Disruptions in Transport Markets with Remote Search

Navid Safahieh

Department of Economics, Purdue University

nsafahie@purdue.edu

November 15, 2026

Click here for the latest version

Abstract

I develop a dynamic spatial directed search model with remote matching, in which customers post fees by carrier location and choose how long to wait, and carriers choose where to wait. Fees, waiting times and the distribution of carriers are determined together. I use the model to study disruptions at chokepoints, and find that they have network effects: they move fees on routes that do not use the chokepoint. Closing Bab el-Mandeb raises the fee on the Persian Gulf–Asia route, which does not use the passage, when the detour to Europe is worth taking, and lowers it otherwise. A toll for leaving the Persian Gulf through the Strait of Hormuz lowers the fee on the West Africa–Asia route, which avoids the strait. Customers bear a larger share of the toll under remote matching than under local matching, because a carrier already in the Gulf must pay the toll to leave, while one hired outside can avoid it by working other routes. In a two-location example, I compare the model with one in which a carrier can be hired only once it reaches the cargo. The two coincide when trips are of little value. Once trips are valuable enough, with remote matching customers hire carriers that are farther away and wait longer for them, and carriers sign a contract first and then sail to the cargo, instead of sailing there without one, consistent with tankers being hired farther from loading ports when freight markets are strong. Carriers are then worth more at both locations, and the premium for being at the cargo shrinks.

1 Introduction

The continuous operation of global shipping is both essential to everyday life and vulnerable to disruption. Ships carry about 80 percent of world trade by volume (UNCTAD, 2023) and about 76 percent of the world's oil.[1] Much of the fuel, food and other goods people use every day reaches them by sea. Yet this trade depends on a handful of narrow straits and canals,[2] and a conflict near any one of them can disrupt this trade. The Strait of Hormuz, only about 21 nautical miles wide, carries more than a quarter of the oil traded by sea. Since February 2026, the war between Iran and the United States has disrupted the strait: in the first days of the war, transits through the strait fell by about 92 percent, and 8.5 percent of the very large crude carriers were trapped in the Persian Gulf.[3] Bab el-Mandeb, at the southern end of the Red Sea, was disrupted before it: after the Houthis began attacking commercial ships in the Red Sea in November 2023, many tankers went around the Cape of Good Hope instead, and oil flows through the strait fell from 8.7 million barrels per day in 2023 to 4.0 million in the first eight months of 2024.[4] Global shipping is therefore indispensable to everyday economic activity, yet its reliance on a limited number of strategic waterways makes it particularly vulnerable to disruption.

This paper asks what such a disruption does to the transport market. I study the effects on fees and delivery times jointly, because they are interconnected: they affect each other, and both depend on the distribution of carriers. A customer that finds a carrier nearby can pay more for that, or hire one that is far away and wait while it sails over, and this choice also affects carriers strategies in positioning themselves. How does a disruption change the fees customers pay and how long they wait for their carrier, on the routes that use the chokepoint and on those that do not? And through what mechanism: when do carriers sail to the cargo before they are hired, and when do customers instead hire carriers that

[1] EIA, World Oil Transit Chokepoints.

[2] For oil, the main ones are the Strait of Hormuz, the Strait of Malacca, the Suez Canal, Bab el-Mandeb, the Turkish Straits, the Danish Straits and the Panama Canal; see EIA, World Oil Transit Chokepoints.

[3] Kpler, 12 March 2026.

[4] EIA, Today in Energy, 11 October 2024.

are still far away?

Building on Brancaccio et al. (2020), I develop a dynamic spatial directed search model with remote matching, in which customers post fees by carrier location and carriers apply from wherever they are. Customers are differentiated by their origin and destination, and carriers by their current location. Each period, every customer posts a fee for each location a carrier may come from and a ranking of those locations; every open carrier observes the fees and rankings and applies to at most one customer; and each customer hires its highest-ranked applicant. A customer that hires a distant carrier pays a storage cost until the carrier arrives, so customers choose their waiting times. Carriers that are not hired either stay or sail empty to another location. As is standard under directed search with posted prices, the equilibrium is efficient, so it solves a concave planning problem; this shows that the equilibrium exists and is unique, and lets me solve the model exactly, including the path after a shock. The model nests local matching, so comparing it with an otherwise identical model with local matching only isolates what remote matching does.

I use directed search because I want to study fees. Contracts are negotiated bilaterally through brokers, ship by ship, so the fee can reflect where a ship is, as the fee for each location does in the model. Because customers post fees and carriers choose where to apply, a fee has a simple form: it pays a carrier its gain from being able to apply, divided by its chance of being hired, plus what being hired costs it. Both parts depend on how many carriers are at each location and on what each location is worth, so the model shows directly how the distribution of carriers moves fees. I use remote matching because I want to study waiting, and because this is how tankers are hired: very large crude carriers fixed to load in the Persian Gulf are on average almost 2,900 miles from the loading area when the contract is signed, and about 3,700 nautical miles in strong markets against 2,800 in weak ones (Prochazka et al., 2019; Regli and Nomikos, 2019). With remote matching, a customer that hires a carrier still far away waits while it sails over, and it chooses how long a wait to accept, so a disruption shows up in how long customers wait as well as in fees.

I first study two disruptions in small models with three locations, meant to show the mechanism rather than to match the data. Both have network effects: they move fees on routes that do not use the chokepoint, because carriers move between routes. Closing Bab el-Mandeb lengthens voyages from the Persian Gulf to Europe and leaves voyages from the Gulf to Asia unchanged (Section 3.1). When the detour is worth taking, customers shipping to Europe keep shipping, the longer voyages tie up carriers, and the route to Asia takes most of the loss: shipments to Asia fall by 25 percent and shipments to Europe by 1 percent, and the fee for a carrier in the Gulf rises by 12 percent for cargo to Asia and by 80 percent for cargo to Europe. Customers shipping to Europe also wait longer for their carrier, about 10 days instead of 8, because they hire more carriers that are still far from the Gulf. When the detour is not worth taking, carriers move to the Asia route and its fee falls.

In the second example, I study what would happen if ships leaving the Persian Gulf through the Strait of Hormuz had to pay a toll (Section 3.2). Crude is loaded in the Gulf and in West Africa, both for Asia, and voyages from West Africa go around the Cape of Good Hope, so they avoid the strait. Unlike the closure, which takes carriers' time, the toll takes their money, and it has the opposite effect on the route that avoids the strait: with a toll of one million dollars per cargo, Gulf customers ship somewhat less, carriers spend the time on cargo from West Africa, and the fee there falls by 14 percent. The toll also changes where carriers wait and how long customers wait for them. With no toll, a fifth of the carriers that Gulf customers hire are still in Asia, and a Gulf customer waits about eight days for its carrier. With the toll, the carriers in Asia turn to West Africa: Gulf customers hire almost only carriers already in the Gulf and wait less than six days, while West African customers start hiring carriers in Asia and wait almost nine days instead of five. As the toll rises further, more carriers wait in West Africa, so the wait there falls back to about five days and the fee keeps falling. At a toll of about 4.4 million dollars per cargo, Gulf customers stop shipping.

Finally, I show the mechanism behind these choices in a two-location example in which

cargo moves in one direction (Section 4). As trips become more valuable, customers hire carriers that are farther away and wait longer for them, and carriers stop sailing to the cargo before they are hired: they sign a contract first, then sail. This matches the fact that tankers are hired farther from the loading area in strong markets. Remote matching changes outcomes only when trips are valuable enough: below a cutoff in the value of a trip, remote and local matching coincide; above it, remote matching raises carrier values at both locations and shrinks the premium for being at the cargo.

The paper contributes first to the literature on search and matching in transport markets. Work on these markets shows that fees depend on where carriers are: Brancaccio et al. (2020) show that the transport sector creates network effects in trade costs, because carriers move across locations, and Regli and Nomikos (2019) show that the number of tankers near the loading area moves freight rates. Search models of transport markets, however, assume local matching (Lagos, 2000; Fréchette et al., 2019; Buchholz, 2022; Brancaccio et al., 2020, 2023): a carrier can be hired only where it is, so it must sail to the cargo without a contract. It sinks the cost of the trip for a cargo it may not win, and once there it can contract only with customers at that location. A hired carrier is then already at the cargo, so customers never wait for their carrier, and these models cannot say how a disruption changes delivery times. My model nests local matching and lets customers hire carriers wherever they are, so fees, delivery times and the distribution of carriers are determined together.

The paper also contributes to the literature on directed search (Moen, 1997; Burdett et al., 2001; Menzio and Shi, 2011). As in Shi (2002) and Shimer (2005), both sides of the market are heterogeneous: customers differ in their origin and destination, and carriers in their location, which determines which cargo they can reach quickly and how customers rank them. My model is a dynamic version in which matches change carriers' types: a hired carrier ends up at its customer's destination, and an unmatched one can relocate, so the distribution of carriers over locations evolves with the matches, and where carriers

wait and how long customers wait for them are equilibrium outcomes.

Finally, it relates to maritime economics, which studies how freight rates form and documents that tankers are hired far from the cargo (Tvedt, 2011; Prochazka et al., 2019; Regli and Nomikos, 2019); I provide an equilibrium model in which these distances, and the fees that go with them, are outcomes.

The rest of the paper is organized as follows. Section 2 presents the model: customers post fees by carrier location, carriers apply from wherever they are, and unmatched carriers stay or sail empty. It defines the equilibrium and shows that it is efficient, exists and is unique. Section 3 studies disruptions in two examples with three locations: closing Bab el-Mandeb (Section 3.1), and a toll for leaving the Persian Gulf through the Strait of Hormuz, and how it changes fees, where carriers wait and how long customers wait (Section 3.2). Section 4 solves the model with two locations, gives the cutoffs in the value of a trip that separate its regimes, and compares it with an otherwise identical model with local matching. Section 5 concludes.

2 Model

This section sets out the environment and the meeting process, and then studies the market with an exogenous inflow of customers (Section 2.1). Table 1 collects the notation, and Table 2 lists what each agent chooses.

Environment. Time is discrete and the horizon is infinite. There are l locations (ports), $\mathcal{L} = \{1, \ldots, l\}$, and two kinds of agents, customers and carriers, who are risk neutral and discount the future at $\beta \in (0, 1]$. A customer (i, j) has one cargo at its origin i bound for its destination j; let $\mathcal{M}$ be the set of customer types. Carriers are physically identical and differ only in where they are. Travel takes time: a carrier on the leg $i \to j$ arrives at j in the next period with probability $d_{ij} \in (0, 1]$, so the leg lasts $1/d_{ij}$ periods on average. Durations satisfy the triangle inequality, $1/d_{ij} \leq 1/d_{ik} + 1/d_{kj}$ for every k: no detour shortens a trip.

Every period, a mass $x_{ij} > 0$ of customers (i, j) enters, each with one cargo to move; the inflow is exogenous and constant over time. A customer obtains w_{ij} from shipping its cargo. Once hired, the carrier must be at i to load the cargo, and the customer pays a storage cost c_i per period until it is aboard: for one period if the carrier is already at i, and for every period of the carrier's leg to i if it is hired from elsewhere. A customer that is not matched in the period it enters leaves with a payoff of zero.

There is a fixed fleet of carriers of mass v. A carrier begins each period in one of three states: *open* at a location k, able to apply for cargo; *in transit* from k to j, empty or laden; or *fixed*, hired at k by a customer (i, j) and sailing to i, where it loads on arrival and starts the leg $i \to j$. During the period, an open carrier is either hired, and becomes fixed, or not hired, and then stays at k or sails empty to another location; carriers in transit and fixed carriers reach the end of their leg with the probabilities d. Staying at k is treated as the trip $k \to k$: it costs $c_{kk} \geq 0$ and ends next period, $d_{kk} = 1$. A carrier in transit $k \to j$ incurs a per-period sailing cost $c_{kj} \geq 0$. I denote the masses of carriers in the three states at the start of a period by $v^{\mathcal{O}}_k$, v^{T}_{kj}, and $v^{\mathrm{F}}_{k,ij}$. The fleet starts from any distribution $\mathbf{v}_0$ over these states.

Carriers make two choices. First, an open carrier at k applies for at most one cargo per period, and it may apply to *any* customer type (i, j), not only to those at k; carriers in transit do not apply. Second, a carrier that is not hired either stays at k or starts a leg to some $j \neq k$. A carrier's location is therefore its type: it determines which cargo the carrier can reach quickly and how it ranks among the applicants for it. Unlike a worker's skill in Shimer (2005), this type is chosen, by relocating, and the distribution of carriers over locations is the model's endogenous state; the customer side is exogenous, as the distribution of skills is there.

Within a period, customers and open carriers play a three-stage game. First, each customer (i, j) posts, at the same time, a fee menu $\phi_{ij} = (\phi_{k,ij})_{k \in \mathcal{L}}$, where $\phi_{k,ij}$ is the fee it pays a carrier hired from location k, and a ranking $\succ_{ij}$ of locations: whom it hires first. Then each open carrier observes all menus and rankings, and so its chance of being hired,

and applies to at most one customer. Finally, each customer that receives at least one application from a location it ranks hires its highest-ranked applicant and pays the posted fee. A ranking may be random, and it may leave locations out: a customer never hires an applicant from a location it does not rank. Following Shimer (2005), fees and hiring may depend on a carrier's location but not on its identity, and carriers at the same location use the same application strategy: a share $\alpha_{k,ij}$ of the open carriers at k applies to customers (i,j) and a share α_k does not apply, with $\alpha_k + \sum_{(i,j)\in\mathcal{M}} \alpha_{k,ij} = 1$. Among the carriers that are not hired at k, a share σ_{kj} sails to j and a share σ_{kk} stays, with $\sum_j \sigma_{kj} = 1$.

Since carriers at the same location cannot coordinate where they apply, a customer in (i,j) receives a Poisson number of applications from location k, independently across k, with mean equal to the *queue*

$$\theta_{k,ij} = \frac{\alpha_{k,ij}\, v_k^{\mathcal{O}}}{x_{ij}}. \tag{1}$$

Some customers therefore receive no applicant while others receive several, of which all but one are turned down. Which applicant a customer hires is set by the ranking it posts, and Lemma A.5 in Appendix A.1 shows that a customer ranks applicants by the joint surplus of a match,

$$\Delta_{k,ij} = w_{ij} + X^{\mathrm{F}}_{k,ij} + V^{\mathrm{F}}_{k,ij} - V^{\mathrm{U}}_{k}, \tag{2}$$

where $X^{\mathrm{F}}_{k,ij}$ is the customer's value of waiting for a carrier from k, $V^{\mathrm{F}}_{k,ij}$ the carrier's value of being hired from k, and V^{U}_{k} its value of not being hired at k. The fee is a transfer and cancels from $\Delta_{k,ij}$. The customer's waiting value follows from the timing: it pays c_i every period until the carrier reaches i, so that

$$X^{\mathrm{F}}_{k,ij} = -c_i + \beta(1-d_{ki})X^{\mathrm{F}}_{k,ij} = -\frac{c_i}{1-\beta(1-d_{ki})}, \qquad X^{\mathrm{F}}_{i,ij} = -c_i. \tag{3}$$

Hiring a carrier from another location widens the customer's pool of applicants but makes it wait while the carrier sails over.

Order the locations that apply to (i,j) by surplus, as Lemma A.5 justifies, $\Delta_{(1),ij} \geq$

$\Delta_{(2),ij} \geq \ldots$, and let $\Theta_{\leq r} = \sum_{s \leq r} \theta_{(s)}$ be the queue of applicants of rank r or better. The probability that a customer hires a carrier of rank r, because no better applicant arrives and at least one of rank r does, and the probability that a given applicant of rank r is hired are

$$\pi^{X}_{(r)} = e^{-\Theta_{\leq r-1}}\left(1 - e^{-\theta_{(r)}}\right), \qquad \pi^{V}_{(r)} = \frac{\pi^{X}_{(r)}}{\theta_{(r)}}. \tag{4}$$

A customer is matched with probability $q_{ij} = \sum_r \pi^{X}_{(r)} = 1 - e^{-\Theta_{\leq R}}$.

The meeting process is the main departure from Brancaccio et al. (2020) and Brancaccio et al. (2023), where ships and exporters meet through a matching function within each region and bargain over the price. Here search is directed and not local: customers post fees, carriers direct their applications, and a carrier at k is a legitimate applicant for a cargo at $i \neq k$, typically ranked below the carriers already at i but hired when none of them applies.

Discussion of the assumptions. Four assumptions shape the results. First, matching is remote: a customer can hire a carrier wherever it is, not only one already at the cargo. This nests local matching, the assumption of existing search models of transport markets: a carrier from another location is typically ranked below the carriers already at the cargo and hired only when none of them applies, and when a remote hire has no surplus, no customer makes one and the two models coincide (Section 4.5). Remote matching lets a carrier look for a contract without first sinking the cost of sailing to a cargo it may not win. It gives the customer a choice between paying for a carrier nearby and waiting for one that is farther away, and both the fee and the wait depend on how many carriers are at each location. It matters most when carriers near the cargo are scarce: when they are plentiful, customers hire them and local matching is a good approximation; when they are scarce, local matching rules out the hires that keep cargo moving. It is also how tankers are hired: very large crude carriers fixed to load in the Persian Gulf are on average almost 2,900 miles from the loading area when the contract is signed (Prochazka et al., 2019; Regli

and Nomikos, 2019). Second, a customer's fee may depend on where the carrier comes from. A customer cares about when its cargo is picked up, not about where the carrier is. In the model, carriers sail at the same speed, so a carrier's distance from the customer stands for the wait the customer expects: a carrier at k reaches the cargo at i in $1/d_{ki}$ periods on average, and a fee by location is a fee by expected pickup time. Tanker fixtures are negotiated bilaterally, ship by ship, through brokers, and nothing ties the freight rate agreed with a ship far from the loading port to the rate agreed with one already there; a posted menu by origin is the simplest way to let the fee reflect a carrier's position. It is one of the features that make the market efficient (Section 2.1.1). Third, a customer that is not matched in the period it enters leaves. A cargo has a loading window set by the schedules of the seller and the refinery, and a charterer that cannot find a ship for that window must make other arrangements. The assumption keeps the customer side stationary, and it does not rule out waiting: a customer that hires a remote carrier waits for it, so waiting is a choice of whom to hire rather than of when to look. Fourth, carriers at a location cannot coordinate where they apply, so the number of applicants a customer receives is random, as in Burdett et al. (2001) and Shimer (2005). Ownership is fragmented — the five largest tanker owners hold less than 13 percent of the fleet (Frittelli, 2024) — and owners decide through their own brokers, so some cargoes attract several ships and others none.

2.1 Exogenous customers

This section keeps the inflow of customers x_{ij} exogenous. It defines the equilibrium and shows that it is efficient. The steady state can depend on where the fleet starts, so the results hold from any starting distribution $\mathbf{v}_0$.

2.1.1 Equilibrium

Stocks. The state at the start of date t is $\mathbf{v}_t = (v_t^{\mathcal{O}}, v_t^{\mathrm{T}}, v_t^{\mathrm{F}})$: open, in transit, and fixed carriers. Write $a_{k,ij,t} = \alpha_{k,ij,t} v_{k,t}^{\mathcal{O}}$ for the applications from k to customers (i,j), so that

$\theta_{k,ij,t} = a_{k,ij,t}/x_{ij}$, $m_{k,ij,t}$ for the hires from k by customers (i,j), and $u_{k,t} = v^{\mathcal{O}}_{k,t} - \sum_{(i,j)} m_{k,ij,t}$ for the unmatched carriers at k. The state moves by

$$\begin{aligned} v^{\mathcal{O}}_{k,t+1} &= \underbrace{\textstyle\sum_h d_{hk}\left(\sigma_{hk,t}u_{h,t} + v^{\mathrm{T}}_{hk,t}\right)}_{\text{arrive at } k}, \\ v^{\mathrm{T}}_{kj,t+1} &= \underbrace{(1-d_{kj})\left(\sigma_{kj,t}u_{k,t} + v^{\mathrm{T}}_{kj,t}\right)}_{\text{on the trip } k\to j\text{, not yet arrived}} + \underbrace{\textstyle\sum_h d_{hk}\left(m_{h,kj,t} + v^{\mathrm{F}}_{h,kj,t}\right)}_{\text{load cargo at } k \text{ bound for } j}, \\ v^{\mathrm{F}}_{k,ij,t+1} &= \underbrace{(1-d_{ki})\left(m_{k,ij,t} + v^{\mathrm{F}}_{k,ij,t}\right)}_{\text{still on the way to the cargo at } i}, \end{aligned} \tag{5}$$

written $\mathbf{v}_{t+1} = F(\mathbf{v}_t; \alpha_t, \sigma_t, m_t)$.

- *Open carriers at k* next period are those arriving at k: unmatched carriers that set off for k ($\sigma_{hk,t}u_{h,t}$, including those that stayed, $h = k$, $d_{kk} = 1$) and carriers already sailing to k ($v^{\mathrm{T}}_{hk,t}$), each arriving with probability d_{hk}.

- *Carriers in transit $k \to j$* come from two places: those already on that trip, empty or laden, that have not arrived (probability $1 - d_{kj}$), and those that load cargo at k bound for j. The latter were hired by customers (k,j) from some origin h, this period ($m_{h,kj,t}$) or earlier ($v^{\mathrm{F}}_{h,kj,t}$), and reach k with probability d_{hk}; a local hire ($h = k$) loads at once. If no customers ship from k to j, this term is zero.

- *Fixed carriers* hired at k by customers (i,j) are those, hired this period or earlier, that have not yet reached the cargo at i (probability $1 - d_{ki}$); a local hire is never fixed for more than a period, since $d_{ii} = 1$.

Carrier values. Along an equilibrium, the values of carriers at date t are as follows. A carrier in transit $k \to j$, a carrier hired at k by a customer (i,j) and sailing to i, a carrier that

is not hired at k, and an open carrier at k are worth

$$V^{\mathrm{T}}_{kj,t} = -c_{kj} + \beta d_{kj} V^{\mathcal{O}}_{j,t+1} + \beta(1-d_{kj}) V^{\mathrm{T}}_{kj,t+1}, \tag{6}$$

$$V^{\mathrm{F}}_{k,ij,t} = -c_{ki} + \beta d_{ki} V^{\mathrm{T}}_{ij,t+1} + \beta(1-d_{ki}) V^{\mathrm{F}}_{k,ij,t+1}, \tag{7}$$

$$V^{\mathrm{U}}_{k,t} = \textstyle\sum_j \sigma_{kj,t} V^{\mathrm{T}}_{kj,t}, \qquad V^{\mathcal{O}}_{k,t} = V^{\mathrm{U}}_{k,t} + S_{k,t},$$

where $S_{k,t}$ is the rent an applicant from k earns at date t. In a steady state the values do not depend on t, and I drop the subscript. Staying at k is the trip $k \to k$: it costs c_{kk} and ends next period, $d_{kk} = 1$; a carrier hired at the cargo's own location ($k = i$) spends one period there before the leg $i \to j$ starts. Appendix A.1 rewrites these values relative to a carrier that stays put forever, which is how the two-location example of Section 4 reads them.

Next, I define the equilibrium, show that it is efficient (Proposition 1), and use this to show that it exists and is unique (Proposition 2). Both rest on how customers rank applicants: a customer that posts its ranking with its fees ranks them by the surplus Δ (Lemma A.5). This order need not follow distance: a closer carrier has a cheaper pickup but can lose more by being hired.

Equilibrium. Customers post fees and rankings, carriers apply and relocate, and the stocks follow carriers' choices. Write $\Theta_{\succ k,ij,t} = \sum_{h \succ_{ij,t} k} \theta_{h,ij,t}$ for the queue of applicants that a customer (i,j) ranks above location k.

