading to a specific destination. In a model of directed search, profitable destinations attract more carriers, thereby reducing their matching probabilities compared to less profitable ones. In equilibrium carriers are indifferent across destinations and matching rates and prices are the equilibrating mechanism.⁴² We first provide suggestive evidence that a directed search model does not provide a good fit for the data. In particular, we show that matching rates from a given origin do not vary by destination; therefore they do not reflect the attractiveness of different destinations, as would be the case in a world of directed search. Next, we test directly the key implication of directed

⁴⁰If heterogeneity were an important driver of ships' ballasting decisions, we would expect ships to choose diverse destinations from a given origin. Yet ballast choices are similar across ships from a given origin (the CR_3 measure for the chosen destinations, i.e. the concentration ratio of the top 3 destinations, is higher than 70% in most regions). Moreover, home-ports are not an important consideration for shipowners, as the crew flies to their home country every 6-8 months.

⁴¹Practitioners explain that contracts tend to be signed with ships that are nearby because "a ship is not a train" and it cannot promise exact arrival times far in advance, due to weather conditions and port congestion. These delays are costly for exporters.

⁴²Attractive destinations attract many searching carriers, but as more carriers enter, matching rates (and prices) fall, reducing the attractiveness of the destination; carriers continue to enter, until matching rates are sufficiently low (and in equilibrium all destinations are equally profitable for the carriers).

search, namely that carriers are indifferent across searching for different destinations, by perturbing our model to allow for directed search. Our results reject the indifference prediction in the overwhelming majority of regions. See Appendix B.2 for the results.

4 Model estimation and results

In this section we estimate the main parameters of the model: in Section 4.1 we estimate the matching functions; in Section 4.2 the ship costs; and in Section 4.3 the exporter parameters. We estimate the model under the assumption of a steady state; in a steady state equilibrium, ships and exporters respond optimally to their expectations of the endogenous market variables, which are constant over time. We can thus omit the dependence on t . We discuss the implications of this assumption at the end of this Section. This Section is self-contained and can be omitted by the reader if so desired.

4.1 Matching Function Estimation

A sizable literature has estimated matching functions in several different contexts (e.g. labor markets, marriage markets, taxicabs). For instance, in labor markets, one can use data on unemployed workers, job vacancies and matches to recover the underlying matching function. In our data, we observe ships and matches, but not searching exporters. As in BKP, we simultaneously recover both exporters, as well as a nonparametric matching function. This approach extends the literature by avoiding parametric restrictions on the matching function; this is important, since as shown in Theorem 1, in frictional markets, the conditions for constrained efficiency depend crucially on the elasticity of the matching function.

Our estimation approach follows BKP; we thus provide but a brief overview and then present the implications for search frictions. Appendix C.1 expands on the approach for completeness. We require that $m(s, e)$ is increasing in s and e , that it exhibits constant returns to scale, and that an instrument that shifts the number of ships exists (the weather shocks employed in Section 3.2).⁴³ The methodology

⁴³As explained in Appendix C.1, the methodology requires a restriction in order to distinguish monotonic transformations of the matching function and the exporters. We impose constant returns to scale (CRS) motivated by both the labor and the urban transport literature that has found strong evidence supporting this assumption. For instance, Petrongolo and Pissarides (2001) note that “Testing for homogeneity, or constant returns to scale, has been one of the preoccupations of the empirical literature” and yet “divergences from constant returns are only mild and rare”. In addition, increasing returns (IRS) do not arise naturally in most models (Petrongolo and Pissarides, 2001 and Duranton and Puga, 2004); in fact Lagos (2003), who microfound a matching function for taxis, proves that it is CRS. One notable exception is the “stock-flow”

delivers exporters point-wise and the matching function of each location i nonparametrically.

Figure 7 in Appendix C.1 presents contour plots of the estimated matching function, as well as the elasticity of the matching function with respect to the number of exporters. The estimates imply that the matching function is steep in the number of exporters, with an average elasticity of 0.87. Therefore, an exporter’s decision to enter the market has a substantial positive externality on matching rates. Vice versa, the matching function is relatively inelastic with respect to the number of ships searching, with an average elasticity of 0.13. This result has important implications for thick market and congestion externalities, as we show below.

Figure 6 in Appendix C.1 presents the estimated search frictions, computed as the average percentage of weekly “unrealized” matches; i.e. $(\min\{s_i, e_i\} - m_i) / \min\{s_i, e_i\}$ in every region. Search frictions are heterogeneous over space and may be somewhat sizable, with up to 20% of potential matches “unrealized” weekly in regions like South and Central America, and Europe. On average, 13.5% of potential matches are “unrealized”.⁴⁴

Finally, estimated search frictions are positively correlated with the observed within-region price dispersion (0.47), another indicator of search frictions. Moreover, frictions are negatively correlated with the Herfindahl-Hirschman Index of charterers reported in the Clarksons contract data (-0.31); this suggests that when the clientele is disperse, frictions are higher. Finally, we estimate the matching function separately for Capesize (biggest size) and Handysize (smallest size) vessels; as expected, for Capesize, where the market is thinner, search frictions are lower.

4.2 Ship costs

To estimate the ship cost parameters parameters, $\{c_{ij}^s, c_i^s\}$, as well as the scale parameter of the logit shocks σ , we use a Nested Fixed Point algorithm (Rust, 1987) as in BKP: at every guess of the parameters

model (see Coles, 1994 and Coles and Smith, 1998), whose key implication is that the matching probability is high when an agent first enters the search pool, but then declines and remains constant. When we test this in our data we find that matching rates do not exhibit this pattern, but instead are fairly flat with respect to the ship’s time already searching.

That said, IRS could be relevant in the context of urban transportation, especially in the case of sparser cities. As Frechette et al. (2019) demonstrate, if the number of agents increases on the same geography, density increases and IRS could materialize. We do not believe that density plays a similar role in the case of shipping. In particular, the number of exporters and ships per unit of space does not directly affect the meeting process, as it does in the case of taxis, as those meetings take place “offline” through the use of brokers.

⁴⁴It is worth noting that this does not imply that in the absence of search frictions there would be 13.5% more matches, as we would need to take into account the optimal response of ships and exporters. This is simply a measure of the severity of search frictions in different regions.

$\{c_{ij}^s, c_i^s, \sigma\}$, we employ a fixed point algorithm to solve for the ship value functions V_i^s, V_{ij}^s, U_i^s , for all ij from equations (1)-(3), using the observed average prices for each route ij and the observed meeting probability λ_i^s (which is set equal to the average m_i/s_i). We then match the ship ballast choices predicted by our model and given by the logit choice probabilities, $\mathbb{P}_{ij}^b = \exp(V_{ij}^s/\sigma) / \sum_l \exp(V_{il}^s/\sigma)$, to the observed ballast choices. We maximize over the parameters via Maximum Likelihood. See BKP for further details on identification and estimation. The ship cost estimates are reported in Table 9 in Appendix D.⁴⁵

4.3 Exporter Parameters

We now turn to the exporter parameters. To proceed, we impose a specific pricing mechanism, Nash bargaining, and we denote by γ_i the ship bargaining coefficient in market i . We also assume that exporter valuations satisfy $w_{ij}(q) = w_{ij}$, for all ij , and that the exporter wait costs, c_{ij}^e do not vary by destination. The parameters of interest here are the exporter valuations w_{ij} , the waiting costs c_i^e , the bargaining coefficients γ_i and the exporters' average entry costs k_{ij} , for all ij .⁴⁶

The valuations w_{ij} are the revenues of exporters in i from selling their goods to destination j . We compute them using trade data from Comtrade, which reports product-level export values and quantities by country pair; the data details are provided in Appendix C.2.

Next, we turn to c_i^e and γ_i , which we estimate from observed shipping prices. Given the estimated ship value functions and the observed prices, we can compute the ship surplus. However, since we do not observe the exporter waiting cost, we do not know the total surplus, and therefore we cannot directly compute the bargaining coefficients. Nonetheless, we can use differences in prices across destinations to identify them, while exporter wait costs can be identified from the level of prices.

Formally, the Nash bargaining surplus sharing condition, $(1 - \gamma_i) (p_{ij} + V_{ij}^s - U_i^s) = \gamma_i [w_{ij} - p_{ij} - \beta V_{ij}^e]$

⁴⁵We construct seven groups for the sailing cost, c_{ij}^s , roughly based on the continent and coast of origin; we estimate port wait costs c_i^s , for all i . Noting that the major component of c_{ij}^s is the cost of fuel, we “normalize” one of the seven groups to equal the average weekly fuel price (\$40,000). Moreover, since the fuel cost is paid by the exporter when the ship is loaded, we add it to the observed prices.

⁴⁶Here, we diverge from the estimation procedure of BKP. The main reason for doing so is to allow the bargaining coefficients γ_i to vary by region, given the importance of these parameters. Moreover, we introduce the Comtrade data to obtain w_{ij} ; this permits us to include exporter waiting costs c_i^e which are not present in BKP.

implies that prices are given by,⁴⁷

$$p_{ij} = \gamma_i \frac{c_i^e + (1 - \beta)(1 - \lambda_i^e)w_{ij}}{1 - \beta(1 - \lambda_i^e) - \gamma_i \lambda_i^e} + \frac{(1 - \gamma_i)[1 - \beta(1 - \lambda_i^e)]}{1 - \beta(1 - \lambda_i^e) - \gamma_i \lambda_i^e} (U_i^s - V_{ij}^s) \quad (33)$$

In this equation, the only unknowns are γ_i and c_i^e , for all i ; indeed, note that λ_i^e is known from the matching function (since it equals m_i/e_i); U_i^s and V_{ij}^s are known once the ship cost parameters have been estimated; w_{ij} is obtained from Comtrade data as described above; and β is calibrated to 0.995. This expression can be rearranged as follows,

$$Y_{ij} = \gamma_i X_{ij} + \gamma_i c_i^e$$

where $Y_{ij} = [1 - \beta(1 - \lambda_i^e)][p_{ij} - (U_i^s - V_{ij}^s)]$ and $X_{ij} = \lambda_i^e p_{ij} + (1 - \beta)(1 - \lambda_i^e)w_{ij} - [1 - \beta(1 - \lambda_i^e)](U_i^s - V_{ij}^s)$, which are known. Taking differences across j we can identify γ_i . Once γ_i has been pinned down, we can directly identify the waiting cost, c_i^e from the average level of prices from (33).

Finally, exporters' average entry costs k_{ij} are estimated using the exporter entry probabilities, which are given by $\mathbb{P}_{ij}^e = \exp(V_{ij}^e - k_{ij})/[1 + \sum_{l \neq i} \exp(V_{il}^e - k_{il})]$, where $V_{ii}^e = k_{ii} \equiv 0$. Then, $\ln \mathbb{P}_{ij}^e - \ln \mathbb{P}_{ii}^e = V_{ij}^e - k_{ij}$, where k_{ij} is the only unknown.⁴⁸

The exporter parameters are presented in Table 8 in Appendix D. The wait costs, c_i^e , are equal to about 3% of the exporters' valuation on average, but there is substantial heterogeneity over space; the estimated costs are highest in Central and South America, as well as parts of Africa. These parameters capture inventory expenditures, delay costs, risk of damage or theft etc. Consistent with this interpretation, exporter wait costs are positively correlated with the recovered wait costs for ships (0.34), and negatively correlated with the World Bank index of quality of port infrastructure (-0.50), as shown in Figure 9 in Appendix D. Finally, the estimates for the bargaining coefficients suggest that exporters get a somewhat larger share of the surplus in almost all regions.

⁴⁷To obtain this, substitute the exporter value βV_{ij}^e from its steady state value $\beta V_{ij}^e = [-c_i^e + \lambda_i^e(w_{ij} - p_{ij})]/[1 - \beta(1 - \lambda_i^e)]$.

⁴⁸To recover \mathbb{P}_{ii}^e , the share of the "outside good", corresponding to the choice of not exporting, we use the total production of the relevant commodities for each region i .

5 Efficiency in dry bulk shipping

In this section we exploit our theoretical results to explore welfare in bulk shipping. In Section 5.1 we check whether the efficiency conditions hold. Section 5.2 presents our main welfare analysis. In Section 5.3 we discuss recent interventions that have affected the industry.

5.1 Is dry bulk shipping efficient?

Efficiency requires that the following conditions are met: (i) the elasticity of the matching function must equal the corresponding share of the surplus, which here is given by the bargaining coefficient (thick market and congestion externalities); (ii) the ship surplus must equalize across destinations (composition externalities). We test each of these conditions in the data.

Figure 2 examines whether the thick market and congestion externalities are internalized. For each region, the left panel presents the average estimated matching function elasticity with respect to exporters, η_i^e , as well as the estimated exporter bargaining coefficient, $\gamma_i^e \equiv 1 - \gamma_i$. For several regions, as shown in the right panel, we reject that the two are equal, so that the “Hosios conditions” (22) and (27) are not satisfied. Although the “knife-edge” nature of these conditions implies that this finding is not particularly surprising, it is worth noting that the difference between the elasticity and the bargaining coefficient is often large. Moreover, γ_i^e tends to be lower than η_i^e , suggesting that the planner would like to see an increase in the share of the surplus accruing to the exporter.

Figure 3 checks whether the composition externalities are internalized. For each region i , it plots the coefficient of variation of the ship surplus from matching with exporters headed to different destinations $j \neq i$. When composition externalities are internalized, the coefficient of variation should be equal to zero, since the ship is indifferent across destinations. Figure 3 demonstrates that this is not the case. In all regions the coefficient of variation is significantly different from zero, and larger than 20%, while in several regions it is substantially higher. It is worth noting that since the ship surplus only depends on the ship cost parameters and observed matching rates, this test does not rely on the Nash bargaining assumption on pricing (nor on assumptions about the matching functions). We conclude that the market has not internalized neither externality.

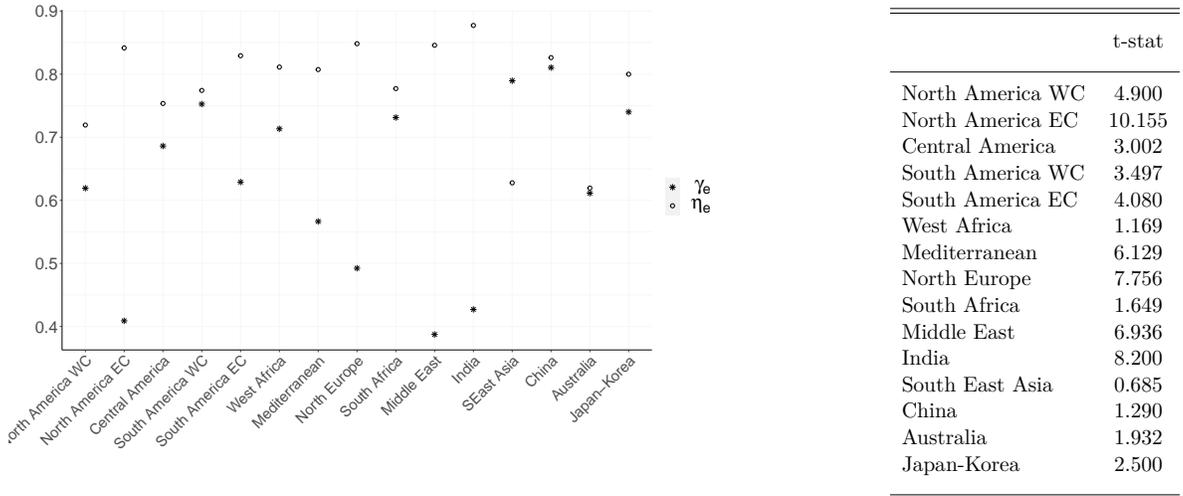


Figure 2: The left panel compares the exporter bargaining coefficient γ_i^e and the average elasticity of the matching function with respect to exporters, estimated nonparametrically. The right panel presents the t-statistic for the null that γ_i^e coincides with the average elasticity of the matching function with respect to exporters.

5.2 Welfare loss

We now come to our main welfare analysis. We begin by comparing the market equilibrium to the frictionless equilibrium (first-best), i.e. a perfectly competitive benchmark where search frictions are not present, so that $m = \min\{s, e\}$, and prices are set to clear demand and supply for trips on every route. To fix ideas, we can think of a model in which, at each location, carriers can offer transportation towards specific destinations, and a Walrasian auctioneer sets the prices to clear supply and customer demand for trips on each route. In this model the first welfare theorem holds as there are no further distortions; therefore, decentralized equilibria always reach the first best, and there is no role for intervention.⁴⁹ This two-way comparison, presented in column I of Table 2, serves as a benchmark, and demonstrates that frictions have a substantial impact on trade and welfare. The frictionless equilibrium welfare is 18% higher compared to the baseline equilibrium, which amounts to almost \$100 million per week. Similarly, trade volume is 4.5% higher, while net trade value (i.e. $w_{ij} - k_{ij}$) is 3.6% higher under the first-best.

We next examine whether a centralizing platform can bring us closer to the first-best. In several setups, from urban transportation to real estate, centralizing platforms have emerged as a market-based solution to search frictions (for instance, Uber/Lyft in the taxi market, Uber Freight and other entrants in the trucking industry, Airbnb in the short-term rental housing market). By centralizing transactions, these

⁴⁹We present the algorithm we use in Appendix C.4.

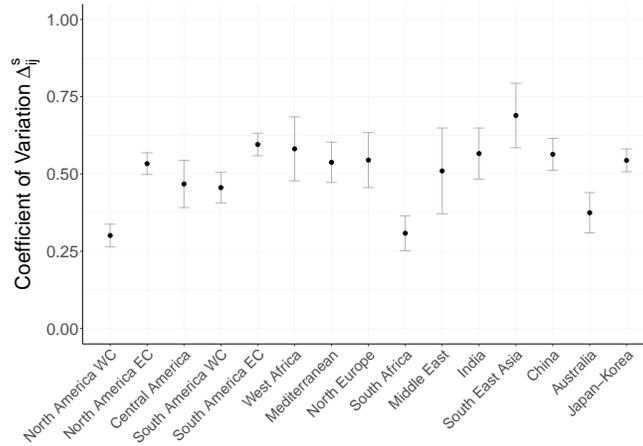


Figure 3: For each region i , we plot the coefficient of variation (standard deviation over mean) of ship surplus for all destinations $j \neq i$. When composition externalities are internalized, the coefficient of variation should be zero.

platforms facilitate meetings and can improve market efficiency. Ideally, centralization would coincide with the frictionless equilibrium discussed above (and indeed the literature has often modeled centralization as the eradication of search frictions taking the regulated prices as given, e.g. Frechette et al. (2019); Buchholz (2021)). However, a centralizing platform does not act like a benevolent planner. Instead, it chooses prices to maximize its profits, thus substituting one friction (search frictions) for another (market power).

To assess the trade-off between search frictions and market power, we consider the optimal pricing problem of a monopolist platform that owns all ships. This is a complex problem, as it entails choosing prices for all possible origin-destination pairs in order to maximize the monopolist profits, taking into account the dynamic paths of ships. To solve this optimization problem, we build an algorithm relying on Rosaia (2020a); see Appendix C.4 for details. Column II of Table 2 compares the market equilibrium to the centralizing monopolist solution. As expected, the increase in welfare is substantially lower than what is achieved in the frictionless equilibrium (around 9%). As a result of the platform’s monopoly power, both trade volume and value drop by 4.5% and 3.5% respectively compared to the baseline. Importantly, the welfare gains are appropriated to a large extent by the centralizing monopolist in the form of profits. The key source of welfare gains is a decline in realized exporter wait costs, driven by both the elimination of search frictions, as well as the increase in shipping prices, which curbs exporter entry.

