Pricing the Right to Renege in Search Markets: Evidence from Trucking

Richard Faltings *

October 15, 2026

Abstract

In many search markets, advance contracts include a right to renege, letting one party leave a match for a better one found later. Contract law would price this right at compensatory damages, the loss to the party left behind, but a broken match may have positive spillovers on other market participants. I study the pricing of the right to renege at a trucking broker that allocates shipments to carriers through dynamic auctions ahead of pickup and penalizes late cancellations through carriers' reputation. I estimate a dynamic model of carriers' search and bidding from data on bids and cancellations. Compensatory damages are more than ten times what a cancellation costs the market, since the broker recovers part of its loss on other shipments and other carriers gain most of the rest. Even a penalty equal to the market's cost lowers welfare, because carriers price it into their bids and fewer matches last to pickup. Welfare is highest with no penalty, and so is the broker's profit when the penalty is reputational, but a broker able to charge a monetary penalty may set it at a positive level in order to extract carriers' rents from outside opportunities.

Keywords: market design; timing frictions; search markets; dynamic auctions; limited commitment; breach of contract; trucking
JEL Codes: D47, D83, K12, L14

*The University of California San Diego, rfaltings@ucsd.edu. Special thanks go to Eugenio Miravete, Dan Ackerberg, Nick Buchholz, and Victoria Marone for their exceptional mentorship and support. I also thank Bob Town, Jorge Balat, Jackson Dorsey, Andrey Ordin, Lauri Kytömaa, Hayden Parsley, Andrey Fradkin, Brett Hollenbeck, Adam Kapor, Yingkang Xie, and Alexander McKay for helpful conversations and comments. This paper also benefited from comments by seminar participants at UT Austin and conference participants at WoPA 2025, IIOC 2025, ACM EC 2025, and SITE 2025 (Empirical Market Design). This research was made possible through data provided by an anonymous company. I wish to thank JL, AC, FK, and numerous others at the company for their generous support and feedback on the project. The company provided no financial compensation for this research and limited its review to ensuring the protection of proprietary information.

1 Introduction

In many decentralized markets, participants search sequentially over alternatives on the other side of the market. When faced with a potential transaction partner, a market participant must make a decision without knowing future matching opportunities. This creates a timing friction: accepting or rejecting matches without complete information can lead to suboptimal outcomes. Centralized clearing mechanisms offer a well-documented solution to these frictions, yet their success is frequently compromised by private incentives to circumvent the market (Kapor, Karnani, and Neilson 2024), occasionally resulting in a complete breakdown (Roth and Xing 1994; Niederle and Roth 2003). Motivated by the fragility of centralized clearing, this paper analyzes a decentralized contractual remedy: a *right to renege* that allows participants to form temporary matches while preserving the option value of continued search.

This option value typically carries a price in the form of a penalty. When reneging occurs on the demand side, such as for airline and hotel bookings, the penalty is often a monetary cancellation fee. In contrast, when reneging occurs on the supply side, such as for job searchers, penalties are often non-monetary in nature—in tight-knit professional circles, reneging on an accepted offer may lead to costly reputational penalties (Niederle and Roth 2009). This asymmetry is particularly salient in gig economy platforms, where consumers are typically subject to monetary fees, while workers are instead subject to pseudo-reputational penalties in the form of quality metrics that must be maintained to certain standards to receive greater match probabilities or to simply remain active on the platform.[1] One exception is the home-services platform Handy, which charged its workers fees for cancelling jobs. A 2023 settlement under California's worker-classification law limited these fees in California to a short schedule of small fixed amounts, charged only within 48 hours of the job (People v. Handy Technologies, Inc. 2023). This paper studies both the welfare implications of different penalty levels and types, and the profit-maximizing penalty design from the perspective of the firms setting them.[2]

A natural benchmark for the level of such a penalty comes from contract law, which prices a broken agreement at compensatory damages, the loss to the party left behind. Diamond and Maskin (1979) compare damage rules in a market where matched parties keep searching, so that a broken match also affects others. On a platform, the worker who cancels may take an-

1. For example, drivers working for Uber Eats can earn priority access to higher-paying deliveries by maintaining cancellation rates below 8% (Uber 2025).

2. While in principle, the right to renege can itself be subject to bilateral negotiations, I focus on "posted" penalties set by larger firms on a standardized basis.

other of the platform's jobs, and the job left behind may go to another worker, so the platform's loss on the job can differ widely from what the cancellation costs the market. A penalty also affects whether matches form at all, since a worker who may have to bear a penalty for leaving a match values the match less. I measure each of these effects for a large freight broker, and compare the penalties they imply with the penalties that maximize welfare and the broker's profit.

I focus on the penalty design problem of a large brokerage firm in the trucking industry, a pillar of the U.S. economy: in 2024, the combined value of all domestic truck shipments exceeded $11 trillion, or 74% of all domestic freight (USDOT 2024). The market is also characterized by significant search frictions due to its high fragmentation on both the demand and supply sides, with 91% of carriers operating fleets of ten or fewer trucks (FMCSA 2023). Furthermore, reneging is a widespread phenomenon in the industry.[3]

The brokerage first negotiates rates with shippers on either a contract or spot basis, and subsequently matches tendered shipments via an online auctions platform. The dynamic auction format allows carriers to renege at two stages. First, unlike standard auctions, no commitment is required at the time of bidding. Winning carriers are instead given a right of first refusal, requiring their explicit confirmation before the match is finalized. Second, even after confirming the match, carriers can still unilaterally cancel, but at the risk of a reputational penalty on future bids on the platform.[4] Whereas most economic data only observe final, realized transactions, the data provided by the brokerage firm record every stage of the matching process, making them uniquely suited to studying reneging behavior. Reneging is indeed frequent: over half of winning bids are reneged on, while 13.8% of confirmed matches are eventually canceled before pickup. As for the motives behind reneging, two empirical patterns are relevant. First, the cancellation rate is decreasing in the match price, suggesting that carriers cancel opportunistically when they find a better deal. Second, the cancellation rate is increasing in the time interval between match confirmation and pickup. Both these patterns are consistent with carriers continuing a search process in the style of McCall (1970) even after they have won and confirmed a match on the platform. Carriers also respond to the penalty: the hazard of a discretionary cancellation drops by a third at the 48-hour mark, when the penalty starts to apply.

3. To illustrate, Caplice and Sheffi (2006) discuss how reneging made combinatorial auctions infeasible: "So, even if another carrier was assigned to the outbound lane in the strategic bidding process, the shipper may choose to tender the load to an alternative carrier for a specific load. Not only is this accepted behavior for shippers, most analysts and transportation management software packages consider such opportunistic continuous move optimization a key capability."

4. Auction winners are selected according to the sum of their bids and an algorithmic quality score, which is computed on the basis of several past performance metrics, including cancellation rates.

The data also measure the compensatory damages of a cancellation. For cancellations in the last 48 hours, the broker ends up paying 0.141 of the spot rate more to have the load moved than the price that the carrier who cancelled had agreed to. The broker does not always lose the carrier who cancels, however. Among owner-operators on loads of 300 miles or more, 15.2% of cancellations end with the carrier completing another of the broker's loads with the same pickup date.

Guided by the empirical evidence, I build a dynamic model of carrier search and bidding behavior under two types of penalties: a monetary penalty that acts as a transfer to the broker, or a reputational penalty that solely reduces the carrier's utility. The key mechanism in the model is that carriers price in the foregone option value induced by the penalty, raising their bids in response. This creates a direct trade-off for the firm: while higher penalties reduce potentially costly cancellations, they also increase the cost of successful matches and reduce their number.[5] How many cancellations a penalty prevents depends on the upper tail of carriers' outside offers. If offers better than a carrier's match become rare quickly as their value rises, a penalty prevents many cancellations. If they remain common, it prevents few, and a monetary penalty mostly collects money from the cancellations that still happen. Under monetary penalties, the firm can thus profit by raising the penalty level in order to extract rents from carriers' outside offers. This sets up a key tension: while monetary penalties are a more efficient punishment mechanism at any given penalty level, the accompanying rent-extraction motive can lead to a more socially inefficient outcome than reputational penalties, which nobody receives.