The carriers' problem. A carrier takes the posted fees, the rankings and the queues as given. By (4), an applicant from k to a customer (i,j) is hired with probability

$$\pi^{\mathrm{V}}_{k,ij,t} = e^{-\Theta_{\succ k,ij,t}} \frac{1 - e^{-\theta_{k,ij,t}}}{\theta_{k,ij,t}}, \tag{8}$$

which is $e^{-\Theta_{\succ k,ij,t}}$ if $\theta_{k,ij,t} = 0$. For a random ranking, this and the hiring probabilities below are averages over rankings, and they are zero for a location the customer leaves out. If hired, it receives the fee $\phi_{k,ij,t}$ and becomes a fixed carrier; if not, it chooses a trip. An open

carrier at k therefore solves

$$V^{\mathcal{O}}_{k,t} = \max\Big\{ V^{\mathrm{U}}_{k,t},\ \max_{(i,j)} \Big[\pi^{\mathrm{V}}_{k,ij,t}\big(\phi_{k,ij,t} + V^{\mathrm{F}}_{k,ij,t}\big) + \big(1 - \pi^{\mathrm{V}}_{k,ij,t}\big) V^{\mathrm{U}}_{k,t} \Big] \Big\},$$
$$V^{\mathrm{U}}_{k,t} = \max_{j} V^{\mathrm{T}}_{kj,t}, \tag{9}$$

with V^{T} and V^{F} given by (6) and (7). Its application choice $\alpha_{k,t}$ puts weight only on the options that attain the first maximum, not applying ($\alpha_{k,t}$) or applying to a customer (i,j) ($\alpha_{k,ij,t}$), and its relocation choice $\sigma_{k,t}$ only on the trips that attain the second, $\sigma_{kj,t} > 0 \Rightarrow j \in \arg\max_{j'} V^{\mathrm{T}}_{kj',t}$. The rent of an applicant from k is $S_{k,t} = V^{\mathcal{O}}_{k,t} - V^{\mathrm{U}}_{k,t} \geq 0$.

The customers' expected payoff. A customer (i,j) that hires a carrier from k obtains w_{ij}, pays the fee $\phi_{k,ij,t}$, and waits for the carrier, which is worth $X^{\mathrm{F}}_{k,ij}$ by (3); a customer that hires no one gets zero. By (4), it hires from k with probability $\pi^{\mathrm{X}}_{k,ij,t} = e^{-\Theta_{\succ k,ij,t}}\big(1 - e^{-\theta_{k,ij,t}}\big)$, so its expected payoff from posting $(\phi_{ij}, \succ_{ij})$ is

$$P_{ij,t}(\phi_{ij}, \succ_{ij}) = \sum_{k} \pi^{\mathrm{X}}_{k,ij,t}\big(w_{ij} + X^{\mathrm{F}}_{k,ij} - \phi_{k,ij}\big). \tag{10}$$

The queues it faces depend on what it posts. As in Shimer (2005), a customer expects carriers from k to apply until an applicant earns the rent $S_{k,t}$ it can get elsewhere, and none to apply if the customer offers less:

$$\pi^{\mathrm{V}}_{k,ij,t}\big(\phi_{k,ij} + V^{\mathrm{F}}_{k,ij,t} - V^{\mathrm{U}}_{k,t}\big) \leq S_{k,t}, \quad \text{with equality if } \theta_{k,ij,t} > 0. \tag{11}$$

A customer chooses its fees and ranking, and so its queues, to maximize (10) subject to (11).

Definition 1 (Equilibrium). *An equilibrium from* $\mathbf{v}_0$ *is a path of carrier values* $(V^{\mathcal{O}}, V^{\mathrm{U}}, V^{\mathrm{T}}, V^{\mathrm{F}})_t$, *bounded over time, carrier choices* (α_t, σ_t), *posted fees and rankings* $(\phi_t, \succ_t)$, *and stocks* $\mathbf{v}_t$ *such that, at every date:*

(i) given the fees, rankings and queues, the values satisfy (6), (7) *and* (9), *and* (α_t, σ_t) *attain the*

maxima in (9);

(ii) given carriers' rents $S_{k,t}$*, each customer's fees and ranking maximize its expected payoff* (10) *subject to* (11);

(iii) the queues are consistent with carriers' choices, $\theta_{k,ij,t} = \alpha_{k,ij,t} v^{\mathcal{O}}_{k,t} / x_{ij}$;

(iv) the stocks follow carriers' choices, $\mathbf{v}_{t+1} = F(\mathbf{v}_t; \alpha_t, \sigma_t, m_t)$*, with* m_t *the hires implied by the rankings.*

Efficiency. The total surplus of customers and carriers at date t is

$$\Pi_t = \sum_{k,(i,j)} m_{k,ij,t} \left(w_{ij} + X^{\mathrm{F}}_{k,ij} - c_{ki} \right) - \sum_{k,j} c_{kj} \left(\sigma_{kj,t} u_{k,t} + v^{\mathrm{T}}_{kj,t} \right) - \sum_{k,(i,j)} c_{ki}\, v^{\mathrm{F}}_{k,ij,t}; \tag{12}$$

fees are transfers and cancel. Hires are feasible if and only if, for every customer type (i,j) and every group K of origins,

$$m_{k,ij,t} \geq 0, \qquad \sum_{k \in K} m_{k,ij,t} \;\leq\; x_{ij} \left(1 - e^{-\sum_{k \in K} \theta_{k,ij,t}} \right) : \tag{13}$$

customers hiring from K need an applicant from K. Every ranking of applicants produces hires in this set, and every point of the set is produced by one. The equilibrium is efficient if no feasible path gives more total surplus, that is, if it attains

$$W(\mathbf{v}_0) = \max_{\{\alpha_t, \sigma_t, m_t\}} \sum_{t \geq 0} \beta^t \, \Pi_t \quad \text{s.t. (13) and (5).} \tag{14}$$

Proposition 1 (Efficiency)**.** *Let* $\beta < 1$*. From any starting distribution* $\mathbf{v}_0$*, every equilibrium path attains* (14)*, and every path that attains it is an equilibrium path: the market is efficient. The fees that support it are*

$$\phi_{k,ij,t} = \frac{\theta_{k,ij,t}\, e^{\Theta_{\succ k,ij,t}}}{1 - e^{-\theta_{k,ij,t}}} \, S_{k,t} + V^{\mathrm{U}}_{k,t} - V^{\mathrm{F}}_{k,ij,t}. \tag{15}$$

The proof is in Appendix A.3.

The fees internalize the externalities that carriers impose on one another. An extra applicant from k to a customer (i, j) is hired only if no applicant ranked at or above k is there, and it then adds the surplus $\Delta_{k,ij}$; but it also displaces the lower-ranked applicants who would otherwise have been hired. In equilibrium, an applicant's rent is this net contribution, its marginal product:

$$\begin{gathered} S_{k,t} \;\geq\; e^{-\Theta_{\succeq k,ij,t}}\Delta^{+}_{k,ij,t} - \textstyle\sum_{h\prec k} e^{-\Theta_{\succ h,ij,t}}\left(1-e^{-\theta_{h,ij,t}}\right)\Delta^{+}_{h,ij,t}, \quad \text{with equality if } \alpha_{k,ij,t}>0, \\ S_{k,t} \geq 0, \quad \text{with equality if } \alpha_{k,t}>0, \end{gathered} \tag{16}$$

where $\Delta^{+} \equiv \max\{\Delta, 0\}$. Carriers therefore apply where they contribute most, and only if the rent is nonnegative.

Efficiency rests on three features. Customers post fees by carrier origin, so every carrier earns its marginal product, S_k, and the same rent at every customer type it applies to. Customers post their ranking with their fees, so the ranking that attracts carriers is also the one they use. And the cost of a queue does not depend on the ranking, so customers choose the ranking that maximizes surplus, by Δ (Shi, 2002; Shimer, 2005).

Proposition 1 turns the equilibrium, a dynamic game among a continuum of customers and carriers, into the concave problem (14). Existence and uniqueness follow from it, as in Shimer (2005).

Proposition 2 (Existence and uniqueness)**.** *Let $\beta < 1$. From any starting distribution $\mathbf{v}_0$, an equilibrium exists. If along it no surplus $\Delta_{k,ij,t}$ is zero and no two surpluses of a customer type tie, all equilibria have the same queues and hires, apart from applications from an origin with $\Delta_{k,ij,t} < 0$ and $S_{k,t} = 0$, which are never accepted. If, in addition, no two trip values $V^{\mathrm{T}}_{kj,t}$ from a location tie, all equilibria have the same path of stocks, applications, hires and relocations, apart from those applications.*

The proof is in Appendix A.4.

The start matters for the path, and it can matter for the steady state, but not for efficiency.

In the two-location example of Section 4, the steady state does not depend on the start, except when there is no trade and carriers stay wherever they start. The path does: carriers that start at the cargo's origin deliver once even when no trade survives in the long run. Whether the path from every start converges to the steady state is not established here; in the example, holding carriers' steady-state behavior fixed, the fleet spread evenly over open carriers is within 10^{-6} of the steady state after 13 to 29 periods.

3 Disruptions

This section shows that a disruption on one route has network effects: it moves fees on routes that do not use the chokepoint, because carriers move between routes. I use two small examples with three locations, meant to show the mechanism rather than to match the data. In the first, closing Bab el-Mandeb lengthens voyages from the Persian Gulf to Europe, and the fee on the Gulf–Asia route, which does not use the passage, rises or falls depending on whether the detour is worth taking (Section 3.1). In the second, a toll for leaving the Persian Gulf through the Strait of Hormuz lowers the fee on the route from West Africa to Asia, which avoids the strait, and Gulf fees rise by less than the toll (Section 3.2). Both examples are solved exactly with the model of Section 2.1, using the efficiency of the equilibrium (Proposition 1).

3.1 Closing Bab el-Mandeb

When Bab el-Mandeb is closed to shipping, voyages between the Persian Gulf and Europe go around the Cape of Good Hope instead of through the Red Sea and the Suez Canal. Voyages between the Gulf and East Asia do not use the passage. Does the closure change what happens on them? This subsection answers the question in a small example with three locations, meant to show the mechanism rather than to match the data.

3.1.1 Setup

The locations are the Persian Gulf, Europe and East Asia (Figure 1). Cargo is loaded in the Gulf, for Europe or for Asia, so there are two customer types, served by one fleet. Customers are exogenous, as in Section 2.1. The closure lengthens every trip between the Gulf and Europe and between Europe and Asia, laden or empty, and leaves trips between the Gulf and Asia unchanged. It is unanticipated and permanent: the economy starts in the steady state with the passage open, the trip lengths change, and it moves to the steady state with the passage closed. Table 3 gives the parameters. They are rough and only illustrative: trip lengths follow from voyage distances at about 13 knots, a period is five days, and money is in millions of dollars. A trip to Europe is worth more to a customer than a trip to Asia; how much more turns out to matter, and Section 3.1.3 varies it.

3.1.2 The closure in the steady state

Table 4 compares the steady states with the passage open and closed. Longer voyages tie up the fleet: the share of carriers at sea rises from 72 to 79 percent, and the carriers waiting in the Gulf fall by a third. Customers shipping to Europe, whose trips are valuable, keep shipping, but the fee for a carrier rises by 80 percent. Customers shipping to Asia, whose route is unchanged, ship a quarter less, and their fee rises by 12 percent: the longer voyages leave fewer carriers waiting in the Gulf for them. Carriers gain at every location, because the fleet is scarcer; the cost falls on customers. Remote matching shifts as well. With fewer carriers in the Gulf, customers shipping to Europe hire more carriers that are still elsewhere, while customers shipping to Asia stop hiring them.

3.1.3 When the closure spills over

Whether the closure raises the fee on the Asia route depends on how much customers value a trip to Europe. Figure 2 repeats the comparison for values of a trip to Europe from 5 to 20. When a trip to Europe is worth little more than a trip to Asia, the closure makes

the longer voyage not worth serving fully: fewer European cargoes ship, carriers move to the Asia route, and its fee falls. When a trip to Europe is worth enough, European cargoes keep shipping, the longer voyages tie up the fleet, and the fee on the Asia route rises. The switch is at a value of about 8.4; near 9.5, European customers go from partly to fully served, which is why the curves jump. A closure therefore spills over to the routes that do not use the passage, and the sign of the spillover depends on whether demand on the disrupted route is strong enough to keep it shipping.

3.2 A toll for leaving the Gulf

In July 2026, the United States proposed charging ships that pass the Strait of Hormuz a toll, and Iran replied that it, not the United States, should collect one.[5] This subsection asks a hypothetical question: what would happen if such a toll were imposed, and what would its network effects be, on a route that does not pass the strait?

3.2.1 Setup

The locations are the Persian Gulf, West Africa and East Asia (Figure 3). Crude is loaded in the Gulf and in West Africa, both for Asia, so there are two customer types. Voyages from the Gulf to Asia pass the strait; voyages from West Africa go around the Cape of Good Hope and do not. A carrier pays a toll τ every time it leaves the Gulf, loaded or empty; in the steady states below, no carrier leaves the Gulf empty, so the toll is paid once per Gulf cargo. The toll leaves the market. Nothing else changes: trip lengths, costs, the values of a trip, the inflow of customers and the fleet are those of Table 5. As in Section 3.1, the parameters are rough and only illustrative: a period is five days, and money is in millions of dollars, so $\tau = 1$ is a toll of one million dollars per cargo, a sixth of the value of a trip from the Gulf.

[5] The proposal was 20 percent of the value of the cargo; Iran's foreign minister answered that "20% is of course too much. We will be fair." The proposal was dropped the next day. CNBC, 13 July 2026.

In the model, the toll is a cost on the trip from the Gulf. A carrier hired by a Gulf customer pays it when it leaves the Gulf, so the value of being hired becomes

$$V^{\mathrm{F}}_{k,ij} = -c_{ki} + \beta\Big[d_{ki}\big(V^{\mathrm{T}}_{ij} - \tau\big) + (1 - d_{ki})V^{\mathrm{F}}_{k,ij}\Big], \qquad i = \text{Gulf},$$

and an unmatched carrier in the Gulf compares $V^{\mathrm{T}}_{\mathrm{Gulf},j} - \tau$ across $j \neq$ Gulf with staying. The equilibrium is still efficient, since the toll is a cost like the sailing costs, and I solve it through the planner's problem as before.[6]

3.2.2 What the toll does

Figure 4 shows the steady state at every toll from 0 to 4.5: what happens to Gulf customers, to West African customers and to carriers. Table 6 gives the numbers at four tolls. Two tolls, the dotted lines in the figure, divide it into three ranges. At $\tau \approx 0.81$, carriers waiting in Asia start to apply to West African customers instead of Gulf customers, and at $\tau \approx 4.42$, Gulf customers stop shipping.

Gulf customers. With a small toll, Gulf customers keep hiring carriers that are still in Asia, about a fifth of their hires, and wait about eight days for their carrier. The fee they pay rises from 4.35 with no toll to 4.75 at $\tau = 0.8$, and their chance of shipping falls a little, from 0.91 to 0.86. As carriers in Asia turn to West Africa, Gulf customers ship with probability 0.80 rather than 0.86, and between $\tau \approx 0.65$ and $\tau \approx 1.2$ they stop hiring carriers in Asia and their wait falls to five days. From then on, they hire only carriers already in the Gulf, and the toll changes little for them until it is large: their chance of shipping falls only from 0.77 to 0.70, and their fee rises from 4.80 to 5.49, while the toll rises from 1 to 4.4. Their expected payoff falls to zero, and at $\tau \approx 4.42$ they stop shipping.

[6] The equilibrium conditions hold to within 10^{-3}, except for tolls between 0.65 and 0.81, where carriers in Asia are indifferent about applying to Gulf customers: the planner then leaves their queue undetermined, and the hiring condition is off by up to 0.04. Hires, values and fees are unaffected.

West African customers. West African customers gain at every toll. Their fee falls from 7.47 with no toll to 6.92 at $\tau = 0.8$, and their chance of shipping rises from 0.78 to 0.85. At $\tau \approx 0.81$, the carriers in Asia that stop applying to Gulf customers apply to them: West African customers then ship almost with certainty, 14 percent of their hires are carriers still in Asia, and their wait for a carrier rises from five days to nine. As the toll rises further, more carriers wait in West Africa, 0.40 at $\tau = 1$ and 1.10 at $\tau = 4.4$, so fewer hires are remote, the wait falls back to about five days, and the fee keeps falling, to 2.03 when Gulf customers stop shipping.

Carriers. Carriers lose at every location, by about the same amount: the value of an open carrier falls by about a quarter at $\tau = 1$, and keeps falling as the toll rises. The toll lowers what a carrier earns from a Gulf cargo, and since a carrier can serve either route, it lowers what carriers earn everywhere. Carriers stay about as busy: the fleet is fixed, and West African customers, who ship almost with certainty from $\tau \approx 0.81$, cannot use more carriers, so carriers keep carrying Gulf cargo although they keep less of the fee. The share of the fleet at sea falls only from 0.75 to 0.69, and more of the carriers that wait do so in West Africa. When Gulf customers stop shipping, the carriers that served them have nothing to do: the share of the fleet at sea falls to 0.28, 5.4 of the 9 carriers sit idle in Asia, and an open carrier is worth about what staying idle forever is worth, -10.

The fee for West African cargo falls at every toll until Gulf customers stop shipping, and it falls at every toll from 0 to 3 in each of ten other parameter sets I tried, which vary the fleet from 6 to 12 carriers, the value of a trip from the Gulf from 3 to 10 and from West Africa from 9 to 15, and the inflows. A toll takes money from carriers, not their time: Gulf customers ship less, so carriers are no busier, and carriers are worth less at every location. Unlike the closure of Section 3.1, which lengthens voyages that are still worth making, the toll does not make carriers scarcer on the route that avoids the strait, in any of these cases.

4 Trip value: two locations

This section solves the model of Section 2.1, with exogenous customers, on the smallest network in which remote matching matters: cargo appears at one location and is delivered to the other. It shows what happens when trips are worth more. Customers hire carriers that are farther away and wait longer for them, and carriers sign a contract first, then sail to the cargo. Compared with an otherwise identical model with local matching, once trips are valuable enough, remote matching raises carrier values at both locations and shrinks the premium for being at the cargo. Closed-form cutoffs in the value of a trip separate the regimes the economy passes through (Section 4.2); the customer's storage cost decides which of them occur (Section 4.3); the carriers' values explain why carriers stop repositioning toward the cargo when trade is most valuable (Section 4.4); and the model with local matching shows which of these results depend on remote matching (Section 4.5). The application rents and the fees that local and remote carriers earn are in Appendix B.

4.1 Setup

Take two locations $\{1, 2\}$ and one customer type, cargo at 1 bound for 2, with inflow x and value of a trip $w \equiv w_{12}$; every loaded trip ends at 2. A customer hires a local carrier, from 1, or a remote one, from 2, and pays the storage cost c_1 for one period with a local carrier and until it arrives with a remote one. Write θ_1, θ_2 for the local and remote queues, $\Theta = \theta_1 + \theta_2$, and Δ_1, Δ_2 for the surpluses of a local and a remote hire. I read carriers' values relative to staying idle, as in Appendix A.1: y_k is what an open carrier at k is worth above staying idle at k forever, ρ_{kj} is the discount over a trip $k \to j$, and κ_{kj} is the cost of sailing empty from k to j and staying idle there forever, over staying idle at k forever.

I assume $\kappa_{12}, \kappa_{21} > 0$: staying idle forever where a carrier is beats sailing empty to the other location and staying idle there. So if no carrier could ever be hired, no carrier would

have an incentive to relocate.[7] By Lemma A.2, an unmatched carrier at k gains

$$V^{\mathrm{T}}_{kj} - V^{\mathrm{T}}_{kk} = \rho_{kj} y_j - \beta y_k - \kappa_{kj}$$

from sailing empty to j instead of staying: it sails only if being open at j is worth enough more than being open at k to cover κ_{kj}.

I do not assume which applicant ranks first. The ranking is an outcome of the equilibrium values (Lemma A.5 in Appendix A.1), not of distance: hiring a local carrier uses up a carrier already at the valuable location, so when being at 1 is valuable enough, a remote applicant can rank first. I study the case in which local applicants rank first, $\Delta_1 > \Delta_2$, and check in the solution that it holds throughout. Unmatched carriers at 1 then always stay.

The surpluses of a local and a remote hire, the carriers' values, and the queues are derived in Appendix A.5 (Lemmas A.7 and A.8). Given what carriers do, the queues come from fleet accounting and do not depend on w. I index them by the regimes of Table 7. In local hiring (regime 3), where only carriers at 1 apply and carriers at 2 sail back, the local queue $\theta^{(3)}$ solves $v/x = \theta^{(3)} + (1/d_{12} + 1/d_{21})(1 - e^{-\theta^{(3)}})$. In mixed matching (regime 4), where all carriers at 2 apply and those not hired sail back, the total queue $\Theta^{(4)}$ solves $v/x = \Theta^{(4)} + (1/d_{12} + 1/d_{21} - 1)(1 - e^{-\Theta^{(4)}})$, split into a remote queue $\theta^{(4)}_2 = 1 - e^{-\Theta^{(4)}}$ and a local queue $\theta^{(4)}_1 = \Theta^{(4)} - \theta^{(4)}_2$. In remote hiring (regime 5), the queues are $(0, \Theta^{(4)})$.

4.2 Regimes

As w rises, carriers at 2 change what they do in three steps: they start sailing back to serve customers at 1, they start applying remotely, and they stop sailing back. Each step happens over an interval of w, along which one indifference condition holds while the local queue θ_1 moves. Proposition 3 gives the ends of these intervals in closed form, and Table 7 lists the regimes they separate.

[7] An open carrier would then be worth only its best trip, $y_k = \max\{\beta y_k, \rho_{kj} y_j - \kappa_{kj}\}$ by Lemma A.2. At the location with the largest y, both terms are below that y unless it is zero, so $y = 0$ and no carrier would move.

Proposition 3 (Cutoffs). *There are seven cutoffs $\bar{w}_1, \ldots, \bar{w}_7$ in the value of a trip, given in closed form in Appendix A.5, such that:*

(i) a carrier at 1 *applies iff $w > \bar{w}_1 \equiv c_1 + \beta\kappa_{12}$, when carriers at* 2 *never come back;*

(ii) carriers at 2 *are indifferent between staying and sailing back, with no rent, at $\bar{w}_2 = \bar{w}_1 + \kappa_{21}/\rho_{21}$, with no carrier at* 1*, and at $\bar{w}_3 = \bar{w}_1 + \left(\beta + (1-\beta)e^{\theta^{(3)}}\right)\kappa_{21}/\rho_{21}$, with the local queue of local hiring;*

(iii) a remote hire has zero surplus, while carriers at 2 *sail back, at $\bar{w}_4$, with the local queue $\theta^{(3)}$ of local hiring, and at $\bar{w}_5$, with the local queue $\theta_1^{(4)}$ of mixed hiring;*

(iv) unmatched carriers at 2 *are indifferent between staying and sailing back, while all carriers at* 2 *apply, at $\bar{w}_6$, with the local queue $\theta_1^{(4)}$, and at $\bar{w}_7$, with no carrier at* 1*.*

Every cutoff is affine in the storage cost c_1.

The proof is in Appendix A.5.

In (ii) to (iv), carriers at 2 are indifferent over an interval of w, between two consecutive cutoffs, and a growing share of them switches: the share that switches is what keeps them indifferent, and given w it is unique. These intervals are short, and I count each as the start of the regime it leads to. When $\bar{w}_3 < \bar{w}_4$ and $\bar{w}_5 < \bar{w}_6$, the long run passes through five regimes as w rises (Table 7). Regime 2 differs from regime 1 only along the path from a start with carriers at 1: they deliver once and then stay at 2.

Example. With $\beta = 0.95$, $d_{12} = d_{21} = 0.5$, $c_{11} = 0.3$, $c_{22} = 0.295$, $c_{12} = c_{21} = 0.4$, $c_1 = 1$, $x = 0.5$, $v = 1$, idling is slightly cheaper at 2, and $\kappa_{12} = 0.100$, $\kappa_{21} = 0.290$. The queues are $\theta^{(3)} = 0.479$, $\Theta^{(4)} = 0.618$ and $(\theta_1^{(4)}, \theta_2^{(4)}) = (0.157, 0.461)$, and the cutoffs are $\bar{w}_1 = 1.095$, $\bar{w}_2 = 1.416$, $\bar{w}_3 = 1.426$, $\bar{w}_4 = 3.873$, $\bar{w}_5 = 4.419$, $\bar{w}_6 = 22.750$ and $\bar{w}_7 = 23.776$. Matches rise from 0.190 to 0.230 once carriers at 2 apply, and the remote share of matches rises $0 \to 68.5\% \to 100\%$.