	Frictionless benchmark	Constrained Efficient	Monopolist
Δ welfare (%)	18.01	8.45	9.10
Δ welfare (\$100,000)	978	459	494.41
Δ Trade (%)	4.54	3.88	-4.59
Δ Trade value (net) (%)	3.58	3.30	-3.48

Table 2: Welfare Analysis. The first column presents the frictionless benchmark, i.e., a perfectly competitive benchmark where search frictions are not present, so that $m = \min\{s, e\}$, and prices are set to clear demand and supply for trips on every route. The second column presents the constrained efficient allocation, i.e. the market equilibrium under the optimal taxes. The third columns present the allocation under a centralizing platform. All comparisons are to the observed equilibrium.

Next, following Section 2.3, we study a policy-based solution to search frictions. A policymaker who is not able to directly affect the meeting process can still levy the optimal taxes and subsidies on ships and exporters, derived in Corollary 1 to achieve constrained efficiency. Column III of Table 2 compares the observed equilibrium to the constrained efficient outcome.⁵⁰ The increase in welfare is again about half of what is achieved in the frictionless equilibrium (8.5%). Although this gain is approximately the same as in the case of the centralizing monopolist, the distributional consequences of the two allocations are very different. Optimal policy encourages exporter entry, leading to an increase in both trade volume (almost 4%) and trade value (3.5%); realized exporter wait costs however increase by 23%. In contrast to the centralizing monopolist, now the welfare gains are redistributed to the market agents.

A burgeoning literature has studied search frictions and welfare through different types of counterfactuals. The novelty of the machinery developed here is that it allows us to explicitly compute the efficient spatial allocation, and therefore quantify the loss from frictions. Moreover, it enables us to consider the trade-off between search frictions and market power, which is a first-order issue in understanding the impact of platforms. Our findings suggest that it is important to distinguish the platform from the first-best (competitive benchmark), or the second-best (policymaker who is unable to directly affect the search environment but can still levy the optimal taxes/subsidies) as the corresponding allocations can indeed be very different.

⁵⁰To compute the constrained efficient outcome, we compute the equilibrium under the optimal taxes/subsidies derived in Corollary 1; in Appendix C.3 we provide the algorithm used, which we found to be very robust and always delivers a unique solution. Alternatively we can impose the efficient prices given in equation (31) of Corollary 2. The resulting allocation is the same.

Robustness with respect to the steady state assumption The model is estimated under the assumption of a steady state. Here we focus on short-run decisions that shipowners make (shipowners’ *weekly* decisions of where to sail). Although there is uncertainty stemming from business cycle variation, it is not clear that this uncertainty is pervasive for these decisions. These reallocation decisions are very short-lived and do not have long-lasting implications, given that trips last only a few weeks at most (the average is three weeks). Contrast this with e.g. models of firm entry or investment where a firm’s capital decision has implications that last for many years, and sometimes decades. As a result, we expect that when a shock hits, transitions to the new steady state are very fast.

Our intuition is confirmed when we recompute our welfare analysis to allow for transition dynamics from an initial steady state (the market equilibrium) to a new steady state (the constrained efficient outcome). As shown in the right panel of Figure 10 in Appendix D, convergence to the steady state is very fast, with almost all of the welfare gains from the imposition of optimal taxes/subsidies realized within two weeks of the policy change. The same is true when we look at changes in prices, as shown in the left panel of Figure 10.⁵¹ In addition, the transition is monotonic: welfare keeps rising at each step. This is important from a practical standpoint, as it suggests that we do not have to be concerned about welfare loss during the transition.⁵²

5.3 Efficiency and Policies

The results presented in Section 2.3 propose policy instruments to obtain efficiency; these can be directly implemented in markets like taxis and trucking. In addition, our results allow us to examine what is the impact of different policy inventions on efficiency, even when that is not their stated goal. To illustrate,

⁵¹In our setup, since agents are small, an individual agent’s action does not affect the aggregate state. Hence, agents keep track only of their own state (their location or route) as well as the aggregate state (which in steady state is constant). When considering transition dynamics, the economy starts at a steady state and is hit by a one-time shock; in this case, the social planner imposes the optimal tariffs and subsidies. As agents can perfectly forecast the entire transition path from the initial steady state to the final one, agents’ strategies depend only on their own state and on time t .

We briefly describe the algorithm to compute transition dynamics which is similar to that used in Benkard et al. (2021); in our case the approximation is exact and there is no additional assumption made. We first compute the initial (market equilibrium) and the new steady state (which will prove to be the efficient equilibrium). Next, suppose the economy has converged by T steps, where we set T to be a large number so that it is not binding. Begin by setting the strategy equal to the strategy under the efficient outcome. Under this strategy, compute the aggregate state for each period $t = 0, \dots, T$. Now, given this transition (which is known by the agents), solve for the new strategy and value functions using backwards induction from $T, \dots, 0$. Update the strategy and keep iterating until convergence.

⁵²As an additional robustness check, we have re-estimated the model using only the first or only the second half of our sample. The results are robust, with the welfare loss equal to 9.6% during 2012-2014, and 6.9% during 2014-2016.

we consider two recent examples that have heavily affected the shipping industry: China’s Belt and Road Initiative, and the International Maritime Organization’s 2020 rule that limits sulphur in the fuel oil used by ships (“IMO 2020”), focusing on the recent compliance of African ports. We demonstrate that both interventions substantially affect the externalities and thus the ensuing welfare and trade. This suggests that taking into account search externalities is germane for any proposed intervention. These examples also serve as an illustration of how the two externalities operate; thick market and congestion externalities depress trade by affecting the numbers of agents, while composition externalities misallocate carriers and distort trade flows.

5.3.1 China’s Belt and Road Initiative

China’s Belt and Road Initiative, henceforth BRI, is a “development strategy proposed by China that focuses on connectivity and cooperation on a trans-continental scale” (The World Bank, 2018). A key aspect of the BRI relates to considerable investments in transportation infrastructure aiming to stimulate trade and growth in Asia, Europe and Eastern Africa. To that end, China has supplied substantial funds to ports in these regions, among other projects. We explore the impact of this intervention and its interaction with the search externalities, by considering a reduction of port costs for ships and exporters in the relevant regions (China, India, Southeast Asia, Middle East, and South and West Africa).⁵³

We present the results in Table 3. Column I computes the impact of the BRI under the baseline prices, while column II under the efficient prices; comparing the two columns allows us to discern how the BRI interacts with the search externalities. As expected, the BRI leads to an increase in trade volume, and thus welfare: this is because the investment in infrastructure reduces wait costs and thus acts as a subsidy that encourages exporter entry. This increase is higher under the baseline prices, suggesting that search externalities amplify the impact of the investment in port infrastructure. Why?

The answer relates to the thick market and congestion externalities, which distort the number of searching agents and, therefore, the total number of matches formed. As shown in 2.3, thick market and

⁵³In particular, to capture the upgrading of existing ports we follow the The World Bank (2018) and assume “that the time it takes to handle the merchandise in the port decreases to the minimum time estimates in the region”. To execute this we employ the World Bank index of quality of port infrastructure (PQI). We impose that the PQI in countries where China invested, increases to the maximum PQI in the region. We then compute the associated change to the average PQI in the region, and assume that the wait costs change by the same percentage. Recall that as mentioned in 4 the PQI is highly correlated with our estimates of port costs c_i^e and c_i^s .

	Impact	Impact under efficient prices
Δ welfare (%)	8.91	8.45
Δ Trade (%)	3.89	1.87
Δ Trade value (net) (%)	1.95	0.96

Table 3: The first column of the table reports the simulated impact of China’s Belt and Road Initiative under bargained prices. The second column reports the impact of this policy under the efficient prices.

congestion externalities can be corrected via a tax/subsidy on searching exporters (or ships). The BRI acts like such a subsidy: recall that based on our estimates, the elasticity of the matching function with respect to exporters, η_i^e , is larger than the exporter bargaining coefficient in most regions (see Figure 2). Therefore, an additional exporter has a substantial positive externality on matching rates, but the baseline shipping prices are too high to appropriately encourage the entry of exporters; the planner corrects this imbalance by lowering prices and increasing exporter entry. In a similar vein, the BRI reduces port costs, incentivizing exporter entry, thus indirectly correcting partially for the thick market and congestion externalities; as a consequence, welfare, as well as trade volume and value, increase substantially more than when we do not correct for these externalities.

5.3.2 Compliance to IMO 2020

As of January 2020, a regulation known as IMO 2020, enforced new emissions standards designed to significantly curb pollution produced by the world’s ships. The mandate commanded a drop in the global sulphur limit in marine fuels. Although practically all affected countries have signed on to the mandate, compliance has been an issue, especially for African ports. Recently, the Port Management Association for West and Central Africa (PMAWCA), which is the regional umbrella organization for port authorities, has began supporting its member ports in their implementation capacity through international cooperation (e.g. through transfer of know-how on monitoring, reporting and waste management).⁵⁴

Here, we explore how the compliance of these ports to IMO 2020 interacts with the search externalities. We do so by considering an increase in the ship cost of traveling to West African ports. We present the

⁵⁴See for example here: <https://euroshore.com/blog/news/2021-02-23/strengthening-marpol-implementation-african-ports>

results in Table 4. Column I computes the impact of the policy under the baseline prices, while column II under the efficient prices. Since we do not consider the environmental impact here, the regulation leads to a welfare reduction, as it increases ship costs. This decline is smaller in the observed equilibrium than under the efficient prices, suggesting that this policy too indirectly corrects partially for the search externalities. Indeed, the environmental regulation acts as a tax to African ports as destinations, and as such it directly impacts the composition externalities.

Why does the planner want to tax African ports? As discussed in Section 2.3, composition externalities distort the allocation of ships over space. Generally speaking, the planner subsidizes “socially valuable” destinations to correct composition externalities. This involves a few different characteristics of destinations. First, the planner subsidizes routes ij with high w_{ij} . Note that the Nash bargained price depends positively on w_{ij} , so that high value exporters pay more to transport their goods. This is because search frictions create a hold-up problem, whereby ships “take advantage” of high value exporters and charge them a higher price, which in turn inefficiently reduces entry of high value exporters. To counteract this, the planner precisely promotes trips with high w_{ij} . West Africa however has the second lowest value imports.

In addition, for every route ij , the planner also cares about the value of the ship’s subsequent trips, which the exporter does not fully internalize. Therefore, the planner subsidizes routes ij whose destination j is itself a high-value exporter. Naturally, the recursion continues, as the planner optimizes the entire dynamic path of ships’ travels. The planner therefore also subsidizes routes ij , for which location j ’s exports are directed to locations that themselves have high value exports, and so on. Indeed, the planner subsidizes regions that are “central”; i.e. regions with high value imports and which further export to high value locations, and so on. Consistent with this, the planner taxes West African ports because not only do they have low value imports, but they are also provide poor reloading options to ships: indeed, West Africa predominantly exports to unattractive destinations (e.g. China and the Middle East).

6 Conclusion

Transport markets are the engine of economic activity. Yet, little is known about their efficiency properties. In this paper we contribute to a nascent but growing literature studying optimal transport networks, by

	Impact	Impact under efficient prices
Δ welfare (%)	-0.29	-0.33
Δ Trade (%)	-0.34	-0.5
Δ Trade value (net) (%)	-0.21	-0.34

Table 4: The first column of the table reports the simulated impact on world trade of IMO 2020 enforcement in West African countries under bargained prices. The second column reports the impact of this policy under the efficient prices.

focusing on decentralized transport markets, such as taxis, trucks and bulk shipping. We show that in these markets, search frictions distort the transportation network and the dynamic allocation of carriers over space, and we derive explicit and intuitive conditions for efficiency, which lead naturally to efficient pricing rules. The latter imply that destination-based pricing is essential to attain efficiency. Then, using data from dry bulk shipping, we demonstrate that search frictions lead to a sizeable social loss and substantial misallocation of ships over space. Optimal policy can eliminate about half of the welfare loss. Can a centralizing platform, often arising as a market-based solution to search frictions, do better? Interestingly, the answer is no; although the platform eradicates frictions, it exerts market power thus eroding the welfare gains. Finally, we use two recent interventions in the industry (China’s Belt and Road Initiative and the environmental initiative IMO 2020) to demonstrate that taking into account the efficiency properties of transport markets is germane for any proposed policy.

References

- ACEMOGLU, D. (2001): “Good Jobs versus Bad Jobs,” *Journal of Labor Economics*, 19, 1–21.
- ACEMOGLU, D. AND R. SHIMER (1999): “Holdups and Efficiency with Search Frictions,” *International Economic Review*, 40, 827–849.
- ALBRECHT, J., P. A. GAUTIER, AND S. VROMAN (2016): “Directed search in the housing market,” *Review of Economic Dynamics*, 19, 218–231.
- ALLEN, J., R. CLARK, AND J.-F. HOUDE (2018): “Search frictions and market power in negotiated price markets,” *forthcoming, Journal of Political Economy*.

- ALLEN, T. (2014): “Information Frictions in Trade,” *Econometrica*, 82, 2041–2083.
- ALLEN, T. AND C. ARKOLAKIS (2021): “The Welfare Effects of Transportation Infrastructure Improvements,” *mimeo*, Yale University.
- ANTRAS, P. AND A. COSTINOT (2011): “Intermediated Trade,” *Quarterly Journal of Economics*, 126, 1319–1374.
- ARCIDIACONO, P. AND R. MILLER (2011): “Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity,” *Econometrica*, 79, 1823–1867.
- ASPLUND, E. (1968): “Fréchet Differentiability of Convex Functions,” *Acta Mathematica*, 121, 31–47.
- ASTURIAS, J. (2020): “Endogenous Transportation Costs,” *European Economic Review*, 123.
- ATKESON, A. G., A. L. EISFELDT, AND P.-O. WEILL (2015): “Entry and Exit in OTC Derivatives Markets,” *Econometrica*, 83, 2231–2292.
- BELLONI, A., D. CHEN, V. CHERNOZHUKOV, AND C. HANSEN (2012): “Sparse Models and Methods for Optimal Instruments with an Application to Eminent Domain,” *Econometrica*, 80, 2369–2429.
- BENKARD, C. L., P. JEZIORSKI, AND G. Y. WEINTRAUB (2021): “Transitional Market Dynamics in Complex Environments,” *mimeo*, Stanford University.
- BIAN, B. (2019): “Search Frictions, Network Effects and Spatial Competition: Taxis versus Uber,” *mimeo*, Shanghai University of Finance and Economics.
- BILAL, A. (2022): “The Geography of Unemployment,” *mimeo*, Harvard University.
- BRANCACCIO, G., M. KALOUPSTIDI, AND T. PAPAGEORGIOU (2018): “Geography, Search Frictions and Endogenous Trade Costs,” *NBER Working Paper*, 23581.
- (2020a): “Geography, Transportation, and Endogenous Trade Costs,” *Econometrica*, 88, 657–691.
- (2020b): “A Guide to Estimating Matching Functions in Spatial Models,” *International Journal of Industrial Organization*, (Special Issue EARIE 2018), 70.
- BRANCACCIO, G., D. LI, AND N. SCHUERHOFF (2020c): “Learning by Trading: The Case of the US Market for Municipal Bonds,” *mimeo*, Cornell University.
- BROOKS, L., N. GENDRON-CARRIER, AND G. RUA (2021): “The Local Impact of Containerization,” *Journal of Urban Economics*, 126.

- BUCHHOLZ, N. (2021): “Spatial Equilibrium, Search Frictions and Dynamic Efficiency in the Taxi Industry,” *forthcoming, The Review of Economic Studies*.
- CAO, G., G. Z. JIN, X. WENG, AND L.-A. ZHOU (2021): “Market Expanding or Market Stealing? Competition with Network Effects in Bike-Sharing,” *The RAND Journal of Economics*, 52, 778–814.
- CASTILLO, J. C. (2020): “Who Benefits from Surge Pricing?” *mimeo, University of Pennsylvania*.
- CHIONG, K. X., A. GALICHON, AND M. SHUM (2016): “Duality in dynamic discrete-choice models,” *Quantitative Economics*, 7, 83–115.
- COLES, M. (1994): “Understanding the matching function: The role of newspapers and job agencies,” Tech. rep.
- COLES, M. G. AND E. SMITH (1998): “Marketplaces and Matching,” *International Economic Review*, 239–254.
- COSAR, A. K. AND B. DEMIR (2018): “Shipping inside the Box: Containerization and Trade,” *Journal of International Economics*, 114, 331–345.
- DAVIS, S. J. (2001): “The Quality Distribution of Jobs and the Structure of Wages in Search Equilibrium,” *NBER Working Paper*, 8434.
- DECREUSE, B. (2008): “Choosy Search and the Mismatch of Talents,” *International Economic Review*, 49, 1067–1089.
- DEMSETZ, H. (1968): “The Cost of Transacting,” *The Quarterly Journal of Economics*, 82, 33–53.
- DIAMOND, P. A. (1982): “Wage Determination and Efficiency in Search Equilibrium,” *Review of Economic Studies*, 49, 217–227.
- DUCRUET, C., R. JUHASZ, D. K. NAGY, AND C. STEINWENDER (2022): “All Aboard: The Effects of Port Development,” *mimeo, Columbia University*.
- DUFFIE, D., N. GÂRLEANU, AND L. H. PEDERSEN (2005): “Over-the-counter markets,” *Econometrica*, 73, 1815–1847.
- DUPAS, P., J. ROBINSON, AND S. SAAVEDRA (2020): “The daily grind: Cash needs and labor supply,” *Journal of Economic Behavior & Organization*, 177, 399–414.
- DURANTON, G. AND D. PUGA (2004): “Micro-foundations of urban agglomeration economies,” in *Handbook of Regional and Urban Economics*, ed. by J. V. Henderson and J. F. Thisse, Elsevier, vol. 4.

- EATON, J., D. JINKINS, J. TYBOUT, AND D. XU (2022): “Two-Sided Search in International Markets,” *NBER Working Paper*, 29684.
- ERICSON, R. AND A. PAKES (1995): “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *The Review of Economic Studies*, 62, 53–82.
- FAJGELBAUM, P. D. AND E. SCHAAL (2020): “Optimal Transport Networks in Spatial Equilibrium,” *Econometrica*, 88, 1411–1452.
- FARBOODI, M., G. JAROSCH, AND R. SHIMER (2021): “The Emergence of Market Structure,” *forthcoming, The Review of Economic Studies*.
- FRECHETTE, G. R., A. LIZZERI, AND T. SALZ (2019): “Frictions in a Competitive, Regulated Market Evidence from Taxis,” *American Economic Review*, 109, 2954–2992.
- GAVAZZA, A. (2011): “The Role of Trading Frictions in Real Asset Markets,” *American Economic Review*, 101, 1106–1143.
- (2016): “An Empirical Equilibrium Model Of a Decentralized Asset Market,” *Econometrica*, 84, 1755–1798.
- GHILI, S. AND V. KUMAR (2021): “Spatial Distribution of Supply and the Role of Market Thickness: Theory and Evidence from Ride Sharing,” *mimeo, Yale University*.
- HOLMES, T. J. AND E. SINGER (2018): “Indivisibilities in Distribution,” *NBER Working Paper*, 24525.
- HOPENHAYN, H. A. (1992): “Entry, Exit, and Firm Dynamics in Long Run Equilibrium,” *Econometrica*, 60, 1127–1150.
- HOSIOS, A. J. (1990): “On the Efficiency of Matching and Related Models of Search and Unemployment,” *The Review of Economic Studies*, 57, 279–298.
- HOTZ, V. J. AND R. A. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60, 497–529.
- HUMMELS, D. AND S. SKIBA (2004): “Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture,” *Journal of Political Economy*, 112, 1384–1402.
- IMBENS, G. W. AND W. K. NEWEY (2009): “Identification and Estimation of Triangular Simultaneous Equations Models Without Additivity,” *Econometrica*, 77, 1481–1512.