To take the model to the data, I extend the baseline model to allow stochastic carrier and shipment arrivals, multi-round multi-unit auctions, and imperfect attention to win notifications, which reconciles the relatively low cancellation rate with the much higher bid-reneging rate. I then estimate the process for outside offers, the attention parameters, and carriers' subjective valuation of the platform's reputational penalty, separately for two lanes, Seattle to San Francisco and Los Angeles to Phoenix. The penalty is identified by the drop in cancellations at the 48-hour mark, given how quickly better outside offers become rare, which the response of cancellations to the posted Accept-Now price reveals. Carriers act as if a cancellation in the last 48 hours costs them about a tenth of the spot rate, 9.8% on Seattle to San Francisco and 8.1% on Los Angeles to Phoenix. The two lanes differ in their outside offers, however. On Seattle to San Francisco, offers well above the value of a typical match remain common, so the penalty prevents a smaller share of cancellations. I separately estimate arrival and on-platform search dynamics using directly observed moments.

5. This incentive extends to firms that can extract part of the option value of reneging by raising posted prices.

I then simulate the market at every level of the penalty, and compare each cancellation with a twin market in which it does not happen. On Seattle to San Francisco, a penalty equal to the *bilateral damages* of a cancellation, the broker's loss on the load left and the compensatory damages of contract law, would be 0.151 of the spot rate. A penalty equal to its *market-wide cost*, the cost to the broker and every other carrier, would be only 0.012. With no penalty, about a third of the broker's loss on the load left is recovered on its other loads, mostly from carriers who cancel to move another of them, and more than half is gained by other carriers. The estimated penalty lies between the two benchmarks, well above the market-wide cost.

Even a penalty equal to the market-wide cost lowers welfare, because a penalty also taxes matching. A carrier who may face a penalty for leaving a match values the match less, so it bids higher and takes fewer loads at the Accept-Now price. Cancellations fall, but fewer matches are made, and fewer last to pickup. A completed match is worth more to the market than to the carrier who chooses it, since the broker's margin and the cost of a load left unmatched lie outside the carrier's choice. Welfare is therefore highest with no penalty on both lanes, whether the penalty is reputational or monetary. At the estimated reputational penalty, welfare changes by -7.1% relative to no penalty on Seattle to San Francisco and by -5.3% on Los Angeles to Phoenix.

The broker's own incentives depend on the type of penalty. With a reputational penalty, the broker collects nothing, and its profit is also highest with no penalty. On Seattle to San Francisco, a broker that could charge a monetary penalty would charge one, because at its current posted prices carriers bear most of the penalty rather than pricing it into their bids. The profit-maximizing monetary penalty there is well above the market-wide cost, and the broker would have no reason to confine it to the last 24 hours, the only ones in which cancellations cost the market, since it gains about as much from applying it earlier. On Los Angeles to Phoenix, where carriers pass more of the penalty into their bids, the profit-maximizing penalty is zero. A limit on when a monetary penalty applies is undone by a higher penalty, while a cap on its amount near the market-wide cost keeps welfare close to its peak. The broker would also not charge a monetary penalty if it put about a third of its weight on carriers' surplus, or if collecting each payment cost about two percent of the spot rate.[6]

Literature Review Classic work on unraveling in entry-level labor markets shows how decentralized timing can induce inefficiently early commitment (Roth and Xing 1994; 1997; Li

6. Platforms in the early stages of growth may value supplier retention more. For example, AirBnB imposed small fees of $50 to $100 for host cancellations until August 21, 2022, after which cancellation fees increased dramatically to scale with the reservation amount, up to $1000 (Airbnb 2022).

and Rosen 1998; Li and Suen 2000; Kagel and Roth 2000). In many such markets, salaries are rigid, so the timing of offers becomes a key strategic lever and the cost of commitment cannot be priced in, leading to market unraveling. Consistent with this, Niederle and Roth (2009) show experimentally that non-binding agreements can improve efficiency but are not sustained in competitive equilibrium. Consequently, the main proposed remedy is a centralized clearing mechanism. In practice, even well-designed centralized mechanisms are vulnerable to strategic offers made outside the mechanism, which can reduce efficiency (Kapor, Karnani, and Neilson 2024) or sometimes lead to complete collapse (McKinney, Niederle, and Roth 2005). This paper contributes to this literature by showing that a welfare-enhancing right to renege can emerge in a decentralized competitive market with endogenous prices. In that sense, this paper is close in spirit to De Groote et al. (2025), who show how a multi-round college admissions mechanism can re-allocate slots that students have reneged on in favor of off-platform options, though I shift the focus to how endogenous prices govern both the initial match and the decision to renege.

This paper also bridges two literatures studying dynamic allocation problems under different mechanisms. The dynamic market design literature has focused on how non-price mechanisms such as waitlists shape forward-looking incentives in the allocation of donor organs (Agarwal et al. 2018; Agarwal et al. 2021), housing (Waldinger 2021) and the reallocation of physicians (Huitfeldt, Marone, and Waldinger 2024). In contrast, the dynamic auction literature has focused on the forward-looking bidding strategies that emerge in sequential multi-unit auctions (Jofre-Bonet and Pesendorfer 2003; Zeithammer 2006; Backus and Lewis 2024) or in auction formats. The mechanism here combines these two elements. A carrier's bid determines its position in the platform's sequence of provisional offers, while the right to renege determines the value of keeping search alive after an offer arrives or even after a match is confirmed. Evaluating cancellation penalties therefore requires a model of both forward-looking acceptance decisions and forward-looking bidding. The mechanism is also connected to work on auction formats allowing ex-post bid modifications or withdrawals (Bajari, Houghton, and Tadelis 2014; Cramton, Ellermeyer, and Katzman 2015; Haberman and Jagadeesan 2023), with the key difference being that flexibility in this setting is not triggered in response to new information about the value or cost of the transaction, but rather in response to the arrival of new outside offers.

Finally, this paper contributes to the literature on the industrial organization of trucking. Classic work studies monitoring, contractual form, and firm boundaries in trucking (Hubbard 2000; 2001; Baker and Hubbard 2001; 2003), while Yang (2022) emphasizes spatial frictions generated by drivers' preferences to return home. The most closely related paper on the value

of flexibility in the market is Harris and Nguyen (2025). Their flexibility margin is tender rejection under long-term fixed-rate contracts; mine is cancellation after acceptance under spot-negotiated rates. The corresponding allocative questions are therefore different and complementary. Harris and Nguyen (2025) study the boundary between relationships and the spot market, where relationships exploit match-specific efficiencies but can thin the spot market. I instead study the efficient real-time allocation of particular carriers to particular loads within the spot market itself, taking into account the frictions induced by the decentralized timing of offers. A final difference is that Harris and Nguyen (2025) take the degree of flexibility allowed by long-term contracts as given and analyze its equilibrium consequences, whereas this paper studies the endogenous choice of flexibility and its consequences for welfare.

The rest of this paper is structured as follows. In Section 2, I describe the platform and the data, and measure what a cancellation costs the broker on the load left. In Section 3, I present a stylized model of the carrier's bid and cancellation decisions, show how carriers respond to the penalty, and define the two types of penalty. I extend the model to a dynamic setting in Section 4, estimate it in Section 5, and use it to compare penalties in Section 6. I conclude with Section 7.

2 Setting and data

Data for this study were directly provided by a trucking broker and spans 2021 and 2022. It is an early pioneer of the digital brokerage business model, seeking to automate the traditionally labor-intensive brokering process via an online auction platform. The majority of the broker's upstream demand came through long-term contracts with large shippers (such as major retailers and manufacturers), which would regularly tender shipments to the platform at the contracted rate. The broker's margin on a shipment is the difference between the price paid by the shipper and the price paid to the carrier and varies from shipment to shipment. This section describes the auction platform's design and cancellation policy, summarizes the data, and presents in-sample estimates of the impact of a carrier cancellation on a shipment's final matching outcomes and the broker's margin.