The customer's tradeoff. A customer at 1 chooses its queues, θ_1 local and θ_2 remote applicants, and with local applicants ranked first it ends in one of three ways (Table 8). A remote applicant is a backup: it is hired only when no local carrier applies. A longer remote queue lets the customer ship instead of leaving, which gains w, but makes it wait for the carrier to sail over, which costs the extra storage $c_1/(1-\beta(1-d_{21}))-c_1$ by (3). That cost does not depend on w.

Waiting times. As the value of a trip rises, customers are willing to wait longer for a carrier: the cost of waiting does not change, but the trip it secures is worth more, so customers attract a longer remote queue and wait while remote carriers sail over. A local carrier loads within the period, while a remote one takes $1/d_{21}$ periods on average to arrive, so the average wait of a hired customer is one period plus the remote share of hires times $1/d_{21}-1$. In the example it rises from one period in local hiring, where all hires are local, to about 1.7 periods in mixed hiring and two periods in remote hiring. It rises as carriers at 2 start applying and as unhired carriers there start staying, while the remote share grows, and is constant otherwise, where the queues do not change.

4.3 Storage cost

The storage cost c_1 is the second parameter I vary. It sets the cost of a remote hire relative to a local one: a customer pays c_1 for one period with a local carrier, but for the whole pickup trip with a remote one. The queues do not depend on c_1, and every cutoff is affine in it. The cutoffs $\bar{w}_1$ to $\bar{w}_3$ rise one for one with c_1, while those at which remote hiring matters rise faster, in the example with slopes of about 4 for $\bar{w}_4$ and $\bar{w}_5$ and 24 for $\bar{w}_6$ and $\bar{w}_7$: a higher storage cost delays remote hiring. Figure 5 shows the long-run regimes in the (c_1, w) plane, drawn from these cutoffs. There are two thresholds on the storage cost. Let $\bar{c}$ be the storage cost at which $\bar{w}_3 = \bar{w}_4$. Above it, the economy passes through the five

regimes of Table 7 as w rises. Let

$$w_D \equiv -X^{\mathrm{F}}_{2,12} + \kappa_{21} + \rho_{21}\kappa_{12}$$

be the value of a trip above which a remote hire has positive surplus without rents: the value of a trip, net of the customer's wait for a remote carrier, $X^{\mathrm{F}}_{2,12}$, covers the extra cost of the pickup and delivery trips. Let $\underline{c} < \bar{c}$ be the storage cost at which $w_D = \bar{w}_2$. Below it, remote hires pay before carriers at 2 would sail back, and carriers at 2 never do: trade moves from regime 2 straight to regime 5 at $w = w_D$. In the example $\underline{c} = 0.039$ and $\bar{c} = 0.050$; both are in closed form in Appendix A.5. Between them, the regimes in which carriers at 2 sail back shrink and disappear one by one.[8] When waiting is cheap, customers wait for remote carriers instead of carriers repositioning toward the cargo.

4.4 Why carriers stay at the destination

In regime 1 the value of an open carrier at each location equals the value of staying idle forever there, $\bar{V}_k$, and since idling is cheaper at 2, $V^{\mathcal{O}}_2 > V^{\mathcal{O}}_1$. In regime 2 a carrier at 1 can be hired, and $V^{\mathcal{O}}_1$ overtakes $V^{\mathcal{O}}_2$ at $w \approx 1.19$. From then on being at the cargo's origin is worth more and the gap grows with w, yet in regime 5 unmatched carriers at 2 stay. By Lemma A.2, an unmatched carrier at 2 sails back if and only if

$$\underbrace{\rho_{21}(y_1 - y_2)}_{\text{gain}} > \underbrace{\kappa_{21} + (\beta - \rho_{21})\, y_2}_{\text{cost}},$$

where $y_k = V^{\mathcal{O}}_k - \bar{V}_k$. The gain is the premium of being at 1, discounted for the trip. The cost is the trip, κ_{21}, and the time lost sailing instead of applying at 2, $(\beta - \rho_{21})y_2$, which grows with the value of being open at 2. Figure 6 plots both. In regimes 1 and 2 there is

[8] In this narrow range the cutoffs are not ordered, and the figure clips them from above so that they are. A numerical solution of the model from its primitives, without the cutoffs, gives the same regime at all but 4 of 83,200 points of a grid, all in this range (Appendix A.5).

nothing to gain at 1 and the trip alone keeps carriers at 2. From local hiring through mixed hiring the gain exceeds the cost. Both then grow roughly in proportion to y_2, but the cost faster: the lines cross as remote hiring begins, and in remote hiring two periods without applying at 2 cost more than being one pickup trip closer to the cargo is worth.

4.5 Remote versus local matching

To see which of these results depend on remote matching, I solve an otherwise identical model with local matching, as in the existing models of transport markets: a customer at 1 can hire only carriers at 1. Everything else, including the fleet, the costs and the inflow of customers, is unchanged. Both models are solved exactly, using the efficiency of the equilibrium (Proposition 1), and the values, rents and fees are read off as in the rest of this section and Appendix B. Table 9 summarizes the comparison, and Figures 7 to 11 show it for every value of a trip; Appendix B compares the fees.

Regimes. With local matching, carriers at 2 cannot apply, so only regimes 1 to 3 exist. Their cutoffs are those of Proposition 3, $\bar{w}_1$, $\bar{w}_2$ and $\bar{w}_3$: below $\bar{w}_4 \approx 3.9$ no customer hires a remote carrier, and the two models coincide. Above it, the model with local matching stays in local hiring however valuable trips become, while the model with remote matching passes through mixed hiring and remote hiring. Every cutoff of the model with local matching is c_1 plus a constant, so the storage cost only shifts the value of a trip at which trade starts; it never moves the economy out of local hiring (Figure 7). The result of Section 4.3, that customers wait for remote carriers when waiting is cheap, has no counterpart.

Repositioning, waiting and the fleet. With local matching, sailing back is the only way for a carrier at 2 to be hired, so every unmatched carrier at 2 sails back, 0.190 per period, and a hired customer always waits one period. With remote matching, empty sailings fall

to 0.073 once all carriers at 2 apply and stop in remote hiring, and the wait rises to 1.7 and then two periods. Customers attract a longer queue, $\Theta^{(4)} = 0.618$ against $\theta^{(3)} = 0.479$, and ship with probability 0.461 against 0.380. The fleet waits elsewhere: with local matching 0.239 carriers are open at 1 and 0.190 at 2; with remote matching, 0.078 at 1 and 0.230 at 2 in mixed hiring, and none at 1 and 0.309 at 2 in remote hiring (Figure 8). Because the inflow of customers is exogenous and the fleet is fixed, these quantities depend only on the regime: they are constant within each regime in both models, and with local matching they do not change with w at all once every unmatched carrier at 2 sails back.

Carrier values and rents. Above $\bar{w}_4$, carriers are worth more at both locations with remote matching: from $w = 5$ on, about 15 to 18 percent more at 1 and 22 percent more at 2 (Figure 9). The location without cargo gains more, so local matching overstates the premium of being near the cargo: at $w = 10$ an open carrier at 1 is worth 2.5 more than one at 2 with remote matching and 3.5 more with local matching; at $w = 30$, 6.2 against 11.0. With local matching, a carrier at 2 earns no application rent, $S_2 = 0$, since it cannot apply; its only way to a contract is to sail back. With remote matching it earns a rent from mixed hiring on, 1.39 at $w = 10$ and 6.18 at $w = 30$. The rent of a carrier at 1 is also higher, 1.93 against 1.71 at $w = 10$ (Figure 10).

Customers and welfare. Remote matching raises welfare, the discounted surplus from a start with all carriers open, half at each location, by 11 percent at $w = 5$, 20 percent at $w = 10$ and 24 percent at $w = 30$. The gain grows with w, because the time a carrier would spend sailing without a contract is worth more when trips are more valuable. Customers do not always share in it. For w between 3.9 and 6.7, their expected payoff is lower with remote matching, 0.04 against 0.10 at $w = 5$: a remote hire has little surplus just above $\bar{w}_4$, and since fewer carriers wait at 1, a customer hires a local carrier with probability 0.145 rather than 0.380, and pays it more. Above 6.7 customers gain, and at $w = 30$ their expected payoff is about twice as high, 1.46 against 0.77 (Figure 11).

5 Conclusion

This paper develops a dynamic spatial model of directed search in which carriers and customers match remotely. Customers post fees by carrier location and rankings of applicants; carriers apply from anywhere, and choose where to wait and whether to sail empty. A remote hire makes the customer wait, so fees, waiting times and the distribution of carriers are determined together.

Two results follow. First, disruptions at chokepoints have network effects: they move fees on routes that do not use the chokepoint, because carriers move between routes. Closing Bab el-Mandeb raises the fee on the Gulf–Asia route when the detour to Europe is worth taking and lowers it otherwise. A toll for leaving the Gulf through the Strait of Hormuz lowers the fee on the West Africa–Asia route; Gulf customers stop hiring carriers that are still in Asia and wait less for their carrier, and those carriers serve West Africa instead. Second, when trips are worth more, customers hire carriers that are farther away and wait longer for them, and carriers sign a contract first and then sail to the cargo; once trips are valuable enough, remote matching raises carrier values everywhere and shrinks the premium for being at the cargo.

References

Brancaccio, G., Kalouptsidi, M., and Papageorgiou, T. (2020). Geography, transportation, and endogenous trade costs. *Econometrica*, 88(2):657–691.

Brancaccio, G., Kalouptsidi, M., Papageorgiou, T., and Rosaia, N. (2023). Search frictions and efficiency in decentralized transport markets. *Quarterly Journal of Economics*, 138(4):2451–2503.

Buchholz, N. (2022). Spatial equilibrium, search frictions, and dynamic efficiency in the taxi industry. *Review of Economic Studies*, 89(2):556–591.

Burdett, K., Shi, S., and Wright, R. (2001). Pricing and matching with frictions. *Journal of Political Economy*, 109(5):1060–1085.

Fréchette, G. R., Lizzeri, A., and Salz, T. (2019). Frictions in a competitive, regulated market: Evidence from taxis. *American Economic Review*, 109(8):2954–2992.

Frittelli, J. (2024). The global oil tanker market: An overview as it relates to sanctions. Technical Report R47962, Congressional Research Service.

Huangfu, Q. and Hall, J. A. J. (2018). Parallelizing the dual revised simplex method. *Mathematical Programming Computation*, 10(1):119–142.

Kelley, James E., J. (1960). The cutting-plane method for solving convex programs. *Journal of the Society for Industrial and Applied Mathematics*, 8(4):703–712.

Lagos, R. (2000). An alternative approach to search frictions. *Journal of Political Economy*, 108(5):851–873.

McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38.

McKelvey, R. D. and Palfrey, T. R. (1998). Quantal response equilibria for extensive form games. *Experimental Economics*, 1(1):9–41.

Menzio, G. and Shi, S. (2011). Efficient search on the job and the business cycle. *Journal of Political Economy*, 119(3):468–510.

Moen, E. R. (1997). Competitive search equilibrium. *Journal of Political Economy*, 105(2):385–411.

Prochazka, V., Adland, R., and Wolff, F.-C. (2019). Contracting decisions in the crude oil transportation market: Evidence from fixtures matched with AIS data. *Transportation Research Part A: Policy and Practice*, 130:37–53.

Regli, F. and Nomikos, N. K. (2019). The eye in the sky: Freight rate effects of tanker supply. *Transportation Research Part E: Logistics and Transportation Review*, 125:402–424.

Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press, Princeton, NJ.

Shi, S. (2002). A directed search model of inequality with heterogeneous skills and skill-biased technology. *Review of Economic Studies*, 69(2):467–491.

Shimer, R. (2005). The assignment of workers to jobs in an economy with coordination frictions. *Journal of Political Economy*, 113(5):996–1025.

Turocy, T. L. (2005). A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence. *Games and Economic Behavior*, 51(2):243–263.

Turocy, T. L. (2010). Computing sequential equilibria using agent quantal response equilibria. *Economic Theory*, 42(1):255–269.

Tvedt, J. (2011). Short-run freight rate formation in the VLCC market: A theoretical framework. *Maritime Economics & Logistics*, 13(4):442–455.

UNCTAD (2023). Review of maritime transport 2023.

A Lemmas for Section 2

A.1 Carrier values relative to staying put

The steady-state values of Section 2.1.1 are easiest to read relative to a carrier that stays at k forever. Let

$$D_{kj} \equiv 1-\beta(1-d_{kj}), \qquad \rho_{kj} \equiv \frac{\beta d_{kj}}{D_{kj}}, \qquad \bar{V}_k \equiv -\frac{c_{kk}}{1-\beta}, \qquad y_k \equiv V_k^{\mathcal{O}} - \bar{V}_k,$$

$$\kappa_{kj} \equiv \frac{c_{kj}-c_{kk}}{D_{kj}} + \rho_{kj}\big(\bar{V}_k - \bar{V}_j\big).$$

$1/D_{kj}$ is the expected discounted length of a trip $k \to j$ and ρ_{kj} the expected discount factor at its end; y_k is what an open carrier at k is worth above staying put; κ_{kj} is the cost of a trip $k \to j$ that ends idle at j, over staying idle at k. For staying, $D_{kk} = 1$, $\rho_{kk} = \beta$ and $\kappa_{kk} = 0$.

Lemma A.1 (Staying put, split at a trip's end)**.** *For every trip $k \to j$,*

$$\bar{V}_k = -\frac{c_{kk}}{D_{kj}} + \rho_{kj}\,\bar{V}_k.$$

Proof. $D_{kj} - \beta d_{kj} = 1-\beta$, so $1-\rho_{kj} = (1-\beta)/D_{kj}$, and $-c_{kk}/D_{kj} = \bar{V}_k(1-\beta)/D_{kj} = (1-\rho_{kj})\bar{V}_k$. □

Staying idle at k forever is the same as staying idle for as long as a trip to j would last, then staying idle forever after.

Lemma A.2 (Trips)**.** *For all k and j, including $j = k$,*

$$V_{kj}^{\mathrm{T}} = \bar{V}_k - \kappa_{kj} + \rho_{kj}\,y_j, \qquad \textit{so} \qquad V_{kk}^{\mathrm{T}} = \bar{V}_k + \beta\,y_k.$$

Proof. Solving (6) in a steady state, $V_{kj}^{\mathrm{T}} = -c_{kj}/D_{kj} + \rho_{kj}(\bar{V}_j + y_j)$. Subtract Lemma A.1 and collect the definition of κ_{kj}. □

A trip is staying put, minus the extra cost of the trip, plus what the carrier is worth at the destination, discounted for the trip.

Lemma A.3 (Hired carriers). *A carrier hired at k by a customer (i, j) is worth*

$$V^{\mathrm{F}}_{k,ij} = \bar{V}_k - \kappa_{ki} - \rho_{ki}\,\kappa_{ij} + \rho_{ki}\rho_{ij}\,y_j.$$

In particular, a local hire is worth $V^{\mathrm{F}}_{i,ij} = \bar{V}_i - \beta\kappa_{ij} + \beta\rho_{ij}\,y_j$.

Proof. Solving (7) in a steady state, $V^{\mathrm{F}}_{k,ij} = -c_{ki}/D_{ki} + \rho_{ki}V^{\mathrm{T}}_{ij}$. By Lemma A.2 with $y_i = 0$, $-c_{ki}/D_{ki} + \rho_{ki}\bar{V}_i = \bar{V}_k - \kappa_{ki}$; substitute $V^{\mathrm{T}}_{ij} = \bar{V}_i - \kappa_{ij} + \rho_{ij}y_j$. For $k = i$, use $\kappa_{ii} = 0$ and $\rho_{ii} = \beta$. □

A hired carrier is worth staying put at k, minus the extra cost of the pickup trip $k \to i$, minus the extra cost of the delivery $i \to j$ (discounted for the pickup), plus what it is worth once open at j (discounted for both trips). $\bar{V}_k$ is only a benchmark: during the pickup the carrier pays c_{ki}, not c_{kk}, because κ_{ki} swaps one for the other. A local hire pays c_{ii} for the loading period, the same as staying put, so only the delivery adds an extra cost.

Lemma A.4 (Surplus of a match). *Let $u_k \equiv V^{\mathrm{U}}_k - \bar{V}_k = \sum_j \sigma_{kj}\big(\rho_{kj}y_j - \kappa_{kj}\big)$. Then*

$$\Delta_{k,ij} = w_{ij} + X^{\mathrm{F}}_{k,ij} - \kappa_{ki} - \rho_{ki}\,\kappa_{ij} + \rho_{ki}\rho_{ij}\,y_j - u_k, \qquad y_k = u_k + S_k.$$

Proof. Substitute Lemma A.3 and $V^{\mathrm{U}}_k = \sum_j \sigma_{kj}V^{\mathrm{T}}_{kj}$ (Lemma A.2) into (2); the benchmark $\bar{V}_k$ cancels. The second claim is $V^{\mathcal{O}}_k = V^{\mathrm{U}}_k + S_k$. □

The surplus of a match depends on staying costs only through the trips' extra costs κ, and on carriers' values only through what they are worth above staying put.

Lemma A.5 (Who gets hired). *Fix the queues $\theta_{\cdot,ij}$ of a customer type (i, j).*

(i) Among hires $m_{k,ij} \geq 0$ with $\sum_{k\in K} m_{k,ij} \leq x_{ij}\big(1 - e^{-\sum_{k\in K}\theta_{k,ij}}\big)$ for every group K of origins (hiring from K needs an applicant from K), those that maximize $\sum_k m_{k,ij}\Delta_{k,ij}$ hire in order of

Δ: *with* $h \succ k$ *iff* $\Delta_{h,ij} > \Delta_{k,ij}$,

$$m_{k,ij} = x_{ij}\, e^{-\Theta_{\succ k,ij}} \left(1 - e^{-\theta_{k,ij}}\right) \text{ if } \Delta_{k,ij} > 0, \qquad m_{k,ij} = 0 \text{ otherwise.}$$

(ii) A customer posting its fees and ranking maximizes its expected payoff by posting the Δ *ranking: among the applicants who arrive, it hires the one with the highest* Δ.

Proof. (i) Every ranking of applicants, deterministic or random, gives hires in this set, and every point of it is reached by one. The set function $K \mapsto x_{ij}\left(1 - e^{-\sum_K \theta_{k,ij}}\right)$ is submodular because $1 - e^{-x}$ is concave, so the set is a polymatroid, on which a linear objective is maximized greedily: fill the constraints in order of the weights Δ, highest first, and stop at $\Delta \leq 0$. The greedy hires of the top r origins total $x_{ij}(1 - e^{-\Theta_{\leq r}})$; their differences are the $m_{k,ij}$ above. (ii) Substituting the fees, the customer's payoff is $R_{ij}(\theta) - \sum_k \theta_{k,ij} S_k$, where R_{ij} is the expected surplus of the applicant it hires (Shimer, 2005). The cost of the queues does not depend on the ranking, and the Δ ranking hires the best applicant in every realization, hence gives the largest R_{ij}. □

The ranking of applicants is therefore chosen, by customers and by any allocation that maximizes surplus alike, and it is the same ranking. Because customers post it with their fees, carriers know their chance of being hired when they apply. It need not follow distance: by Lemma A.4, a closer carrier has a smaller pickup cost but can have a higher value of not being hired, u_k.

A.2 The surplus-maximizing paths

The proofs of Propositions 1 and 2 use the following characterization of the paths that attain (14).

Lemma A.6 (The surplus-maximizing paths)**.** *Let* $\beta < 1$. *Problem* (14) *is a concave program. A feasible path from* $\mathbf{v}_0$ *is optimal if it has bounded shadow values* V_t *that satisfy the value*

equations of Section 2.1.1 date by date, and if its applications satisfy (16), its relocations satisfy $\sigma_{kj,t} > 0 \Rightarrow j \in \arg\max_{j'} V^{\mathrm{T}}_{kj',t}$*, and its hires are those of Lemma A.5(i). Suppose that along this path no surplus* $\Delta_{k,ij,t}$ *is zero and no two surpluses of a customer type tie. Then every optimal path has the same queues and hires, apart from applications from an origin with* $\Delta_{k,ij,t} < 0$ *and* $S_{k,t} = 0$*, which are never accepted. If, in addition, no two trip values* $V^{\mathrm{T}}_{kj,t}$ *from a location tie, the optimal path is unique, apart from those applications.*

Proof of Lemma A.6. Write the date-t controls as flows: applications $a_{k,ij,t} = \alpha_{k,ij,t} v^{\mathcal{O}}_{k,t}$, hires $m_{k,ij,t}$, and relocations $s_{kj,t} = \sigma_{kj,t} u_{k,t}$, all nonnegative, with

$$\textstyle\sum_{(i,j)} a_{k,ij,t} \le v^{\mathcal{O}}_{k,t}, \qquad \sum_j s_{kj,t} = v^{\mathcal{O}}_{k,t} - \sum_{(i,j)} m_{k,ij,t}, \tag{A.1}$$

and $\theta_{k,ij,t} = a_{k,ij,t}/x_{ij}$ in (13). Every choice of $(\alpha_t, \sigma_t, m_t)$ gives such flows and every such flows come from a choice of $(\alpha_t, \sigma_t, m_t)$. Constraint (13) with $K = \{k\}$ gives $m_{k,ij,t} \le x_{ij}(1 - e^{-\theta_{k,ij,t}}) \le a_{k,ij,t}$, so no more carriers are hired at k than are open there.

Step 1: concavity. In the variables $(\mathbf{v}_t, a_t, m_t, s_t)_{t\ge 0}$, the surplus (12), the laws of motion (5), and (A.1) are linear. Each constraint in (13) reads $\sum_{k\in K} m_{k,ij,t} \le g_{ij}\big(\sum_{k\in K} a_{k,ij,t}\big)$ with $g_{ij}(z) = x_{ij}(1 - e^{-z/x_{ij}})$ concave, so the pairs (a_t, m_t) that satisfy it form a convex set. Problem (14) maximizes a linear objective over a convex set.

Step 2: the value of a path. Let V_t be the shadow values of the candidate path: they satisfy (6) and (7), and $V^{\mathcal{O}}_{k,t} = V^{\mathrm{U}}_{k,t} + S_{k,t}$ with $V^{\mathrm{U}}_{k,t} = \sum_j \sigma_{kj,t} V^{\mathrm{T}}_{kj,t} = \max_j V^{\mathrm{T}}_{kj,t}$ by the condition on relocations, and let $\Delta_{k,ij,t}$ be given by (2) at these values. Write $V_t \cdot \mathbf{v}_t \equiv \sum_k V^{\mathcal{O}}_{k,t} v^{\mathcal{O}}_{k,t} + \sum_{k,j} V^{\mathrm{T}}_{kj,t} v^{\mathrm{T}}_{kj,t} + \sum_{k,(i,j)} V^{\mathrm{F}}_{k,ij,t} v^{\mathrm{F}}_{k,ij,t}$ and, for any feasible path, $G_t \equiv \Pi_t + \beta V_{t+1} \cdot \mathbf{v}_{t+1} - V_t \cdot \mathbf{v}_t$. Then

$$\textstyle\sum_{t=0}^{T} \beta^t \Pi_t = V_0 \cdot \mathbf{v}_0 - \beta^{T+1} V_{T+1} \cdot \mathbf{v}_{T+1} + \sum_{t=0}^{T} \beta^t G_t.$$

The stocks add up to the fleet v and the values are bounded, so the middle term vanishes as $T \to \infty$, and

$$\textstyle\sum_{t\ge 0} \beta^t \Pi_t = V_0 \cdot \mathbf{v}_0 + \sum_{t\ge 0} \beta^t G_t. \tag{A.2}$$

All paths start from $\mathbf{v}_0$, so they differ only in $\sum_t \beta^t G_t$.