- KALOUPSIDI, M. (2014): “Time to Build and Fluctuations in Bulk Shipping,” *American Economic Review*, 104, 564–608.
- KENISTON, D. (2011): “Bargaining and welfare: A dynamic structural analysis of the autorickshaw market,” *Unpublished manuscript, Yale University*.
- KOOPMANS, T. C. (1949): “Optimum Utilization of the Transportation System,” *Econometrica*, 17, 136–146.
- KREINDLER, G. (2020): “Peak-Hour Road Congestion Pricing: Experimental Evidence and Equilibrium Implications,” *Mimeo, Harvard University*.
- KROLIKOWSKI, P. M. AND A. H. MCCALLUM (2021): “Goods-market frictions and international trade,” *Journal of International Economics*, 129, 103411.
- LAGOS, R. (2000): “An Alternative Approach to Search Frictions,” *Journal of Political Economy*, 108, 851–873.
- (2003): “An Analysis of the Market for Taxicab Rides in New York City,” *International Economic Review*, 44, 423–434.
- LANGE, F. AND T. PAPAGEORGIOU (2020): “Beyond Cobb-Douglas: Flexibly Estimating Matching Functions with Unobserved Matching Efficiency,” *NBER Working Paper*, 26972.
- LEE, J., S.-J. WEI, AND J. XU (2020): “The Welfare Cost of the Current Account Imbalance: A "Clean" Effect,” *NBER Working Paper*, 27276.
- LIU, T. X., Z. WAN, AND C. YANG (2019): “The Efficiency of A Dynamic Decentralized Two-sided Matching Market,” *mimeo, University of Rochester*.
- LUENBERGER, D. G. (1997): *Optimization by Vector Space Methods*, Cambridge, MA: John Wiley & Sons.
- MA, H., F. FANG, AND D. C. PARKES (2022): “Spatio-Temporal Pricing for Ridesharing Platforms,” *Operations Research*, 70, 1025–1041.
- MATZKIN, R. L. (2003): “Nonparametric Estimation of Nonadditive Random Functions,” *Econometrica*, 71, 1339–1375.
- MOEN, E. R. (1997): “Competitive Search Equilibrium,” *Journal of Political Economy*, 105, 385–411.
- MORTENSEN, D. T. (1982): “The Matching Process as a Noncooperative Bargaining Game,” in *The Economics of Information and Uncertainty*, ed. by J. J. McCall, University of Chicago Press.

- (2003): *Wage Dispersion: Why Are Similar Workers Paid Differently?*, Cambridge, MA: MIT Press.
- OSTROVSKY, M. AND M. SCHWARZ (2018): “Carpooling and the Economics of Self-Driving Cars,” *Mimeo, Stanford University*, 25487.
- PANAYIDES, P. M. (2016): *Principles of Chartering*, CreateSpace Independent Publishing Platform, 2nd ed.
- PETRONGOLO, B. AND C. A. PISSARIDES (2001): “Looking Into the Black Box: A Survey of the Matching Function,” *Journal of Economic Literature*, 39, 390–431.
- PISSARIDES, C. A. (1985): “Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages,” *American Economic Review*, 75, 676–690.
- PLAKANTONAKI, C., ed. (2010): *An Introduction to the World of Shipping Through the Eyes of a Ship Management Company*, Thenamaris (Ships Management) Inc.
- ROSAIA, N. (2020a): “Competing Platforms and Transport Equilibrium: Evidence from New York City,” *mimeo, Harvard University*.
- (2020b): “Duality and Estimation of Undiscounted Markov Decision Processes,” *mimeo, Harvard University*.
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999–1033.
- SHAPIRO, M. H. (2018): “Density of Demand and the Benefit of Uber,” *mimeo, Singapore Management University*.
- SHIMER, R. AND L. SMITH (2001): “Matching, Search, and Heterogeneity,” *The B.E. Journal of Macroeconomics*, 1, 153460131010.
- STOPFORD, M. (2009): *Maritime Economics*, London and New York: Routledge, 3rd ed.
- THE WORLD BANK (2018): “How Much Will the Belt and Road Initiative Reduce Trade Costs?” Tech. rep., The World Bank, Washington, DC.
- WONG, W. F. (2020): “The Round Trip Effect: Endogenous Transport Costs and International Trade,” *forthcoming, American Economic Journal: Applied Economics*.
- WTO (2015): “International Trade Statistics,” Tech. rep., World Trade Organization, Geneva, Switzerland.
- YANG, R. (2022): “(Don’t) Take Me Home: Home Bias and the Effect of Self-Driving Trucks on Interstate Trade,” *mimeo, Harvard University*.

Online Appendix

A Proofs

A.1 Preliminaries

A.1.1 Sequence spaces and optimization

This Section introduces notation and reviews basic concepts related to infinite-dimensional optimization that will be used later in the proofs.

We work with the spaces of bounded sequences and of β -summable sequences, defined by $l^\infty \equiv \{\mathbf{a} \in \mathbb{R}^{\mathbb{N}} : \sup_{t \geq 0} |a(t)| < \infty\}$ and $l^{1,\beta} \equiv \{\mathbf{a} \in \mathbb{R}^{\mathbb{N}} : \sum_{t=0}^{\infty} \beta^t |a(t)| < \infty\}$, respectively. In particular, we consider sequences of vectors whose coordinates are either bounded, or β -summable. That is, vector-valued sequences in $(l^\infty)^N$ and $(l^{1,\beta})^N$ for some $N \in \mathbb{N}$, equipped with the usual topologies.

First, we introduce some standard notation. For generic sequences $\mathbf{a}, \mathbf{b} \in \mathbb{R}^{\mathbb{N}}$, we denote by \mathbf{ab} the product sequence defined by $(\mathbf{ab})(t) = a(t)b(t) \forall t$, we write $\mathbf{a} \geq \mathbf{b}$ (respectively $\mathbf{a} > \mathbf{b}$) if $a(t) \geq b(t)$ (respectively $a(t) > b(t)$) $\forall t$, and we denote by \mathbf{a}^{+1} the sequence obtained by shifting \mathbf{a} one period forward, that is, the sequence defined by $a^{+1}(t) = a(t+1)$, for all t . Moreover, we denote by $\mathbf{1}$ and $\mathbf{0}$ the sequences defined by $1(t) = 1$ and $0(t) = 0$, for all t . For sequences $\mathbf{a} \in (l^\infty)^N$ and $\mathbf{b} \in (l^{1,\beta})^N$, we denote their discounted inner product by $\mathbf{a} \cdot \mathbf{b} = \sum_{t=0}^{\infty} \beta^t \sum_n a_n(t) b_n(t)$.

Second, we introduce some standard terminology. If $f : (l^\infty)^N \rightarrow [-\infty, +\infty)$ is a function mapping bounded vector sequences into the extended real line, we say that: (i) its domain, denoted by $\text{Dom} f$, is the set of all points at which f is finite, that is, $\text{Dom} f = \{\mathbf{a} \in (l^\infty)^N : f(\mathbf{a}) > -\infty\}$; (ii) $\nabla \in (l^{1,\beta})^N$ is a super-gradient of f at $\mathbf{a} \in \text{Dom} f$ if $f(\mathbf{a}') - f(\mathbf{a}) \leq \nabla \cdot [\mathbf{a}' - \mathbf{a}]$ for every $\mathbf{a}' \in \text{Dom} f$; (iii) the Gateaux (directional) derivative of f at a point \mathbf{a} in the interior of $\text{Dom} f$, if it exists, is the map $\partial f(\mathbf{a})$ defined by $\partial f(\mathbf{a}; \mathbf{a}') = \lim_{\alpha \rightarrow 0} [f(\mathbf{a} + \alpha \mathbf{a}') - f(\mathbf{a})] / \alpha$ for every $\mathbf{a}' \in (l^\infty)^N$; (iv) f is Gateaux differentiable at \mathbf{a} in the interior of $\text{Dom} f$ if $\partial f(\mathbf{a})$ exists.

Finally, we use the following standard results. We provide proofs here for completeness.

Lemma 1. (i) If f is Gateaux differentiable at \mathbf{a} and ∇ is a super-gradient of f at \mathbf{a} , then $\partial f(\mathbf{a}; \mathbf{a}') = \nabla \cdot \mathbf{a}'$ for every $\mathbf{a}' \in (l^\infty)^N$. (ii) If $D \subseteq (l^\infty)^N$ and ∇ is a super-gradient of f at $\mathbf{a} \in D$, a sufficient

condition for \mathbf{a} to maximize f over D is that $\nabla \cdot (\mathbf{a}' - \mathbf{a}) \leq \mathbf{0}$ for every $\mathbf{a}' \in D$. (iii) If D is a convex subset of $(l^\infty)^N$, f is Gateaux differentiable at $\mathbf{a} \in D$, and ∇ is a super-gradient of f at \mathbf{a} , a necessary and sufficient condition for \mathbf{a} to maximize f over D is that $\nabla \cdot (\mathbf{a}' - \mathbf{a}) \leq \mathbf{0}$ for every $\mathbf{a}' \in D$

Proof. To prove (i), take $\mathbf{a}' \in (l^\infty)^N$. By the definition of super-gradients, for every α we have $f(\mathbf{a} + \alpha\mathbf{a}') - f(\mathbf{a}) \leq \alpha \nabla \cdot \mathbf{a}'$. This implies $\lim_{\alpha \rightarrow 0^+} [f(\mathbf{a} + \alpha\mathbf{a}') - f(\mathbf{a})]/\alpha \leq \nabla \cdot \mathbf{a}'$ and $\lim_{\alpha \rightarrow 0^-} [f(\mathbf{a} + \alpha\mathbf{a}') - f(\mathbf{a})]/\alpha \geq \nabla \cdot \mathbf{a}'$. Hence, by the definition of Gateaux derivative, we must have $\partial f(\mathbf{a}; \mathbf{a}') = \nabla \cdot \mathbf{a}'$. (ii) follows by definition of super-gradient. To prove (iii), note that sufficiency follows from (ii). To prove necessity, suppose that $\nabla \cdot (\mathbf{a}' - \mathbf{a}) > \mathbf{0}$ for some $\mathbf{a}' \in D$. Take $\alpha, \epsilon > 0$ small enough so that, $\epsilon < \nabla \cdot \mathbf{a}'$ and $|[f(\mathbf{a} + \alpha(\mathbf{a}' - \mathbf{a})) - f(\mathbf{a})]/\alpha - \partial f(\mathbf{a}; \mathbf{a}' - \mathbf{a})| < \epsilon$. Since $\partial f(\mathbf{a}; \mathbf{a}' - \mathbf{a}) = \nabla \cdot (\mathbf{a}' - \mathbf{a})$ by (i), this implies $0 < \alpha[\nabla \cdot (\mathbf{a}' - \mathbf{a}) - \epsilon] < f(\mathbf{a} + \alpha(\mathbf{a}' - \mathbf{a})) - f(\mathbf{a})$. Since D is convex, $\mathbf{a} + \alpha(\mathbf{a}' - \mathbf{a}) \in D$, hence \mathbf{a} does not maximize f over D . \square

Lemma 2. Let $(f(\cdot; t))_{t=0}^\infty$ be a sequence of concave and differentiable functions $f(\cdot; t) : \mathbb{R}^N \rightarrow \mathbb{R}$, and $g(\mathbf{a}') = \sum_{t=0}^\infty \beta^t f(\mathbf{a}'(t); t)$ for all $\mathbf{a}' \in (l^\infty)^N$. Take $\mathbf{a} \in (l^\infty)^N$, and define $\nabla_n(t) = df(\mathbf{a}(t); t)/da_n(t) \forall n, t$. Then $\nabla \in (l^{1,\beta})^N$, and ∇ is a super-gradient of g at \mathbf{a} .

Proof. Take $\mathbf{a}' \in (l^\infty)^N$. By concavity, for every $t \geq 0$ we have $f(\mathbf{a}'(t); t) - f(\mathbf{a}(t); t) \leq \sum_n \nabla_n(t)[a'_n(t) - a_n(t)] \leq f(\mathbf{a}(t); t) - f(2\mathbf{a}(t) - \mathbf{a}'(t); t)$. Since both the left and right hand sides belong to $l^{1,\beta}$, it follows that $\sum_{t=0}^T \beta^t \sum_n \nabla_n(t)[a'_n(t) - a_n(t)]$ converges to $\nabla \cdot (\mathbf{a}' - \mathbf{a})$ as $T \rightarrow \infty$, and $g(\mathbf{a}') - g(\mathbf{a}) \leq \nabla \cdot [\mathbf{a}' - \mathbf{a}]$. Moreover, for every n , taking $a'_m(t) = a_m(t) + \text{sign} \nabla_m(t)$, if $m = n$ and $a'_m(t) = 0$ otherwise for all t , $\sum_{t=0}^T \beta^t |\nabla_n(t)|$ converges as $T \rightarrow \infty$, hence $\nabla \in (l^{1,\beta})^N$.

A.1.2 Customers' entry costs and carriers' idiosyncratic shocks

Recall that $K(n)$ and $\mathcal{E}(b)$ were informally defined in Section 2.3 as the sum of customers' random entry costs when their optimal choices yield inflow of new entrants n_{ij} , and the sum of carriers' random shocks when their optimal choices yield the ballasting flows b_{ij} , respectively. This Section formally describes these functions, and states two related results used in the proofs. For every n, b , define the corresponding entry and ballasting choice probabilities \mathbb{P}^e and \mathbb{P}^b by $\mathbb{P}_{ij}^e = n_{ij}/N_i$, for all $i, j \neq i$ and $\mathbb{P}_{ij}^b = b_{ij}/\sum_k b_{ik}$, for all ij .

Customers' entry costs Starting from customers' entry, recall that the optimal choice probabilities are a function of the vector $V^e = (V_{ij}^e)_{i,j \neq i}$ of search values, given by

$$\mathbb{P}_{ij}^e(V^e) \equiv \Pr[V_{ij}^e - \kappa_{ij} \geq V_{ik}^e - \kappa_{ik}, \forall k]$$

for all $j \neq i$, where we normalize $V_{ii}^e \equiv 0$. It can be shown that function $\mathbb{P}^e(\cdot)$ is invertible.⁵⁵ That is, for every vector \mathbb{P}^e of entry probabilities, there exists a unique vector of valuations $V^e(\mathbb{P}^e) \in \mathbb{R}^{I(I-1)}$ such that $\mathbb{P}^e(V^e(\mathbb{P}^e)) = \mathbb{P}^e$.

If n_{ij} new customers enter on every $i, j \neq i$, the search values must be such that n is the optimal vector of entry flows, i.e. they must satisfy $V_{ij}^e = V_{ij}^e(\mathbb{P}^e)$, for all $i, j \neq i$. Hence, when customers' entry decisions are described by $n = (n_{ij})_{i,j \neq i} \in \mathbb{R}^{I(I-1)}$, the sum of customers' random entry costs is given by

$$K(n) \equiv \sum_{ij} n_{ij} \mathbb{E}[\kappa_{ij} | V_{ij}^e(\mathbb{P}^e) - \kappa_{ij} \geq V_{ik}^e(\mathbb{P}^e) - \kappa_{ik} \forall k] \quad (34)$$

where we denote $V_{ii}^e(\mathbb{P}^e) \equiv 0$ and $n_{ii} \equiv N_i - \sum_{j \neq i} n_{ij}$. In words, the total entry costs are obtained as the sum of the expected customers' entry costs conditional on entering in every route, weighted by the number of customers entering.⁵⁶

This function has been studied extensively in the discrete choice literature. Here, we state a well-known result used in the proofs.⁵⁷

Lemma 3. *$K(n)$ is strictly convex and differentiable. Moreover, for every pair of vectors of search value functions $V^e \in \mathbb{R}^{I(I-1)}$ and entry flows $n \in \mathbb{R}^{I(I-1)}$, the following are equivalent: (i) $n_{ij} = N_i \Pr[V_{ij}^e - \kappa_{ij} \geq V_{ik}^e - \kappa_{ik} \forall k]$ for every $i, j \neq i$; (ii) For every $i, j \neq i$,*

$$V_{ij}^e = dK(n)/dn_{ij} \quad (35)$$

In other words, customers' inverse entry curve coincides with the marginal entry cost. Hence, intu-

⁵⁵For instance, this follows from the results of Hotz and Miller (1993) and the restriction that V_{ii}^e is equal to 0 (in general, in discrete choice models, choice probabilities identify choice-specific payoffs up to a constant).

⁵⁶In particular, note that we allow customers to pay an "entry cost" κ_{ii} if they choose not to enter, so that this framework nests our empirical application, where $\kappa_{ij} = k_{ij} + \epsilon_{ij}$, $\kappa_{ii} = \epsilon_{ii}$ and ϵ is a logit shock.

⁵⁷This follows, from instance, by Chiong, Galichon, and Shum (2016, Theorem 1), after normalizing the payoff of one option to zero.

itively, the total entry costs at n can be obtained by integrating the inverse entry curve below n . This can be seen as a generalization of the classic definition of welfare in economies with a single good, where total welfare can be defined as the area below the inverse demand curve, and prices coincide with the marginal welfare.

Carriers' random ballasting shocks The definition of $\mathcal{E}(\cdot)$ is similar. Again, recall that the optimal relocation choice probabilities are a function of the vector $V^s = (V_{ij}^s)_{ij}$ of traveling values, given by $\mathbb{P}_{ij}^b(V_{ij}^s) = \Pr[V_{ij}^s + \epsilon_{ij} \geq V_{ik}^e + \epsilon_{ik} \ \forall k] \ \forall j$. Hence carriers' payoff from shocks corresponding to relocation flows b can be defined by

$$\mathcal{E}(b) \equiv \sum_{ij} b_{ij} \mathbb{E}[\epsilon_{ij} | V_{ij}^s(\mathbb{P}^b) + \epsilon_{ij} \geq V_{ik}^e(\mathbb{P}^b) + \epsilon_{ik}, \ \forall k]$$

where $V^s(\cdot) : [0, 1]^{I^2} \rightarrow \mathbb{R}^{I^2}$ is any function such that $\mathbb{P}^b(V^s(\mathbb{P}^b)) = \mathbb{P}^b$ for every system of choice probabilities \mathbb{P}^b .⁵⁸

Function $\mathcal{E}(\cdot)$ is also not new to the discrete choice literature. One result that we used in the proofs is the following.⁵⁹

Lemma 4. *$\mathcal{E}(b)$ is strictly concave and differentiable. Moreover, for every system of value functions $U^s \in \mathbb{R}^I$, $V^s \in \mathbb{R}^{I^2}$ and relocation choices $b \in \mathbb{R}_+^{I^2}$, the following are equivalent: (i) $U_i^s = E \max_j (V_{ij}^s + \epsilon_j)$ and $b_{ij} / \sum_k b_{ik} = \Pr[V_{ij}^s + \epsilon_j \geq V_{ik}^s + \epsilon_k \ \forall k]$ for every ij ; (ii) For every ij ,*

$$U_i^s = V_{ij}^s + \frac{d\mathcal{E}(b)}{db_{ij}} \tag{36}$$

Hence, similarly to the entry costs, the gradient of $\mathcal{E}(b)$ pins down the differences between continuation values across destinations at each origin i . However, differently from the entry costs, since we do not normalize the continuation value of one option to 0, this does not pin down the vector $(V_{ij}^s)_j$ uniquely, but only up to a constant, U_i^s .