2.1 Auction Platform Design

Carriers can discover loads by either searching for specific criteria or browsing a personalized feed. Once a shipment appears on their feed or search results, the carrier can click on it to obtain further details (referred to as a detailed *view*), place a bid, or be immediately matched via an "Accept-Now" feature with a posted price. During the study years of 2021 and 2022, the firm implemented its "timed auction" format. This format consists of several rounds, with the reserve prices non-decreasing across rounds. The majority of shipments are posted to the platform between one and two weeks before their pickup time. Auction rounds begin to clear five days before pickup, with one to two rounds per day. If a shipment is unmatched 24 hours before pickup, the auction format changes to update the reserve price more frequently, with additional intervention from human brokers. When choosing whether to bid or Accept Now on a shipment, the carriers can only see the current round's closing time and the Accept Now price, in addition to the shipment characteristics. No information is provided about other bidders and their bid amounts, and the reserve price is hidden.[7] Additionally, carriers are never restricted from bidding on additional shipments, even if these conflict with existing confirmed appointments on the platform.

Each bid is also assigned an additive quality adjustment factor based on a set of performance metrics computed over past matched shipments, including the carriers' past cancellations. When an auction reaches one of its pre-determined closing times, if all bids exceed the reserve price, the auction proceeds to the next round. If there are bids below the reserve price, the platform accepts the lowest adjusted bid and notifies the bidder, who has 10 minutes to confirm. If the bidder does not respond, the process is repeated for the next lowest bid. If no bidders confirm, the auction proceeds to the next round. I use the terminology of the platform *accepting* bids and carriers *confirming* them throughout the paper. Based on these rules, it is clear that bidding entails practically no commitment on the side of the carrier.

After being matched to a shipment, whether through an accepted bid or the Accept-Now feature, the carrier may cancel the match at any time, but with potential consequences for their future bid quality adjustment factor if the cancellation occurs in the last 48 hours before pickup.[8] Carriers' cancellation behavior and consequences for the bid quality adjustment factor are discussed in more detail in section 3.

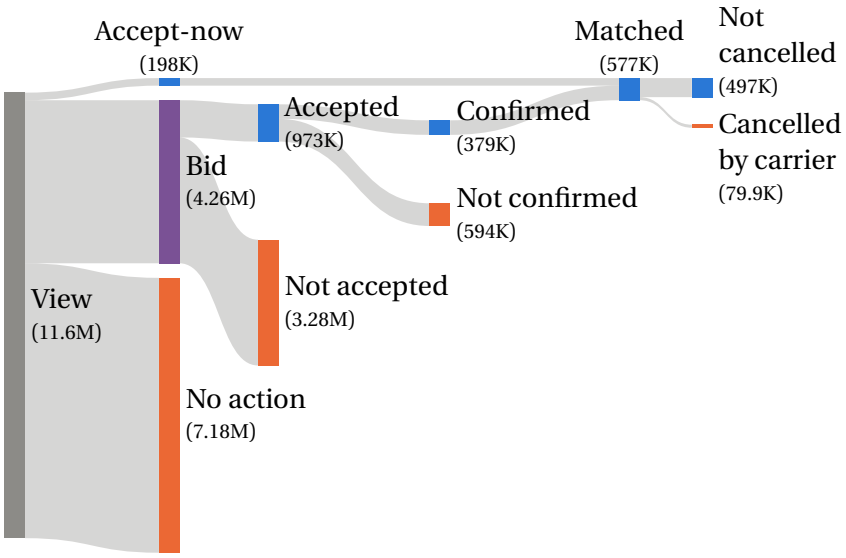
7. Reserve prices were initially publicly displayed at the launch of the auction format, but were hidden soon after launch, as carrier bids tended to bunch at the reserve price.

8. Until early 2021, carriers had to contact a firm representative to cancel. The policy was then changed to self-service cancellations with a disclaimer that "if the request occurs within 48 hours of pickup time, it may negatively affect [the carrier's] scorecard."

2.2 Data

The design of the platform yields a rich dataset on the endogenous decisions by carriers at the multiple stages of the matching process, including viewing loads, placing bids, bid acceptances and confirmations, and cancellations. While the full dataset covers the universe of auctions run by the platform, I exclude all local shipments (within a metropolitan area) from the analysis, due to their high-frequency nature allowing carriers to match with multiple shipments in a single day. The full matching funnel described above is visualized for the sample of 559,379 non-local shipments in Figure 1.

Figure 1: Funnel from views to matches



NOTE: Carriers' decisions, starting from a shipment view on the sample of all non-local shipments in 2021-2022 offered on the timed auctions platform. Accepted refers to bids that were accepted by the platform (winning bids). Confirmed refers to bids that were subsequently confirmed by the carrier.

The structural model is developed for a specific lane market and is thus estimated on two lane subsets: Seattle to San Francisco and Los Angeles to Phoenix. These subsets are compared to the full non-local sample in Table 1. Overall, both lanes are fairly representative of the broader non-local market. The Seattle to San Francisco differentiates itself by having a longer shipment distance and a higher propensity to bid or accept-now conditional on a view.

Freight prices move with the season and with aggregate shocks such as fuel prices. To compare loads across lanes and dates, I divide every price in the paper (bids, accept-now prices, reserves and final carrier prices) by the load's *spot rate*: an ex-post average spot rate for the lane from a third-party data aggregator.

Table 1: Summary statistics

	All non-local	Seattle to San Francisco	Los Angeles to Phoenix
Loads			
Loads	559,379	9,487	16,487
Mean distance (miles)	371	758	349
Mean spot rate ($)	1,387	1,440	1,587
Carriers			
Carriers	88,439	4,733	10,338
Platform miles per driver and year, median	906	3,487	2,214
90th percentile	5,004	16,730	12,009
Carrier decisions			
Views	11,629,188	92,590	364,174
Bids	4,255,680	54,756	159,134
Accept-nows	198,304	3,509	4,519
Views without a bid or accept-now (%)	61.7	37.1	55.1
Bids accepted (%)	22.9	25.4	23.5
Accepted bids confirmed (%)	38.9	38.2	34.9
Bid / spot rate, median	1.106	1.111	1.051
25th percentile	0.976	1.001	0.953
75th percentile	1.274	1.259	1.164
Matches and cancellations			
Matches	615,027	9,270	18,645
Cancelled by the carrier (%)	13.8	12.3	13.8
Cancelled by the shipper (%)	8.2	16.8	9.9

NOTE: Non-local loads are all loads with different pickup and dropoff metropolitan areas with a pickup date in 2021–22 that was offered in the timed auctions format. Carrier counts include all carriers with at least one decision on the column's loads. Platform miles per driver and year are computed over each carrier's drivers that have hauled at least one of the columns loads. Accepted bids confirmed: share of bids accepted by the platform that the carrier confirmed. Spot rate: ex-post average spot rate for the lane from a third-party data aggregator. Matches: confirmed bids and accept-nows, including eventual cancellations.

Shippers cancel matched loads too: 8.2% of matches end in a cancellation by the load's shipper (Table 1). These cancellations are beyond the scope of this paper, because the industry already has a well-established standard for them: when a shipper cancels close to pickup, it pays the carrier a fee known as "truck order not used" (TONU).

Carriers haul only a small part of their miles on the platform. Among the drivers who hauled at least one load on the platform, the median driver hauls 906 miles a year on the platform, and 90% haul fewer than 5,004 miles (Table 1). Given that owner-operators drive over 100,000

miles per year on average, most carriers must find shipments off the platform for a majority of their mileage (OOIDA 2022). Given this degree of multihoming, a key focus of the empirical strategy will be to recover a distribution of outside offers.

2.3 What a Cancellation Costs on the Load

Overall, 13.8% of matches end in a cancellation by the carrier. In the following, I examine what a carrier cancellation does to the rest of the load's life: whether the shipper cancels the shipment altogether, whether it is rescheduled, and the price at which it is eventually matched. Using the full non-local sample from Table 1, I compare loads whose match was cancelled with similar loads whose match was not. The comparisons condition on lead time, the number of hours left until a load's scheduled pickup, which I discretize into four-hour windows over the last four days before pickup, with each window indexed by w. For expositional purposes I further aggregate up to three longer periods indexed by r: 0–24, 24–48 or 48–96 hours before pickup.