Step 3: computing G_t. Substitute (5) into $\beta V_{t+1} \cdot \mathbf{v}_{t+1}$ and collect terms. Carriers on the trip $k \to j$, $s_{kj,t} + v^{\mathrm{T}}_{kj,t}$, are next period open at j with probability d_{kj} and still on the trip otherwise; carriers hired at k by (i,j), $m_{k,ij,t} + v^{\mathrm{F}}_{k,ij,t}$, are next period on the leg $i \to j$ with probability d_{ki} and still fixed otherwise. By the value equations,

$$\beta V_{t+1} \cdot \mathbf{v}_{t+1} = \textstyle\sum_{k,j} \left(V^{\mathrm{T}}_{kj,t} + c_{kj}\right)\left(s_{kj,t} + v^{\mathrm{T}}_{kj,t}\right) + \sum_{k,(i,j)} \left(V^{\mathrm{F}}_{k,ij,t} + c_{ki}\right)\left(m_{k,ij,t} + v^{\mathrm{F}}_{k,ij,t}\right).$$

Add (12): the sailing costs cancel. Subtract $V_t \cdot \mathbf{v}_t$: the terms in v^{T}_t and v^{F}_t cancel. So

$$G_t = \textstyle\sum_{k,(i,j)} m_{k,ij,t}\left(w_{ij} + X^{\mathrm{F}}_{k,ij} + V^{\mathrm{F}}_{k,ij,t}\right) + \sum_{k,j} V^{\mathrm{T}}_{kj,t} s_{kj,t} - \sum_k V^{\mathcal{O}}_{k,t} v^{\mathcal{O}}_{k,t}.$$

Step 4: a bound on G_t. Drop the index t. Three inequalities hold for any feasible flows.

(a) Relocations: $\sum_j V^{\mathrm{T}}_{kj} s_{kj} \le V^{\mathrm{U}}_k \sum_j s_{kj} = V^{\mathrm{U}}_k \left(v^{\mathcal{O}}_k - \sum_{(i,j)} m_{k,ij}\right)$, with equality if and only if $s_{kj} > 0$ only for the best trips. With (2) and $V^{\mathcal{O}}_k = V^{\mathrm{U}}_k + S_k$, this gives $G \le \sum_{k,(i,j)} m_{k,ij}\Delta_{k,ij} - \sum_k S_k v^{\mathcal{O}}_k$.

(b) Applications: since $S_k \ge 0$ and $\sum_{(i,j)} a_{k,ij} \le v^{\mathcal{O}}_k$, $-S_k v^{\mathcal{O}}_k \le -S_k \sum_{(i,j)} x_{ij}\theta_{k,ij}$, with equality if and only if $S_k = 0$ or every open carrier at k applies.

(c) Hires: fix a customer type (i,j) and number the origins with $\Delta_{k,ij} > 0$ as $(1), \ldots, (R)$, with $\Delta_{(1)} \ge \cdots \ge \Delta_{(R)} > 0 \equiv \Delta_{(R+1)}$. Let $M_r = \sum_{s \le r} m_{(s),ij}$, the hires from the top r origins. Then

$$\begin{aligned}\textstyle\sum_k m_{k,ij}\Delta_{k,ij} &\le \textstyle\sum_{r=1}^{R} m_{(r),ij}\Delta_{(r)} = \sum_{r=1}^{R}\left(\Delta_{(r)} - \Delta_{(r+1)}\right)M_r \\ &\le \textstyle\sum_{r=1}^{R}\left(\Delta_{(r)} - \Delta_{(r+1)}\right)x_{ij}\left(1 - e^{-\Theta_{\le r}}\right) \equiv Y_{ij}(\theta_{\cdot,ij}).\end{aligned}$$

The first inequality drops the origins with $\Delta_{k,ij} \le 0$, the equality sums by parts, and the second inequality is (13) for the group of the top r origins, whose weights

$\Delta_{(r)} - \Delta_{(r+1)}$ are nonnegative. The hires of Lemma A.5(i) attain both: they hire no one with $\Delta_{k,ij} < 0$ and set $M_r = x_{ij}(1 - e^{-\Theta_{\leq r}})$ for every r.

Together, for any feasible path,

$$G_t \leq \sum_{(i,j)} H_{ij,t}(\theta_{\cdot,ij,t}), \qquad H_{ij,t}(\theta) \equiv Y_{ij,t}(\theta) - x_{ij} \sum_k S_{k,t}\theta_k. \tag{A.3}$$

$H_{ij,t}$ depends on the path only through its queues at date t.

Step 5: the candidate's queues maximize $H_{ij,t}$. Summing by parts again, $Y_{ij} = x_{ij}\big[\Delta_{(1)} - \sum_{r=1}^{R}(\Delta_{(r)} - \Delta_{(r+1)})e^{-\Theta_{\leq r}}\big] = x_{ij}\sum_{r=1}^{R} e^{-\Theta_{\leq r-1}}(1 - e^{-\theta_{(r)}})\Delta_{(r)}$. The first form is a sum of concave functions $-e^{-\Theta_{\leq r}}$ with nonnegative weights, so H_{ij} is concave. The second form gives its derivatives, for every origin k,

$$\frac{1}{x_{ij}}\frac{\partial H_{ij}}{\partial \theta_k} = e^{-\Theta_{\succeq k,ij}}\Delta^{+}_{k,ij} - \sum_{h \prec k} e^{-\Theta_{\succ h,ij}}\left(1 - e^{-\theta_{h,ij}}\right)\Delta^{+}_{h,ij} - S_k :$$

raising θ_k adds hires from k and takes them from every origin ranked below k. Condition (16) says that, at the candidate's queues θ^*, $\partial H_{ij}/\partial\theta_k \leq 0$, with equality if $\theta^*_k > 0$ (which requires $\alpha_{k,ij} > 0$). By concavity, for every $\theta \geq 0$,

$$H_{ij}(\theta) \leq H_{ij}(\theta^*) + \sum_k \frac{\partial H_{ij}}{\partial \theta_k}(\theta^*)\left(\theta_k - \theta^*_k\right) = H_{ij}(\theta^*) + \sum_k \frac{\partial H_{ij}}{\partial \theta_k}(\theta^*)\,\theta_k \leq H_{ij}(\theta^*). \tag{A.4}$$

Step 6: optimality. Along the candidate path, (a) holds with equality by the condition on relocations, (b) by the second line of (16), and (c) because its hires are those of Lemma A.5(i). So $G^*_t = \sum_{(i,j)} H_{ij,t}(\theta^*_{\cdot,ij,t})$, and for any feasible path, by (A.3) and (A.4), $G_t \leq \sum_{(i,j)} H_{ij,t}(\theta_{\cdot,ij,t}) \leq G^*_t$ at every date. By (A.2), no feasible path does better than the candidate.

Step 7: uniqueness. Let another path be optimal. By (A.2), $\sum_t \beta^t(G_t - G^*_t) = 0$, and every term is nonpositive, so $G_t = G^*_t$ at every date: every inequality in (A.3) and (A.4) holds with equality.

- *Queues.* With no ties and no zero surpluses, $\Delta_{(r)} - \Delta_{(r+1)} > 0$ for every r. The map from $(\theta_{(1)}, \ldots, \theta_{(R)})$ to $(\Theta_{\leq 1}, \ldots, \Theta_{\leq R})$ is linear and one to one, and $-e^{-z}$ is strictly concave, so Y_{ij} is strictly concave in the queues of the origins with $\Delta_{k,ij} > 0$, and H_{ij} is this plus a linear function. If two maximizers of H_{ij} differed in these queues, their midpoint would do strictly better; so these queues are those of the candidate. For an origin with $\Delta_{k,ij} < 0$, $\partial H_{ij}/\partial \theta_k = -x_{ij} S_k$, and equality in (A.4) requires $S_k \theta_k = 0$: its queue is zero unless $S_k = 0$, and it is never hired.

- *Hires.* Equality in (c) requires $M_r = x_{ij}(1 - e^{-\Theta_{\leq r}})$ for every r, since every weight $\Delta_{(r)} - \Delta_{(r+1)}$ is positive, and no hires from origins with $\Delta_{k,ij} < 0$. The queues of the origins with $\Delta_{k,ij} > 0$ are the candidate's, so the hires are the candidate's.

- *Path.* If no two trip values from a location tie, equality in (a) sends every unmatched carrier at k on the one best trip. Suppose the two paths have the same stocks at date t, as they do at $t = 0$. They have the same hires, so the same unmatched carriers $u_{k,t}$, and the same relocations; by (5) they have the same stocks at $t + 1$. By induction they have the same stocks, hires, and relocations at every date. □

A.3 Proof of Proposition 1

Keep the notation of the proof of Lemma A.6. The date-t flows are $f_t = (a_t, m_t, s_t)$; $\Phi(\mathbf{v})$ is the set of flows that satisfy (A.1) and (13) at stocks $\mathbf{v}$; and $\Pi(\mathbf{v}, f)$ and $F(\mathbf{v}, f)$ are the right sides of (12) and (5), both linear in $(\mathbf{v}, f)$. A path is a sequence $(\mathbf{v}_t, f_t)_{t \geq 0}$ with $f_t \in \Phi(\mathbf{v}_t)$ and $\mathbf{v}_{t+1} = F(\mathbf{v}_t, f_t)$. Given surpluses Δ and rents S, Y_{ij} and H_{ij} are defined as in (A.3).

Proof of Proposition 1. Step 0: three facts about a customer type. Fix (i, j) and drop it from the notation.

(F1) *The best applicant.* Let $\Theta(z) = \sum_{h:\Delta_h > z} \theta_h$ be the queue of applicants whose surplus exceeds z. On $[\Delta_{(r+1)}, \Delta_{(r)})$, $\Theta(z) = \Theta_{\leq r}$, so the first form of Y in Step 5 of the proof

of Lemma A.6 is

$$Y(\theta) = x \int_0^\infty \left(1 - e^{-\Theta(z)}\right) dz, \qquad \frac{1}{x}\frac{\partial Y}{\partial \theta_k} = \int_0^{\Delta_k^+} e^{-\Theta(z)}\, dz.$$

Y/x is the expected positive surplus of the best applicant, and the derivative is the one in Step 5. So Y, $\Theta(z)$ for $z \geq 0$, and $\partial Y/\partial\theta_k$ for an origin with $\theta_k > 0$ depend on the surpluses only through Δ_h^+ of the origins with $\theta_h > 0$. If $\theta_k = 0$, raising Δ_k by $\varepsilon \geq 0$ raises $\partial Y/\partial\theta_k$ by at most $x\varepsilon$, because $e^{-\Theta} \leq 1$.

(F2) *Who is hired.* Let m satisfy (13) at queues θ, with $\sum_k m_k \Delta_k = Y(\theta)$. Then $m_k = 0$ if $\Delta_k < 0$, and $m_k > 0$ if $\theta_k > 0$ and $\Delta_k > 0$. Indeed, both inequalities in Step 4(c) hold with equality. The first gives $m_k = 0$ if $\Delta_k < 0$. If $\Delta_k > 0$, the origins T with $\Delta_h \geq \Delta_k$ are the top r origins for an r with $\Delta_{(r)} - \Delta_{(r+1)} > 0$, so the second gives $\sum_{h\in T} m_h = x(1 - e^{-\sum_T \theta_h})$, while (13) for $T \setminus \{k\}$ gives $\sum_{h\in T\setminus\{k\}} m_h \leq x(1 - e^{-\sum_{T\setminus\{k\}} \theta_h})$. Subtracting, $m_k \geq x\, e^{-\sum_{T\setminus\{k\}} \theta_h}(1 - e^{-\theta_k}) > 0$.

(F3) *A customer's payoff.* Let values, with $V_k^{\mathrm{U}} \leq V_k^{\mathcal{O}}$, be given, and let $S_k = V_k^{\mathcal{O}} - V_k^{\mathrm{U}}$ and Δ_k be given by (2). If a customer posts $(\phi, \succ)$ and its queues θ satisfy (11), its payoff (10) is

$$P = \textstyle\sum_k \pi_k^{\mathrm{X}} \Delta_k - \sum_k \theta_k S_k \;\leq\; H(\theta)/x.$$

For $\theta_k > 0$, multiply the equality in (11) by θ_k and use $\pi_k^{\mathrm{X}} = \theta_k \pi_k^{\mathrm{V}}$; for $\theta_k = 0$, $\pi_k^{\mathrm{X}} = 0$. The inequality is Step 4(c), since the hires $x\pi^{\mathrm{X}}$ of a ranking satisfy (13). Conversely, every $\theta \geq 0$ is matched by a posting that pays at least $H(\theta)/x$. Set to zero the queues of the origins with $\Delta_k \leq 0$, which leaves Y unchanged by (F1) and does not lower $-\sum_k \theta_k S_k$. Rank the other origins by Δ and leave the rest out, so that the hires are those of Lemma A.5(i) and $\sum_k \pi_k^{\mathrm{X}} \Delta_k = Y/x$. Post $\phi_k = S_k/\pi_k^{\mathrm{V}} + V_k^{\mathrm{U}} - V_k^{\mathrm{F}}$ if $\theta_k > 0$ (then k is ranked and $\pi_k^{\mathrm{V}} > 0$), and $\phi_k = V_k^{\mathrm{U}} - V_k^{\mathrm{F}}$ if $\theta_k = 0$; then (11) holds. Hence a customer's best payoff is $\sup_\theta H(\theta)/x$, and a posting is optimal if and only if its

queues maximize H and $\sum_k \pi^{\mathrm{X}}_k \Delta_k = Y(\theta)/x$.

Step 1: what an equilibrium is. Take a feasible path with bounded values $(V^{\mathcal{O}}, V^{\mathrm{U}}, V^{\mathrm{T}}, V^{\mathrm{F}})_t$, with $S_t = V^{\mathcal{O}}_t - V^{\mathrm{U}}_t$, with carrier choices (α_t, σ_t) that give its flows, and with Δ_t given by (2). Consider the conditions

(E1) the values satisfy (6) and (7), $V^{\mathrm{U}}_{k,t} = \max_j V^{\mathrm{T}}_{kj,t}$, $S_{k,t} \geq 0$, and $\sigma_{kj,t} > 0$ only if $j \in \arg\max_{j'} V^{\mathrm{T}}_{kj',t}$;

(E2) $S_{k,t} = 0$ if $\sum_{(i,j)} a_{k,ij,t} < v^{\mathcal{O}}_{k,t}$;

(E3) for every (i,j), the queues $\theta_{\cdot,ij,t}$ maximize $H_{ij,t}$ over $\theta \geq 0$, and $\sum_k m_{k,ij,t}\Delta_{k,ij,t} = Y_{ij,t}(\theta_{\cdot,ij,t})$.

An equilibrium satisfies (E1)–(E3). (E1) is condition (i) of Definition 1. For (E2), if some open carriers at k do not apply, $\alpha_{k,t} > 0$, so not applying attains the maximum in (9), and $S_{k,t} = 0$. For (E3), each customer's posting is optimal by condition (ii), and its hires are m_t by condition (iv); (F3) then gives (E3).

A path with values that satisfy (E1)–(E3) is an equilibrium. The hires $m_{\cdot,ij,t}$ satisfy (13), so they are the hires of some ranking. Move every origin with $\theta_{k,ij,t} = 0$ to the top of it; this leaves the hires unchanged. Post the fees as follows.

- If $\theta_{k,ij,t} > 0$ and $\pi^{\mathrm{V}}_{k,ij,t} > 0$, post $\phi_{k,ij,t} = S_{k,t}/\pi^{\mathrm{V}}_{k,ij,t} + V^{\mathrm{U}}_{k,t} - V^{\mathrm{F}}_{k,ij,t}$.
- If $\theta_{k,ij,t} > 0$ and $\pi^{\mathrm{V}}_{k,ij,t} = 0$, then $m_{k,ij,t} = 0$, so $\Delta_{k,ij,t} \leq 0$ by (F2), and $\partial H_{ij,t}/\partial\theta_k = -x_{ij}S_{k,t}$ by (F1). Since $\theta_{k,ij,t} > 0$ maximizes $H_{ij,t}$ in this coordinate, $S_{k,t} = 0$, and any fee will do.
- If $\theta_{k,ij,t} = 0$, an applicant from k would be at the top, so $\pi^{\mathrm{V}}_{k,ij,t} = 1$; post $\phi_{k,ij,t} = S_{k,t} + V^{\mathrm{U}}_{k,t} - V^{\mathrm{F}}_{k,ij,t}$.

Then (11) holds with equality for every k and (i,j). By (E3) and (F3), every customer's posting is optimal: condition (ii). Applying to any customer is worth $V^{\mathrm{U}}_{k,t} + S_{k,t} = V^{\mathcal{O}}_{k,t}$ to

an open carrier at k, and not applying is worth $V^{\mathrm{U}}_{k,t} \leq V^{\mathcal{O}}_{k,t}$. So the values solve (9), every application attains its maximum, and by (E2) not applying attains it whenever some open carriers do not apply (if $v^{\mathcal{O}}_{k,t} = 0$, set $\alpha_{k,t} = 0$). With (E1), this is condition (i). Conditions (iii) and (iv) hold because the path is feasible and its hires are those of the rankings.

Step 2: an equilibrium path is optimal. Steps 2–4 and 6 of the proof of Lemma A.6 use that lemma's conditions only through (E1)–(E3): Step 2 uses bounded values that satisfy the value equations with $V^{\mathrm{U}}_{k,t} = \max_j V^{\mathrm{T}}_{kj,t}$; Step 4 uses $S \geq 0$; and Step 6 uses equality in (a), (b) and (c), which is the relocation condition in (E1), (E2), and the second half of (E3), and that the queues maximize $H_{ij,t}$, which is the first half of (E3). By Step 1, an equilibrium path satisfies (E1)–(E3), so no feasible path does better.

Step 3: an optimal path is an equilibrium path. Let $(\mathbf{v}_t, f_t)$ be optimal from $\mathbf{v}_0$. By Step 1 it suffices to find bounded values that satisfy (E1)–(E3) along it.

(3a) The value function. For stocks $\mathbf{v} \in \mathbb{R}^N_+$ of any fleet size, let $W(\mathbf{v})$ be the value of (14) from $\mathbf{v}$. An optimal path exists from every $\mathbf{v}$. At every date, stocks and flows lie in a compact set, since they are bounded by the fleet, and (A.1), (13) and (5) are closed conditions. So the feasible paths form a compact set in the product topology. The objective is continuous on it, because $|\Pi_t|$ is bounded and $\beta < 1$. W is concave: (13) is convex in (a, m) because $1 - e^{-z}$ is concave, the other constraints are linear, so a convex combination of feasible paths from $\mathbf{v}$ and $\mathbf{v}'$ is feasible from the same combination of $\mathbf{v}$ and $\mathbf{v}'$; and Π is linear. By the principle of optimality, $W(\mathbf{v}) \geq \Pi(\mathbf{v}, f) + \beta W(F(\mathbf{v}, f))$ for every $f \in \Phi(\mathbf{v})$, with equality at $(\mathbf{v}_t, f_t)$ for every t.

(3b) W is Lipschitz. Let $\bar{w} = \max_{ij} w_{ij}$, $\bar{c} = \max_{k,j} c_{kj}$ and $L = \max\{\bar{w}, \bar{c}\}/(1-\beta)$. Adding carriers $\delta \geq 0$ costs at most $\bar{c}|\delta|_1/(1-\beta)$. Follow an optimal path from $\mathbf{v}$, and let the added carriers never apply: open ones stay put; ones in transit finish their trip, and fixed ones their pickup and delivery, and then stay put. They leave (13) unchanged and cost at most $\bar{c}$ each per period. Removing carriers $0 \leq \delta \leq \mathbf{v}$ costs at most $\bar{w}|\delta|_1/(1-\beta)$. Follow an optimal path from $\mathbf{v}$ and remove a cohort that starts as δ. At each date, give

the cohort the share $\lambda_{k,t}$ of the open carriers at k that it holds, and the same share of their flows, so that the remaining flows are $(1-\lambda_{k,t})$ times the original ones. Scaling one origin's applications and hires by $1-\lambda \in [0,1]$ preserves (13): for a group $K \ni k$, the slack $x\left(1-e^{-\sum_K \theta_h + \lambda\theta_k}\right) - \sum_K m_h + \lambda m_k$ is concave in λ, nonnegative at $\lambda = 0$, and nonnegative at $\lambda = 1$, which is the constraint for $K \setminus \{k\}$. Scale the origins one at a time. The other constraints are linear. By linearity of (5), the cohort's stocks follow (5) under its flows and keep the mass $|\delta|_1$, and the remaining stocks are the difference. Π is linear, and the cohort earns at most $\bar{w}|\delta|_1$ per date: it hires at most its open mass, and each hire earns at most $\bar{w}$, since $X^{\mathrm{F}} \leq 0$ and costs are nonnegative. So for any $\mathbf{v}, \mathbf{v}' \geq 0$, with $\delta = (\mathbf{v}-\mathbf{v}')^+$ and $\delta' = (\mathbf{v}'-\mathbf{v})^+$, $W(\mathbf{v}') \geq W(\mathbf{v}) - L|\mathbf{v}-\mathbf{v}'|_1$, and symmetrically.

(3c) Supergradients. Let $\tilde{W}(\mathbf{z}) = \sup_{\mathbf{y}\geq 0}\left[W(\mathbf{y}) - L|\mathbf{z}-\mathbf{y}|_1\right]$ for $\mathbf{z} \in \mathbb{R}^N$. It is concave, being a partial maximum of a jointly concave function; it is finite, L-Lipschitz, and equal to W on $\mathbb{R}^N_+$. So at every $\mathbf{z}$ the set $\partial\tilde{W}(\mathbf{z})$ of supergradients, the p with $\tilde{W}(\mathbf{z}') \leq \tilde{W}(\mathbf{z}) + p \cdot (\mathbf{z}'-\mathbf{z})$ for all $\mathbf{z}'$, is nonempty, and each of them has $|p|_\infty \leq L$ (Rockafellar, 1970, Theorem 23.4).

(3d) Shadow values. Let $G = \{(\mathbf{v}, f) : \mathbf{v} \geq 0,\ f \in \Phi(\mathbf{v})\}$, a convex set. I construct $q_t \in \partial\tilde{W}(\mathbf{v}_t)$ such that, for every t, $(\mathbf{v}_t, f_t)$ maximizes

$$\Lambda_t(\mathbf{v}, f) = \Pi(\mathbf{v}, f) + \beta q_{t+1} \cdot F(\mathbf{v}, f) - q_t \cdot \mathbf{v} \quad \text{over } G. \tag{A.5}$$

Take any $q_0 \in \partial\tilde{W}(\mathbf{v}_0)$, and let $q_t \in \partial\tilde{W}(\mathbf{v}_t)$. For $(\mathbf{v}, f) \in G$, $F(\mathbf{v}, f) \geq 0$, so by (3a)

$$\Pi(\mathbf{v}, f) + \beta\tilde{W}(F(\mathbf{v}, f)) - q_t \cdot \mathbf{v} \leq W(\mathbf{v}) - q_t \cdot \mathbf{v} \leq W(\mathbf{v}_t) - q_t \cdot \mathbf{v}_t,$$

with equality at $(\mathbf{v}_t, f_t)$. So $(\mathbf{v}_t, f_t)$ maximizes the concave function on the left over G. Since $\tilde{W}$ is finite everywhere and F is linear, there is then a $p \in \partial\tilde{W}(F(\mathbf{v}_t, f_t)) = \partial\tilde{W}(\mathbf{v}_{t+1})$ such that $(\mathbf{v}_t, f_t)$ maximizes the linear function $\Pi + \beta p \cdot F - q_t \cdot \mathbf{v}$ over G (Rockafellar, 1970, Theorems 23.8, 23.9 and 27.4). Set $q_{t+1} = p$.