⁵⁸Note that in this case we do not normalize the utility of an option to zero, hence the inverse $V^s(\mathbb{P}^b)$ is not uniquely defined, as choice probabilities identify choice-specific payoffs only up to a constant. However this is irrelevant for defining $\mathcal{E}(\cdot)$. Indeed, notice that all that matters for defining $\mathcal{E}(\cdot)$ are the differences $V_{ij}^s(\mathbb{P}^b) - V_{ik}^s(\mathbb{P}^b)$, which are pinned down uniquely by \mathbb{P}^b .

⁵⁹For instance, this follows from the results in Rosaia (2020b).

Logit specifications Recall that, in our empirical specification, potential customers at i solve $\max_j (V_{ij}^e - k_{ij} + \epsilon_{ij})$, where ϵ is i.i.d. according to a Type I extreme value distribution, and $V_{ii}^e = k_{ii}^e \equiv 0$. Hence the optimal entry probabilities are given by $\mathbb{P}_{ij}^e = \exp V_{ij}^e / [1 + \sum_{k \neq i} \exp V_{ik}^e]$, for all $i, j \neq i$.

In this case, it can be shown that K takes the closed form expression

$$K(n) = \sum_{i,j \neq i} n_{ij} \log n_{ij} + \sum_i (N_i - \sum_{j \neq i} n_{ij}) \log (N_i - \sum_{k \neq i} n_{ik}) - \sum_i N_i \log N_i \quad (37)$$

Note that $dK(n)/dn_{ij} = \log n_{ij} / (N_i - \sum_{k \neq i} n_{ik})$. After some algebra, it is easy to see that

$$V_{ij}^e = \log \frac{n_{ij}}{N_i - \sum_{k \neq i} n_{ik}} \Leftrightarrow n_{ij} = N_i \frac{\exp V_{ij}^e}{1 + \sum_{k \neq i} \exp V_{ik}^e}$$

as Lemma 3 indicates.

Similarly, in our empirical specification, carriers unmatched at i solve $\max_j (V_{ij}^s + \epsilon_j)$, where ϵ is i.i.d. according to a Type I extreme value distribution with scale parameter σ . Hence the optimal choice probabilities are given by $\mathbb{P}_{ij}^b = \exp(V_{ij}^s/\sigma) / \sum_k \exp(V_{ik}^s/\sigma)$, and the inclusive value of the relocation choice by $U_i^s = \sigma \log \sum_j \exp(V_{ij}^s/\sigma)$.

In this case, it can be shown that \mathcal{E} takes the closed form expression

$$\mathcal{E}(b) = -\sigma \sum_{i,j} b_{ij} \log b_{ij} + \sigma \sum_{i,j} b_{ij} \log \sum_k b_{ik} \quad (38)$$

Note that $d\mathcal{E}(b)/db_{ij} = -\sigma \log b_{ij} / \sum_k b_{ik}$. After some algebra, it is easy to see that

$$U_i^s = V_{ij}^s - \sigma \log \frac{b_{ij}}{\sum_k b_{ik}} \Leftrightarrow U_i^s = \sigma \log \sum_j \exp(V_{ij}^s/\sigma) \text{ and } \frac{b_{ij}}{\sum_k b_{ik}} = \frac{\exp(V_{ij}^s/\sigma)}{\sum_k \exp(V_{ik}^s/\sigma)}$$

as Lemma 3 indicates.

A.2 Proof of Proposition 1

Let $\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}'$ denote generic (not necessarily feasible) allocations such that $\mathbf{s}'_i, \mathbf{e}'_i, \mathbf{G}'_{ij}, \mathbf{q}'_{ij}, \mathbf{n}'_{ij}, \mathbf{b}'_{ij} \in l^\infty$, for all ij . For every allocation $\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}'$ and every ij , we denote by $\tilde{\mathbf{s}}'_{ij}$ the sequence of carrier state variables associated with $\mathbf{q}', \mathbf{n}', \mathbf{b}'$, and by $\tilde{\mathbf{e}}'_{ij}$ the sequence of customer state variables associated

with $\mathbf{e}', \mathbf{G}', \mathbf{q}'$. That is, for all t , $\tilde{s}'_{ij}(t)$ and $\tilde{e}'_{ij}(t)$ denote the measures of carriers traveling on ij at the beginning of time t and customers searching on ij at the beginning of time t - i.e. those customers who remained unmatched at $t - 1$. Formally, $\tilde{s}'(0) = \tilde{s}(0)$, $\tilde{e}'(0) = \tilde{e}(0)$, $\tilde{e}'_{ij}{}^{+1} = \mathbf{e}'_i \mathbf{G}'_{ij} - \mathbf{q}'_{ij}$ and $\tilde{s}'_{ij}{}^{+1} = (1 - d_{ij})\tilde{\mathbf{x}}'_{ij} + \mathbf{q}'_{ij} + \mathbf{b}'_{ij}$.

Also, $\mathbf{V}'_i, \mathbf{U}'_i, \mathbf{\Delta}'_{ij}, \mathbf{V}'_{ij}$ denote generic sequences of feasible discounted Lagrange multipliers associated with constraints (9)-(11) and (13), respectively. That is, sequences satisfying $\mathbf{\Delta}'_{ij} \geq \mathbf{0}$ and $\mathbf{V}'_i, \mathbf{V}'_{ij}, \mathbf{U}'_i, \mathbf{\Delta}'_{ij} \in l^{1,\beta}$, for all ij .

With this notation, the Lagrangian of Problem 15 given $\mathbf{s}', \mathbf{e}', \mathbf{G}'$ can be written as

$$\begin{aligned} L(\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}', \mathbf{V}'^s, \mathbf{V}'^e, \mathbf{U}'^s, \mathbf{\Delta}') &= u(\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}') + \sum_i \mathbf{V}'^s_i \cdot (\sum_j d_{ji} \tilde{s}'_{ji} - \mathbf{s}'_i) \quad (39) \\ &+ \sum_i \mathbf{U}'^s_i \cdot [\mathbf{s}'_i - \sum_j (\mathbf{q}'_{ij} + \mathbf{b}'_{ij})] + \sum_{ij} \mathbf{V}'^e_{ij} \cdot (\tilde{e}'_{ij} + \mathbf{n}'_{ij} - \mathbf{e}'_i \mathbf{G}'_{ij}) + \sum_{ij} \mathbf{\Delta}'_{ij} \cdot (\mathbf{m}'_i \mathbf{G}'_{ij} - \mathbf{q}'_{ij}) \end{aligned}$$

L is well defined - that is, the infinite discounted sums converge and are finite - whenever the multipliers are feasible (since $\mathbf{V}'_i, \mathbf{V}'_{ij}, \mathbf{U}'_i, \mathbf{\Delta}'_{ij} \in l^{1,\beta}$, for all ij).

With this notation, to prove Proposition 1, let $\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}, \mathbf{p}$ be an equilibrium, $\mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s$ and $\mathbf{\Delta}^s, \mathbf{\Delta}^e$ denote the corresponding sequences of value functions and surplus shares, and $\mathbf{\Delta} = \mathbf{\Delta}^s + \mathbf{\Delta}^e$ denote the sequence of joint match surpluses. Optimality of customers' entry choices (condition 8) and Lemma 3 imply that

$$V^e_{ij} = dK(n)/dn_{ij} \quad (40)$$

Moreover, optimality of carriers' relocation decisions (conditions 3 and 5) and Lemma 4 imply that

$$U^s_i(t) = d\mathcal{E}(b(t))/db_{ij}(t) + V^s_{ij}(t) \quad (41)$$

Finally, by the definition of $\mathbf{\Delta}^s$ and $\mathbf{\Delta}^e$, and the equilibrium conditions (4) and (7), we have

$$\begin{aligned} \Delta_{ij}(t) &\geq 0, \text{ with equality if } q_{ij}(t) < m_i(t)G_{ij}(t) \text{ and} \quad (42) \\ w_{ij}(q(t)) + V^s_{ij}(t) - U^s_i(t) - \beta V^e_{ij}(t+1) - \Delta_{ij}(t) &\leq 0, \text{ with equality if } q_{ij}(t) > 0 \end{aligned}$$

Note also that, substituting $\tilde{s}'_{ij}(t) = (1 - d_{ij})^t \tilde{s}'_{ij}(0) + \sum_{s=0}^{t-1} (1 - d_{ij})^{t-s-1} [q'_{ij}(s) + b'_{ij}(s)]$ and $\tilde{e}'_{ij}(t+1) =$

$e'_i(t)G'_{ij}(t) - q'_{ij}(t)$, using $V_{ij}^s(t) = \sum_{s=0}^{\infty} (1 - d_{ij})^s \beta^s [d_{ij} \beta V_{ij}^s(t + s + 1) - c_{ij}^s]$, and rearranging terms, we can write

$$\begin{aligned} L(\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}', \mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta}) &= \sum_{ij} \tilde{s}_{ij}(0) [d_{ij} V_{ij}^s(0) + (1 - d_{ij}) V_{ij}^s(0)] + \sum_{ij} \tilde{e}_{ij}(0) V_{ij}^e(0) \\ &+ \sum_{t=0}^{\infty} \beta^t [W(q'(t)) - K(n'(t)) + \mathcal{E}(b'(t))] + \sum_{ij} \mathbf{q}'_{ij} \cdot (\mathbf{V}_{ij}^s - \mathbf{U}_i^s - \beta \mathbf{V}_{ij}^{e,+1} - \mathbf{\Delta}_{ij}) + \sum_{ij} \mathbf{n}'_{ij} \cdot \mathbf{V}_{ij}^e \\ &+ \sum_{ij} \mathbf{b}'_{ij} \cdot (\mathbf{V}_{ij}^s - \mathbf{U}_i^s) + \sum_i \mathbf{s}'_i \cdot (-c_i^s + \mathbf{U}_i^s - \mathbf{V}_i^s) + \sum_{ij} \mathbf{e}'_i \mathbf{G}'_{ij} \cdot (-c_{ij}^e + \beta \mathbf{V}_{ij}^{e,+1} - \mathbf{V}_{ij}^e) + \sum_{ij} \mathbf{\Delta}_{ij} \cdot \mathbf{m}'_i \mathbf{G}'_{ij} \end{aligned}$$

Hence, defining the sequences $\nabla^q, \nabla^n, \nabla^b$ by

$$\begin{aligned} \nabla_{ij}^q(t) &= w_{ij}(q(t)) + V_{ij}^s(t) - U_i^s(t) - U_{ij}^e(t) - \Delta_{ij}(t) \\ \nabla_{ij}^n(t) &= V_{ij}^e(t) - dK(n(t))/dn_{ij}(t) \text{ and } \nabla_{ij}^b(t) = d\mathcal{E}(b(t))/db_{ij}(t) + V_{ij}^s(t) - U_i^s(t), \end{aligned}$$

it is easy to see that $\nabla = (\nabla^q, \nabla^n, \nabla^b)$ is a super-gradient of $L(\mathbf{s}, \mathbf{e}, \mathbf{G}, \cdot, \cdot, \cdot, \mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta})$ at $\mathbf{q}, \mathbf{n}, \mathbf{b}$ - in particular, this follows from Lemma 2. Since $\nabla^n, \nabla^b = \mathbf{0}$ and $\nabla_{ij}^q(t) \leq 0$ with equality if $q_{ij}(t) > 0$, for all ij, t by conditions (40)-(42), it follows that $\mathbf{q}, \mathbf{n}, \mathbf{b}$ maximizes $L(\mathbf{s}, \mathbf{e}, \mathbf{G}, \cdot, \cdot, \cdot, \mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta})$.

Conversely, it is easy to see that, since $\sum_{ij} \mathbf{\Delta}_{ij} \cdot (\mathbf{m}'_i \mathbf{G}'_{ij} - \mathbf{q}'_{ij}) = \mathbf{0}$ by condition (42), $\mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta}$ minimizes $L(\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}, \cdot, \cdot, \cdot)$ over all feasible Lagrange multipliers.

In summary, $\mathbf{q}, \mathbf{n}, \mathbf{b}, \mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta}$ is a saddle point of $L(\mathbf{s}, \mathbf{e}, \mathbf{G}, \cdot, \cdot, \cdot, \cdot, \cdot, \cdot)$. By standard duality arguments (see e.g. Luenberger 1997, Theorem 2, p. 221), it follows that $\mathbf{q}, \mathbf{n}, \mathbf{b}$ solves Problem 14, and

$$\forall \mathbf{q}', \mathbf{n}', \mathbf{b}' \geq \mathbf{0} : v(\mathbf{s}, \mathbf{e}, \mathbf{G}) = L(\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}, \mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta}) \geq L(\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}', \mathbf{n}', \mathbf{b}', \mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta}) \quad (43)$$

A.3 Proof of Theorem 1

A.3.1 Proof outline

For every allocation $\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}'$ and every ij , we denote by $\mathbf{m}'_i, d\mathbf{m}'_i/d\mathbf{s}'_i$ and $d\mathbf{m}'_i/d\mathbf{e}'_i$ the sequences defined by $m'_i(t) = m_i(s'_i(t), e'_i(t))$, $(dm'_i/ds'_i)(t) = dm_i(s'_i(t), e'_i(t))/ds'_i(t)$ and $(dm'_i/de'_i)(t) = dm_i(s'_i(t), e'_i(t))/de'_i(t)$, respectively.

Also, if $\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}'$ satisfies constraints (9)-(13) we write $\mathbf{s}', \mathbf{e}', \mathbf{G}' \in \mathcal{S}$, meaning that $\mathbf{s}', \mathbf{e}', \mathbf{G}'$

is feasible, since it can be supported by the agents' decisions represented by $\mathbf{q}', \mathbf{n}', \mathbf{b}'$. Finally, we denote by \mathcal{S}^s , \mathcal{S}^e and \mathcal{S}^G the sets of sequences \mathbf{s}' , \mathbf{e}' and \mathbf{G}' that are feasible conditional of (\mathbf{e}, \mathbf{G}) , (\mathbf{s}, \mathbf{G}) and (\mathbf{s}, \mathbf{e}) , respectively. Formally:

$$\begin{aligned}\mathcal{S} &= \{\mathbf{s}', \mathbf{e}', \mathbf{G}' \geq \mathbf{0} : \exists \mathbf{q}', \mathbf{n}', \mathbf{b}' \geq \mathbf{0} \text{ s.t. } \mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}' \text{ satisfies (9)-(13)}\} \\ \mathcal{S}^s &= \{\mathbf{s}' \geq \mathbf{0} : \mathbf{s}', \mathbf{e}, \mathbf{G} \in \mathcal{S}\}, \mathcal{S}^e = \{\mathbf{e}' \geq \mathbf{0} : \mathbf{s}, \mathbf{e}', \mathbf{G} \in \mathcal{S}\}, \mathcal{S}^G = \{\mathbf{G}' \geq \mathbf{0} : \mathbf{s}, \mathbf{e}, \mathbf{G}' \in \mathcal{S}\}\end{aligned}$$

As usual, we set $v(\mathbf{s}', \mathbf{e}', \mathbf{G}') = -\infty$, if $\mathbf{s}', \mathbf{e}', \mathbf{G}' \notin \mathcal{S}$. Hence the planner's problem 16 can be written as

$$\sup_{\mathbf{s}', \mathbf{e}', \mathbf{G}' \in \mathcal{S}} v(\mathbf{s}', \mathbf{e}', \mathbf{G}') \text{ s.t. } \sum_j \mathbf{G}'_{ij} = 1, \text{ for all } i$$

and we have $\text{Dom}v = \mathcal{S}$, $\text{Dom}v(\cdot, \mathbf{e}, \mathbf{G}) = \mathcal{S}^s$, $\text{Dom}v(\mathbf{s}, \cdot, \mathbf{G}) = \mathcal{S}^e$ and $\text{Dom}v(\mathbf{s}, \mathbf{e}, \cdot) = \mathcal{S}^G$.

Define the sequences, $\nabla^s, \nabla^e, \nabla^G$ by

$$\begin{aligned}\nabla_{ij}^s(t) &= -c_i^s + dm_i(t)/ds_i(t)\sum_j G_{ij}(t)\Delta_{ij}(t) + U_i^s(t) - V_i^s(t) \\ \nabla_{ij}^e(t) &= \sum_j G_{ij}(t)[-c_{ij}^e + dm_i(t)/de_i(t)\Delta_{ij}(t) + \beta V_{ij}^e(t+1) - V_{ij}^e(t)] \\ \nabla_{ij}^G(t) &= m_i(t)\Delta_{ij}(t) + e_i(t)[-c_{ij}^e + \beta V_{ij}^e(t+1) - V_{ij}^e(t)]\end{aligned}$$

for all ij , and notice that we have: (i) $\nabla^s = \mathbf{0}$, if and only if condition (22) holds; (ii) $\nabla^e = \mathbf{0}$, if and only if condition (27) holds; (iii) $\nabla_{ij}^G = \nabla_{ik}^G$, for every ijk , if and only if condition (25) holds.⁶⁰

The proof proceeds as follows. First, we show that the following holds⁶¹

$$\begin{aligned}v(\mathbf{s}', \mathbf{e}, \mathbf{G}) - v(\mathbf{s}, \mathbf{e}, \mathbf{G}) &\leq \nabla^s \cdot (\mathbf{s}' - \mathbf{s}), \quad v(\mathbf{s}, \mathbf{e}', \mathbf{G}) - v(\mathbf{s}, \mathbf{e}, \mathbf{G}) \leq \nabla^e \cdot (\mathbf{e}' - \mathbf{e}) \\ \text{and } v(\mathbf{s}, \mathbf{e}, \mathbf{G}') - v(\mathbf{s}, \mathbf{e}, \mathbf{G}) &\leq \nabla^G \cdot (\mathbf{G}' - \mathbf{G}) \quad \forall \mathbf{s}' \in \mathcal{S}^s, \mathbf{e}' \in \mathcal{S}^e, \mathbf{G}' \in \mathcal{S}^G\end{aligned}\tag{44}$$

In words, ∇^s , ∇^e and ∇^G are super-gradients of $v(\cdot, \mathbf{e}, \mathbf{G})$ at \mathbf{s} , of $v(\mathbf{s}, \cdot, \mathbf{G})$ at \mathbf{e} and of $v(\mathbf{s}, \mathbf{e}, \cdot)$ at \mathbf{G} , respectively. By Lemma 1, this implies that: (i) \mathbf{s} maximizes $v(\cdot, \mathbf{e}, \mathbf{G})$ over \mathcal{S}^s if $\nabla^s = \mathbf{0}$; (ii) \mathbf{e} maximizes

⁶⁰To see the latter statement, recall that we have $-c_{ij}^e + \beta V_{ij}^e(t+1) - V_{ij}^e(t) = -\lambda_i^e(t)\Delta_{ij}^e(t)$, by equation (20). Hence $\nabla_{ij}^G(t) = m_i(t)[\Delta_{ij}(t) - \Delta_{ij}^e(t)] = m_i(t)\Delta_{ij}^s(t)$, and $\nabla_{ij}^G = \nabla_{ik}^G$ if and only if $\Delta_{ij}^s(t) = \Delta_{ik}^s(t)$.