For each window w, I take the loads that were still active when the window opened, some of which are then cancelled by the carrier during the window and the remainder which serve as comparison loads. A single load can appear in multiple windows as long as it remains active and has not been cancelled by the shipper or a carrier. Aggregating over all windows w, I estimate

$$y_{iw} = \sum_{r=1}^{3} \beta_r \, D_{iw}^{r} + \alpha_{w,g(i)} + \lambda_{\text{lane}(i)} + \mu_{\text{week}(i)} + \varepsilon_{iw}, \tag{1}$$

where D_{iw}^{r} equals one if the carrier of load i cancels during window w and the window lies in period r. Three sets of fixed effects are included: $\alpha_{w,g(i)}$ controls for the current window and the initial lead time window of the shipment $g(i)$, λ and μ control for lane and calendar week, respectively. The outcomes y_{iw} are whether the shipper cancels the load, whether it is later rescheduled to another pickup date, and its final carrier price as a share of the spot rate. The coefficient β_r is the excess outcome of a cancellation in period r.

The second comparison uses nearly the same estimating equation (dropping the dependence of α_w on the intitial lead time $g(i)$) but modifies the comparison set of shipments. Each window w now includes all shipments that were currently unmatched at the start of the window, as well as all shipments that became unmatched during the window due to a carrier cancellation. The only outcome y_{iw} included in this second comparison is the price. Table 2 reports the results of all four comparison regressions.

Table 2: What happens to a load after a cancellation

	Active loads			Unmatched loads
Carrier cancellation, hours before pickup	Shipper cancellation (pp)	Reschedule (pp)	Final price (% of spot rate)	Final price (% of spot rate)
0–24 hours	+2.4 (0.3)	+4.6 (0.2)	+13.1 (0.3)	+1.9 (0.2)
24–48 hours	−0.3 (0.2)	+4.0 (0.3)	+5.0 (0.2)	+0.4 (0.2)
48–96 hours	−0.6 (0.2)	+3.2 (0.2)	+3.0 (0.2)	+0.7 (0.2)

NOTE: Estimates of β_r in (1), one regression per column; rows are the periods r. Comparison loads are restricted to loads viewed by at least one carrier. Shipper cancellation and reschedule in percentage points (comparison means 5.8% and 6.1%). Final price in percent of the spot rate, picked-up loads only. Standard errors clustered by lane in parentheses.

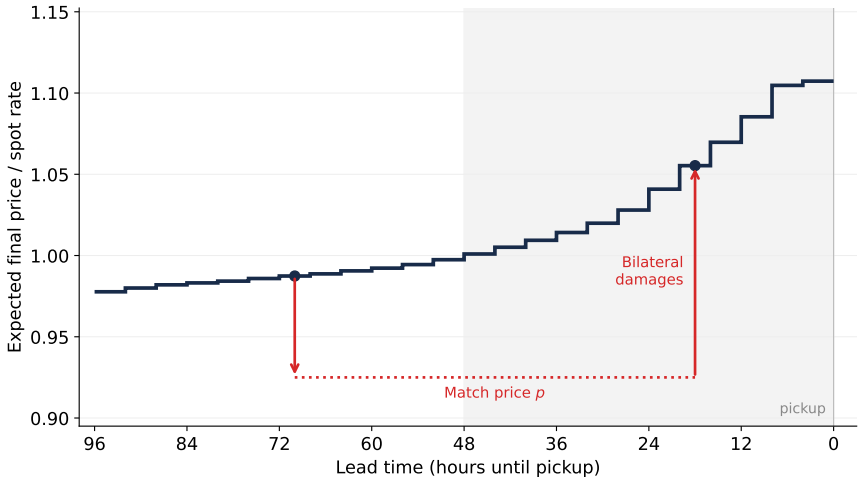
The key findings from the comparison exercise are that, relative to shipments with the same initial lead time, cancellations cause a modest increase in shipment cancellations when they occur close to the pickup date[9], a modest increase in reschedules at all lead times, and an increase in the final matching price that is most pronounced for last-minute cancellations. The second comparison with unmatched loads shows that price effect is largely driven by an effective loss of shipment lead time, which pushes up the expected final matching price.

Figure 2 illustrates the effect of lead time on price, by taking the estimated four-hour window effects α_w from the second comparison regression and plotting them against the lead time windows w, creating an expected price curve conditional on the remaining lead time of an unmatched load. The figure also illustrates the impact of a typical shipment match and cancellation on a shipment's expected final price: initially the match lowers the expected price to the matching price, and a subsequent cancellation raises it back up, now at the remaining lead time. This demonstrates why later cancellations are more costly and, in the legal doctrine of compensatory damages, would warrant higher penalties to be imposed in order to prevent breaches that hurt the platform more than they benefit the carrier. This view counts the platform's loss on the load left, and on that load only.

The rise from the match price back up to the curve is the platform's loss on the load left. I call this loss *bilateral damages*: the loss to the party the carrier contracted with, which is what con-

9. The low impact on shipper cancellation rates suggests minimal disintermediation, unlike findings in Xie and Zhu (2022), which documented substantial disintermediation on a Chinese intra-city cargo delivery platform. This difference likely stems from U.S. shippers' preference for long-term contracts with brokers or large carriers who can reliably source drivers, rather than engaging with small carriers for one-off shipments at a lower price.

Figure 2: The expected price of a load still unmatched, by time left before pickup



NOTE: Expected final carrier price, as a share of the spot rate, of a load still unmatched in each four-hour window before pickup: the window effects α_w of the price regression in the last column of Table 2, at the sample's average lane and week effects. Shaded: the last 48 hours. Right arrow: the bilateral damages of a cancellation 18 hours before pickup.

tract law's compensatory damages make good.[10] Shippers pay their contracted rate whatever the platform pays the carrier, so every dollar of the rise comes out of the platform's margin. For cancellations in the last 48 hours, bilateral damages in the data average 0.141 of the spot rate: the platform finally pays that much more to have the load moved than the price the carrier who cancelled had agreed to.

The load is rarely lost. The premium goes to whoever moves the load instead; whether the market loses it is a question for the model (Section 6). Nor does the platform always lose the carrier who cancels. Among owner-operators on loads of 300 miles or more, 15.2% of cancellations end with the carrier completing another of the platform's loads with the same pickup date. When that happens, the platform's loss on the load left overstates its loss over all its loads. The data cannot show what the carrier who moves the load instead gave up to do so, nor whether the cancelling carrier's other load would have found another carrier anyway. Section 6 uses the model to answer both.

Beyond the load left, the penalty could also matter to the platform through which carriers

10. In contract law, these are expectation damages measured by the cost of cover: the price of a substitute performance less the contract price (ALI 1981, § 347). U.C.C. § 2-712 gives the same measure to a buyer of goods; it applies to a haul only by analogy. Section 6 also subtracts what the platform recovers on its other loads, giving platform-wide damages. That is the platform's own view of being made whole, not a legal measure: a court would likely not reduce the damages by the platform's gain on another load.

it matches with. A worse quality adjustment after a cancellation in the last 48 hours steers future loads away from carriers who cancel. That cuts cancellations, a gain for the platform. But carriers who cancel often also match at lower prices, so steering loads away from them raises what the platform pays, a loss. Appendix A.1 finds that the two roughly offset. Among owner-operators with at least 10 matches, the carriers who cancel most hold a tenth of the matches. Excluding them lowers the cancellation rate from 14.7% to 11.5%. Yet the platform's margin, charging each cancellation the cost of not matching the load, rises only from 8.0% to 8.1% of what the shipper pays. I therefore treat the penalty as neither a gain nor a loss to the platform in itself: it matters to the platform only through how carriers respond to it.

3 Cancellation Penalty and Carrier Response

Recall that carriers are subject to a penalty for cancelling a load within 48 hours of the scheduled pickup time, and this penalty takes the form of additional adjustments to their future bids, making them less likely to win subsequent loads or requiring lower bids to secure them. To develop intuition, this section first introduces a stylized repeated-game model of the quality adjustment. The model captures both how the quality adjustment induces an effective penalty that discourages cancellations and how carriers react to the adjustment and the penalty in their optimal bidding behavior. The section then introduces some empirical evidence that the penalty affects carriers' cancellation behavior but that carriers do not react to variation in their penalty state, which will justify the use of a structural model that treats the reputational penalty on cancellations as a fixed perceived utility shock rather than a fully endogenous outcome of the repeated game.