(3e) Reading (A.5). Write $q_t = (q_t^{\mathcal{O}}, q_t^{\mathrm{T}}, q_t^{\mathrm{F}})$, and let

$$Q_{kj,t}^{\mathrm{T}} = -c_{kj} + \beta d_{kj} q_{j,t+1}^{\mathcal{O}} + \beta(1-d_{kj}) q_{kj,t+1}^{\mathrm{T}}, \qquad Q_{k,t}^{\mathrm{U}} = \max_j Q_{kj,t}^{\mathrm{T}},$$

$$Q_{k,ij,t}^{\mathrm{F}} = -c_{ki} + \beta d_{ki} q_{ij,t+1}^{\mathrm{T}} + \beta(1-d_{ki}) q_{k,ij,t+1}^{\mathrm{F}}, \qquad \Delta_{k,ij,t}^{\mathrm{Q}} = w_{ij} + X_{k,ij}^{\mathrm{F}} + Q_{k,ij,t}^{\mathrm{F}} - Q_{k,t}^{\mathrm{U}}.$$

The computation in Step 3 of the proof of Lemma A.6, with q_{t+1} in place of V_{t+1}, gives

$$\begin{aligned}\Lambda_t = {} & \textstyle\sum_{k,j} (Q_{kj,t}^{\mathrm{T}} - q_{kj,t}^{\mathrm{T}}) v_{kj}^{\mathrm{T}} + \sum_{k,(i,j)} (Q_{k,ij,t}^{\mathrm{F}} - q_{k,ij,t}^{\mathrm{F}}) v_{k,ij}^{\mathrm{F}} \\ & + \textstyle\sum_{k,(i,j)} m_{k,ij} (w_{ij} + X_{k,ij}^{\mathrm{F}} + Q_{k,ij,t}^{\mathrm{F}}) + \sum_{k,j} Q_{kj,t}^{\mathrm{T}} s_{kj} - \sum_k q_{k,t}^{\mathcal{O}} v_k^{\mathcal{O}}.\end{aligned}$$

(a) Only $v^{\mathrm{T}} \geq 0$ and $v^{\mathrm{F}} \geq 0$ constrain the stocks in transit and fixed, so $Q_{kj,t}^{\mathrm{T}} \leq q_{kj,t}^{\mathrm{T}}$, with equality if $v_{kj,t}^{\mathrm{T}} > 0$, and $Q_{k,ij,t}^{\mathrm{F}} \leq q_{k,ij,t}^{\mathrm{F}}$, with equality if $v_{k,ij,t}^{\mathrm{F}} > 0$.

(b) Given $v^{\mathcal{O}}$ and m, the unmatched carriers go only to trips in $\arg\max_j Q_{kj,t}^{\mathrm{T}}$, and the open part of Λ_t is $\sum_k (Q_{k,t}^{\mathrm{U}} - q_{k,t}^{\mathcal{O}}) v_k^{\mathcal{O}} + \sum_{k,(i,j)} m_{k,ij} \Delta_{k,ij,t}^{\mathrm{Q}}$. Given the applications, the best hires earn $\sum_{(i,j)} Y_{ij,t}^{\mathrm{Q}}(\theta_{\cdot,ij})$ by Step 4(c), where Y^{Q} is Y at the surpluses Δ^{Q}. So m_t attains $Y_{ij,t}^{\mathrm{Q}}(\theta_{\cdot,ij,t})$ for every (i,j), and $(v_t^{\mathcal{O}}, a_t)$ maximizes $\sum_k (Q_{k,t}^{\mathrm{U}} - q_{k,t}^{\mathcal{O}}) v_k^{\mathcal{O}} + \sum_{(i,j)} Y_{ij,t}^{\mathrm{Q}}(a_{\cdot,ij}/x_{ij})$ subject to $v^{\mathcal{O}} \geq 0$, $a \geq 0$ and $\sum_{(i,j)} a_{k,ij} \leq v_k^{\mathcal{O}}$. This objective is concave and differentiable, by (F1), and the constraints are linear, so there are multipliers $S_{k,t}^{\mathrm{Q}} \geq 0$ with (Rockafellar, 1970, Corollary 28.2.2 and Theorem 28.3)

$$\frac{1}{x_{ij}} \frac{\partial Y_{ij,t}^{\mathrm{Q}}}{\partial \theta_k} \leq S_{k,t}^{\mathrm{Q}} \text{ (equality if } a_{k,ij,t} > 0), \qquad Q_{k,t}^{\mathrm{U}} + S_{k,t}^{\mathrm{Q}} \leq q_{k,t}^{\mathcal{O}} \text{ (equality if } v_{k,t}^{\mathcal{O}} > 0),$$

and $S_{k,t}^{\mathrm{Q}} = 0$ if $\sum_{(i,j)} a_{k,ij,t} < v_{k,t}^{\mathcal{O}}$.

(c) If $v_{k,t}^{\mathcal{O}} > 0$, some carriers at k are not hired: by (13) for $K = \{k\}$, $m_{k,ij,t} \leq x_{ij}(1 - e^{-\theta_{k,ij,t}}) < a_{k,ij,t}$ whenever $a_{k,ij,t} > 0$, so $\sum_{(i,j)} m_{k,ij,t} < v_{k,t}^{\mathcal{O}}$.

(3f) The values. Let $V_t^{\mathcal{O}} = q_t^{\mathcal{O}}$; let V^{T} and V^{F} be the bounded solutions of (6) and (7) given $V^{\mathcal{O}}$; and let $V_{k,t}^{\mathrm{U}} = \max_j V_{kj,t}^{\mathrm{T}}$ and $S_t = V_t^{\mathcal{O}} - V_t^{\mathrm{U}}$. All are bounded. Compare them

with the q's.

(i) $q^{\mathrm{T}} \geq V^{\mathrm{T}}$ and $q^{\mathrm{F}} \geq V^{\mathrm{F}}$. By (a), $q^{\mathrm{T}}_{kj,t} - V^{\mathrm{T}}_{kj,t} \geq \beta(1-d_{kj})\big(q^{\mathrm{T}}_{kj,t+1} - V^{\mathrm{T}}_{kj,t+1}\big)$. Iterating, the right side vanishes, since both sequences are bounded and $\beta < 1$. Likewise $q^{\mathrm{F}}_{k,ij,t} - V^{\mathrm{F}}_{k,ij,t} \geq \beta d_{ki}\big(q^{\mathrm{T}}_{ij,t+1} - V^{\mathrm{T}}_{ij,t+1}\big) + \beta(1-d_{ki})\big(q^{\mathrm{F}}_{k,ij,t+1} - V^{\mathrm{F}}_{k,ij,t+1}\big)$, whose first term is nonnegative. Hence $Q^{\mathrm{T}} \geq V^{\mathrm{T}}$, $Q^{\mathrm{F}} \geq V^{\mathrm{F}}$, and $Q^{\mathrm{U}} \geq V^{\mathrm{U}}$.

(ii) The two coincide where carriers are. If $v^{\mathrm{T}}_{kj,t} > 0$, the inequality in (i) is an equality; if $d_{kj} < 1$, then $v^{\mathrm{T}}_{kj,t+1} \geq (1-d_{kj})v^{\mathrm{T}}_{kj,t} > 0$. Iterating, $q^{\mathrm{T}}_{kj,t} = V^{\mathrm{T}}_{kj,t}$. The same argument gives $q^{\mathrm{F}}_{k,ij,t} = V^{\mathrm{F}}_{k,ij,t}$ if $v^{\mathrm{F}}_{k,ij,t} > 0$, because then $v^{\mathrm{T}}_{ij,t+1} \geq d_{ki}v^{\mathrm{F}}_{k,ij,t} > 0$. Consequently, $Q^{\mathrm{T}}_{kj,t} = V^{\mathrm{T}}_{kj,t}$ if $s_{kj,t} > 0$, because then $v^{\mathrm{T}}_{kj,t+1} \geq (1-d_{kj})s_{kj,t}$; and $Q^{\mathrm{F}}_{k,ij,t} = V^{\mathrm{F}}_{k,ij,t}$ if $m_{k,ij,t} > 0$, because then $v^{\mathrm{T}}_{ij,t+1} \geq d_{ki}m_{k,ij,t} > 0$ and $v^{\mathrm{F}}_{k,ij,t+1} \geq (1-d_{ki})m_{k,ij,t}$.

(iii) Open locations. Let $v^{\mathcal{O}}_{k,t} > 0$. By (c), some $s_{kj,t} > 0$, with $j \in \arg\max Q^{\mathrm{T}}_{k\cdot,t}$ by (b). By (ii) and (i), $V^{\mathrm{T}}_{kj,t} = Q^{\mathrm{T}}_{kj,t} = Q^{\mathrm{U}}_{k,t} \geq V^{\mathrm{U}}_{k,t} \geq V^{\mathrm{T}}_{kj,t}$. So $Q^{\mathrm{U}}_{k,t} = V^{\mathrm{U}}_{k,t}$, every trip in use maximizes $V^{\mathrm{T}}_{kj,t}$, and, by (b), $S_{k,t} = S^{Q}_{k,t}$.

(iv) Rents. By (b) and (i), $S_{k,t} = q^{\mathcal{O}}_{k,t} - V^{\mathrm{U}}_{k,t} \geq S^{Q}_{k,t} + \big(Q^{\mathrm{U}}_{k,t} - V^{\mathrm{U}}_{k,t}\big) \geq S^{Q}_{k,t} \geq 0$.

(v) Surpluses. By (i), $\Delta - \Delta^{Q} = (V^{\mathrm{F}} - Q^{\mathrm{F}}) - (V^{\mathrm{U}} - Q^{\mathrm{U}}) \leq Q^{\mathrm{U}}_{k,t} - V^{\mathrm{U}}_{k,t}$. If $m_{k,ij,t} > 0$, then $v^{\mathcal{O}}_{k,t} > 0$, and by (ii) and (iii) $\Delta_{k,ij,t} = \Delta^{Q}_{k,ij,t}$. Let $\theta_{k,ij,t} > 0$. If $\Delta^{Q}_{k,ij,t} > 0$, then $m_{k,ij,t} > 0$ by (F2) applied to Δ^{Q}, so $\Delta_{k,ij,t} = \Delta^{Q}_{k,ij,t}$; otherwise $\Delta_{k,ij,t} \leq \Delta^{Q}_{k,ij,t} \leq 0$. Either way $\Delta^{+}_{k,ij,t} = \Delta^{Q+}_{k,ij,t}$.

(3g) (E1)–(E3). (E1): the value equations and $S \geq 0$ hold by (3f), and unmatched carriers exist only where $v^{\mathcal{O}}_{k,t} > 0$, where they use only trips that maximize V^{T}, by (iii); where $v^{\mathcal{O}}_{k,t} = 0$, choose $\sigma_{k,t}$ on $\arg\max_j V^{\mathrm{T}}_{kj,t}$. (E2): if $\sum_{(i,j)} a_{k,ij,t} < v^{\mathcal{O}}_{k,t}$, then $v^{\mathcal{O}}_{k,t} > 0$, and $S_{k,t} = S^{Q}_{k,t} = 0$ by (iii) and (b). (E3): by (v) and (F1), Y and Y^{Q} coincide at $\theta_{\cdot,ij,t}$, and so do their derivatives in the queues that are positive. Since $m_{k,ij,t} > 0$ implies $\Delta_{k,ij,t} = \Delta^{Q}_{k,ij,t}$, $\sum_k m_{k,ij,t}\Delta_{k,ij,t} = \sum_k m_{k,ij,t}\Delta^{Q}_{k,ij,t} = Y^{Q}_{ij,t} = Y_{ij,t}$ at $\theta_{\cdot,ij,t}$. For the queues, $H_{ij,t}$ is concave, so it suffices that $\partial H_{ij,t}/\partial\theta_k \leq 0$ at $\theta_{\cdot,ij,t}$, with equality if $\theta_{k,ij,t} > 0$. If $\theta_{k,ij,t} > 0$, then $a_{k,ij,t} > 0$

and $v^{\mathcal{O}}_{k,t} > 0$, so $\frac{1}{x_{ij}}\partial Y_{ij,t}/\partial\theta_k = \frac{1}{x_{ij}}\partial Y^{Q}_{ij,t}/\partial\theta_k = S^{Q}_{k,t} = S_{k,t}$ by (b) and (iii). If $\theta_{k,ij,t} = 0$, then by (F1), (v), (b) and (iv),

$$\frac{1}{x_{ij}}\frac{\partial Y_{ij,t}}{\partial\theta_k} \le \frac{1}{x_{ij}}\frac{\partial Y^{Q}_{ij,t}}{\partial\theta_k} + \left(Q^{\mathrm{U}}_{k,t} - V^{\mathrm{U}}_{k,t}\right) \le S^{Q}_{k,t} + Q^{\mathrm{U}}_{k,t} - V^{\mathrm{U}}_{k,t} \le S_{k,t}.$$

So every optimal path is an equilibrium path.

Step 4: conclusions. By Steps 2 and 3, the equilibrium paths are the optimal paths, so the market is efficient. Finally, with the Δ ranking, an applicant from k with $\theta_{k,ij,t} > 0$ is hired with the probability (8), and (11) with equality is (15). □

A.4 Proof of Proposition 2

Proof of Proposition 2. By Step (3a) of the proof of Proposition 1, a path that attains (14) exists from every $\mathbf{v}_0$, and by Proposition 1 it is an equilibrium path; so an equilibrium exists. By Proposition 1, every equilibrium path attains (14), so Lemma A.6 applies to all of them: under its no-tie conditions they have the same queues and hires, and, if in addition no two trip values from a location tie, the same path, apart from applications that are never accepted. □

A.5 Two locations

This appendix uses the notation of Section 4 and Appendix A.1, and three more symbols. By (3), a customer's waiting value is $X^{\mathrm{F}}_{1,12} = -c_1$ with a local carrier and $X^{\mathrm{F}}_{2,12} = -c_1/D_{21}$ with a remote one. An unmatched carrier at 2 is worth $u_2 \equiv V^{\mathrm{U}}_2 - \bar{V}_2$ above staying idle; and

$$A \equiv \rho_{12}\rho_{21}, \qquad \Lambda \equiv \kappa_{12} + \rho_{12}\kappa_{21}$$

are the discount and the extra cost of a round trip $1 \to 2 \to 1$.

Lemma A.7 (Rents and values). *Given the queues,*

$$\Delta_1 = w - c_1 - \beta\kappa_{12} + \beta\rho_{12}\, y_2 - \beta y_1, \quad \Delta_2 = w + X^{\mathrm{F}}_{2,12} - \kappa_{21} - \rho_{21}\kappa_{12} + A\, y_2 - u_2,$$
$$(1-\beta)\, y_1 = e^{-\theta_1}\big[\Delta_1 - (1 - e^{-\theta_2})\Delta_2\big], \qquad y_2 = u_2 + e^{-\Theta}\Delta_2 \ \textit{if carriers at 2 apply},$$

with $u_2 = \rho_{21} y_1 - \kappa_{21}$ if unmatched carriers at 2 sail back and $u_2 = \beta y_2$ if they stay.

Proof. The surpluses are Lemma A.4 with $u_1 = \beta y_1$. Substituting the fees, a customer chooses queues to maximize $R(\theta) - \theta_1 S_1 - \theta_2 S_2$ with $R = (1 - e^{-\theta_1})\Delta_1 + e^{-\theta_1}(1 - e^{-\theta_2})\Delta_2$ (Shimer, 2005); the first-order conditions give $S_1 = \partial R/\partial\theta_1$ and $S_2 = \partial R/\partial\theta_2 = e^{-\Theta}\Delta_2$. Lemma A.4 gives $y_k = u_k + S_k$. □

Lemma A.8 (Queues). *In a steady state, given what carriers do, the queues do not depend on w. If only carriers at 1 apply and carriers at 2 sail back (regime 3), the local queue $\theta^{(3)}$ solves*

$$\frac{v}{x} = \theta^{(3)} + \Big(\frac{1}{d_{12}} + \frac{1}{d_{21}}\Big)\big(1 - e^{-\theta^{(3)}}\big).$$

If all carriers at 2 apply, the total queue $\Theta^{(4)}$ solves

$$\frac{v}{x} = \Theta^{(4)} + \Big(\frac{1}{d_{12}} + \frac{1}{d_{21}} - 1\Big)\big(1 - e^{-\Theta^{(4)}}\big),$$

split $\theta_2^{(4)} = 1 - e^{-\Theta^{(4)}}$, $\theta_1^{(4)} = \Theta^{(4)} - \theta_2^{(4)}$ while unmatched carriers at 2 sail back (regime 4), and $(\theta_1, \theta_2) = (0, \Theta^{(4)})$ once they all stay (regime 5).

Proof. Hires are $m = x(1 - e^{-\Theta})$ per period, and a stock that m enter each period and leave at rate d has size m/d. With only local applicants, the fleet is open at 1 ($x\theta^{(3)}$), laden $1 \to 2$ (m/d_{12}), open at 2 for one period (m), and sailing back empty after that period ($(1 - d_{21})m/d_{21}$). Once carriers at 2 apply, the fleet is open ($x\Theta$), laden (m/d_{12}), and empty $2 \to 1$ after the first period ($(1 - d_{21})m/d_{21}$): pickups and returns, m of them starting each period whether unmatched carriers sail back or stay. While they sail back, every arrival at 2 is open there for one period, so $x\theta_2 = m$. □

The cutoffs. Let $K_1 \equiv c_1 + \beta\Lambda$ and $K_2 \equiv -X^{\mathrm{F}}_{2,12} + \rho_{21}\Lambda$, the costs of a local and of a remote hire. The cutoffs of Proposition 3 are $\bar{w}_1 = c_1 + \beta\kappa_{12}$ and, with three functions of the local queue θ_1 — w_T for the trip back to 1, w_R for remote hires, and w_S for staying at 2 —

$$\bar{w}_2 = w_T(0), \quad \bar{w}_3 = w_T(\theta^{(3)}), \quad \bar{w}_4 = w_R(\theta^{(3)}), \quad \bar{w}_5 = w_R(\theta^{(4)}_1), \quad \bar{w}_6 = w_S(\theta^{(4)}_1), \quad \bar{w}_7 = w_S(0),$$

where

$$\begin{aligned}
w_T(\theta_1) &= c_1 + \beta\kappa_{12} + \beta\frac{\kappa_{21}}{\rho_{21}} + e^{\theta_1}(1-\beta)\frac{\kappa_{21}}{\rho_{21}},\\
w_R(\theta_1) &= K_2 + \frac{q(\theta_1)}{1-q(\theta_1)}(K_2 - K_1), \qquad q(\theta_1) \equiv \frac{\rho_{21}(1-A)\,e^{-\theta_1}}{1-\beta+\beta(1-A)\,e^{-\theta_1}},\\
w_S(\theta_1) &= -X^{\mathrm{F}}_{2,12} + \kappa_{21} + \rho_{21}\kappa_{12} + \left[\beta - A + (1-\beta)e^{\Theta^{(4)}}\right] y_2(\theta_1),\\
y_2(\theta_1) &\equiv \frac{G\,e^{-\theta_1} - (1-\beta)\kappa_{21}/\rho_{21}}{(1-\beta)(\beta/\rho_{21}-1) - h\,e^{-\theta_1}},
\end{aligned}$$

with $G \equiv (X^{\mathrm{F}}_{1,12} - X^{\mathrm{F}}_{2,12}) + \kappa_{12}(\rho_{21} - \beta) + \kappa_{21}(1 - \beta/\rho_{21})$ and $h \equiv \beta(1+\rho_{12}) - A - \beta^2/\rho_{21}$.

Proof of Proposition 3. All four use Lemma A.7; given the queues, every condition is linear in (w, y_1, y_2). Parts (ii) to (iv) give w_T, w_R and w_S at any local queue θ_1; the cutoffs evaluate them at the queues of Lemma A.8. (i) No rents, $y_1 = y_2 = 0$: $\Delta_1 = w - c_1 - \beta\kappa_{12}$. (ii) $\theta_2 = 0$, $y_2 = 0$, and $V^{\mathrm{T}}_{21} = V^{\mathrm{T}}_{22}$, i.e. $y_1 = \kappa_{21}/\rho_{21}$ by Lemma A.2; substitute into $(1-\beta)y_1 = e^{-\theta_1}\Delta_1$. (iii) No rent at 2 and sailing back: $y_2 = u_2 = \rho_{21}y_1 - \kappa_{21}$. Then $\Delta_1 = w - K_1 - \beta(1-A)y_1$ and $\Delta_2 = w - K_2 - \rho_{21}(1-A)y_1$. With $\Delta_2 = 0$, $(1-\beta)y_1 = e^{-\theta_1}\Delta_1$ gives $y_1 = e^{-\theta_1}(w - K_1)/\left(1-\beta+\beta(1-A)e^{-\theta_1}\right)$, and $\Delta_2 = 0$ becomes $w - K_2 = q(\theta_1)(w - K_1)$. (iv) $u_2 = \beta y_2 = \rho_{21}y_1 - \kappa_{21}$. From $(1-\beta)y_2 = e^{-\Theta^{(4)}}\Delta_2$ with $\Delta_2 = w + X^{\mathrm{F}}_{2,12} - \kappa_{21} - \rho_{21}\kappa_{12} - (\beta - A)y_2$ follows the expression for w. Since $e^{-\Theta^{(4)}}\Delta_2 = (1-\beta)y_2$, the rent of carriers at 1 reads $(1-\beta)y_1 = e^{-\theta_1}(\Delta_1 - \Delta_2) + (1-\beta)y_2$; with $y_1 = (\beta y_2 + \kappa_{21})/\rho_{21}$, $\Delta_1 - \Delta_2 = G + h\,y_2$, which is one linear equation in y_2. Finally, c_1 enters only through $X^{\mathrm{F}}_{1,12} = -c_1$ and $X^{\mathrm{F}}_{2,12} =$

$-c_1/D_{21}$, linearly in K_1, K_2 and G, and the queues do not depend on it (Lemma A.8); so every cutoff is affine in c_1. □

Storage-cost thresholds. Since every cutoff is affine in c_1, the storage costs at which two cutoffs meet are in closed form. Write $\delta \equiv 1/D_{21} - 1 = \beta(1 - d_{21})/D_{21}$ for the extra discounted periods a customer waits for a remote carrier, and $q_3 \equiv q(\theta^{(3)})$. The five regimes appear in order when $\bar{w}_3 < \bar{w}_4$, that is, when $c_1 > \bar{c}$ with

$$\bar{c} = \frac{(1 - q_3)\left[\beta\kappa_{12} + \left(\beta + (1-\beta)e^{\theta^{(3)}}\right)\kappa_{21}/\rho_{21}\right] - (\rho_{21} - q_3\beta)\Lambda}{\delta}.$$

Trade moves straight from regime 2 to regime 5, at w_D, when remote hires pay before carriers at 2 start sailing back, $w_D < \bar{w}_2$, that is, when $c_1 < \underline{c}$ with

$$\underline{c} = \frac{(\beta - \rho_{21})\kappa_{12} + (1/\rho_{21} - 1)\kappa_{21}}{\delta}.$$

At $c_1 = \underline{c}$, $w_D = \bar{w}_2 = \bar{w}_7$. In the example, $\underline{c} = 0.039$ and $\bar{c} = 0.050$.

A numerical check. I also solve the example numerically from the model's primitives, on a grid of 260×320 values of (c_1, w), without using the cutoffs. At each point, the stocks and queues come from iterating the laws of motion (5), from the fleet spread evenly over open carriers, under a candidate behaviour of carriers at 2; the values come from the steady-state value equations of Section 2.1.1, with rents equal to customers' marginal products of queues; and a candidate is kept if every choice is optimal given those values. Exactly one behavior of carriers survives at every grid point, and the regime it gives agrees with Figure 5, which is drawn from the cutoffs, at all but 4 of the 83,200 points, all with $\underline{c} < c_1 < \bar{c}$, where the cutoffs are not ordered.

B Application rents and fees in the two-location example

This appendix shows the application rents and the fees of the two-location example of Section 4, and compares the fees with those of the model with local matching (Section 4.5).