⁶¹In particular, this shows that $v(\mathbf{s}', \mathbf{e}, \mathbf{G}), v(\mathbf{s}, \mathbf{e}', \mathbf{G}), v(\mathbf{s}, \mathbf{e}, \mathbf{G}') < +\infty$ for every $\mathbf{s}' \in \mathcal{S}^s, \mathbf{e}' \in \mathcal{S}^e, \mathbf{G}' \in \mathcal{S}^G$. Hence the results of Lemma 1 apply to the functions $v(\cdot, \mathbf{e}, \mathbf{G}), v(\mathbf{s}, \cdot, \mathbf{G}), v(\mathbf{s}, \mathbf{e}, \cdot)$.

$v(\mathbf{s}, \cdot, \mathbf{G})$ over \mathcal{S}^e if $\nabla^e = \mathbf{0}$; (iii) \mathbf{G} maximizes $v(\mathbf{s}, \mathbf{e}, \cdot)$ over all $\mathbf{G}' \in \mathcal{S}^G$ such that $\sum_j \mathbf{G}'_{ij} = 1$ for all i , if $\nabla_{ij}^G = \nabla_{ik}^G$ for every ijk ; and (iv) The converse of (i), (ii) and (iii) hold whenever $\mathbf{s}, \mathbf{e}, \mathbf{G}$ lies in the interior of \mathcal{S} , and $v(\cdot, \mathbf{e}, \mathbf{G})$, $v(\mathbf{s}, \cdot, \mathbf{G})$ and $v(\mathbf{s}, \mathbf{e}, \cdot)$ are Gateaux differentiable at \mathbf{s}, \mathbf{e} and \mathbf{G} , respectively.

We conclude the proof by showing that this is the case almost everywhere. That is, we show that $\mathbf{s}, \mathbf{e}, \mathbf{G}$ lies in the interior of \mathcal{S} , and that there exist dense subsets $\mathring{\mathcal{S}}^s \subseteq \mathcal{S}^s$, $\mathring{\mathcal{S}}^e \subseteq \mathcal{S}^e$ and $\mathring{\mathcal{S}}^G \subseteq \mathcal{S}^G$ such that, for every $\mathbf{s}' \in \mathring{\mathcal{S}}^s$, $\mathbf{e}' \in \mathring{\mathcal{S}}^e$ and $\mathbf{G}' \in \mathring{\mathcal{S}}^G$, the maps $v(\cdot, \mathbf{e}, \mathbf{G})$, $v(\mathbf{s}, \cdot, \mathbf{G})$ and $v(\mathbf{s}, \mathbf{e}, \cdot)$ are Gateaux differentiable at \mathbf{s}' , \mathbf{e}' and \mathbf{G}' , respectively.

A.3.2 Proof of (44)

Consider again the Lagrangian L of Problem 14 defined in (39). Using $\tilde{\mathbf{e}}_{ij}^{'+1} = \mathbf{e}'_i \mathbf{G}'_{ij} - \mathbf{q}'_{ij}$ and rearranging terms, for every $\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}'$ satisfying (9)-(13) we can write

$$\begin{aligned} L(\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}', \mathbf{n}', \mathbf{b}', \mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta}) &= L(\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}', \mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta}) \\ &+ \sum_i (\mathbf{s}_i - \mathbf{s}'_i) \cdot (-c_i^s + \mathbf{U}_i^s - \mathbf{V}_i^s) + \sum_i (\mathbf{e}_i - \mathbf{e}'_i) \sum_j \mathbf{G}_{ij} \cdot (-c_{ij}^e + \beta \mathbf{V}_{ij}^{e,+1} - \mathbf{V}_{ij}^e) \\ &+ \sum_{ij} (\mathbf{G}_{ij} - \mathbf{G}'_{ij}) \cdot [\mathbf{e}'_i (-c_{ij}^e + \beta \mathbf{V}_{ij}^{e,+1} - \mathbf{V}_{ij}^e) + \mathbf{m}'_i \mathbf{\Delta}_{ij}] + \sum_i (\mathbf{m}_i - \mathbf{m}'_i) \cdot \sum_j \mathbf{G}_{ij} \mathbf{\Delta}_{ij} \end{aligned}$$

Since $L(\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}', \mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta}) \geq u(\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}')$ for all $\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}'$ satisfying (9)-(13), and $v(\mathbf{s}, \mathbf{e}, \mathbf{G}) \geq L(\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}', \mathbf{n}', \mathbf{b}', \mathbf{V}^s, \mathbf{V}^e, \mathbf{U}^s, \mathbf{\Delta})$ by (43), this implies

$$\begin{aligned} v(\mathbf{s}', \mathbf{e}', \mathbf{G}') &\leq v(\mathbf{s}, \mathbf{e}, \mathbf{G}) + \sum_i (\mathbf{s}'_i - \mathbf{s}_i) \cdot (-c_i^s + \mathbf{U}_i^s - \mathbf{V}_i^s) + \sum_i (\mathbf{e}'_i - \mathbf{e}_i) \sum_j \mathbf{G}_{ij} \cdot (-c_{ij}^e + \beta \mathbf{V}_{ij}^{e,+1} - \mathbf{V}_{ij}^e) \\ &+ \sum_{ij} (\mathbf{G}'_{ij} - \mathbf{G}_{ij}) \cdot [\mathbf{e}'_i (-c_{ij}^e + \beta \mathbf{V}_{ij}^{e,+1} - \mathbf{V}_{ij}^e) + \mathbf{m}'_i \mathbf{\Delta}_{ij}] + \sum_i (\mathbf{m}'_i - \mathbf{m}_i) \cdot \sum_j \mathbf{G}_{ij} \mathbf{\Delta}_{ij}. \end{aligned}$$

Lemma 2 applied to the function $g(\mathbf{e}', \mathbf{s}') = \sum_i (\mathbf{m}'_i - \mathbf{m}_i) \cdot \sum_j \mathbf{G}_{ij} \mathbf{\Delta}_{ij}$ then yields

$$\begin{aligned} v(\mathbf{s}', \mathbf{e}', \mathbf{G}') &\leq v(\mathbf{s}, \mathbf{e}, \mathbf{G}) + \sum_i (\mathbf{s}'_i - \mathbf{s}_i) \cdot (-c_i^s + d\mathbf{m}_i/d\mathbf{s}_i \sum_j \mathbf{G}_{ij} \mathbf{\Delta}_{ij} + \mathbf{U}_i^s - \mathbf{V}_i^s) \\ &+ \sum_i (\mathbf{e}'_i - \mathbf{e}_i) \sum_j \mathbf{G}_{ij} \cdot (-c_{ij}^e + d\mathbf{m}_i/d\mathbf{e}_i \mathbf{\Delta}_{ij} + \beta \mathbf{V}_{ij}^{e,+1} - \mathbf{V}_{ij}^e) \\ &+ \sum_{ij} (\mathbf{G}'_{ij} - \mathbf{G}_{ij}) \cdot [\mathbf{e}'_i (-c_{ij}^e + \beta \mathbf{V}_{ij}^{e,+1} - \mathbf{V}_{ij}^e) + \mathbf{m}'_i \mathbf{\Delta}_{ij}] \end{aligned}$$

which implies the conditions in (44).

A.3.3 Differentiability almost everywhere

By concavity of Problem 14, it follows that the value functions $v(\cdot, \mathbf{e}, \mathbf{G})$, $v(\mathbf{s}, \cdot, \mathbf{G})$ and $v(\mathbf{s}, \mathbf{e}, \cdot)$ are concave. To see this, for every $\mathbf{s}', \mathbf{e}', \mathbf{G}' \in \mathcal{S}$, denote

$$\mathcal{D}(\mathbf{s}', \mathbf{e}', \mathbf{G}') = \{\mathbf{q}', \mathbf{n}', \mathbf{b}' \geq \mathbf{0} : \mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}' \in \text{satisfies (9)-(13)}\}$$

and consider for instance $v(\cdot, \mathbf{e}, \mathbf{G})$. Take $\mathbf{s}^1, \mathbf{s}^2$ such that $\mathbf{s}^1, \mathbf{e}, \mathbf{G} \in \mathcal{S}$ and $\mathbf{s}^2, \mathbf{e}, \mathbf{G} \in \mathcal{S}$. Take two sequences $(\mathbf{q}^{1,n}, \mathbf{n}^{1,n}, \mathbf{b}^{1,n})_n \subset \mathcal{D}(\mathbf{s}^1, \mathbf{e}, \mathbf{G})$ and $(\mathbf{q}^{2,n}, \mathbf{n}^{2,n}, \mathbf{b}^{2,n})_n \subset \mathcal{D}(\mathbf{s}^2, \mathbf{e}, \mathbf{G})$ such that $u(\mathbf{q}^{1,n}, \mathbf{n}^{1,n}, \mathbf{b}^{1,n}, \mathbf{s}^1, \mathbf{e}, \mathbf{G}) \rightarrow v(\mathbf{s}^1, \mathbf{e}, \mathbf{G})$ and $u(\mathbf{q}^{2,n}, \mathbf{n}^{2,n}, \mathbf{b}^{2,n}, \mathbf{s}^2, \mathbf{e}, \mathbf{G}) \rightarrow v(\mathbf{s}^2, \mathbf{e}, \mathbf{G})$ as $n \rightarrow \infty$. By concavity of the matching functions it follows that, for every $\alpha \in [0, 1]$, we have $\alpha(\mathbf{q}^{1,n}, \mathbf{n}^{1,n}, \mathbf{b}^{1,n}) + (1 - \alpha)(\mathbf{q}^{2,n}, \mathbf{n}^{2,n}, \mathbf{b}^{2,n}) \in \mathcal{D}(\alpha\mathbf{s}^1 + (1 - \alpha)\mathbf{s}^2, \mathbf{e}, \mathbf{G})$ for all n . Hence, by concavity of $u(\cdot, \mathbf{e}, \mathbf{G}, \cdot, \cdot, \cdot)$, we have

$$\begin{aligned} v(\alpha\mathbf{s}^1 + (1 - \alpha)\mathbf{s}^2, \mathbf{e}, \mathbf{G}) &\geq u(\alpha\mathbf{s}^1 + (1 - \alpha)\mathbf{s}^2, \mathbf{e}, \mathbf{G}, \alpha(\mathbf{q}^{1,n}, \mathbf{n}^{1,n}, \mathbf{b}^{1,n}) + (1 - \alpha)(\mathbf{q}^{2,n}, \mathbf{n}^{2,n}, \mathbf{b}^{2,n})) \\ &\geq \alpha u(\mathbf{s}^1, \mathbf{e}, \mathbf{G}, \mathbf{q}^{1,n}, \mathbf{n}^{1,n}, \mathbf{b}^{1,n}) + (1 - \alpha)u(\mathbf{s}^2, \mathbf{e}, \mathbf{G}, \mathbf{q}^{2,n}, \mathbf{n}^{2,n}, \mathbf{b}^{2,n}) \rightarrow \alpha v(\mathbf{s}^1, \mathbf{e}, \mathbf{G}) + (1 - \alpha)v(\mathbf{s}^2, \mathbf{e}, \mathbf{G}). \end{aligned}$$

Hence $v(\cdot, \mathbf{e}, \mathbf{G})$ is concave. Concavity of $v(\mathbf{s}, \cdot, \mathbf{G})$ and $v(\mathbf{s}, \mathbf{e}, \cdot)$ can be proven similarly.

It is well known that concave functions are differentiable almost everywhere in the interior of their domain (Asplund 1968). Formally, there exist dense subsets $\mathring{\mathcal{S}}^s \subseteq \mathcal{S}^s$, $\mathring{\mathcal{S}}^e \subseteq \mathcal{S}^e$ and $\mathring{\mathcal{S}}^G \subseteq \mathcal{S}^G$ such that, for every $\mathbf{s}' \in \mathring{\mathcal{S}}^s$, $\mathbf{e}' \in \mathring{\mathcal{S}}^e$ and $\mathbf{G}' \in \mathring{\mathcal{S}}^G$, the maps $v(\cdot, \mathbf{e}, \mathbf{G})$, $v(\mathbf{s}, \cdot, \mathbf{G})$ and $v(\mathbf{s}, \mathbf{e}, \cdot)$ are Gateaux differentiable at \mathbf{s}' , \mathbf{e}' and \mathbf{G}' , respectively, as we wanted to show.

A.3.4 $\mathbf{s}, \mathbf{e}, \mathbf{G}$ lies in the interior of \mathcal{S}

Notice that, in equilibrium, we have $\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{n}, \mathbf{b} > \mathbf{0}$. Indeed, $\mathbf{n} > \mathbf{0}$ by full support of the random entry costs, hence $\mathbf{e}, \mathbf{G} > \mathbf{0}$. On the other hand, $\mathbf{b} > \mathbf{0}$ follows from carriers' random shocks having full support and hence for every i there is always a positive measure of carriers relocating towards i , hence $\mathbf{s} > \mathbf{0}$.

We show that, if $\mathbf{s}, \mathbf{e}, \mathbf{G} > \mathbf{0}$ is such that $\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}$ satisfies (9)-(13) for some $\mathbf{q}, \mathbf{n}, \mathbf{b}$ such that $\mathbf{n}, \mathbf{b} > \mathbf{0}$, then $\mathbf{s}, \mathbf{e}, \mathbf{G}$ lies in the interior of \mathcal{S} . Take such a triplet $\mathbf{s}, \mathbf{e}, \mathbf{G}$. First notice that, without loss of generality, we can take \mathbf{q} such that

$$q_{ij}(t) < m_i(t)G_{ij}(t) \tag{45}$$

Indeed, if $q_{ij}(t) = m_i(t)G_{ij}(t)$ for some ij then, for some small ϵ , decreasing $q_{ij}(t)$ and $n_{ij}(t+1)$ by ϵ and increasing $b_{ij}(t)$ by ϵ leaves all the feasibility constraints (9)-(13) satisfied.

We consider in turn perturbations leading to changes in $s_i(t), e_{ij}(t)$ and $G_{ij}(t)$. Formally, we denote a perturbation of $\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}$ by

$$\mathbf{d} = (\mathbf{ds}, \mathbf{de}, \mathbf{dG}, \mathbf{dq}, \mathbf{dn}, \mathbf{db}) = (ds(t), de(t), dG(t), dq(t), dn(t), db(t))_{t=0}^{\infty}$$

and say that a perturbation is feasible if $\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}'$ satisfies (9)-(13), where $\mathbf{s}', \mathbf{e}', \mathbf{G}', \mathbf{q}', \mathbf{n}', \mathbf{b}'$ is the perturbed allocation defined by

$$\begin{aligned} s'_i(t), e'_{ij}(t), G'_{ij}(t), q'_{ij}(t), n'_{ij}(t), b'_{ij}(t) = & s_i(t) + ds_i(t), e_i(t) + de_i(t), G_{ij}(t) + dG_{ij}(t), \\ & q_{ij}(t) + dq_{ij}(t), n_{ij}(t) + dn_{ij}(t), b_{ij}(t) + db_{ij}(t) \end{aligned}$$

Note that, when all components of \mathbf{d} are small enough so that the positivity constraints do not bind, \mathbf{d} is feasible if and only if it satisfies

$$ds_i(t) = \sum_j d_{ji} d\tilde{s}_{ji}(t) = \sum_j [dq_{ij}(t) + db_{ij}(t)], \quad de_{ij}(t) = d\tilde{e}_{ij}(t) + dn_{ij}(t) \quad (46)$$

$$\text{and } dq_{ij}(t) \leq m_i(s_i(t) + ds_i(t), e_i(t) + de_i(t)) [G_{ij}(t) + dG_{ij}(t)]$$

where $\mathbf{d}\tilde{\mathbf{s}}$ and $\mathbf{d}\tilde{\mathbf{e}}$ are the sequences defined recursively by $d\tilde{s}(0), d\tilde{e}(0) = 0$, $d\tilde{s}_{ij}(t+1) = d\tilde{s}_{ij}(t)(1 - d_{ij}) + dq_{ij}(t) + db_{ij}(t)$ ij , and $d\tilde{e}_{ij}(t+1) = de_{ij}(t) - dq_{ij}(t)$, for all ij .

i) Perturbing $G_{ij}(t)$ Fix an arbitrary ij such that $j \neq i$ and $t \geq 0$. We show that, for small enough $|\epsilon|$, there exists a feasible perturbation \mathbf{d} such that all the components of \mathbf{ds} and \mathbf{de} are zero, and $dG_{ij}(t) = \epsilon$ is the only non-zero component of \mathbf{dG} . To see this, it is enough to consider perturbations of $n_{ij}(t)$ and $n_{ij}(t+1)$, so we let all other components of \mathbf{d} be zero. Given this, the constraints in (46) reduce to

$$e_i(t)dG_{ij}(t) = dn_{ij}(t), \quad dn_{ij}(t+1) = -d\tilde{e}_{ij}(t+1) \text{ and } q_{ij}(t) \leq m_i(t)[G_{ij}(t) + dG_{ij}(t)]$$

$$\text{where } d\tilde{e}_{ij}(t+1) = e_i(t)dG_{ij}(t)$$

Hence we can take $dG_{ij}(t) = \epsilon$, $dn_{ij}(t) = \epsilon e_i(t)$, $dn_{ij}(t+1) = -\epsilon e_i(t)$, and $|\epsilon|$ small enough so that $|\epsilon| m_i(t) < m_i(t)G_{ij}(t) - q_{ij}(t)$.

ii) Perturbing $e_i(t)$ Fix an arbitrary it such that $t \geq 0$. We show that, for small enough $|\epsilon|$, there exists a feasible perturbation \mathbf{d} such that all the components of \mathbf{ds} and \mathbf{dG} are zero, and $de_i(t) = \epsilon$ is the only non-zero component of \mathbf{de} . To see this, it is enough to consider perturbations of $n_{ij}(t)$ and $n_{ij}(t+1)$, for all j , so we let all other components of \mathbf{d} be zero. Given this, the constraints in (46) reduce to:

$$de_i(t)G_{ij}(t) = dn_{ij}(t), \quad dn_{ij}(t+1) = -d\tilde{e}_{ij}(t+1) \quad \text{and} \quad q_{ij}(t) \leq m_i(s_i(t), e_i(t) + de_i(t))G_{ij}(t), \quad \text{for all } j$$

where $d\tilde{e}_{ij}(t+1) = de_i(t)G_{ij}(t)$, for all j

Hence we can take $de_i(t) = \epsilon$, $dn_{ij}(t) = \epsilon G_{ij}(t) \forall j$, $dn_{ij}(t+1) = -\epsilon G_{ij}(t)$, for all j , so that, by continuity of m_i and condition (45), the above constraints are always satisfied for $|\epsilon|$ small enough.

ii) Perturbing $s_i(t)$ Fix $t \geq 1$. We show that, for all i , for small enough $|\epsilon|$, there exists a feasible perturbation \mathbf{d} such that all the components of \mathbf{de} and \mathbf{dG} are zero, and $ds_i(t) = \epsilon$ is the only non-zero component of \mathbf{ds} .