3.1 Model of Bid Adjustment and Reputational Penalties

A carrier with discount factor $\delta \in (0,1)$ bids b for a shipment costing c, with win probability $\gamma(b+q)$ given its current dollar quality adjustment q and $\gamma'(b+q) < 0$. After the auction, whether matched or unmatched, it observes a net outside payoff $\pi \sim G$. A matched carrier chooses between completion for $b-c$ and cancellation for π. An unmatched carrier's only choice is between the outside offer and a normalized payoff of zero.

The current quality adjustment q is state-dependent: given a cancellation, it will increase. Let $\mathcal{V}(q)$ denote the carrier's continuation value when it enters the next auction with quality ad-

justment q. Let q^C denote the adjustment after completion or remaining unmatched, and let $q^X \geq q^C$ denote the adjustment after cancellation, so that $\mathcal{V}(q^X) \leq \mathcal{V}(q^C)$. A successful completion may also lower the adjustment, so that $q^C < q$. Conditional on the current state, these values do not depend on the current bid or outside offer. We can then define:

$$\kappa = \delta \left[\mathcal{V}(q^C) - \mathcal{V}(q^X) \right] \tag{2}$$

as the carrier's effective (state-dependent) cancellation penalty. The matched carrier's choice when faced with the outside offer can then be written as

$$\max\{b - c + \delta\mathcal{V}(q^C),\ \pi + \delta\mathcal{V}(q^X)\} = \delta\mathcal{V}(q^C) + \max\{b - c,\ \pi - \kappa\}.$$

Thus the carrier will cancel if and only if the outside offer exceeds the effective reservation value $r^* = b - c + \kappa$. The cancellation probability is then simply $1 - G(b - c + \kappa)$, so that κ directly affects the likelihood of cancellation. An unmatched carrier's payoff is simply $\delta\mathcal{V}(q^C) + \max\{0, \pi\}$.

The carrier picks the bid b^* that maximizes its expected payoff. Appendix A.6 gives the derivation. Let $r^* = b^* - c + \kappa$ be the outside-offer cutoff at the optimal bid. At an interior optimum with $G(r^*) > 0$, the optimal bid satisfies

$$b^*(c) = c + \frac{\kappa[1 - G(r^*)] + \int_0^{r^*} \pi\, dG(\pi)}{G(r^*)} - \frac{\gamma(b^* + q)}{\gamma'(b^* + q)} \tag{3}$$

The bid consists of cost, compensation for outside opportunities forgone and cancellation costs incurred after an award, and a positive markup. Dividing the opportunity-cost terms by $G(r^*)$ expresses them per completed shipment, since payment is received only upon completion. Equation (3) is implicit: the cutoff and markup both depend on the optimal bid.

3.2 Empirical Patterns of Carrier Cancellations

When cancelling a load, carriers provide a reason among a predefined list of options. Under the broker's policy, the stated cancellation reason does not affect the penalty applied to the carrier, reducing misreporting incentives. The breakdown of these reasons is shown in Table 3. These reasons can broadly be classified as discretionary or forced, representing the amount of agency a carrier had in the decision to cancel. "Accepted by accident" is counted as discretionary due to its high prevalence and vague description. It also shares many of the

same patterns with other discretionary cancellations in the analysis that follows. Overall, each type of cancellation represents about half of the total.

Table 3: Stated reasons for carrier cancellations

Reason	Share of cancellations (%)
Discretionary	
Does not want the load	21.4
Accepted by accident	19.5
Overbooked, no truck available	8.0
Found a better load	0.2
Total discretionary	49.1
Forced	
Truck problem	21.2
Delay or problem with a previous load	10.3
Driver illness or personal matter	7.2
Compliance, quality or insurance problem	3.2
Out of legal driving hours	3.0
Other	6.0
Total forced	50.9
All carrier-fault cancellations	100.0

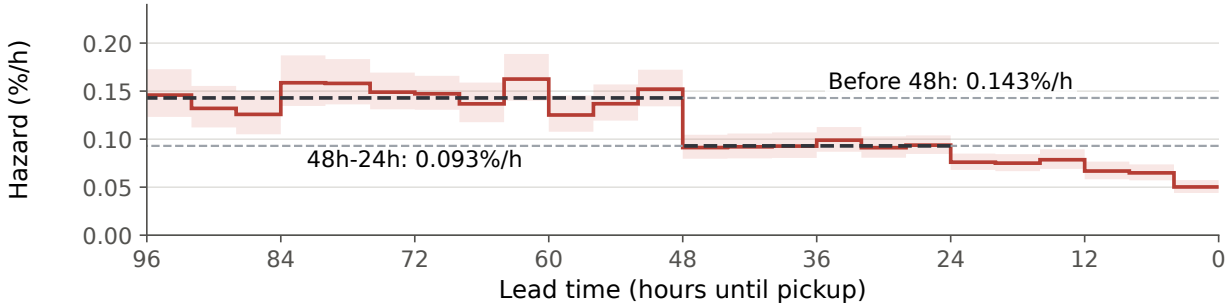
NOTE: Share of the 13.8 of carrier-fault cancellations of matches (confirmed bids and accept-nows) on the sample of 2021-2022 non-local timed auction loads.

Recall the official platform policy from subsection 2.1: a cancellation inside 48 hours of the scheduled pickup "may negatively affect" the carrier's scorecard. To examine the impact of this policy on carrier behavior, I analyze the hazard rate of discretionary carrier cancellations in 4-hour bins of time before the pickup time, including fixed effects for time of day, first hour of match, and separate coefficients for bid amounts and Accept-Now prices. Appendix A.3 provides further details. The resulting hazard patterns are shown in Figure 3. While the cancellation hazard remains mostly flat initially, save for some estimation noise, it drops sharply by a third right at the 48-hour mark, clearly reflecting the bite of the cancellation policy. The same hazard analysis for forced cancellations, depicted in Figure A.3, shows a steady increase towards the pickup time, with no visible discontinuity at the 48-hour mark, suggesting these cancellations are unavoidable and thus unaffected by the policy.

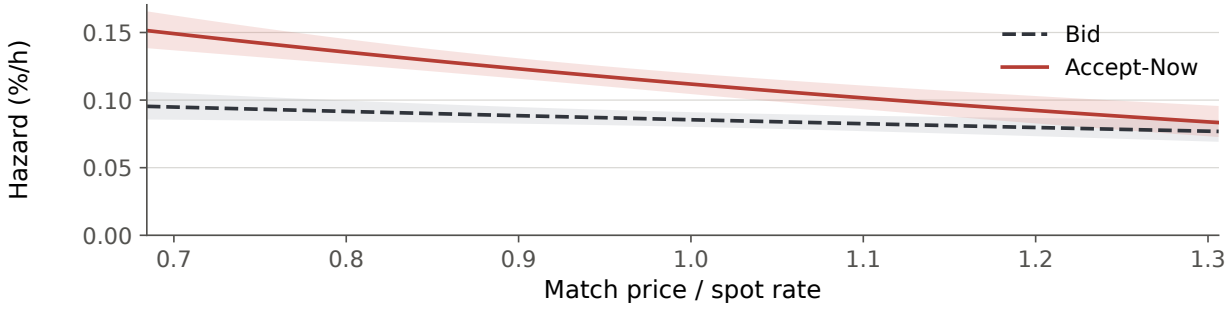
Panel (b) of Figure 3 shows effect of the match price on the discretionary cancellation hazard. Intuitively, a higher match price relative to the spot market rate makes the current match more attractive relative to other opportunities a carrier might encounter in the market. This effect is much more apparent on the Accept-Now price than on bid amounts. This can be explained by

Figure 3: Discretionary cancellations and bid confirmations

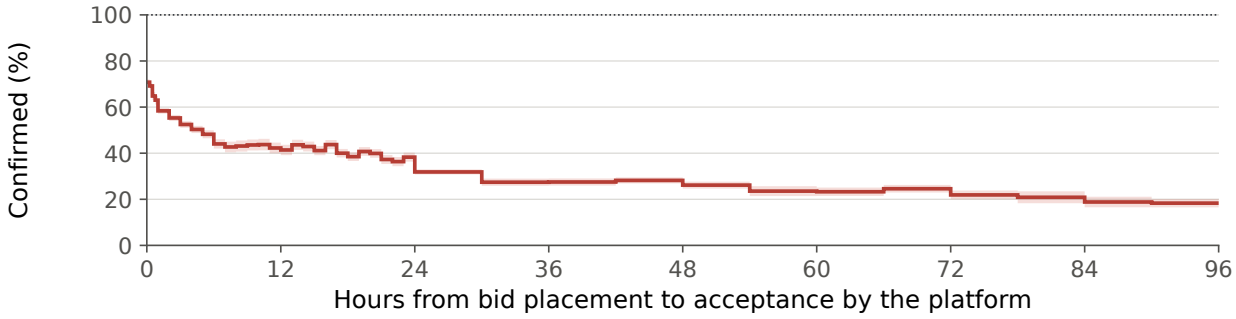
(a) Discretionary cancellation hazard, by time before pickup