B.1 Application rents

An open carrier's application rent, $S_k = V_k^{\mathcal{O}} - V_k^{\mathrm{U}}$, is what it gains from being able to apply. Figure 12 shows it for carriers at each location. Carriers at 2 earn no rent until all of them apply, at $\bar{w}_5$: before customers hire remote carriers there is nothing for them to apply to, and while carriers at 2 start applying, between $\bar{w}_4$ and $\bar{w}_5$, a remote hire has zero surplus, so customers attract remote applicants without paying them a rent. Carriers at 1 earn a rent as soon as trade starts, and always more than carriers at 2, because local applicants rank first; whenever carriers at 2 apply,

$$S_1 - S_2 = e^{-\theta_1}\big(\Delta_1 - \Delta_2\big).$$

Both rents grow with w, to about a fifth of the value of a trip. The gap is largest when remote applications begin, and it shrinks through mixed matching as the local premium in surplus, $\Delta_1 - \Delta_2$, shrinks.

B.2 Fees

Figure 13 shows the fees customers post for a local and for a remote carrier as w rises. By (15), a fee has two parts,

$$\phi_k = \frac{S_k}{\pi_k^{\mathrm{V}}} + \big(V_k^{\mathrm{U}} - V_k^{\mathrm{F}}\big) :$$

the rent of an applicant from k divided by its chance of being hired, and compensation for what being hired costs the carrier. Where no carrier applies from a location, its fee is not pinned down; the figure shows instead the reservation fee, at which the first applicant

from there would just be willing to apply.

Without rents, the compensation is the extra cost of the trips a hire commits the carrier to: $\beta\kappa_{12}$ for a local carrier, which only delivers, and $\kappa_{21} + \rho_{21}\kappa_{12}$ for a remote one, which also sails to the cargo. These are the reservation fees when there is no trade. As w rises, both fees rise faster than w over most of the range, and they end up taking about four fifths of the value of a trip.

Local carriers are paid more than remote ones until $w \approx 13.6$, in mixed matching, even though a remote carrier has to sail to the cargo. The compensation is larger for the local carrier throughout: being hired takes it away from the location where cargo appears. Its rent per hire is also larger at first, because it is ranked first. As w rises, the remote carrier's rent per hire overtakes it. A remote applicant is hired only when no local one arrives, so it must earn more per hire to collect its rent, and as rents grow with w this outweighs the local carrier's premium in surplus, $\Delta_1 - \Delta_2$, which shrinks through mixed hiring (Section B.1). In remote hiring, no carrier waits at 1, and the local fee is only a reservation fee.

B.3 Fees with local matching

With local matching a customer posts only a fee for a local carrier. Remote matching raises that fee by up to 5 percent in remote entry and early mixed hiring, for w between 3.9 and 5.7, and lowers it above: by 6 percent at $w = 10$ and by 12 percent at $w = 30$, where it is a reservation fee (Figure 14). The difference comes mostly from the compensation $V_1^{\mathrm{U}} - V_1^{\mathrm{F}}$ in (15), what being hired costs a carrier at 1: a hired carrier gives up being at 1 and ends its trip at 2, and with local matching being at 1 is worth much more. At $w = 30$ the compensation is 19.9 with local matching and 17.2 with remote matching. With local matching the fee takes 84 percent of the value of a trip at $w = 10$ and 90 percent at $w = 30$; with remote matching, 79 percent.

C Computing the equilibrium

The two-location example has closed-form cutoffs; a general network does not. I compute steady states as the limit of a logit quantal response equilibrium (McKelvey and Palfrey, 1995), following the path-following method of Turocy (2005, 2010), and check them against an exact solution of (14), which gives the equilibrium by Proposition 1. This appendix sets the model up as a game, defines the logit equilibrium, writes it as a system of equations, and describes how the system is solved.

C.1 The game

The steady state is a game in agent form (McKelvey and Palfrey, 1998): each kind of decision in the model is made by its own agent, and the continuum of carriers who face the same decision plays the same mixed strategy.

Agents and strategies. At every location $k \in \mathcal{L}$ there are two carrier agents, and every customer type is an agent.

- An *open carrier at k* chooses whether and where to apply. Its strategies are $S_k^{\mathcal{O}} = \{0\} \cup \mathcal{M}$: not applying, 0, or applying to customer type $m = (i, j) \in \mathcal{M}$; it has $|\mathcal{M}| + 1$ strategies. Its mixed strategy is the vector of application shares $(\alpha_k, (\alpha_{k,m})_{m \in \mathcal{M}})$ of Section 2.

- An *unmatched carrier at k* chooses a trip. Its strategies are $S_k^{\mathrm{U}} = \mathcal{L}$: the trip $k \to j$, with $j = k$ for staying; it has l strategies. Its mixed strategy is the vector of relocation shares $(\sigma_{kj})_{j \in \mathcal{L}}$.

- A *customer of type m* posts fees and chooses whom to hire among the applicants who arrive. Its fees satisfy its first-order conditions, as in the equilibrium; only the hiring choice is noisy.

A strategy profile $\pi = (\alpha_k, \sigma_k)_{k \in \mathcal{L}}$ collects the carriers' mixed strategies.

Payoffs. The expected payoff of a strategy, holding the profile fixed, is the value it leads to. For an open carrier at k,

$$u^{\mathcal{O}}_{k,0}(\pi) = V^{\mathrm{U}}_k, \qquad u^{\mathcal{O}}_{k,m}(\pi) = V^{\mathrm{U}}_k + S_{k,m}, \tag{B.1}$$

where $S_{k,m}$ is the rent a customer of type m offers carriers from k: its marginal product of their queue. For an unmatched carrier at k,

$$u^{\mathrm{U}}_{kj}(\pi) = V^{\mathrm{T}}_{kj}. \tag{B.2}$$

These payoffs depend on the whole profile: the values depend on future choices, and the rents depend on the queues, which depend on the stocks that the profile generates. A customer of type m that hires an applicant from h gets the surplus $\Delta_{h,m}$.

C.2 The logit equilibrium

In a quantal response equilibrium, agents observe noisy payoffs $\tilde{u} = u + \varepsilon$, with the ε drawn independently from an extreme value distribution with precision λ. Each agent then plays each strategy with the logit probability of its payoff:

$$\alpha_{k,m} = \frac{e^{\lambda S_{k,m}}}{1 + \sum_{m'} e^{\lambda S_{k,m'}}}, \qquad \sigma_{kj} = \frac{e^{\lambda V^{\mathrm{T}}_{kj}}}{\sum_{j'} e^{\lambda V^{\mathrm{T}}_{kj'}}}, \tag{B.3}$$

since V^{U}_k is common to every strategy of an open carrier. The value of a choice is the expected maximum of its noisy payoffs:

$$V^{\mathrm{U}}_k = \frac{1}{\lambda} \ln \sum_j e^{\lambda V^{\mathrm{T}}_{kj}}, \qquad V^{\mathcal{O}}_k = V^{\mathrm{U}}_k + \frac{1}{\lambda} \ln\Big(1 + \sum_m e^{\lambda S_{k,m}}\Big), \tag{B.4}$$

with V^{T} and V^{F} given by the value equations of Section 2.1.1. A customer compares each applicant's surplus with noise: an applicant from h counts as better than a level z

with probability $s_h(z) = 1/\left(1 + e^{-\lambda(\Delta_h - z)}\right)$. Replacing the indicator $\mathbf{1}\{\Delta_h > z\}$ of (F1) in Appendix A.3 by $s_h(z)$, a customer's expected surplus is $R(\theta) = \int_0^\infty \left(1 - e^{-\Theta(z)}\right) dz$ with $\Theta(z) = \sum_h \theta_h s_h(z)$, and its rents and hiring probabilities are

$$S_k = \frac{\partial R}{\partial \theta_k} = \int_0^\infty s_k(z)\, e^{-\Theta(z)}\, dz, \qquad \pi_k^{\mathrm{X}} = \frac{\partial R}{\partial \Delta_k} = \int_0^\infty \theta_k \lambda\, s_k(z)\big(1 - s_k(z)\big) e^{-\Theta(z)}\, dz. \quad \text{(B.5)}$$

A *steady-state logit equilibrium* at precision λ is a profile π, values and stocks $\mathbf{v}$ such that the choices are the logit responses (B.3), the values satisfy (B.4) and the value equations, the queues are $\theta_{k,m} = \alpha_{k,m} v_k^{\mathcal{O}} / x_m$, the rents and hires are (B.5), and the stocks are stationary under the laws of motion (5) with hires $x_m \pi_{k,m}^{\mathrm{X}}$. As $\lambda \to \infty$, the logit responses become best responses, the log-sums become maxima, s_h becomes the indicator, and (B.5) becomes the rents and hires of the Δ ranking: these are the conditions of Definition 1 in a steady state. At $\lambda = 0$, every agent mixes uniformly.

The noise serves the computation. The exact equilibrium is full of corners — locations from which no carrier applies, trips that no carrier takes, applicants who are never hired — and of intervals in which carriers are exactly indifferent, where its conditions are not differentiable. At $\lambda < \infty$ every strategy is played with positive probability and every condition is smooth.

C.3 The computational problem

Take not applying as the reference strategy of an open carrier, and staying as the reference strategy of an unmatched carrier. Writing the logit responses as log-odds, a logit equilibrium solves

$$H_{k,m}^{\mathcal{O}} \equiv \ln \alpha_{k,m} - \ln \alpha_k - \lambda\left(u_{k,m}^{\mathcal{O}} - u_{k,0}^{\mathcal{O}}\right) = \ln \alpha_{k,m} - \ln \alpha_k - \lambda S_{k,m} = 0, \quad \text{(B.6)}$$

$$H_{kj}^{\mathrm{U}} \equiv \ln \sigma_{kj} - \ln \sigma_{kk} - \lambda\left(V_{kj}^{\mathrm{T}} - V_{kk}^{\mathrm{T}}\right) = 0, \quad \text{(B.7)}$$

one equation for each strategy other than the reference one, together with the equations that make the payoffs consistent with the profile: the values (B.4), and the stationarity of the stocks under (5), in which one equation is replaced by the fleet size because the laws of motion conserve the fleet. The shares of each agent add up to one by construction, since they are computed from the log-odds.

Two features of the model reduce the system. In a steady state the trip and hired-carrier values are closed-form functions of $V^{\mathcal{O}}$, so (B.7) is solved by the logit formula once $V^{\mathcal{O}}$ is known, and the relocation shares drop out of the unknowns. And stocks must stay positive, so they are carried in logs. The unknowns are therefore

$$z = \Big((V_k^{\mathcal{O}})_k,\ \big(\ln \alpha_{k,m} - \ln \alpha_k\big)_{k,m},\ \ln(\text{stocks}) \Big),$$

where the stocks are the open carriers at each location, the carriers in transit on each trip $k \to j \neq k$, and the remote hires on their way to the cargo; a local hire loads within the period, so its fixed stock is always zero and is left out. The logit equilibrium at precision λ is a zero of a system $H(z, \lambda) = 0$ with as many equations as unknowns.

C.4 Computational steps

I trace the logit equilibria along a path of increasing precision (Turocy, 2005).

1. Start at $\lambda_0 = 0.5$, where choices are close to uniform and the solution is easy to find: application shares equal, the fleet spread evenly over open carriers, and every value at that of staying idle forever.

2. For $t = 1, \ldots, 70$, with λ_t on a geometric grid up to 10^5, solve $H(z, \lambda_t) = 0$ by Powell's hybrid method, starting from the solution at λ_{t-1}. If a step fails, retry it through the geometric midpoint of λ_{t-1} and λ_t, up to four times.

3. Read the equilibrium off the solution at $\lambda = 10^5$, where a difference in payoffs of

10^{-4} already tilts a choice by a factor of e^{10}.

If Powell's method fails at some λ, a policy iteration supplies a new starting point. Given the profile, it evaluates the payoffs — the values, rents and stocks the profile implies — and then improves the profile toward the logit responses to those payoffs, incrementally: the application shares move to $(1-\rho)\alpha + \rho\alpha'$, with α' the logit response and $\rho = 0.05$, and the values and stocks move by the same fraction toward their update.

C.5 The exact steady state, and accuracy

For $\beta < 1$, an exact steady state is computed as a fixed point. Given stocks $\bar{\mathbf{v}}$, solve

$$\max_{\nu} \ \Pi_\beta(\nu) \quad \text{s.t.} \quad \nu = (\mathbf{v}, f) \in G, \qquad \mathbf{v} = (1-\beta)\bar{\mathbf{v}} + \beta F(\nu), \tag{B.8}$$

where ν stands for the discounted average $(1-\beta)\sum_t \beta^t(\mathbf{v}_t, f_t)$ of the stocks and flows along a path from $\bar{\mathbf{v}}$, and G is the set of stocks and flows that satisfy (A.1) and (13). Then replace $\bar{\mathbf{v}}$ by the stocks of the solution, and repeat until they stop changing. At a fixed point, $\mathbf{v}^* = \bar{\mathbf{v}}$, the constraint reads $\mathbf{v}^* = F(\nu^*)$, so ν^* is stationary; and since the discounted average of every feasible path from $\mathbf{v}^*$ satisfies the constraints of (B.8), the stationary path earns at least as much as any of them: it attains (14) from $\mathbf{v}^*$, so by Proposition 1 it is an equilibrium steady state. Problem (B.8) is linear except for the hiring constraints (13), whose right side is concave; they are imposed through tangent cuts (Kelley, 1960), added at the current applications for every violated constraint until none is violated by more than 10^{-7}, and each linear program is solved with HiGHS (Huangfu and Hall, 2018).

Table 10 compares the logit equilibrium at $\lambda = 10^5$ with the exact steady state, in the two-location example of Section 4, at values of w in each regime and as each regime begins. The two agree to four decimals where carriers' choices are strict. As local hiring and remote hiring begin, where carriers at 2 are exactly indifferent between staying and sailing back, the share that sails back differs by less than one percent: the noise matters most where

carriers are indifferent. With no trade, any distribution of idle carriers is a steady state, and the two methods settle on different ones; neither ships any cargo. Both also agree with the regimes of Figure 5, found by a third, direct method.

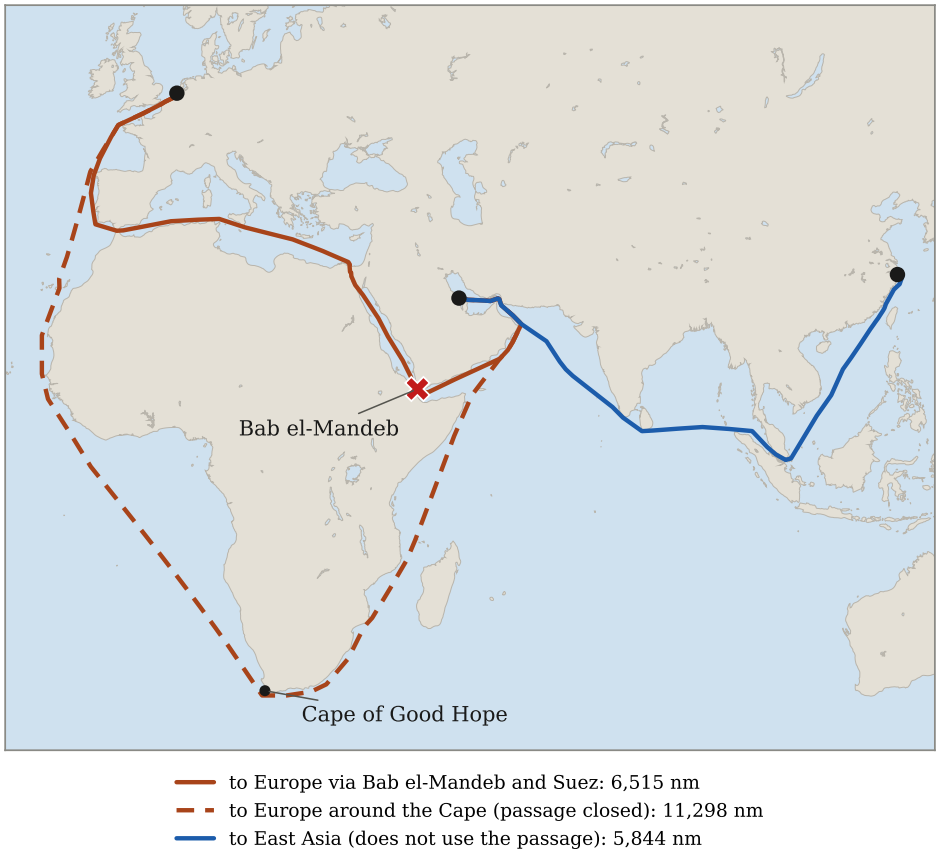


Figure 1: The three locations of the closure example. Voyages from the Persian Gulf to Europe pass Bab el-Mandeb and the Suez Canal, and go around the Cape of Good Hope when the passage is closed; voyages from the Gulf to East Asia do not use the passage.

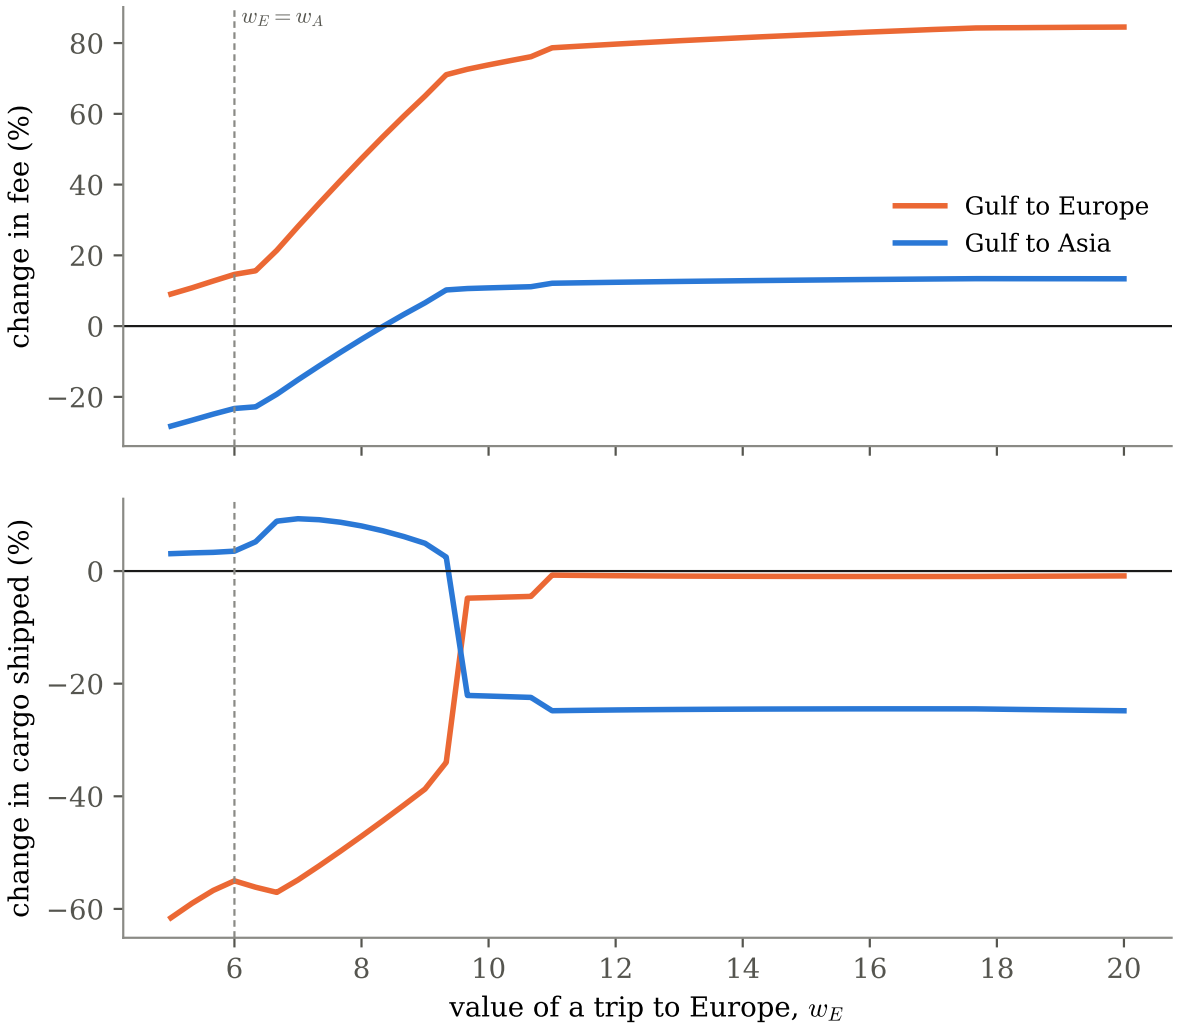


Figure 2: The effect of the closure on the fees for a carrier in the Gulf (top) and on cargo shipped (bottom), in percent of their values with the passage open, by the value of a trip to Europe. The dashed line marks equal values of a trip to Europe and to Asia. Other parameters as in Table 3.

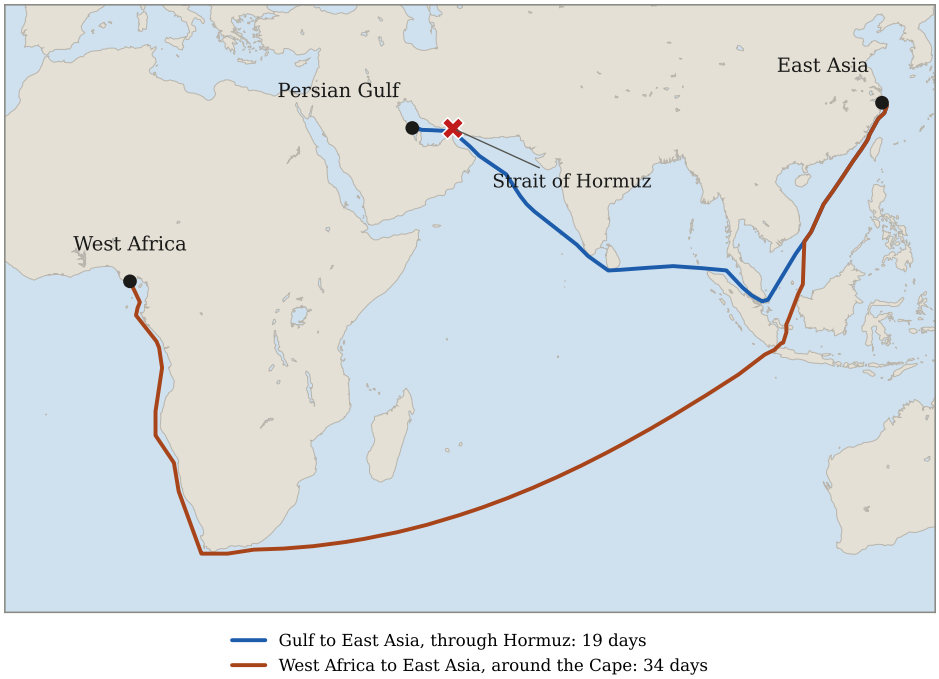


Figure 3: The three locations of the toll example. Voyages from the Persian Gulf to East Asia pass the Strait of Hormuz; voyages from West Africa go around the Cape of Good Hope.