Pick an arbitrary pair ij with $j \neq i$. We first show that, for small enough $|\epsilon^{ij}|$, there exists a feasible perturbation \mathbf{d}^{ij} such that all the components of $\mathbf{d}^{ij}\mathbf{e}$ and $\mathbf{d}^{ij}\mathbf{G}$ are zero, and $d^{ij}s_i(t), d^{ij}s_j(t)$ are the only non-zero component of $\mathbf{d}^{ij}\mathbf{s}$, with

$$d^{ij}s_i(t) = \epsilon^{ij} \quad \text{and} \quad d^{ij}s_j(t) = \epsilon^{ij} \frac{d_{ij}(1 - d_{ji})}{d_{ji}(1 - d_{ij})} \quad (47)$$

To see this, it is enough to consider perturbations of $b_{ii}(t-1), b_{ij}(t-1), b_{jj}(t-1), b_{ji}(t-1), b_{ij}(t), b_{ji}(t)$,

so we let all other components of \mathbf{d}^{ij} be equal to zero. Given this, the constraints in (46) reduce to

$$\mathbf{d}^{ij}b_{ii}(t-1) + \mathbf{d}^{ij}b_{ij}(t-1) = 0, \quad \mathbf{d}^{ij}b_{jj}(t-1) + \mathbf{d}^{ij}b_{ji}(t-1) = 0$$

$$\mathbf{d}^{ij}s_i(t) = \mathbf{d}^{ij}\tilde{s}_{ii}(t) + d_{ji}\mathbf{d}^{ij}\tilde{s}_{ji}(t) = \mathbf{d}^{ij}b_{ij}(t) = -\mathbf{d}^{ij}\tilde{s}_{ij}(t)(1 - d_{ij})$$

$$\mathbf{d}^{ij}s_j(t) = \mathbf{d}^{ij}\tilde{s}_{jj}(t) + d_{ij}\mathbf{d}^{ij}\tilde{s}_{ij}(t) = \mathbf{d}^{ij}b_{ji}(t) = -\mathbf{d}^{ij}\tilde{s}_{ji}(t)(1 - d_{ji})$$

$$q_{ik}(t) \leq m_i(s_i(t) + \mathbf{d}^{ij}s_i(t), e_i(t))G_{ik}(t) \text{ and } q_{ji}(t) \leq m_j(s_j(t) + \mathbf{d}^{ij}s_j(t), e_j(t))G_{jk}(t) \quad \forall k$$

$$\text{where } \mathbf{d}^{ij}\tilde{s}_{ii}(t) = \mathbf{d}^{ij}b_{ii}(t-1), \quad \mathbf{d}^{ij}\tilde{s}_{ij}(t) = \mathbf{d}^{ij}b_{ij}(t-1), \quad \mathbf{d}^{ij}\tilde{s}_{jj}(t) = \mathbf{d}^{ij}b_{jj}(t-1), \quad \mathbf{d}^{ij}\tilde{s}_{ji}(t) = \mathbf{d}^{ij}b_{ji}(t-1)$$

Take $\mathbf{d}^{ij}s_i(t)$ and $\mathbf{d}^{ij}s_j(t)$ as in (47), define ϕ, ψ so that $\epsilon^{ij} = \phi(1 - d_{ij})$ and $\phi d_{ij} = \psi d_{ji}$, and let

$$\mathbf{d}^{ij}b_{ii}(t-1) = \phi, \quad \mathbf{d}^{ij}b_{ij}(t-1) = -\phi, \quad \mathbf{d}^{ij}b_{jj}(t-1) = \psi, \quad \mathbf{d}^{ij}b_{ji}(t-1) = -\psi$$

$$\mathbf{d}^{ij}b_{ij}(t) = \phi(1 - d_{ij}), \quad \mathbf{d}^{ij}b_{ji}(t) = \psi(1 - d_{ji})$$

By continuity of m_i and condition (45), the above constraints are always satisfied for $|\epsilon^{ij}|$ small enough.

Now, pick three arbitrary distinct locations i, j, k . For $|\epsilon|$ small enough, we can choose $\epsilon^{ij}, \epsilon^{jk}, \epsilon^{ki}$ such that

$$\epsilon^{ij} + \epsilon^{ki} \frac{d_{ki}(1 - d_{ik})}{d_{ik}(1 - d_{ki})} = \epsilon, \quad \epsilon^{jk} + \epsilon^{ij} \frac{d_{ij}(1 - d_{ji})}{d_{ji}(1 - d_{ij})} = 0, \quad \epsilon^{jk} \frac{d_{jk}(1 - d_{kj})}{d_{kj}(1 - d_{jk})} + \epsilon^{ki} = 0$$

and the sum $\mathbf{d} = \mathbf{d}^{ij} + \mathbf{d}^{jk} + \mathbf{d}^{ki}$ of the perturbations defined above is feasible. By construction, we have $\mathbf{d}\mathbf{e} = 0$, $\mathbf{d}\mathbf{G} = 0$, and the unique non zero component of $\mathbf{d}\mathbf{s}$ is $\mathbf{d}s_i(t) = \epsilon$. This completes the proof.

A.4 Proof of Corollary 1

For concreteness, we consider the case in which the match tax is paid by customers. The proof for the case in which the tax is paid by carriers is identical.

Note that, under a sequence of taxes $\boldsymbol{\tau} = (\boldsymbol{\tau}^q, \boldsymbol{\tau}^s, \boldsymbol{\tau}^e)$, the agents' value functions are modified as follows. First, equations (1), (3), (4), (5) and (17) - the definitions of traveling and unmatched carriers' value functions, the description of carriers' optimal policy, and the definition of carriers' surplus shares -

are unchanged. The value function of searching carriers are instead now given by

$$V_i^s(t) = -c_i^s - \tau_i^s(t) + \lambda_i^s(t) \Sigma_j G_{ij}(t) \Delta_{ij}^s(t) + U_i^s(t) \quad (48)$$

so that the search costs now include the per-period search tax.

Second, conditions (7) and (8), describing customers' match acceptance/rejection policy and optimal entry decisions are also unchanged, while conditions (18) and (20), describing customers' match surplus and search value functions, are replaced by

$$\begin{aligned} \Delta_{ij}^e(t) &= \max\{w_{ij}(t) - p_{ij}(t) - \beta V_{ij}^e(t+1), 0\} - \tau_{ij}^q(t) \\ V_{ij}^e(t) &= -c_{ij}^e - \tau_i^e(t) + \lambda_i^e(t) \Delta_{ij}^e(t) + \beta V_{ij}^e(t+1) \end{aligned}$$

so that matching costs for customers now include the match tax, while search costs now include the per-period search tax.

Moreover, as mentioned in the main text, the social surplus generated by a match now takes into account the planner's tax revenues:

$$\Delta_{ij}(t) = \Delta_{ij}^s(t) + \Delta_{ij}^e(t) + \tau_{ij}^q(t)$$

Given this, proceeding exactly as in the proof of Theorem 1, we can prove the following.

Lemma 5. *If $\mathbf{s}, \mathbf{e}, \mathbf{G}, \mathbf{q}, \mathbf{n}, \mathbf{b}, \mathbf{p}$ is an equilibrium under taxes $\boldsymbol{\tau}$, then:*

(i) *Carriers internalize thick market and congestion if and only if*

$$-\tau_i^s(t)/\lambda_i^s(t) + \Sigma_j G_{ij}(t) \Delta_{ij}^s(t) = \eta_i^s(t) \Sigma_j G_{ij}(t) \Delta_{ij}(t), \text{ for all } i$$

(ii) *Customers internalize thick market and congestion externalities if and only if*

$$-\tau_i^e(t)/\lambda_i^e(t) + \Sigma_j G_{ij}(t) \Delta_{ij}^e(t) = \eta_i^e(t) \Sigma_j G_{ij}(t) \Delta_{ij}(t), \text{ for all } i$$

(iii) Customers internalize composition externalities if and only if

$$\Delta_{ij}(t) - \Delta_{ij}^e(t) = \Delta_{ik}(t) - \Delta_{ik}^e(t), \text{ for all } ijkt$$

Now recall that, under Nash bargaining, the equilibrium prices must satisfy the additional surplus sharing conditions $\Delta_{ij}^s(t) = \gamma_i[\Delta_{ij}^s(t) + \Delta_{ij}^e(t)] = \gamma_i[\Delta_{ij}(t) - \tau_{ij}^q(t)]$ and $\Delta_{ij}^e(t) = (1 - \gamma_i)[\Delta_{ij}(t) - \tau_{ij}^q(t)]$. Substituting this into (i)-(iii) and re-arranging, implies that, at an equilibrium under Nash bargaining and taxes τ : (i) Carriers internalize thick market and congestion externalities if and only if for all it

$$\tau_i^s(t)/\lambda_i^s(t) + \gamma_i \Sigma_j G_{ij}(t) \tau_{ij}^q(t) = [\gamma_i - \eta_i^s(t)] \Sigma_j G_{ij}(t) \Delta_{ij}(t) \quad (49)$$

(ii) Customers internalize thick market and congestion if and only if for all it

$$\tau_i^e(t)/\lambda_i^e(t) + (1 - \gamma_i) \Sigma_j G_{ij}(t) \tau_{ij}^q(t) = [1 - \gamma_i - \eta_i^e(t)] \Sigma_j G_{ij}(t) \Delta_{ij}(t) \quad (50)$$

(iii) Customers internalize composition externalities if and only if for all it

$$(1 - \gamma_i)[\tau_{ij}^q(t) - \tau_{ik}^q(t)] = \gamma_i[\Delta_{ik}(t) - \Delta_{ij}(t)] \quad (51)$$

The system of equations (49)-(51) fully characterizes the set of efficient taxes τ . Since condition(51) only pins down the differences between the efficient match-destination taxes, it has one degree of freedom for every it . Imposing $\Sigma_j G_{ij}(t) \tau_{ij}^q(t) = 0$, for all it , it yields the efficient tax system (28)-(30).

B Homogeneity and random search in bulk shipping

B.1 Homogeneity

In this section we present additional descriptive evidence that bulk ships are homogeneous. Table 5 regresses shipping prices on ship and shipowner characteristics and fixed effects and shows that they are not are significant.

Next, note that the fit of the ballast discrete choice model is very good (see BKP); this already suggests

	log(price per day)			
	I	II	III	IV
$I \{ \text{orig.} = \text{home country} \}$			0.004 (0.019)	
$I \{ \text{dest.} = \text{home country} \}$			-0.012 (0.015)	
ln(Number Employees)				0.008 (0.007)
ln(Operating Revenues)				0.003 (0.005)
Time FE	Qtr×Yr	Qtr×Yr	Qtr×Yr	Qtr×Yr
Shipowner FE	No	Yes	No	No
Ship characteristics	Yes	Yes	Yes	Yes
Region FE	Orig. & Dest.	Orig. & Dest.	Orig. & Dest.	Orig. & Dest.
Observations	7,263	7,263	7,973	7,973
Adj. R ²	0.530	0.540	0.537	0.537

*p<0.1; **p<0.05; ***p<0.01

Table 5: Regression of shipping prices on shipowner characteristics and fixed effects (Table SI in Supplement to BKP). Shipping prices, ships' characteristics (age and size), and the identity of the shipowner are obtained from Clarksons. Information on shipowner characteristics is obtained from ORBIS. In particular, we match the shipowners in Clarksons to ORBIS; we do so for two reasons: (i) ORBIS allows us to have reliable firm identities, as shipowners may appear under different names in the contract data; (ii) ORBIS reports additional firm characteristics (e.g. number of employees, revenue, headquarters). Here we identify the shipowner with the global ultimate owner (GUO); results are robust to controlling for the identity of the domestic owner (DUO) and the shipowner as reported in Clarksons. Finally, the data used span the period 2010-2016.

that the factors capturing a region's attractiveness considered in our model, can predict behavior, leaving little room for unobserved persistent ship heterogeneity. Moreover, as mentioned in Footnote 40, ships'

ACC CORUS (charterer)
50000/5 MT COAL (cargo size and type)
POLAND TO IMMINGHAM (load/discharge port/range)
23/30 APRIL (laycan)
3 DAYS SHINC, 24 HRS TTIME (loading time)
15000 MT SHINC, 24 HRS TTIME (discharge time)
3.75 TTL (total deductible commission)
PLS INDICATE (OR OFFER) (request for owners' interest)

Figure 4: Sample broker email, from Plakantonaki (2010)

ballast choices are not dispersed from a given origin (i.e. ships tend to ballast to the same regions).

B.2 Random Search

In this section we further investigate the assumption of random search in the model of Section 2. As discussed in Section 3.2, we present an extract from Plakantonaki (2010), which is the internal manual of a large shipping firm, that describes the contracting process. Figure 4 shows an example of an email that the shipowner's broker receives in his inbox (again, this is one out of several thousand emails in a day).⁶² Plakantonaki (2010) writes: "This order means: Please indicate or offer a vessel to be chartered for the account of the charterer CORUS to load a full cargo of bulk coal, but minimum 47,500 mt and maximum 52,500 mt , i.e. 50,000 +/- 5%. The loading port will be in Poland and the discharge port is Immingham. The laydays for loading will commence on the 23rd of April and the canceling date will be on the 30th of April. For loading the laytime is 3 days Sundays Holidays Included, allowing up to 24 hours for delays prior to counting the 3 days allowed. [...] For discharging the time will be calculated on the basis of 15,000 t discharged per day, usually of 24 consecutive hours, Sundays Holidays Included, allowing up to 24 hours for delays prior counting of the time, if a berth is not available. The total commission payable by the owners to the charterer and the competitive brokers will be 3.75% on freight, deadfreight and demurrage earned. [...]"

As soon as the in-house broker receives this information he considers whether he has a ship available,

⁶²We have corroborated in interviews with brokers that this email is indeed representative.

usually in the vicinity, which can safely arrive at the load port during the requested laycan. In the dry cargo example above, the ship should arrive in Poland ready to load on or after the 23rd of April and before the 30th of April.

Assuming for example that the M/V Prince, a Supramax vessel of about 58,000 mt summer deadweight is expected to complete its discharge in nearby Russian St. Petersburg around the 22nd of April, it can reach Polish ports on the 23rd of April or the 24th depending on the port.”

Testing for random search We next test whether search in bulk shipping is random (or undirected). We contrast this with the case of directed search (see Moen, 1997), where carriers choose to search in a specific “market”, i.e. a market for customers heading to a specific destination. First, as discussed in the main text, under directed search, profitable markets attract more carriers, thereby reducing their matching probabilities compared to less profitable markets. We can directly test this implication suggestive of directed search by checking whether in a given origin, i , ships’ waiting time is different across destinations j . Since we use 15 regions, for a given region there are 14 possible destinations; therefore there are $\binom{14}{2} = 91$ such equalities to test for every origin i . Using a simple F-test we are only able to reject the null of no difference for 16% of the equalities.

Second, we directly test the condition that carriers, in equilibrium, are indifferent across searching in different destinations. Consider the following extension to our setup that allows carriers to choose which market to search in. The value of a ship in region i for searching for exporters heading to j is defined as,

$$J_{ij}^s = -c_i^w + \lambda_{ij}^s (p_{ij} + V_{ij}^s) + (1 - \lambda_{ij}^s) U_i^s$$

where as before, $U_i^s = \max\{\beta V_i^s + \sigma \epsilon_i, \max_{i \neq j} V_{ij}^s + \sigma \epsilon_j\}$. Hence, the value of searching in location i is given by, $V_i^s = \max_{j \neq i} \{J_{ij}^s\}$. In equilibrium, carriers are indifferent between searching in different markets, so that, $J_{ij}^s = J_{ik}^s$, for all $k \neq j$. Therefore it must be the case that,

$$\lambda_{ij}^s (p_{ij} + V_{ij}^s - U_i^s) - \lambda_{ik}^s (p_{ik} + V_{ik}^s - U_i^s) = 0$$

Following the dynamic discrete choice literature, (Arcidiacono and Miller, 2011, lemma 1) we have:⁶³

$$V_{ij}^s - U_i^s = \sigma \left(\log P_{ij} - \gamma^{euler} \right)$$

Hence, replacing above, we obtain

$$\lambda_{ik}^s p_{ik} - \lambda_{ij}^s p_{ij} = \left[\lambda_{ij}^s \left(\log P_{ij} - \gamma^{euler} \right) - \lambda_{ik}^s \left(\log P_{ik} - \gamma^{euler} \right) \right] \sigma + \epsilon_{ijk} \quad (52)$$

We test the directed search model by testing whether the residuals of regression (52) are mean zero as follows. Set $y_{ijk} = \lambda_{ik}^s \tau_{ik} - \lambda_{ij}^s \tau_{ij}$ and $x_{ijk} = \lambda_{ij}^s \left(\log P_{ij} - \gamma^{euler} \right) - \lambda_{ik}^s \left(\log P_{ik} - \gamma^{euler} \right)$ and for every market estimate regression (52) with the addition of a constant. The test boils down to checking if the constant is significant.⁶⁴

Figure 5 presents the results. The dots report the value of the statistic, while the black horizontal line is the critical value: anything above this line rejects the null that ships direct their search. For the majority of markets (all but virtually two) we find that indeed directed search is rejected. It is worth noting that even if in some markets there was directed search, unless all of them have directed search the efficiency conditions are violated.

C Estimation and computation details

In this section we discuss the estimation of the model and in particular, the matching function and computation of exporter valuations, as well as the computational algorithms to compute equilibria.

C.1 Matching function estimation

We briefly outline the approach adopted to estimate the matching function in BKP. The estimation draws from the literature on nonparametric identification (Matzkin, 2003) and non-separable instrumental variable techniques (e.g. Imbens and Newey, 2009); see also Brancaccio et al. (2020b) for a guide on the

⁶³This is easy to show, by leveraging the Hotz-Miller inversion, stating that, in the case of logit, $\log P_{ij} - \log P_{ik} = V_{ij}^s - V_{ik}^s$. From here, one can show that for any chosen action j we can write the value function U_i^s as $U_i^s = V_{ij}^s + \psi(P)$ where ψ is a known function, in the case of logit given by $\sigma \log P_{ij} - \sigma \gamma^{euler}$.

⁶⁴We estimate σ_i separately for each market for further flexibility.