(b) Discretionary cancellation hazard and the match price, by channel



(c) Confirmation of accepted bids, by the wait from the bid to the acceptance



NOTE: Sample of owner-operators on loads of 300 miles or more with a pickup date in 2021–22. Panel (a): hazard of a discretionary cancellation, in percent per hour of an active match, by four-hour bin of time before pickup, with 95% confidence intervals clustered by carrier (shaded). Dashed: the hazard more than 48 hours before pickup and from 48 to 24 hours before pickup, from the models with one constant on each side of the deadline and with a third constant for the last 24 hours. Panel (b): the price gradients from the hazard model of panel (a) for matches at the posted accept-now price (solid) and matches won with a bid (dashed). Histogram of matching prices provided in Figure A.4. Panel (c): share of bid acceptances the carrier confirms, by time from the bid to the acceptance, adjusted for time of day and lead time until pickup of the acceptance. Filtered to the first acceptance of a carrier for a given pickup date. 95% confidence intervals clustered by carrier (shaded but small).

the carrier's optimal bidding behavior in Equation 3, where the bid is a direct function of the cost. Thus, depending on the properties of the auction win probability, the markup over a carrier's cost may vary very little with the actual bid amount, producing a nearly flat relationship between the bid and the cancellation hazard.

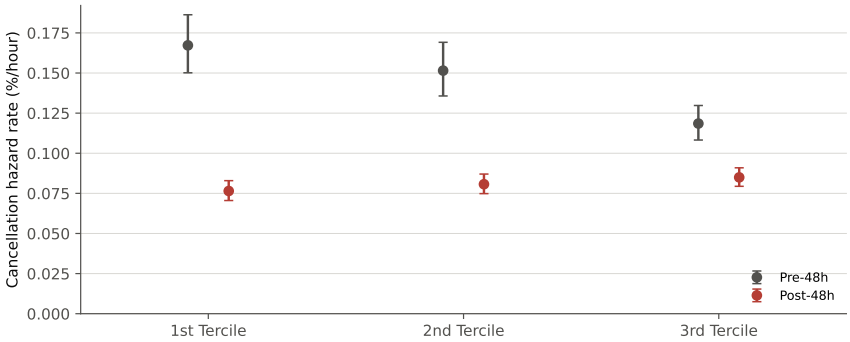
Finally, Panel (c) of Figure 3 shows the confirmation behavior of carriers after accepting a bid conditional on the time interval from when the bid was placed to when it was accepted by the platform. The highly granular time bins show that the confirmation probability starts at less than 100% immediately after the bid is placed and drops rapidly in the first few hours, before stabilizing at a steadier rate of decline over time. This suggests that while there may be a similar rate of outside opportunities that take carriers off the market over time, the rapid initial drop also reflects some carrier inattention, given the limited 10-minute exclusivity window a carrier has before a load is offered to the next best bidder.

3.3 Heterogeneity in the Cancellation Penalty

As the model in subsection 3.1 suggests, the rational, repeated-game reputational penalty could be highly state-dependent. To investigate the empirical relevance of this heterogeneity, I use the empirical relationship between carrier cancellations (and on-time pickup and delivery) to derive the carrier's effective additional quality adjustment on future bids for each of their matches, $q^X - q^C$ in the model of subsection 3.1. The distribution of these incremental penalties is shown in Figure A.2 with a mean of 0.62% of the spot rate for a single cancellation. I then split carrier matches into terciles of the effective penalty and run the cancellation hazard model of Figure 3 but with constant hazard rates before and after the 48-hour mark, separately for each effective penalty tercile. If carriers were responsive to the heterogeneity in the cancellation hazards, we would expect a larger drop in the cancellation hazard for the higher penalty terciles from a common baseline prior to the 48-hour deadline.

Instead, the figure shows a *smaller* drop in the cancellation hazard for the higher penalty terciles, almost entirely driven by a lower baseline cancellation hazard *prior* to the 48-hour deadline. This pattern is not simple selection on carriers' baseline cancellation tendencies, as that would suggest the opposite pattern: higher-penalty terciles would have higher pre-48-hour hazards if they were simply the more cancellation-prone carriers. Instead, this may reflect carriers' tenure on the platform, as the quality adjustment uses a rolling window of recent performance with a Bayesian prior, which naturally makes the effective penalty higher for newer carriers. As an explicit cancellation option is not a norm in the industry, carriers may take time

Figure 4: The response to the deadline by realized-penalty tercile



NOTE: For each tercile of the realized penalty of one cancellation, the hazard two to four days before pickup (Pre-48h, left) and in the last two days (Post-48h, right), using the same specification as in Figure 3a, with 95% confidence intervals clustered by carrier. Mean penalty by tercile: 0.18, 0.34 and 1.34 percent of the spot rate. Full distribution in Figure A.2.

to get accustomed to the ability to cancel loads without facing immediate repercussions. As learning effects are out of the scope of this paper, I model the reputational penalty as having a uniform effect across carriers in the model.

3.4 Reputational and Monetary Penalties

The platform's penalty for a late cancellation is reputational. A cancellation in the last 48 hours worsens the quality adjustment on the carrier's future bids, which lowers its chances of winning future auctions (Section 2.1). The carrier who cancels bears this cost, and nobody receives it. The platform collects no money, and it gains little from steering future loads away from carriers who cancel, since those carriers also match at lower prices (Appendix A.1). Carriers also see neither their quality adjustment nor how it is computed (Appendix A.2).

The alternative is a monetary penalty, a cancellation fee paid to the platform, such as airlines and hotels charge their customers. A carrier perceives a monetary penalty at face value, so a monetary penalty of κ changes its decisions just as a reputational penalty perceived as κ does. The two penalties thus differ only in who receives the payment. A monetary penalty is a transfer from the carrier to the platform, while a reputational penalty is a loss to the carrier that nobody receives. This difference matters for welfare and for the platform's incentive to raise the penalty, which Section 6 studies. Section 6.5 discusses why the platform may not charge monetary penalties.

In conclusion, this section shows that the platform's reputational penalty mechanism has a clear impact on carrier cancellation decisions but that this impact is limited to discretionary cancellations. It also showed that carriers respond somewhat to the shipment price in their cancellation behavior, and that their bids combine both immediate and gradual forms of attrition. Both these patterns are consistent with the simple model of Section 3.1 where carriers respond to the reputational penalty by integrating it into their outside offer reserve price $r^* = b - c + \kappa$. Section 3.3 found that carriers whose cancellation would raise their adjustment more do not respond more to the 48-hour deadline (Figure 4). I therefore treat κ as a single perceived penalty that rationalizes the sudden drop in the cancellation hazard rate. In the next section, I develop the full structural model which rationalizes these observed patterns in carrier behavior.

4 Empirical Model

This section builds the full structural model of carrier decisions and the equilibrium model of the platform's auction dynamics. The model is specified for an individual *market* which is a combination of a lane and a pickup date. Time is continuous, and the model covers the last eight days before pickup. I index days by d, the number of days left before pickup: in the data, $d = 0$ is the last 24 hours, $d = 1$ the 24 hours before that, and so on, so the communicated 48 hour penalty period includes days 0 and 1. The market ends when $d = -1$. In the model, a day does not last exactly 24 hours. It ends at a random time, at rate η, which I normalize to one, so it lasts one day on average. Every other event also arrives at random, at a rate per day that can depend on d. The model is thus a continuous-time Markov chain, simplifying transition dynamics by eliminating multiple different events occurring simultaneously (Doraszelski and Judd 2012). Since the time horizon is very short, I assume both carriers and the platform do not discount their future payoffs.