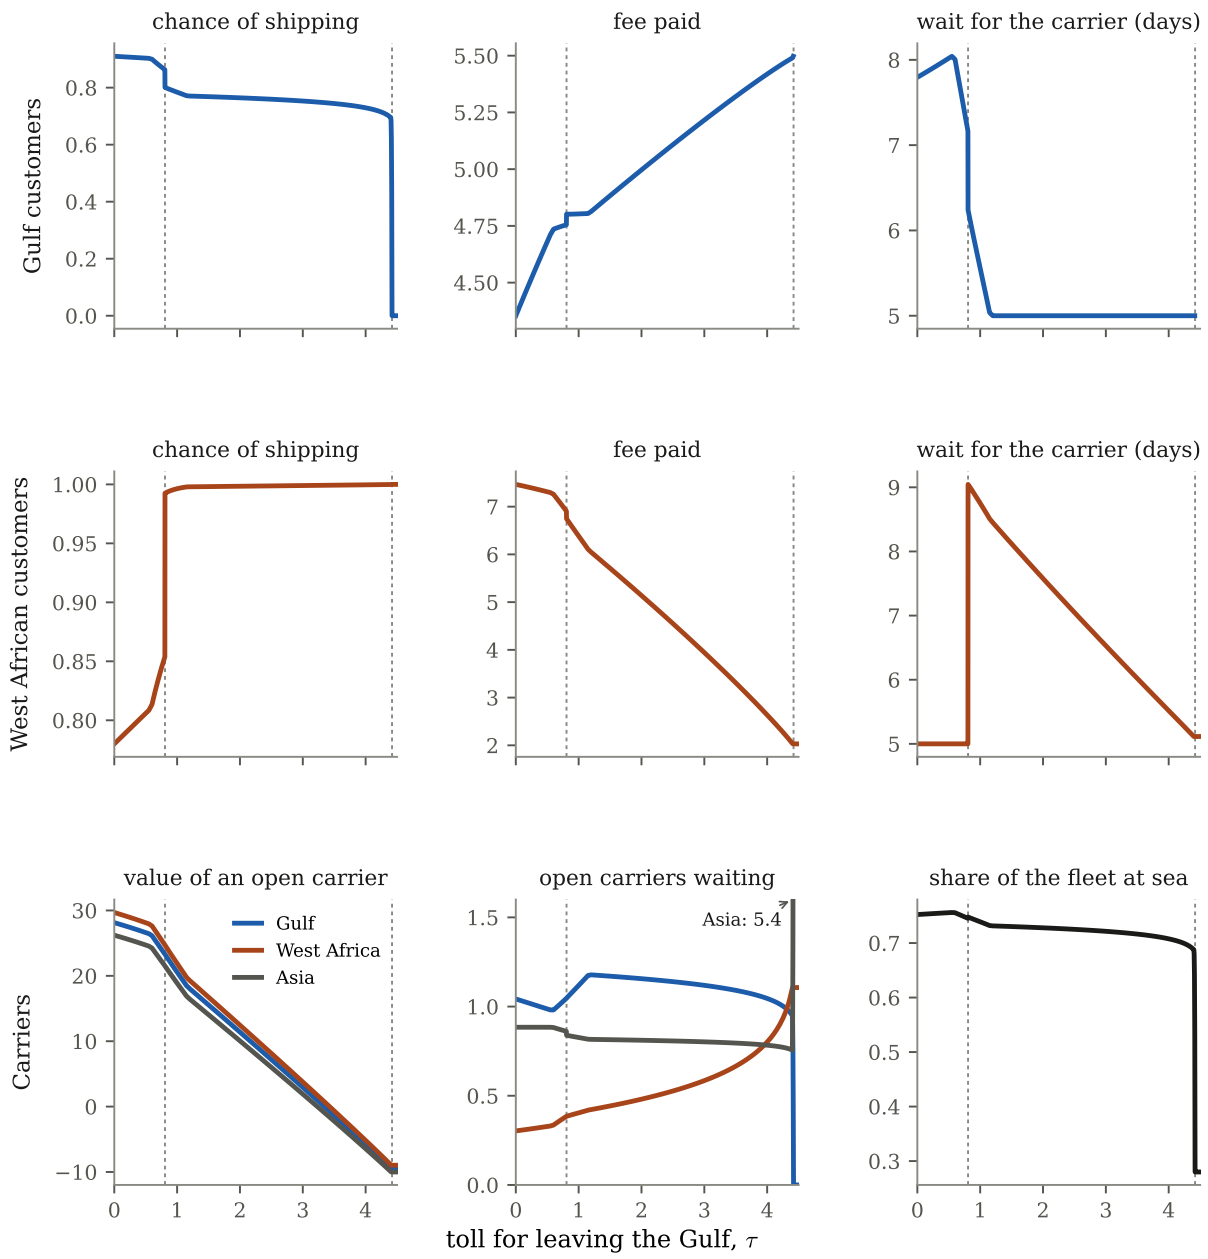


Figure 4: Steady states by the toll for leaving the Gulf, τ. Top and middle: Gulf and West African customers' chances of shipping, the fees they pay and their wait for a carrier. Bottom: the value of an open carrier and the open carriers waiting at each location, and the share of the fleet at sea; once Gulf customers stop shipping, 5.4 carriers wait in Asia. Dotted lines: $\tau \approx 0.81$, where West African customers start hiring carriers in Asia, and $\tau \approx 4.42$, where Gulf customers stop shipping. Parameters as in Table 5.

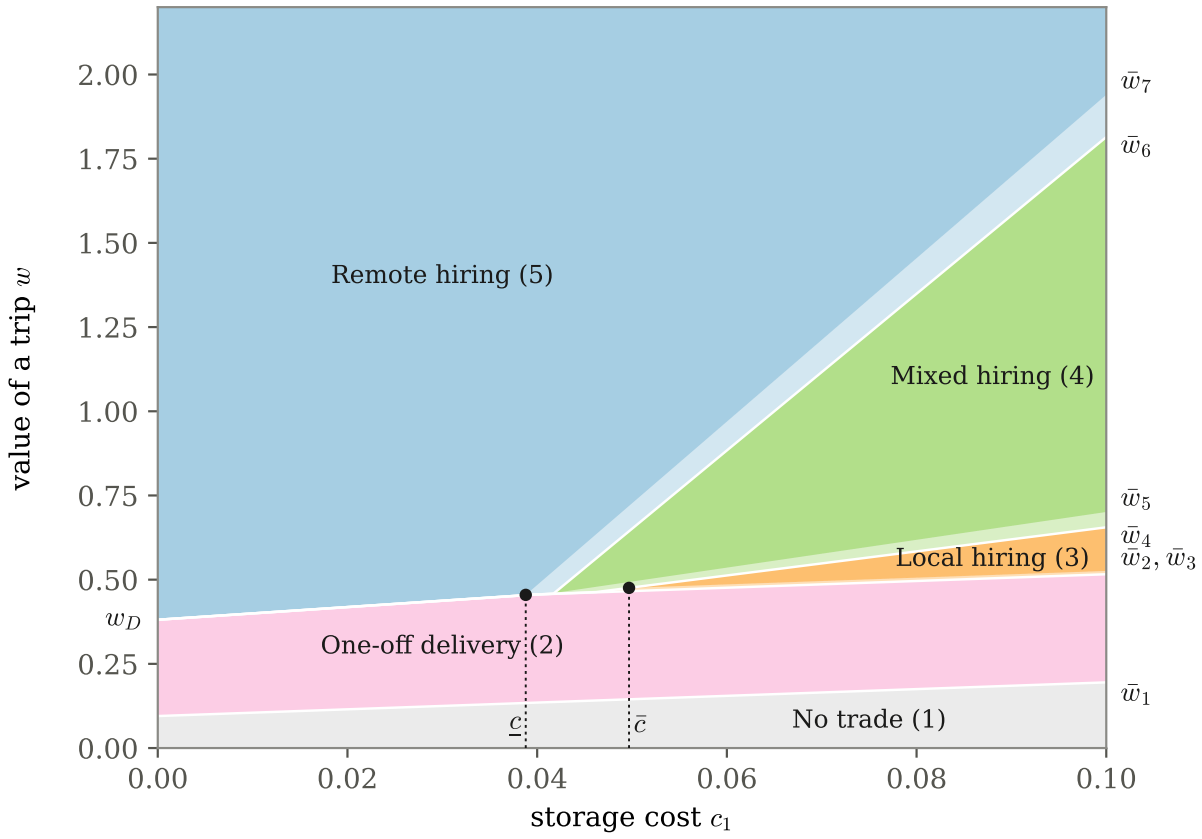


Figure 5: Long-run regimes of the two-location economy by the storage cost c_1 and the value of a trip w, drawn from the cutoffs; the other parameters are those of the example. Regimes are numbered as in Table 7, and the three short intervals in which a growing share of carriers at 2 switches are shown in a lighter shade of the regime they start: partial repositioning, between $\bar{w}_2$ and $\bar{w}_3$, in which a growing share of unmatched carriers at 2 sails back; remote entry, between $\bar{w}_4$ and $\bar{w}_5$, in which a growing share of carriers at 2 applies; and partial staying, between $\bar{w}_6$ and $\bar{w}_7$, in which a growing share of unhired carriers at 2 stays. The cutoffs of Proposition 3 are marked on the right, and w_D on the left; $\underline{c}$ and $\bar{c}$ are the storage-cost thresholds. The dot above $\underline{c}$ is where the regimes between one-off delivery and remote hiring first appear, at $w_D = \bar{w}_2$; the dot above $\bar{c}$ marks $\bar{w}_3 = \bar{w}_4$, above which every carrier at 2 sails back before any is hired remotely.

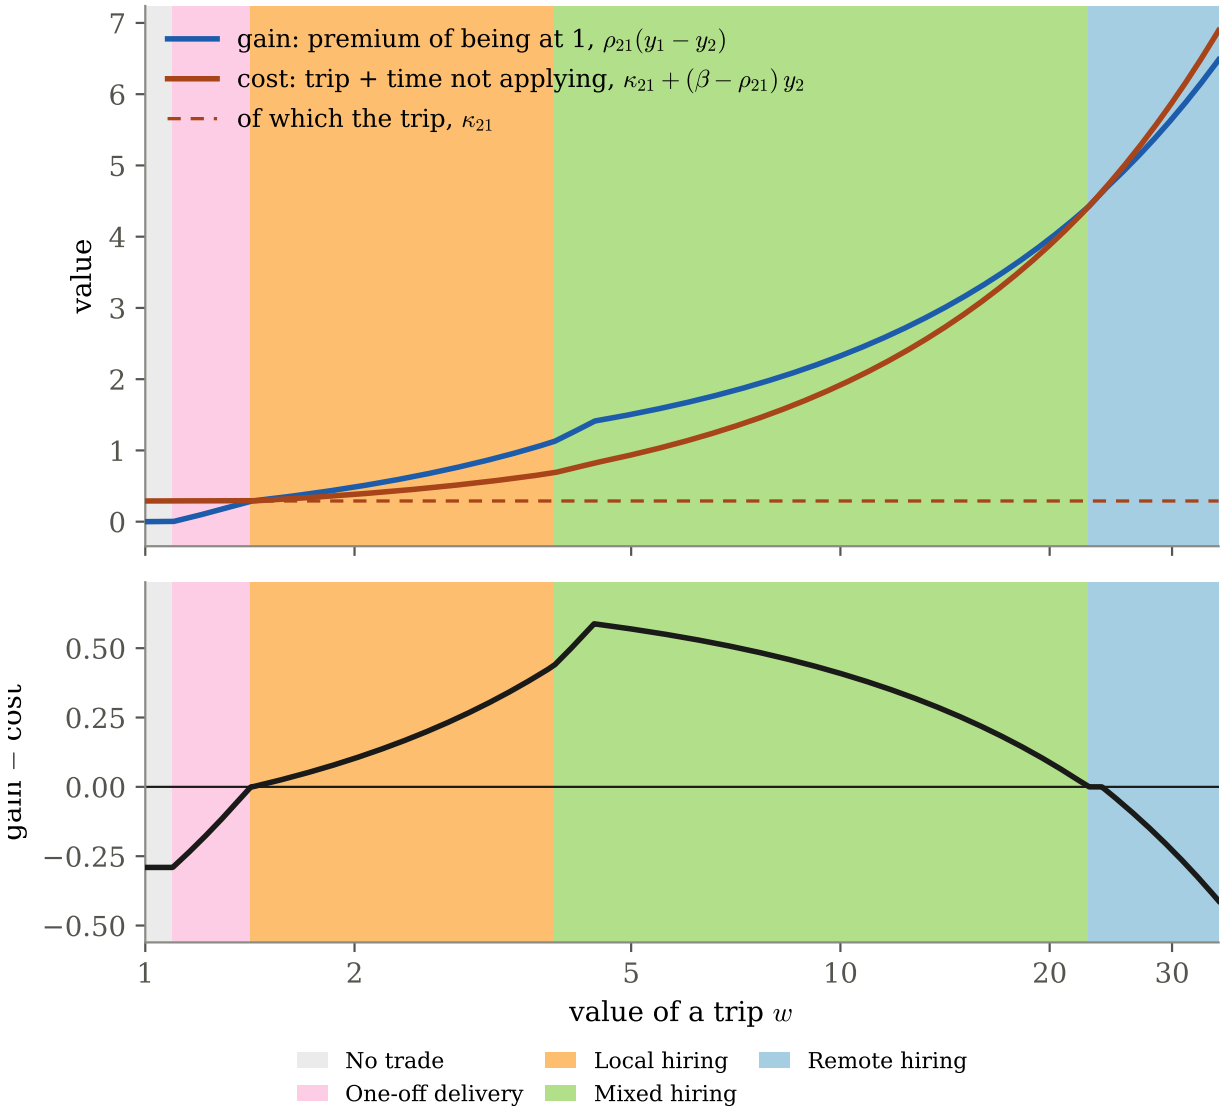


Figure 6: The decision of an unmatched carrier at 2, by the value of a trip w on a log scale, in the example. Top: the gain and the cost of sailing back to 1; the dashed line is the part of the cost due to the trip, κ_{21}. Bottom: gain minus cost, $V_{21}^{\mathrm{T}} - V_{22}^{\mathrm{T}}$; carriers sail back where it is positive. The shading shows the regimes, as in Table 7 and Figure 5.

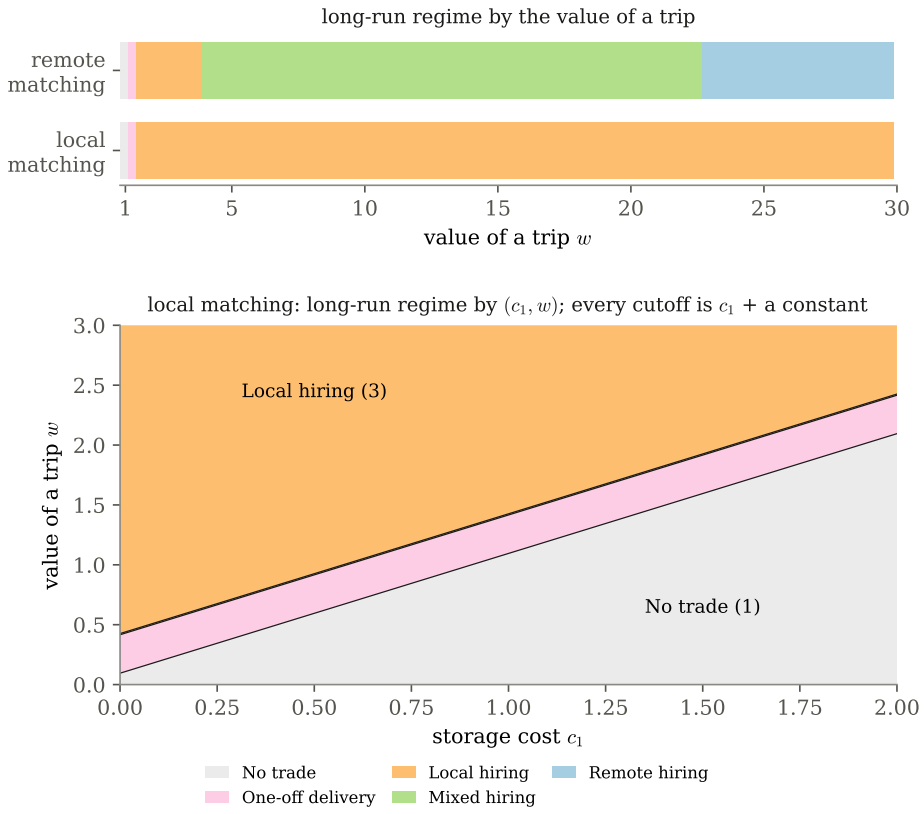


Figure 7: Long-run regimes with remote and with local matching, by the value of a trip w in the example (top), and with local matching by the storage cost c_1 and w (bottom). Regimes are named as in Table 7.

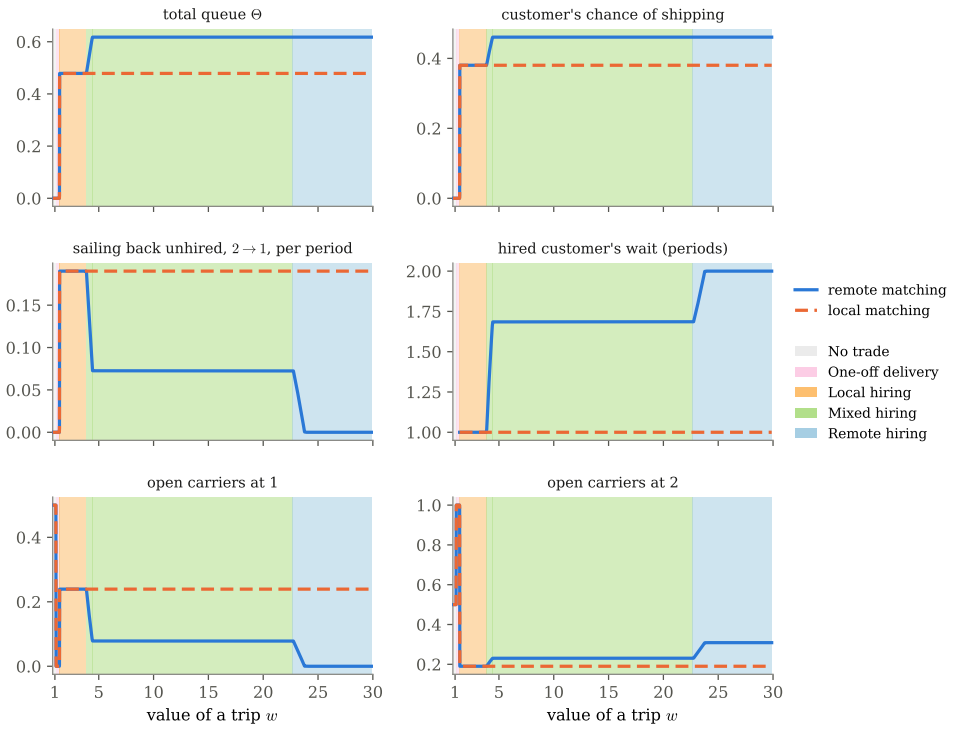


Figure 8: Steady states with remote matching (solid) and with local matching (dashed), by the value of a trip w in the example. The shading shows the regimes with remote matching, as in Table 7.

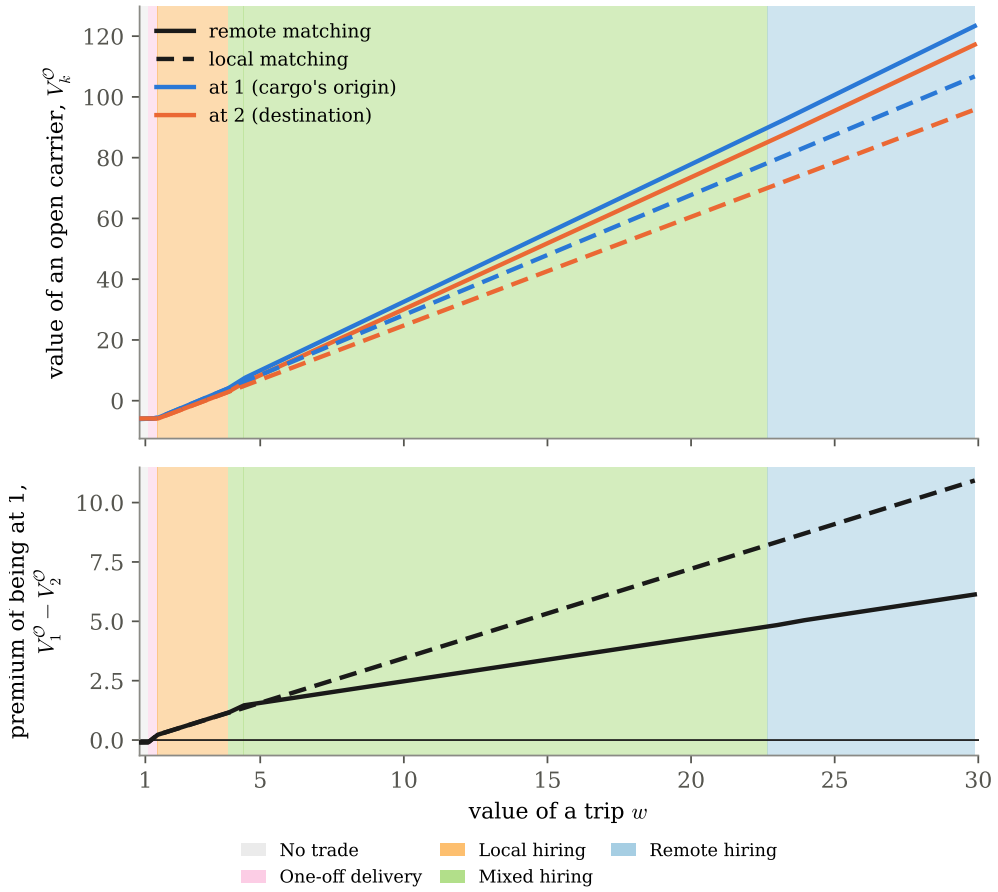


Figure 9: Value of an open carrier at 1 and at 2 (top) and the premium of being at 1 (bottom), with remote matching (solid) and with local matching (dashed), by the value of a trip w in the example. The shading shows the regimes with remote matching.

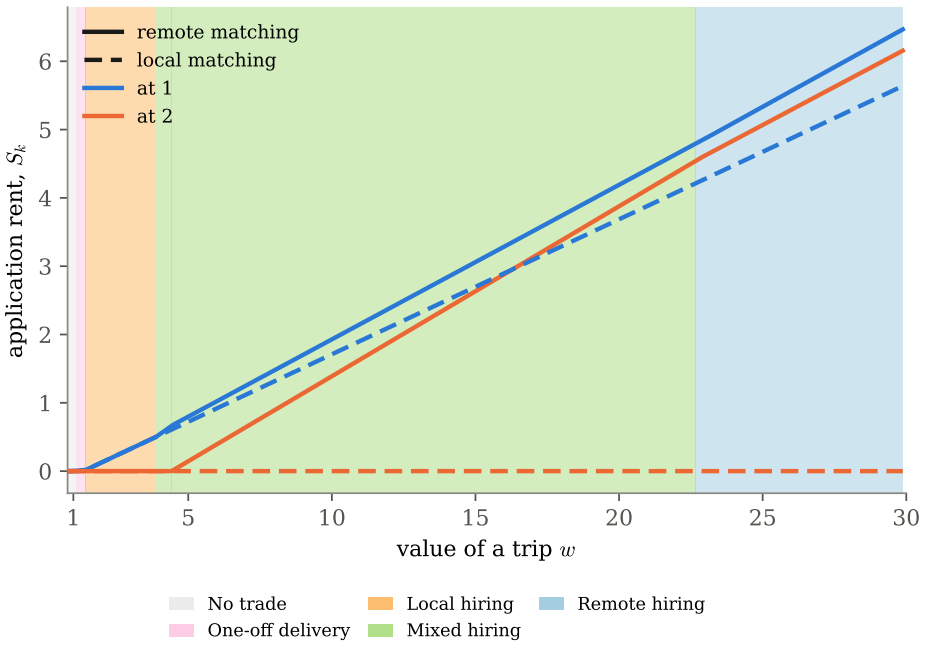


Figure 10: Application rents of carriers at 1 and at 2 with remote matching (solid) and with local matching (dashed), by the value of a trip w in the example. With local matching, $S_2 = 0$. The shading shows the regimes with remote matching.