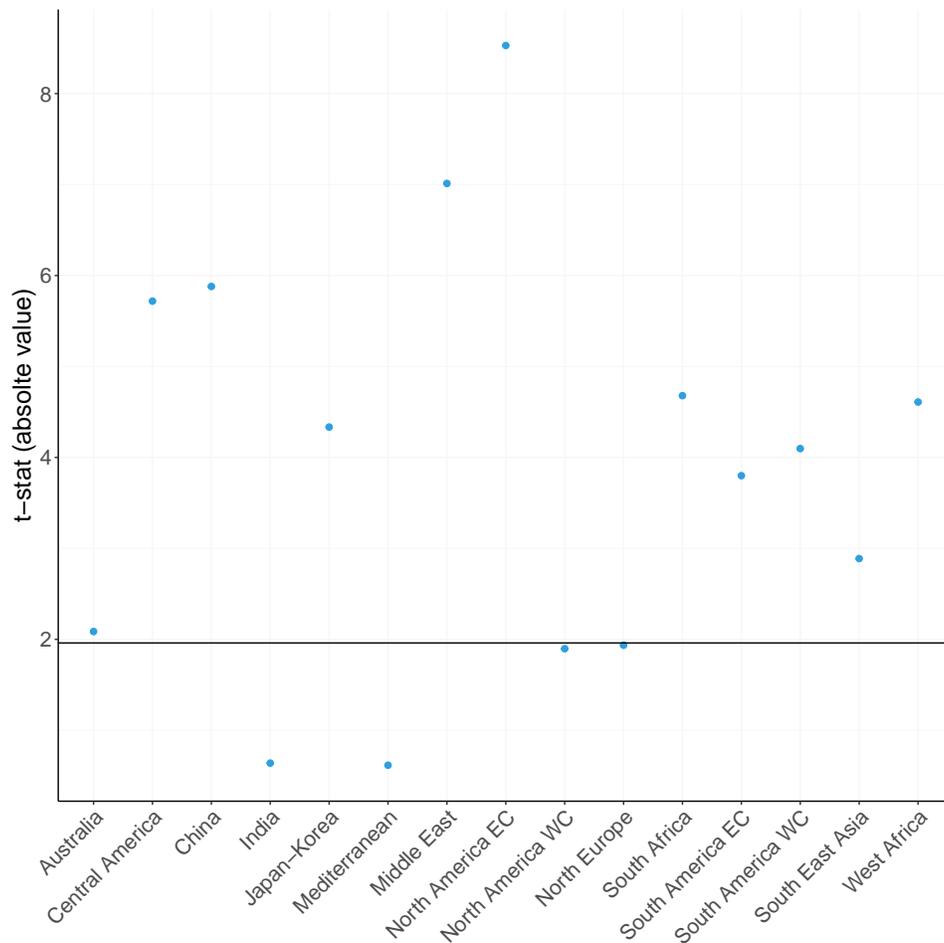


Figure 5: Test of directed search model. For each market i we test regression (52), the main indifference condition satisfied in a directed search model, by testing if its residuals are mean zero. We plot the corresponding t-statistic for each market. The horizontal black line denotes the significance level; anything above the line suggests that the regression is rejected.

implementation of this approach in this and other settings.⁶⁵

To illustrate, assume that s and e are independent. Moreover assume that $m(s, e)$ is continuous and strictly increasing in e , that it exhibits constant returns to scale (CRS), so that $m(as, ae) = am(s, e)$ for all $a > 0$, and that there is a known point $\{\bar{s}, \bar{e}, \bar{m}\}$, such that $\bar{m} = m(\bar{s}, \bar{e})$. Suppose we have a sample $\{s_{it}, m_{it}\}_{t=0}^T$ for each market i . The unknowns of interest are the I matching functions $m_i(\cdot)$ and the exporters e_{it} , for all i, t ; henceforth, we suppress the i subscript to ease notation. Let $F_{m|s}$ denote the distribution of matches conditional on ships, and F_e the distribution of exporters, e . Then at a given point $\{s_t, e_t, m_t\}$ we have:

$$\begin{aligned} F_{m|s=s_t}(m_t|s=s_t) &= \Pr(m(s, e) \leq m_t|s=s_t) = \Pr(e \leq m^{-1}(s, m_t)|s=s_t) \\ &= \Pr(e \leq m^{-1}(s_t, m_t)) = F_e(e_t) \end{aligned}$$

This equation, along with the CRS assumption, allows us to recover the distribution $F_e(e)$, for all e : using the known point $\{\bar{s}, \bar{e}, \bar{m}\}$ and letting $a = e/\bar{e}$, for all e ,

$$F_e(a\bar{e}) = F_{m|s=a\bar{s}}(m(a\bar{s}, a\bar{e})|s=a\bar{s}) = F_{m|s=a\bar{s}}(a\bar{m}|s=a\bar{s})$$

We use this and vary a to trace out $\hat{F}_e(e)$, relying on a kernel density estimator for the conditional distribution $\hat{F}_{m|s=a\bar{s}}(a\bar{m}|s=a\bar{s})$.⁶⁶

Since it is unlikely that s and e are independent, we employ an instrument, which consists of the ocean weather conditions (unpredictable wind at sea) that shift the arrival of ships at a port without affecting the number of exporters (also employed in the search frictions test, see Section 3.2).⁶⁷ Table 6 presents the first stage estimates. Figure 6 reports our estimates for search frictions, along with confidence intervals constructed from 200 bootstrap samples, and Figure 7 shows the matching function contour plot and elasticity.

⁶⁵For an application to labor markets see Lange and Papageorgiou (2020).

⁶⁶We choose the known point, $\{\bar{s}, \bar{e}, \bar{m}\}$, to be of the form $1 = m(\bar{s}, 1)$, so that one exporter is always matched when there are \bar{s} ships. We set \bar{s} iteratively, to be the lowest value such that $m_t \leq e_t$, for all t , thus obtaining a conservative bound on search frictions.

⁶⁷Assume that an instrument z exists such that $s = h(z, \eta)$, with z independent of e, η . The approach now has two steps. In the first step, we recover η using the relationship $s = h(z, \eta)$. In the second step, we repeat the above conditioning on both s (as before) and η , $F_{m|s=s_t, \eta}(m_t|s=s_t, \eta) = F_{e|\eta}(e_t|\eta)$. We recover the unknowns of interest e and $m(\cdot)$, by integrating both sides over η .

	F-stat
North America West Coast	21.132
North America East Coast	18.429
Central America	17.877
South America West Coast	18.671
South America East Coast	16.889
West Africa	16.333
Mediterranean	46.072
North Europe	28.651
South Africa	13.153
Middle East	68.037
India	29.521
South East Asia	34.909
China	28.642
Australia	35.977
Japan-Korea	32.794

Table 6: First Stage, Matching Function Estimation. Regressions of the number of ships in each region on the unpredictable component of weather conditions in the surrounding seas. The table reports the F-statistic. For the construction of the instrument, see Table 1.

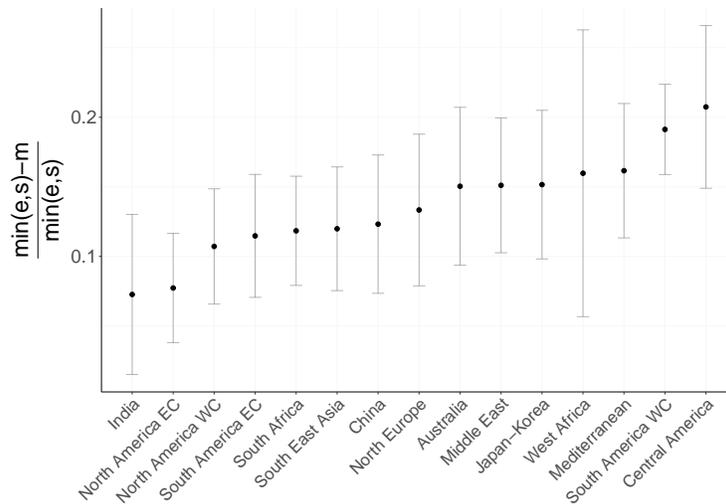


Figure 6: Search Frictions. Average weekly share of unrealized matches, with confidence intervals from 200 bootstrap samples.

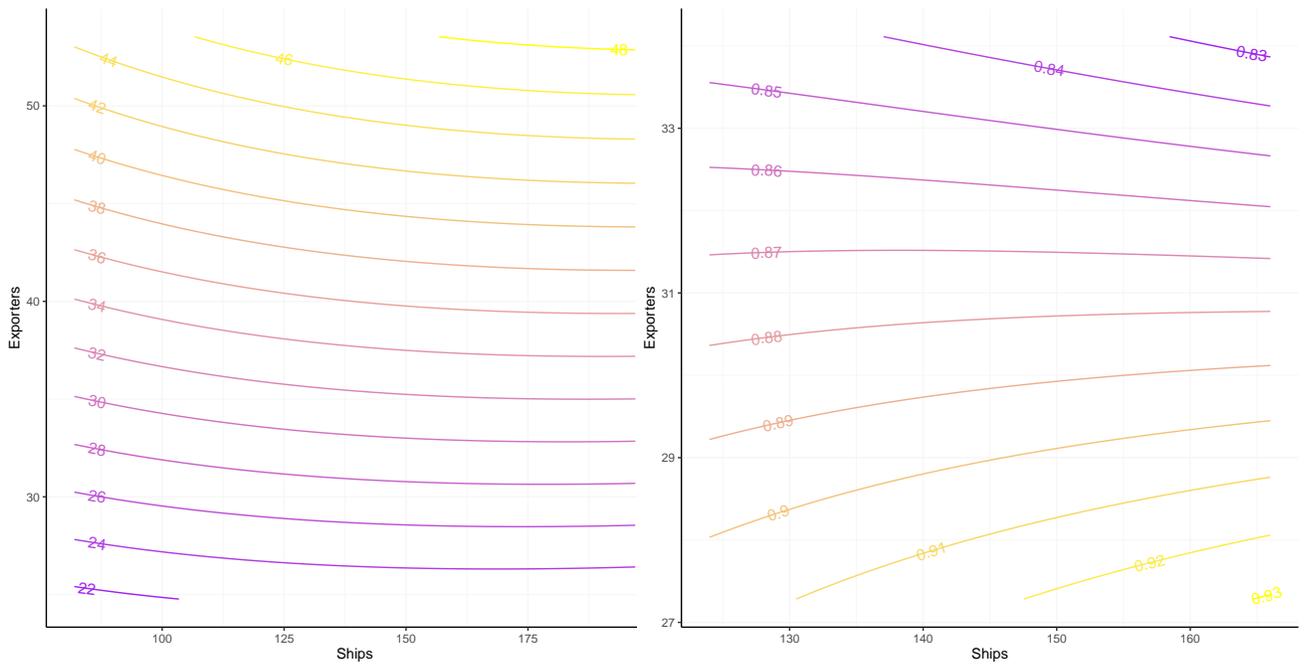


Figure 7: Matching function estimates. The left panel depicts the contour plot for the estimated matching function, averaged across markets. The right panel plots the estimated elasticity with respect to exporters, averaged across markets.

C.2 Exporter valuations

We construct exporter valuations, w_{ij} , from product-level data on export value and quantity by country-pair, obtained from Comtrade. We select bulk commodities among all possible 4-digit HS product codes. The list includes cereals (except rice and barley); oil seeds (which consists of mostly soybeans); cocoa beans; salt and cement; ores; mineral fuels (except petroleum coke); fertilizers; fuel wood and wood pulp; metals; cermets and articles thereof.

To compute the average value of a cargo exported from region i to j , we first compute the average “price” of a ton exported by dividing total export value by total export quantity from i to j . Then, we multiply this price by the average ship tonnage capacity in our sample.⁶⁸

Finally, although most countries belong to one of our regions (depicted in Figure 8), the USA and Canada each belong to two regions (according to the coast). We thus need to split the Comtrade data for the USA and Canada into east and west coast export values. To do so, we employ data on state-level exports from the US Census, as well as on province-level exports from the Canadian International Merchandise Trade Database. In particular, we assign every state (province) to either the east or the west coast and compute, for every product, the share of the total value of trade in that commodity that is exported by east and west coast states (provinces). Then, we compute the total value and quantity of trade for the region East Coast of North America (West Coast of North America) by summing over products the share of the value of east (west) coast trade by the total value of the country’s trade for the USA and Canada. Implicitly, this approach assumes that export values from these two regions are only different due to the composition of products, not their prices.

C.3 Algorithm to compute the constrained efficient allocation

Here, we describe the algorithm employed to compute the steady state of our model. We approximate the matching function with a Cobb Douglas specification in order to avoid computing matching rates nonparametrically in every iteration of the algorithm, which would be computationally intense, as well as in order to smooth the extrapolation.⁶⁹ In particular, for each region we impose $m_{it} = A_i s_{it}^{1-\alpha_i} e_{it}^{\alpha_i}$, and

⁶⁸This is robust to using the average ship tonnage capacity on route ij .

⁶⁹Nonetheless, the nonparametric approach has substantial benefits in the model estimation, as it yields better estimates for the exporters point-wise. Moreover, it provides a “cleaner” test for the thick market and congestion externalities, as we now allow elasticities in a market to vary over time, whereas the Cobb-Douglas specification imposes that they are constant.

select the parameters (A_i, α_i) through non-linear least-squares using the nonparametrically estimated exporters. Recall that τ_{ij}^q , τ_i^e , and τ_i^s denote the tax on loaded trips, on searching exporters, and on searching carriers and that $\sum_j \tau_{ij}^q G_{ij} \equiv 0$.

The algorithm proceeds as follows:

1. Make an initial guess for $\{V_{ij}^{e,0}, p_{ij}^0, s_i^0, e_{ij}^0\}$.
2. At each iteration k , inherit $\{V_{ij}^{e,k-1}, p_{ij}^{k-1}, s_i^{k-1}, e_{ij}^{k-1}\}$. Let G_{ij}^{k-1} , e_i^{k-1} , and q_i^{k-1} denote, respectively, the associated destination shares, searching exporters, and matches.⁷⁰ Moreover, let $\lambda_i^{e,k-1}$ and $\lambda_i^{s,k-1}$ denote the associated matching rates. We update our guess according to the following steps:
 - (a) First, in an inner loop compute the ship optimal policy and value functions implied by $\lambda_i^{s,k-1}$, p_{ij}^{k-1} , and G_{ij}^{k-1} . In particular, after initializing $V_i^{s,0}$, repeat the following steps until convergence:
 - i. At iteration m , compute the value of traveling $V_{ij}^{s,m}$ from $V_{ij}^{s,m} = \frac{-c_{ij}^s + d_{ij} \beta V_i^{s,m-1}}{1 - \beta(1 - d_{ij})}$.
 - ii. Compute the value $U_i^{s,m}$ from:

$$U_i^{s,m} = \sigma \ln \left(\exp \frac{\beta V_i^{s,m-1}}{\sigma} + \sum_{j \neq i} \exp \frac{V_{ij}^{s,m}}{\sigma} \right) + \sigma \gamma^{euler}$$

where γ^{euler} is the Euler constant.⁷¹

- iii. Update $V_i^{s,m}$ from $V_i^{s,m} = -c_i^s - \tau_i^{s,k-1} + (1 - \lambda_i^{s,k-1}) U_i^{s,m} + \lambda_i^{s,k-1} \sum_j G_{ij}^{k-1} (V_{ij}^{s,m} + p_{ij}^{k-1})$
- iv. Upon convergence, we set $V_{ij}^{s,k} = V_{ij}^{s,\infty}$, $V_i^{s,k} = V_i^{s,\infty}$, $U_i^{s,k} = U_i^{s,\infty}$, and compute the ship optimal choice probabilities based on $\mathbb{P}_{ij}^{b,k} = \exp \left(\frac{V_{ij}^{s,k}}{\sigma} \right) / \left(\sum_{l \neq i} \exp \left(\frac{V_{il}^{s,k}}{\sigma} \right) + \exp \left(\frac{\beta V_i^{s,k}}{\sigma} \right) \right)$ for $i \neq j$ and $\mathbb{P}_{ii}^{b,k} = \exp \left(\frac{\beta V_i^{s,k}}{\sigma} \right) / \left(\sum_{l \neq i} \exp \left(\frac{V_{il}^{s,k}}{\sigma} \right) + \exp \left(\frac{\beta V_i^{s,k}}{\sigma} \right) \right)$ for $i = j$.

⁷⁰That is, $e_i^{k-1} = \sum_j e_{ij}^{k-1}$, $G_{ij}^{k-1} = \frac{e_{ij}^{k-1}}{e_i^{k-1}}$, and $q_i^{k-1} = m(s_i^{k-1}, e_i^{k-1})$

⁷¹This is the closed form expression for the expectation of the maximum over multiple choices, and is obtained by integrating U_i^s over the distribution of ϵ .

(b) Update the optimal taxes as follows:

$$\begin{aligned}\tau_{ij}^{q,k} &= \frac{\gamma_i}{1 - \gamma_i} \left(\sum_j G_{ij} \Delta_{ij}^k - \Delta_{ij}^k \right) \\ \tau_i^{s,k} &= \lambda_i^{s,k-1} \left(\gamma_i - \eta_i^{s,k-1} \right) \sum_j G_{ij}^{k-1} \Delta_{ij}^k \\ \tau_i^{e,k} &= \lambda_i^{e,k-1} \left(1 - \gamma_i - \eta_i^{e,k-1} \right) \sum_j G_{ij}^{k-1} \Delta_{ij}^k\end{aligned}$$

where Δ_{ij}^k is defined based on the current iteration's value functions. To update the prices p^k solve the surplus sharing condition,

$$(1 - \gamma_i) \left(p_{ij}^k + V_{ij}^{s,k} - U_i^{s,k} \right) = \gamma_i \left(w_{ij} - p_{ij}^k - \tau_{ij}^{q,k} - \beta V_{ij}^{e,k-1} \right)$$

(b') If the planner is using the efficient prices instead of the optimal taxes/subsidies, then to update the efficient prices p^k use

$$p_{ij}^k = \eta_i^{s,k-1} \sum_j G_{ij}^{k-1} \Delta_{ij}^k + U_i^{s,k} - V_{ij}^{s,k}$$

(b'') If instead the algorithm is used to compute the market equilibrium, compute prices under Nash bargaining using the surplus sharing condition.

(c) Update the exporter value function $V^{e,k}$ based on the prices p_{ij}^k and matching rates $\lambda_i^{e,k-1}$ setting $V_{ij}^{e,k} = \frac{-c_i^e + \lambda_i^{e,k-1} (w_{ij} - p_{ij}^k)}{1 - \beta \delta (1 - \lambda_i^{e,k-1})}$.⁷²

(d) Finally, update the number of ships and exporters searching $\{s_i^k, e_{ij}^k\}$ according to

$$s_i^k = \sum_j \mathbb{P}_{ij}^{b,k} \left(s_j^{k-1} - q_j^{k-1} \right) + \sum q_{ji}^{k-1},$$

⁷²Following BKP we assume that unmatched exporters survive with probability δ so that their effective discount factor is $\beta \delta$ (this is also true in the estimation procedure even though it was omitted there for notational simplicity). We calibrate $\delta = 0.99$. This makes no difference in our theoretical analysis.

and

$$e_{ij}^k = N_i \underbrace{\frac{\exp(V_{ij}^{e,k} - \kappa_{ij})}{1 + \sum_{l \neq i} \exp(V_{il}^{e,k} - \kappa_{il})}}_{\text{new entrants}} + \delta \underbrace{(e_i^{k-1} - q_i^{k-1})}_{\text{unmatched}}$$

3. If $\|s_i^k - s_i^{k-1}\| < \epsilon$, $\|e_{ij}^k - e_{ij}^{k-1}\| < \epsilon$, $\|V_{ij}^{e,k} - V_{ij}^{e,k-1}\| < \epsilon$, and $\|p_{ij}^k - p_{ij}^{k-1}\| < \epsilon$ stop; otherwise go back to point (a).

C.4 Algorithm to solve the centralizing monopolist problem and the competitive benchmark

Here, we begin by describing the algorithm employed to compute the allocation that maximizes the profits of a monopolist who owns the entire fleet of ships and, therefore, can eliminate search frictions through centralization. For this reason, for each region we impose $m_i(t) = \min\{s_i(t), e_i(t)\}$. At the end of the section we detail how to modify the algorithm to compute the competitive benchmark, where there are no search frictions and prices are set competitively on each route.