4.1 Carrier Behavior

Carriers view shipment opportunities on the platform at an exogenous rate. Each time they view a shipment, they draw costs $c \sim N(\mu_d^c, \sigma_d^c)$ from a day-dependent distribution. Given this cost, they can ignore the shipment, match with it immediately at the Accept-Now price p^{AN}, or place a bid b. A pending bid b wins at Poisson rate $\gamma_d(b)$ on day d.[11] Up until the market's pickup

11. The win rate also depends on the day the bid was placed, which I leave implicit until Appendix B.3.

time, a carrier receives outside offers at a constant rate λ. Each outside offer is described by its net payoff $\pi \sim G$. If a carrier is matched to a load, taking an outside offer may subject them to a time-dependent penalty κ_d, the perceived reputational penalty of Section 3.4. Based on the evidence from subsection 3.2, we will assume $\kappa_d = 0, \forall d \geq 2$ and $\kappa_d = \kappa > 0, d = 0, 1$. Carriers may also be subject to *forced* cancellations due to exogenous events like truck breakdowns. These immediately end a carrier's participation in the market and leave them with a normalized payoff of 0 and an eventual penalty κ_d if they were matched.[12] When a bid wins, the carrier has only ten minutes to confirm it. Many carriers are owner-operators who drive their own trucks and cannot always watch their bids. A carrier therefore notices a win in time to confirm it with probability α_d, which I call attention, where d is the day of the win. Figure 3c motivates this feature, as even bids accepted within minutes of being placed often go unconfirmed, which outside offers alone cannot explain.

Carriers can bid on many loads at once, even on loads that conflict with one they already hold (see Section 2.1). Following every load a carrier holds would make its problem very large, so I model each load on its own and restrict the estimation sample to the first decision or commitment a carrier makes to reduce the cross-influence of different shipment opportunities within the platform. On a given load, the carrier is in one of three positions: it is *free*, with no bid or match on the load; it has a *pending bid*; or it is *matched*. A matched carrier is paid the agreed price b at pickup and bears its cost c of hauling the load, so it earns $b - c$. Other platform loads will still be anticipated via the outside offer process λ and G, much as bidders in Backus and Lewis (2024) form expectations about future auctions. In the simulated market of Section 4.4, a carrier can hold bids on several loads of the same lane and pickup date at once, and leaving a match for another of these loads works like leaving it for an outside offer.

4.2 Values and the Cancellation Rule

A carrier's value on a load depends on its position and on the day. I write U_d for the value of being free, $W_d(b, c)$ for the value of a pending bid b, and $V_d(b, c)$ for the value of a match at price b, at cost c. Each value is then an average over what can happen next: an outside offer, a forced cancellation, a win, or the end of the day, each weighted by its rate. Leaving out the

12. These could also be modelled as a very attractive outside offer, such that no reasonable penalty could prevent a carrier from taking it. Since the counterfactual simulation will difference out baseline carrier welfare, the carrier's exit value upon a forced shock is immaterial to the results.

arguments b and c, the values solve

$$(\eta + \lambda + \lambda_d^F)\, U_d = \lambda\, \mathbb{E}_{\pi\sim G} \max\{U_d,\, \pi\} + \eta\, U_{d-1}, \tag{4}$$

$$(\eta + \lambda + \lambda_d^F)\, V_d = \lambda\, \mathbb{E}_{\pi\sim G} \max\{V_d,\, \pi - \kappa_d\} - \lambda_d^F\, \kappa_d + \eta\, V_{d-1}, \tag{5}$$

$$(\eta + \lambda + \lambda_d^F + \gamma_d)\, W_d = \lambda\, \mathbb{E}_{\pi\sim G} \max\{W_d,\, \pi\} + \gamma_d \big[\alpha_d V_d + (1-\alpha_d)\, U_d\big] + \eta\, W_{d-1}, \tag{6}$$

with terminal values at pickup $U_{-1} = W_{-1} = 0$ and $V_{-1} = b - c$. From here on, $\eta = 1$. Appendix B derives the three equations.

In equation (4), a free carrier takes an offer if it beats the value of staying free. A forced shock leaves it with nothing, and when the day ends it moves on to the next one. Equation (5) adds the penalty. A matched carrier who takes an offer bears κ_d, and so does a matched carrier hit by a forced shock. In equation (6), a pending bid also wins at rate $\gamma_d(b)$. The carrier then confirms with probability α_d and becomes matched; otherwise it is free again. Walking away from a bid costs nothing, so a carrier with a pending bid takes any offer that beats $W_d(b, c)$. A forced shock leaves the carrier with nothing from the bid.

A matched carrier cancels for an outside offer when the offer, net of the penalty, beats the value of keeping the match, that is, when π exceeds the cutoff $r_d^*(b, c) = V_d(b, c) + \kappa_d$. This is the cutoff $r^* = b - c + \kappa$ of Section 3.1, with $b - c$ replaced by the value of keeping the match. That value counts the option to cancel for a better offer later and the risk of a forced cancellation, and it equals $b - c$ when neither can happen before pickup. Let $\bar{G} = 1 - G$, so that $\bar{G}(r)$ is the chance that an offer beats r. Offers above the cutoff then arrive at rate $\lambda\bar{G}\big(r_d^*(b, c)\big)$. On a day without a forced cancellation, a matched carrier cancels if such an offer arrives before the day ends. This happens with probability

$$\Pr(\text{cancel on day } d \mid b, c) = \frac{\lambda\bar{G}\big(r_d^*(b, c)\big)}{1 + \lambda\bar{G}\big(r_d^*(b, c)\big)}. \tag{7}$$

When the penalty starts at 48 hours, the cutoff jumps up and the daily chance of a cancellation falls, which is the model's version of the drop in Figure 3a. The same probability applies to a match taken at the platform's posted Accept-Now price p^{AN}, with p^{AN} in place of b. A higher match price raises the value of keeping the match, and so the cutoff. In Figure 3b, cancellations fall with the Accept-Now price much more than with the bid. The platform sets the Accept-Now price, so how cancellations move with it is informative about G (Section 5). A bid, in contrast, moves with the carrier's cost, as equation (3) shows.

4.3 Confirmation, Bidding and the First Choice

A won bid is confirmed if two things happen. First, the carrier must still be there when the bid wins, that is, it must not have left for an offer above $W_d(b,c)$ or been hit by a forced shock since placing the bid. A carrier who has left does not confirm, and the platform moves on to the next bid (Section 2.1). Second, the carrier must notice the win, which it does with probability α_d on the day of the win. For a bid placed on day d^b and accepted by the platform on day d^a, the confirmation probability is

$$\Pr(\text{confirm} \mid b,c) = \alpha_{d^a} \prod_{k=d^a}^{d^b} \frac{1+\gamma_k(b)}{1+\gamma_k(b)+\lambda\bar{G}\big(W_k(b,c)\big)+\lambda_k^F}. \tag{8}$$

The product falls as the wait between the bid and the win grows, as confirmation does in Figure 3c. Attention keeps confirmation below one even when the bid wins soon after it is placed. Because the model tracks the day but not the hour, attention also absorbs the fast fall in confirmation over the first hours after a bid.

A carrier bids once, when it first views the load, and does not revise its bid.[13] A carrier with cost c viewing the load on day d picks the bid $b_d^*(c)$ that maximizes $W_d(b,c)$. Appendix B.3 derives the bid's first-order condition, equation (B.3). It also writes the optimal bid, in equation (B.4), as a weighted average, over the days on which the bid can win, of bids of the form of equation (3). Each of these covers the cost, the outside offers the carrier gives up and the penalties it expects to bear, per completed load, plus a markup. The expected penalties also include those of forced cancellations inside the last 48 hours. For a given bid, the first-order condition gives the cost at which that bid is optimal, and Section 5 uses this to recover each bidder's cost from its bid.