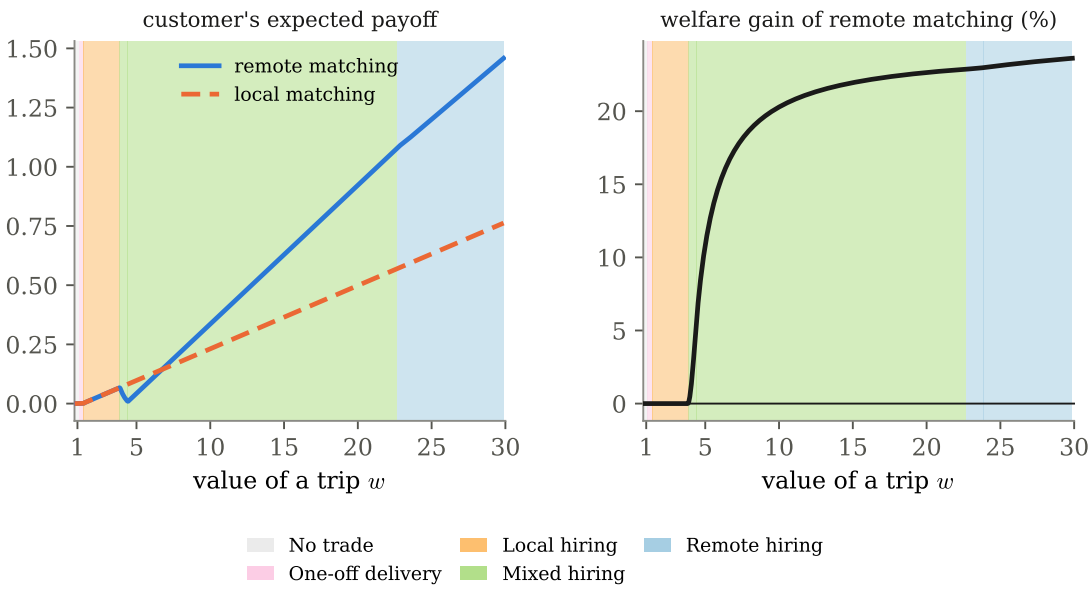


Figure 11: A customer's expected payoff with remote matching (solid) and with local matching (dashed), and the welfare gain of remote matching, by the value of a trip w in the example. Welfare is $(1-\beta)$ times the equilibrium's total discounted surplus (14) from a start with all carriers open, half at each location. The shading shows the regimes with remote matching.

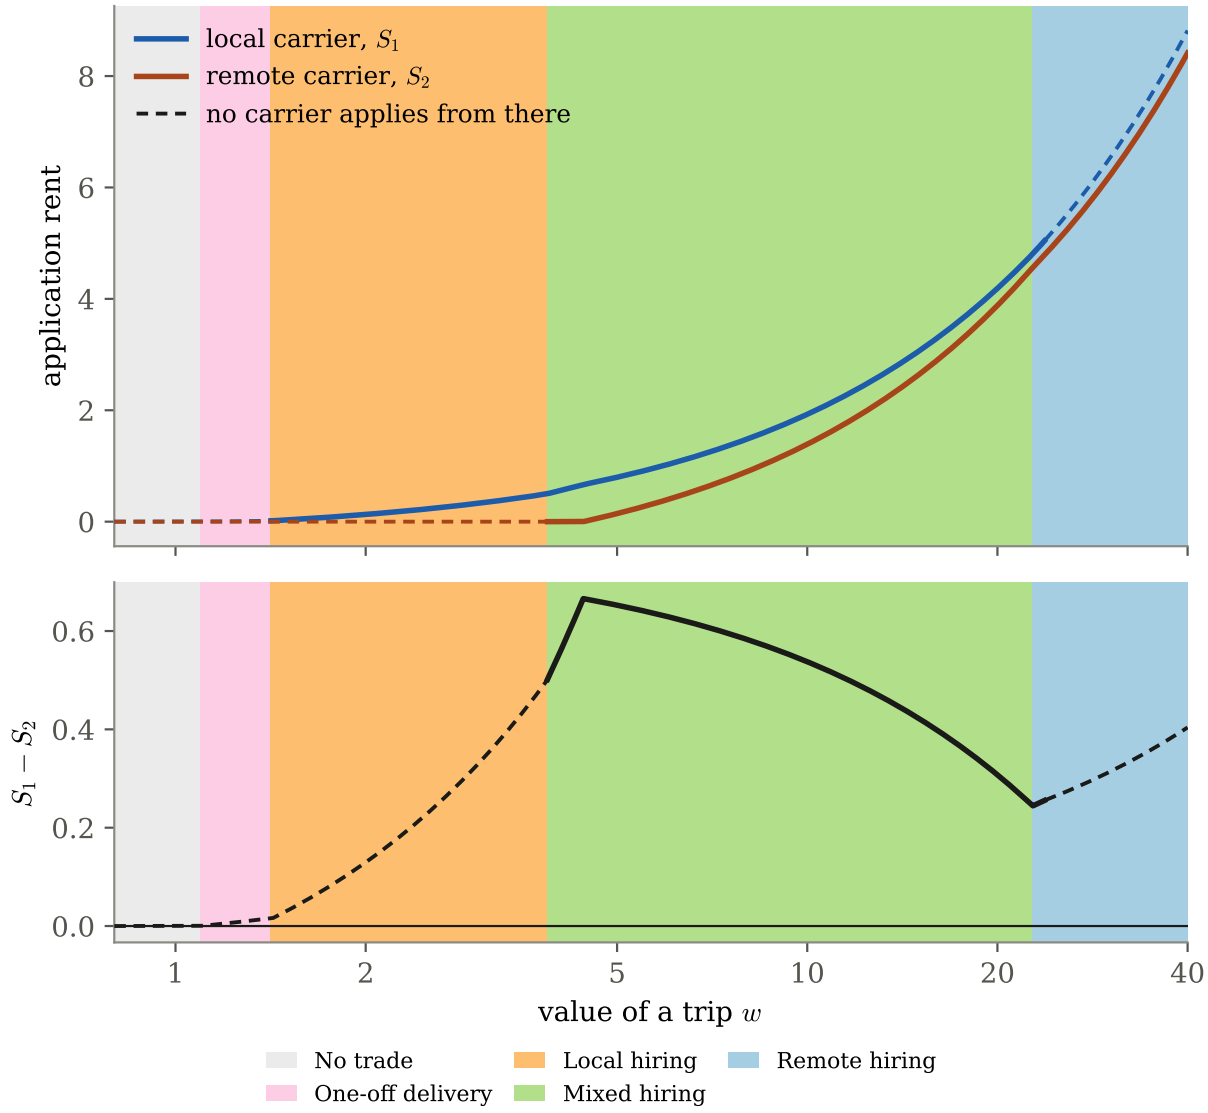


Figure 12: Steady-state application rents of carriers at locations 1 and 2 (top) and their difference (bottom), by the value of a trip w on a log scale, in the example. Dashed: rents at a location from which no carrier applies. The shading shows the regimes, as in Table 7.

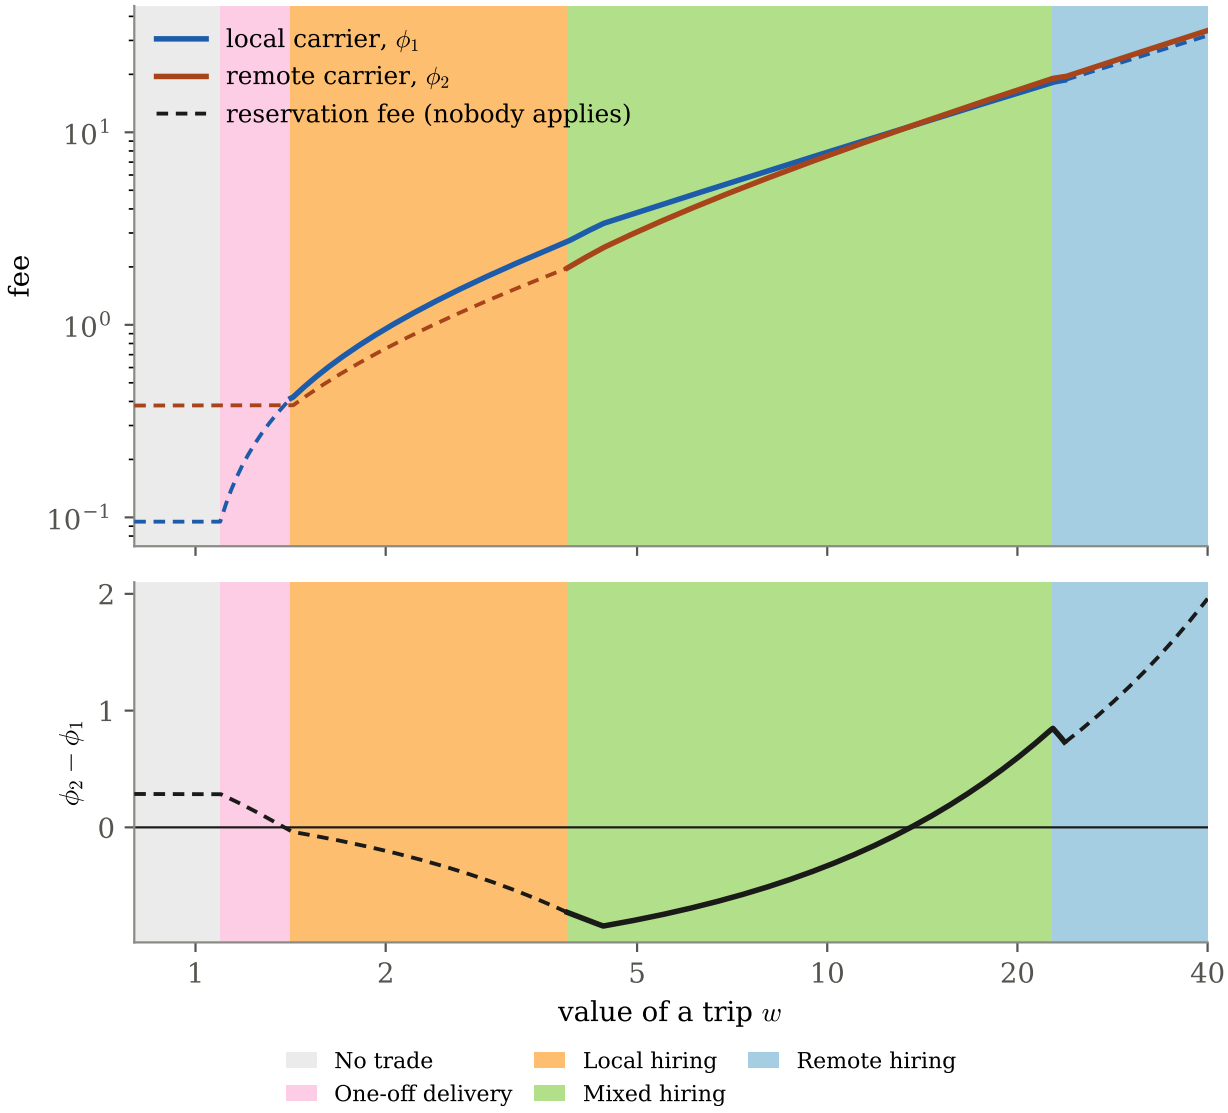


Figure 13: Steady-state fees for a local and a remote carrier (top, log scale) and their difference (bottom), by the value of a trip w on a log scale, in the example. Dashed: reservation fees, where no carrier applies from that location. The shading shows the regimes, as in Table 7.

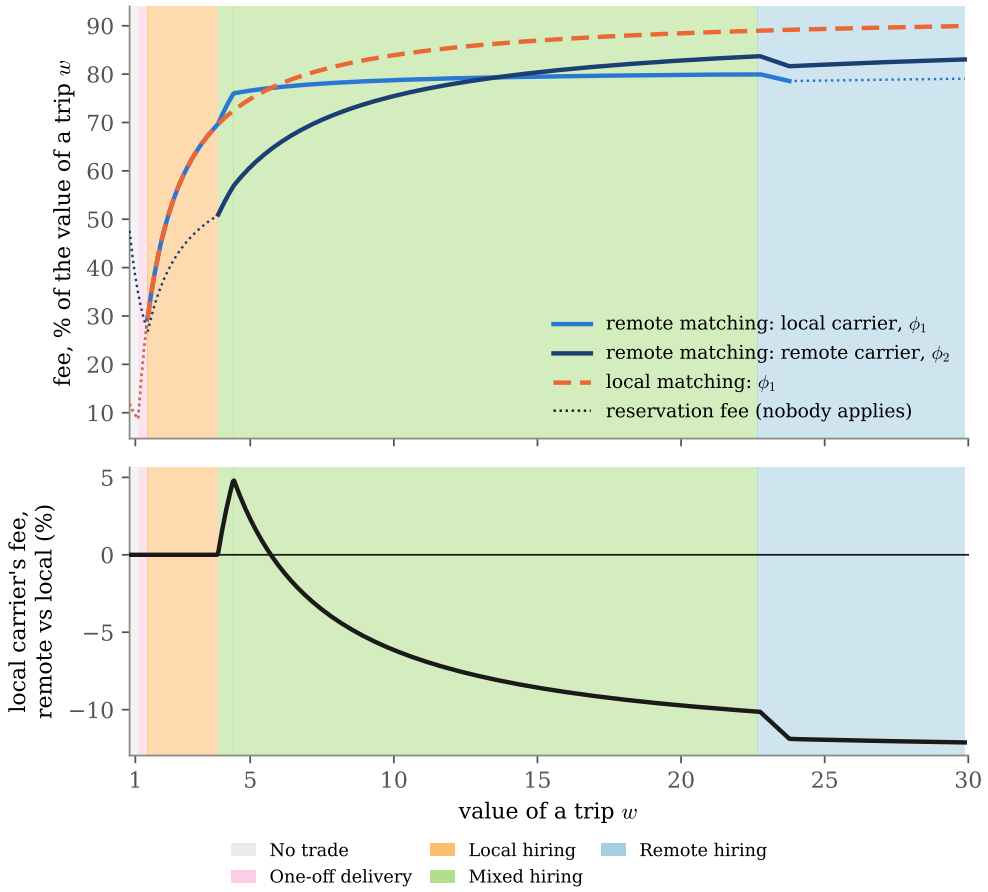


Figure 14: Fees for a local and a remote carrier with remote matching and the fee for a local carrier with local matching, as a percent of the value of a trip (top), and the percent difference in the fee for a local carrier (bottom), by the value of a trip w in the example. Dotted: reservation fees, where no carrier applies from that location. The shading shows the regimes with remote matching.

Table 1: Notation

Symbol	Meaning
Indices	
$k, h \in \mathcal{L}$	locations of carriers
$(i, j) \in \mathcal{M}$	customer type: cargo at origin i bound for destination j
t	date; dropped in a steady state
Primitives	
β	discount factor
d_{kj}	probability that a leg $k \to j$ ends this period; $d_{kk} = 1$
c_{kj}, c_{kk}	per-period cost of sailing $k \to j$; of staying at k
c_i	customer's per-period storage cost at i
w_{ij}	value of shipping a cargo (i, j)
x_{ij}	inflow of customers (i, j)
v	fleet size
Stocks of carriers	
$v_k^{\mathcal{O}}$	open at k
v_{kj}^{T}	in transit $k \to j$, empty or laden
$v_{k,ij}^{\mathrm{F}}$	hired at k by a customer (i, j), sailing to i
Choices and matching	
$\phi_{k,ij}$, $\succ_{ij}$	fee for a carrier from k; ranking of origins
$\alpha_{k,ij}$, α_k	share of open carriers at k applying to (i, j); not applying
σ_{kj}	share of unhired carriers at k sailing to j ($j = k$: staying)
$\theta_{k,ij}$	queue: applicants from k per customer (i, j)
$\Theta_{\succ k,ij}$	queue of applicants ranked above k
$\pi_{k,ij}^{\mathrm{X}}$, $\pi_{k,ij}^{\mathrm{V}}$	probability that a customer hires from k; that an applicant from k is hired
$m_{k,ij}$	hires from k by customers (i, j)
Values	
$V_k^{\mathcal{O}}$, V_k^{U}	carrier open at k; not hired at k
V_{kj}^{T}, $V_{k,ij}^{\mathrm{F}}$	carrier in transit $k \to j$; hired at k by (i, j)
S_k	application rent of a carrier at k, $V_k^{\mathcal{O}} - V_k^{\mathrm{U}}$
$X_{k,ij}^{\mathrm{F}}$, $X_{ij}^{\mathcal{O}}$	customer's value of waiting for a carrier from k; expected payoff
$\Delta_{k,ij}$	joint surplus of a match between a customer (i, j) and a carrier from k

Table 2: What each agent chooses within a period

Stage	Agent	Chooses	Taking as given
1	Customer (i, j)	fee menu ϕ_{ij} and ranking $\succ_{ij}$	carriers' rents S_k
2	Open carrier at k	whether and where to apply, α_k	fees, rankings, queues
3	Customer (i, j)	hires its highest-ranked applicant	its posted ranking
3	Unhired carrier at k	stay or sail to j, σ_k	trip values V^{T}_{kj}
–	Carrier in transit or fixed	nothing; sails on	—

Table 3: Parameters of the three-location example

Parameter	Value	Note
Period, discount factor	5 days, $\beta = 0.99$	
Trip length, Gulf–Europe	21 days, 36 when closed	via Suez; around the Cape
Trip length, Europe–Asia	33.5 days, 44 when closed	via Suez; around the Cape
Trip length, Gulf–Asia	19 days	does not use the passage
Sailing cost, idle cost	0.05, 0.02 per day	c_{kj}, c_{kk}
Storage cost in the Gulf	0.1 per day	c_{Gulf}
Value of a trip	12 to Europe, 6 to Asia	w_{ij}
Customer inflow	0.06 to Europe, 0.16 to Asia per day	x_{ij}
Fleet	9	v

Table 4: Steady states with the passage open and closed

	Open	Closed	Change
Cargo shipped per period, Gulf to Europe	0.295	0.293	−1%
Cargo shipped per period, Gulf to Asia	0.644	0.485	−25%
Fee for a carrier in the Gulf, Europe cargo	4.86	8.74	+80%
Fee for a carrier in the Gulf, Asia cargo	4.48	5.03	+12%
Share of hires that are remote, Europe cargo	0.22	0.30	
Share of hires that are remote, Asia cargo	0.04	0.00	
Carriers waiting in the Gulf	1.62	1.09	−33%
Share of the fleet at sea	0.72	0.79	
Value of an open carrier, Gulf	28.4	36.5	+29%
Value of an open carrier, Europe	26.3	32.4	+23%
Value of an open carrier, Asia	26.5	34.4	+29%

Table 5: Parameters of the toll example

Parameter	Value	Note
Period, discount factor	5 days, $\beta = 0.99$	
Trip length, Gulf–Asia	19 days	through Hormuz
Trip length, West Africa–Asia	34 days	around the Cape
Trip length, Gulf–West Africa	26 days	empty sailings only
Sailing cost, idle cost	0.05, 0.02 per day	c_{kj}, c_{kk}
Storage cost	0.1 per day	c_{Gulf}, c_{WAF}
Value of a trip	6 from the Gulf, 9 from West Africa	w_{ij}
Customer inflow	0.16 from the Gulf, 0.04 from West Africa per day	x_{ij}
Fleet	9	v
Toll for leaving the Gulf	0 to 4.5	τ, paid on departure

Table 6: Steady states by the toll for leaving the Gulf

	$\tau = 0$	$\tau = 1$	$\tau = 2$	$\tau = 4.42$
Fee paid, Gulf cargo	4.35	4.80	5.00	—
Customer's chance of shipping, Gulf	0.91	0.78	0.76	0
Customer's wait for its carrier (days), Gulf	7.8	5.6	5.0	—
Share of Gulf hires that are remote	0.20	0.04	0	—
Customer's expected payoff, Gulf	0.80	0.50	0.38	0
Fee paid, West African cargo	7.47	6.39	5.14	2.03
Customer's chance of shipping, West Africa	0.78	1.00	1.00	1.00
Customer's wait for its carrier (days), West Africa	5.0	8.8	7.6	5.1
Share of West African hires that are remote	0	0.13	0.09	0.00
Customer's expected payoff, West Africa	0.81	1.76	3.12	6.46
Value of an open carrier, Gulf	28.1	20.6	11.4	−9.4
Value of an open carrier, West Africa	29.7	21.9	12.5	−9.0
Value of an open carrier, Asia	26.2	18.9	10.0	−10.0

Fees are averages over the hires from each location. A customer's wait runs from the contract to the carrier's arrival at the cargo. At $\tau = 4.42$, the smallest toll at which no Gulf customer ships, there are no Gulf hires.

Table 7: Regimes in the two-location economy

	Regime	w	Carriers
1	No trade	$w \leq \bar{w}_1$	stay; no trade at any date
2	One-off delivery	$\bar{w}_1 < w < \bar{w}_2$	those starting at 1 deliver once, then stay at 2
3	Local hiring	$\bar{w}_2 \leq w < \bar{w}_4$	at 2: sail back (all from $\bar{w}_3$)
4	Mixed hiring	$\bar{w}_4 \leq w < \bar{w}_6$	at 2: apply (all from $\bar{w}_5$); unhired sail back
5	Remote hiring	$w \geq \bar{w}_6$	at 2: all apply; unhired stay (all from $\bar{w}_7$)

Each of regimes 3 to 5 begins with a short interval, between $\bar{w}_2$ and $\bar{w}_3$, $\bar{w}_4$ and $\bar{w}_5$, or $\bar{w}_6$ and $\bar{w}_7$, in which a growing share of carriers at 2 switches to the new behavior.

Table 8: The outcomes of a customer at 1

Outcome	Probability	Customer gets	Wait (periods)
Local carrier	$1 - e^{-\theta_1}$	$w - c_1 - \phi_1$	1
Remote carrier	$e^{-\theta_1}\left(1 - e^{-\theta_2}\right)$	$w - \dfrac{c_1}{1 - \beta(1 - d_{21})} - \phi_2$	$1/d_{21}$
Leave (no applicant)	$e^{-\theta_1 - \theta_2}$	0	—

Table 9: Steady states with remote and local matching in the two-location example

	$w = 5$		$w = 10$		$w = 30$	
	Remote	Local	Remote	Local	Remote	Local
Regime	4	3	4	3	5	3
Empty sailings from 2 to 1, per period	0.073	0.190	0.073	0.190	0	0.190
Customer's wait for its carrier, periods	1.68	1	1.68	1	2	1
Customer's chance of shipping	0.461	0.380	0.461	0.380	0.461	0.380
Open carriers at 1	0.078	0.239	0.078	0.239	0	0.239
Open carriers at 2	0.230	0.190	0.230	0.190	0.309	0.190
Value of an open carrier at 1, $V_1^{\mathcal{O}}$	9.9	8.4	32.5	28.2	123.9	107.3
Value of an open carrier at 2, $V_2^{\mathcal{O}}$	8.4	6.9	30.1	24.8	117.8	96.3
Premium of being at 1	1.6	1.6	2.5	3.5	6.2	11.0
Application rent at 1, S_1	0.80	0.72	1.93	1.71	6.50	5.66
Application rent at 2, S_2	0.14	0	1.39	0	6.18	0
Fee for a local carrier, ϕ_1	3.83	3.74	7.88	8.39	23.72*	26.99
Fee for a remote carrier, ϕ_2	3.04	—	7.55	—	24.92	—
Customer's expected payoff	0.04	0.10	0.34	0.23	1.46	0.77
Welfare gain of remote matching	11%		20%		24%	

*Reservation fee: in remote hiring no carrier waits at 1.

Table 10: Quantal response equilibrium at $\lambda = 10^5$ and the exact steady state

		Quantal response			Exact		
Regime	w	matches	remote	σ_{21}	matches	remote	σ_{21}
One-off delivery	1.300	0	—	0	0	—	0
Local hiring, as it begins	1.421	0.1174	0	0.2283	0.1178	0	0.2299
Local hiring	2.500	0.1902	0	1	0.1902	0	1
Mixed hiring, as it begins	4.100	0.2081	0.3163	1	0.2081	0.3164	1
Mixed hiring	10.000	0.2304	0.6849	1	0.2304	0.6849	1
Remote hiring, as it begins	23.300	0.2304	0.8451	0.4629	0.2304	0.8453	0.4625
Remote hiring	30.000	0.2304	1	0	0.2304	1	0

Matches are cargoes shipped per period, remote is the share hired from location 2, and σ_{21} is the share of unmatched carriers at 2 that sail back. The example's parameters are those of Section 4.2.