Both in our sample and counterfactual simulations, the number of ships in each region likely far exceeds the number of exporters. For this reason, we present the counterfactual under the assumption that $s_i(t) \geq e_i(t)$. Hence $m_i(t) = e_i(t)$ (and $\lambda_i^e(t) = 1$), and the number of exporters entering every period in route ij , the number of exporters searching, and the number of loaded trips (matches) on that route must coincide $n_{ij}(t) = e_{ij}(t) = q_{ij}(t)$. Moreover, since $\lambda_i^e(t) = 1$, searching customers' value functions are given by $V_{ij}^e(t) = -c_{ij}^e + w_{ij}(t) - p_{ij}(t)$, hence customers' optimal entry can be written as $n_{ij}(t) = N_i \exp[-c_{ij}^e + w_{ij}(t) - p_{ij}(t) - k_{ij}] / \{1 + \sum_{k \neq i} \exp[-c_{ik}^e + w_{ik}(t) - p_{ik}(t) - k_{ik}]\}$. Setting $q_{ij}(t) = n_{ij}(t)$ and re-arranging terms yields

$$\frac{q_{ij}(t)}{N_i - \sum_{k \neq i} q_{ik}(t)} = \exp[-c_{ij}^e + w_{ij}(t) - p_{ij}(t) - k_{ij}] \quad (53)$$

Problem description In our counterfactual, the monopolist directly pays ship traveling and waiting costs, and can choose their movements. Hence we can think of the monopolist as solving two subsequent problems. First, he decides how many loaded trips $q_{ij}(t)$ to supply in every route ij and period t , and how many ships $s_i(t)$ should be available in each market i and period t . From equation (53), note that

to ensure that exactly $q_{ij}(t)$ exporters need transportation every period on route ij , the monopolist must set transportation prices equal to

$$p_{ij}(t) = p_{ij}(q(t)) \equiv -c_{ij}^e + w_{ij}(t) - \kappa_{ij} - \log \frac{q_{ij}(t)}{N_i - \sum_{k \neq i} q_{ik}(t)} \quad (54)$$

Hence the discounted sum of the monopolist's revenues is given by $\sum_{t=0}^{\infty} \beta^t \sum_{ij} q_{ij}(t) p_{ij}(q(t))$. Notice that, by the log-sum inequality, this is a concave and differentiable function of \mathbf{q} .

Second, the monopolist faces the logistics problem of how to organize the ships' movements in order to minimize the transportation costs, while providing trips as described by \mathbf{q} and available ships as described by \mathbf{s} . In particular, the monopolist must choose ships' ballast flows \mathbf{b} so that the feasibility constraints (9)-(13) are satisfied. Under our previous assumption that the matching constraint is not binding (and that $n_{ij}(t) = e_{ij}(t) = q_{ij}(t)$), these constraints reduce to

$$s_i(t) = \sum_j d_{ji} \tilde{s}_{ji}(t) \quad (55)$$

$$\sum_j b_{ij}(t) = s_i(t) - \sum_j q_{ij}(t)$$

$$\text{where } \tilde{s}_{ij}(t+1) = \tilde{s}_{ij}(t)(1 - d_{ij}) + q_{ij}(t) + b_{ij}(t)$$

We introduce a cost function $C(\mathbf{q}, \mathbf{s})$ that describes the minimum traveling and waiting costs (and maximum i.i.d. re-location utility shocks) that the monopolist must pay in order to provide \mathbf{q}, \mathbf{s} . Recall that, under our logit specification, we have equation (38), repeated here for convenience, $\mathcal{E}(b) = -\sigma \sum_{ij} b_{ij} \log b_{ij} + \sigma \sum_{ij} b_{ij} \log \sum_k b_{ik}$, so that,

$$C(\mathbf{q}, \mathbf{s}) = \min_{\mathbf{b} \geq 0} \sum_{t=0}^{\infty} \beta^t \left\{ \sum_{ij} [q_{ij}(t) + b_{ij}(t)] \frac{c_{ij}^s}{1 - \beta(1 - d_{ij})} + \sum_i s_i(t) c_i^s \right. \quad (56)$$

$$\left. + \sigma \sum_{ij} b_{ij}(t) \log b_{ij}(t) - \sigma \sum_{ij} b_{ij}(t) \log \sum_j b_{ij}(t) \right\} \text{ s.t. (55)}$$

With this notation, the monopolist problem can be written as

$$\max_{\mathbf{q}, \mathbf{s} \geq 0} \sum_{t=0}^{\infty} \beta^t \sum_{ij} q_{ij}(t) p_{ij}(q(t)) - C(\mathbf{q}, \mathbf{s}). \quad (57)$$

To compute a profit-maximizing allocation, we exploit the results of Rosaia (2020b), to show the following:

Lemma 6. *The function $C(\cdot)$ is concave and differentiable. Moreover, fix $\mathbf{q}, \mathbf{s}, \mathbf{b}$ satisfying constraints in (55), and let the $\delta, \mathbf{V}^s, \mathbf{U}^s$ be sequences satisfying*

$$V_{ij}^s(t) = c_{ij}^s + (1 - d_{ij})\beta V_{ij}^s(t) + d_{ij}\beta V_j^s(t), \quad V_i^s(t) = \delta_i - c_i^s + \lambda_i^s(t) \sum_j G_{ij} V_{ij}^s(t) + [1 - \lambda_i^s(t)] U_i^s(t)$$

$$U_i^s(t) = \log \sum_j \exp V_{ij}^s(t) \quad \text{and} \quad b_{ij}(t) / \sum_k b_{ik}(t) = \exp V_{ij}^s(t) / \sum_k \exp V_{ik}^s(t).$$

Then

$$\frac{dC(\mathbf{q}, \mathbf{s})}{dq_{ij}(t)} = \beta^t [U_i^s(t) - V_{ij}^s(t)] \quad \text{and} \quad \frac{dC(\mathbf{q}, \mathbf{s})}{ds_i(t)} = \beta^t [c_i^s + V_i^s(t) - U_i^s(t)],$$

and therefore

$$\begin{aligned} \frac{d\text{profits}}{dq_{ij}(t)} &= \beta^t \left[\frac{d\sum_{k \neq i} q_{ik} p_{ik}(q)}{dq_{ij}} - U_i^s(t) - V_{ij}^s(t) \right] \\ \frac{d\text{profits}}{ds_i(t)} &= -\beta^t [c_i^s + V_i^s(t) - U_i^s(t)]. \end{aligned}$$

Intuitively, one can think of $\delta_i(t)$ as the average payoff that carriers searching at i would need to receive to implement the allocation (\mathbf{q}, \mathbf{s}) in a decentralized fashion. Indeed, $\mathbf{V}^s, \mathbf{U}^s$ can be thought of as carriers' optimal value functions under these payoffs, and, as the outcome of carriers' optimal re-location policy, they yield the ballasting behavior \mathbf{b} solving Problem (56). In other words, Lemma 6 states that the monopolist's cost minimization problem can be decentralized by finding carriers' payoffs $\delta_i(t)$ and letting them optimize their movements autonomously.

Then, Lemma 6 states that the marginal cost of an additional trip on ij at t coincides with carriers' opportunity cost of being matched. Similarly, the marginal cost of an additional searching carrier coincides with the opportunity cost of searching one more period.

Solution algorithm We solve problem (57) focusing on steady state allocations q, s, b . From the constraints in (55), in steady state we must have $\tilde{s}_{ij} = (q_{ij} + b_{ij})/d_{ij}$; hence q, s, b must satisfy

$$s_i = \sum_j (q_{ji} + b_{ji}) \quad (58)$$

$$\sum_j b_{ij} = s_i - \sum_j q_{ij}$$

We therefore say that q, s, b is a profit-maximizing steady state allocation if it satisfies the constraints in (58), and the constant sequences $\mathbf{q}, \mathbf{s}, \mathbf{b}$ defined by $q(t), s(t), b(t) = q, s, b$, for all t , are such that \mathbf{b} solves Problem (56) and \mathbf{q}, \mathbf{s} solve Problem (57) under the initial condition $\tilde{s}_{ij}(0) = (q_{ij} + b_{ij})/d_{ij}$ for all ij .

The algorithm to solve Problem (57) is as follows. We start from a steady state equilibrium allocation (q, s) and modify the allocation based on the gradient of the steady-state monopolist's profits in Lemma 6.

With that in mind, the algorithm proceeds as follows:

1. Make an initial guess $\{q^0, s^0\}$
2. At each iteration k , start with an allocation $\{q^k, s^k\}$. Let $G_{ij}^k = q_{ij}^k/q_i^k$ denote the associated destination shares and $\lambda_i^{s,k} = q_i^k/s_i^k$ denote the associated ship matching rates. We update our guess according to the following steps:
 3. In an inner-loop, we find the average payoff δ and ballast probabilities $\mathbb{P}_{ij}^{b,k}$ that implement $\{q^k, s^k\}$. Rosaia (2020a) guarantees that the algorithm converges.
 - (a) Start with a guess δ^0
 - (b) At step m , given $\delta^m, \lambda^{s,k}, G^k$, compute the ship values $(U_i^m, V_i^{s,m}, V_{ij}^{s,m})$ that solve the fixed point

$$U_i^m = \log \sum_j \exp(V_{ij}^{s,m}) + g_{\text{euler}}$$

$$V_i^{s,m} = \delta_i^m - c_i^s + \lambda_i^{s,k} \sum_j G_{ij}^k V_{ij}^m + (1 - \lambda_i^{s,k}) U_i^m$$

$$V_{ij}^{s,m} = c_{ij}^s + (1 - d_{ij}) \beta V_{ij}^m + d_{ij} \beta V_j^m$$

as well as the associated ballast probabilities $\mathbb{P}_{ij}^{b,m}$

$$\mathbb{P}_{ij}^{b,m} = \frac{\exp(V_{ij}^{s,m})}{\sum_{j \neq i} \exp(V_{ij}^{s,m}) + \exp(\beta V_i^{s,m})}$$

- (c) Given $\mathbb{P}_{ij}^{b,m}$, compute the corresponding steady state distribution of searching ships s^m
- (d) If $\|s^m - s^k\| < \epsilon$ set $U_i^k = U_i^m$, $V_{ij}^{s,k} = V_{ij}^{s,m}$, and $V_i^{s,k} = V_i^{s,m}$. Otherwise update the price based on $\delta^{m+1} = \delta^m + \alpha (s^k - s^m)$ for some small constant α and go back to (a).

4. Compute the gradient of the monopolist's objective function as

$$\begin{aligned} \frac{d\text{profits}^k}{dq_{ij}} &= p_{ij}(q_i^k; s_i^k) + \sum_h q_{ih}^k \frac{\partial p_{ih}(q_i^k; s_i^k)}{\partial q_{ij}} V_{ij}^{s,k} - U_i, \\ \frac{d\text{profits}^k}{ds_i} &= -c_{ij}^s - V_i^{s,k} + U_i^k. \end{aligned}$$

5. For some small constant ψ , update

$$\begin{aligned} q_{ij}^{k+1} &= q_{ij}^k + \psi \frac{\partial \text{profits}^k}{\partial q_{ij}} \\ s_i^{k+1} &= s_i^k + \psi \frac{\partial \text{profits}^k}{\partial s_i}. \end{aligned}$$

6. Keep iterating until $\nabla \text{profits}^k \sim 0$.

Competitive benchmark The algorithm above can be easily modified to compute the equilibrium allocation under no search frictions where prices are set competitively on each route (that is, the frictionless benchmark, or first-best). In particular, we simulate the behavior of a ‘‘benevolent platform’’. Such a platform directly pays ship traveling and waiting costs and can choose their movements. However, as opposed to the monopolist, the benevolent platform maximizes welfare rather than profits. Denoting by $C(\mathbf{q}, \mathbf{s})$ the minimum cost of implementing allocation (\mathbf{q}, \mathbf{s}) , the objective function for the benevolent platform is

$$v(\mathbf{q}, \mathbf{s}) = \sum_{t=0}^{\infty} \beta^t \sum_{ij} q_{ij}(t) w_{ij} - C(\mathbf{q}, \mathbf{s})$$

The algorithm above can be easily adapted using

$$\frac{dv(\mathbf{q}, \mathbf{s})}{dq_{ij}(t)} = \beta^t [w_{ij}(t) - U_i^s(t) - V_{ij}^s(t)]$$

$$\frac{dv(\mathbf{q}, \mathbf{s})}{ds_i(t)} = -\beta^t [c_i^s + V_i^s(t) - U_i^s(t)].$$

D Additional figures and tables

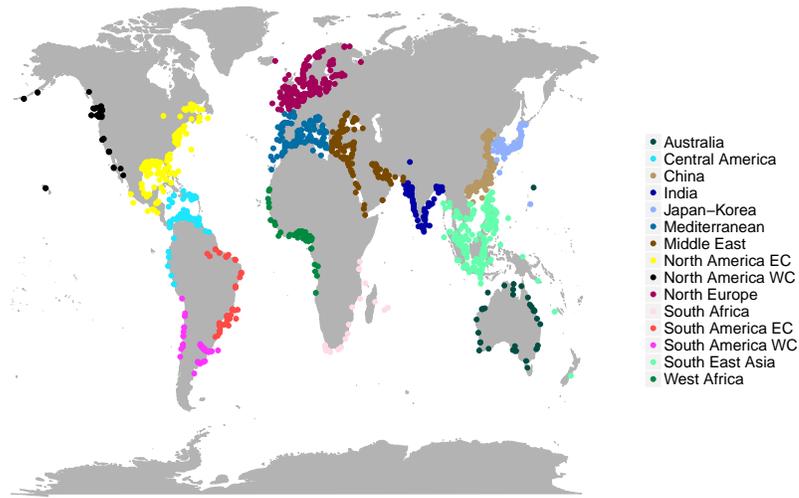


Figure 8: Definition of regions. Each color depicts one of the 15 geographical regions.

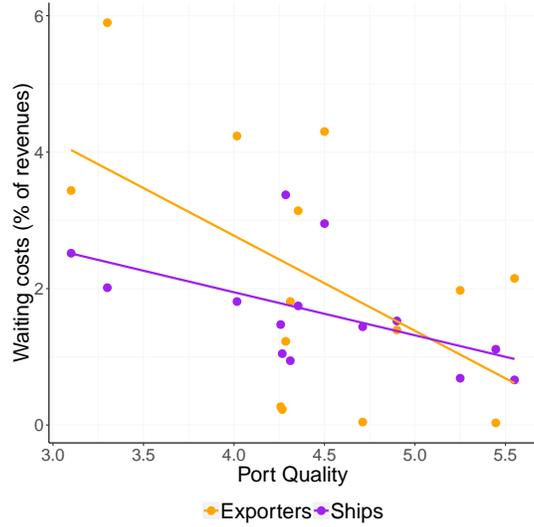


Figure 9: The plot correlates the wait cost estimates for ships and exporters with the World Bank index of port quality.

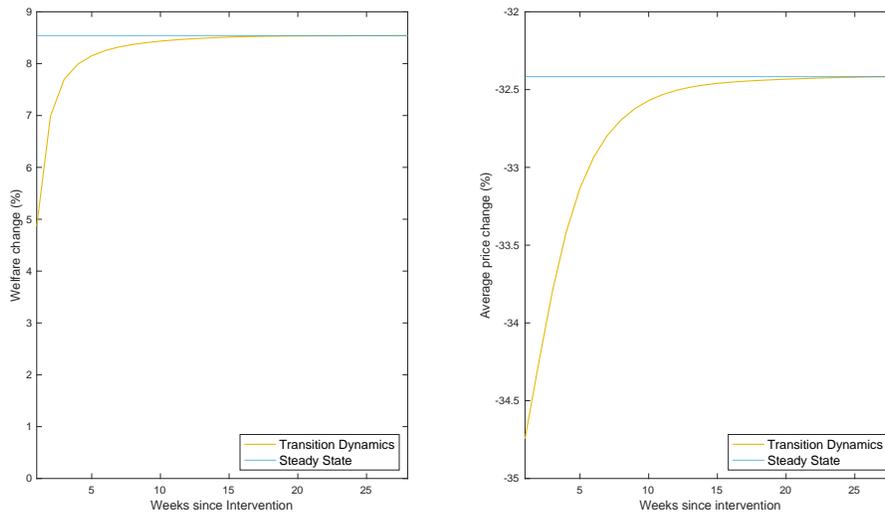


Figure 10: Transition dynamics from the observed equilibrium to the constrained efficient outcome. The left panel depicts welfare change and the right panel depicts average price change.

	I	II	III
	log(price per day)		
Probability of ballast		0.234** (0.030)	0.556** (0.081)
Avg duration of ballast trip (log)		0.166** (0.014)	0.065** (0.032)
Coal			0.088** (0.045)
Fertilizer			0.245** (0.051)
Grain			0.131** (0.048)
Ore			0.124** (0.045)
Steel			0.135** (0.049)
Constant	10.284** (0.103)	9.127** (0.099)	8.915** (0.408)
Destination FE	Yes	No	No
Origin FE	Yes	Yes	Yes
Ship type FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Obs	11,014	11,011	1,662
Adjusted R ²	0.694	0.674	0.664

** $p < 0.05$, * $p < 0.1$

Table 7: Shipping price regressions (Table II in BKP). The dependent variable is the logged price per day in USD. The independent variables include combinations of: the average frequency of ballast traveling after the contract’s destination (Probability of ballast), the average logged duration (in days) of the ballast trip after the contract’s destination, as well as ship type, origin, destination and quarter FEs. The product is reported in only 20% of the sample, so the regression in column III has substantially fewer observations. The omitted product category is cement.

	Exporter wait costs	Ship bargaining coefficient	Average exporter value
	c_i^e	γ_i	\bar{w}_i
North America West Coast	83.49 (10.72)	0.384 (0.018)	13,738
North America East Coast	83.49 (10.72)	0.585 (0.012)	12,192
Central America	302.3 (69.28)	0.344 (0.038)	14,350
South America West Coast	302.3 (69.28)	0.259 (0.017)	20,096
South America East Coast	302.3 (69.28)	0.371 (0.042)	6,971
West Africa	396.9 (512.74)	0.292 (0.078)	4,547
Mediterranean	3.44 (8.60)	0.412 (0.014)	10,508
North Europe	3.44 (8.60)	0.517 (0.014)	14,577
South Africa	396.9 (512.74)	0.24 (0.075)	6,224
Middle East	20.53 (9.04)	0.615 (0.026)	7,160
India	20.53 (9.04)	0.568 (0.022)	6,305
South East Asia	174.01 (44.00)	0.215 (0.036)	4,918
China	282.8 (77.86)	0.194 (0.028)	8,231
Australia	174.01 (44.00)	0.388 (0.037)	12,475
Japan-Korea	282.8 (77.86)	0.265 (0.038)	2,977

Table 8: Average exporter valuation (over destinations), wait costs and bargaining coefficients estimates. All the estimates are in 1,000 USD. To gain power, we restrict exporter wait costs to be constant within a continent. Standard errors computed from 200 bootstrap samples.

	Ship wait costs c_i^s	Ship traveling cost $c_i^s j$	Preference Shock σ
North America West Coast	137.14 (4.75)	45.94 (0.26)	
North America East Coast	181.39 (2.07)	– –	
Central America	138.89 (2.64)	45.94 (0.26)	
South America West Coast	165.07 (2.77)	45.94 (0.26)	
South America East Coast	177.41 (2.09)	– –	
West Africa	116.88 (3.302)	46.95 (0.22)	
Mediterranean	101.66 (1.51)	46.54 (0.23)	
North Europe	105.96 (1.13)	46.54 (0.23)	
South Africa	140.02 (3.08)	46.95 (0.22)	
Middle East	99.66 (1.11)	46.54 (0.23)	
India	83.11 (0.71)	45.39 (0.28)	
South East Asia	83.89 (0.49)	41.23 (0.22)	
China	81.13 (0.722)	41.73 (0.21)	
Australia	147.89 (1.720)	41.23 (0.22)	
Japan-Korea	89.26 (1.17)	41.73 (0.21)	
			13.88 (0.072)

Table 9: Ship cost estimates and scale parameter of their preference ships. All the estimates are in 1,000 USD. Standard errors computed from 200 bootstrap samples. The sailing cost for the East Coast of North and South America is set equal to the weekly fuel cost at 40,000 U.S. dollars.