When a carrier views a load on day d, it chooses among three options. Ignoring the load leaves it free, which is worth U_d. Accepting the load now makes it matched at the Accept-Now price p^{AN}, worth $V_d(p^{AN},c)$. Bidding is worth $W_d\big(b_d^*(c),c\big)$. To fit the observed shares of each choice, I subtract hassle costs c_{AN} and c_{bid} from the values of accepting now and of bidding, and add independent logit shocks with scale σ^{choice}.[14] The probability of option k, for ignoring (I),

13. Carriers can change their bids at any time, but they rarely do.

14. I leave the hassle costs and the logit shocks out of welfare. Counting them would change carriers' payoffs mechanically with every view.

accepting now (A) or bidding (B), is then

$$\Pr(k \mid c) = \frac{\exp\left(u^k/\sigma^{choice}\right)}{\sum_{\ell\in\{I,A,B\}} \exp\left(u^\ell/\sigma^{choice}\right)}, \tag{9}$$
$$u^I = U_d, \qquad u^A = V_d(p^{AN}, c) - c_{AN}, \qquad u^B = W_d\big(b_d^*(c), c\big) - c_{bid}.$$

Section 5 estimates the model from four of its pieces: the cost behind each bid, the confirmation probability (8), the daily cancellation probability (7) for matches won with a bid or taken at the Accept-Now price, and the choice probabilities (9).

4.4 The Simulated Market

For the counterfactuals, I place these carriers in a simulated market for one lane and pickup date. Over the last eight days before pickup, carriers and loads arrive and carriers view loads, at daily rates taken from the data; some loads are already posted when the market opens. At a view, the carrier draws its cost and ignores the load, bids or accepts now, as in Section 4.3. A load's reserve and Accept-Now prices are multiples of the platform's price forecast for the load, with one multiple per day taken from the data (Appendix D).

Each load's auction clears at random times over the last five days. In a clearing, the platform goes up through the bids at or below the reserve price, lowest first. A bidder confirms if it notices the win, with probability α_d, and does not hold a match it values more than the new one net of the penalty for leaving it. The first bidder who confirms wins the load. The bids passed over are lost, and the others stay pending.

A carrier can hold several bids and one match. A matched carrier can leave its match for another load on the platform, by accepting it now or by winning its auction, and it then bears the penalty. A carrier leaves the market when an outside offer beats the value of its best position, plus the penalty if it is matched, or when a forced shock hits, at the penalty if it holds a match. Loads also fall off for reasons outside the carriers' control.

At pickup, each match pays the carrier its price minus its cost, and pays the platform what the shipper pays minus that price. A load still unmatched at pickup costs the platform a share L of what the shipper pays. L is the cost of not matching the shipment, in which case the platform has to ask the shipper to reschedule or cancel it. Since most shipments come through relational contracts with shippers, this also has consequences for the platform's relationship with the shipper. Welfare adds the platform's profit to carriers' payoffs, measured against what

carriers would earn without the platform. A reputational penalty is a loss to the carrier who cancels that nobody receives, while a monetary penalty of the same size is a transfer to the platform (Section 3.4). I let it affect the platform only through carriers' choices, since steering future loads away from carriers who cancel adds little to the platform's margin (Appendix A.1).

In a counterfactual, the win rates change with carriers' behavior, so I solve for them as an equilibrium: given the win rates, carriers bid and confirm as above, and the simulated auctions produce new win rates. I repeat until the two agree (Appendix E).

5 Estimation

I estimate the carrier model of Section 4 separately for the two lanes, Seattle to San Francisco and Los Angeles to Phoenix. Section 5.1 explains which of carriers' decisions the estimation uses, and why. Section 5.2 describes the likelihood and what in the data identifies the penalty and the outside offers. Section 5.3 presents the estimates.

5.1 Which Decisions the Estimation Uses

Since the model treats each load on its own (Section 4), the estimation uses only each carrier's first decisions for each pickup date: the loads it views and ignores up to its first bid or accept-now, together with that bid or accept-now; whether it confirms its first accepted bid; and its daily choices to keep or cancel its first match. A bid or a match on another lane also counts in deciding which decision comes first, since it can take the same truck.

These are the decisions a carrier makes before it is matched to another load for the date, which is the situation the model describes. Later decisions depend on the carrier's other matches and bids in ways the model would misread: a carrier that turns down a win because it is already matched would look inattentive, a backup bid would invert to the wrong cost, and ignoring a load while another bid is pending would look like a high cost. Whether a decision enters the sample depends only on what happened before it, and each load brings a fresh cost while outside offers arrive at the same rate whatever the carrier's other matches and bids, so the decisions kept are as informative about costs and offers as those left out.

The sample covers each lane's loads with pickup dates in 2021–22, all carriers active on them, the last eight days before each pickup, and bids and Accept-Now prices in the usual range rel-

ative to the spot rate. Appendix C.1 gives each rule, the counts and the limits of the argument above.

5.2 Likelihood and Identification

I take two inputs directly from the data, before estimating the rest of the model. The first is the daily rate of forced cancellations, λ_d^F: the number of forced cancellations per day matched without a cancellation, from the first matches of the sample. The second is the chance that a pending bid wins on a given day, as a function of the bid, which gives the win rates $\gamma_d(b)$. I estimate it from all bids on the lane's loads, separately by the day the bid was placed and the day it wins (Appendix C.3). Unlike the likelihood, this uses every bid, since every bid competes in the auction.

The other 21 parameters are estimated jointly by maximum likelihood. They are the penalty κ, five attention probabilities $\alpha_0, \ldots, \alpha_4$, the mean μ^π and standard deviation σ^π of the outside offers, the logit scale σ^{choice} and the hassle costs c_{bid} and c_{AN}. The cost distribution adds a mean and a standard deviation for loads viewed on each of the last four days before pickup and for loads viewed earlier. The net payoffs of outside offers are normal, $G = N(\mu^\pi, \sigma^\pi)$, the same on every day. I normalize the rate at which offers arrive to $\lambda = 1$, one offer a day on average. The rate can be normalized because only offers good enough to take affect behavior, and these arrive at rate $\lambda\bar{G}$, so more frequent offers act much like a fatter tail of G.

The likelihood follows each decision in the sample, beginning with a bid. At given parameters, the bid's first-order condition, equation (B.3), gives the cost at which the bid is optimal, and I invert each bid to that cost. Every part of the likelihood that involves the bid then uses this cost. For a first bid, the cost distribution of the day of the view gives the density of the bid, adjusted for how fast the inverted cost changes with the bid, and equation (9) gives the probability of choosing to bid rather than ignore the load or accept it now. If the bid is the carrier's first acceptance, equation (8) gives the probability that the carrier confirms it. If it becomes the first match, equation (7) gives the daily probability of cancelling it. Since the inverted cost depends on the parameters, the inversion sits inside the likelihood and is redone at every trial value of the parameters.

For a load that the carrier accepts now or ignores, the cost is not revealed. The probability of the choice, together with the daily probabilities of cancelling that follow an accept-now, is therefore integrated over the cost distribution of the day of the view. A forced cancellation ends a match for reasons outside the carrier's choice, so a match enters the likelihood only up

to the day before a forced cancellation. Appendix C.2 states the likelihood in full. Standard errors come from a sandwich estimator, clustered by carrier and pickup date and corrected for the sampling error of the two inputs from the data (Appendix C.5).

Identification. Outside offers are not observed, so the distribution G must be read from carriers' choices. The Accept-Now price is informative about the slope of the tail of G: how fast the chance of a better offer, $\bar{G}$, falls as the value of the carrier's match rises. The platform posts the Accept-Now price, so unlike a bid it does not move with the carrier's cost. A higher Accept-Now price raises the value of a load taken at that price, and with it the cutoff an offer must beat for the carrier to cancel. How carriers' choices to accept now, and their cancellations after accepting now, change with the posted price therefore show how quickly better offers become rare.

Given G, the size of the drop in cancellations at the 48-hour deadline gives the penalty. Take a match whose value of keeping, V, is about the same on both sides of the deadline. Before the deadline, the carrier cancels it for any offer above V. Inside the last 48 hours, it cancels only for an offer above $V+\kappa$. When the daily chance of cancelling is small, equation (7) gives, approximately,

$$\frac{\Pr(\text{cancel on a day in the last 48 hours})}{\Pr(\text{cancel on a day before them})} \approx \frac{\bar{G}(V+\kappa)}{\bar{G}(V)}. \tag{10}$$

The data give the ratio on the left. Given the slope of the tail of G, the drop shows how far the penalty must raise the cutoff, which is κ. The faster the chance of a better offer falls, the smaller the penalty that produces a given drop.

Attention comes mainly from bids that win on the day they are placed. Such a bid has little time to lose its carrier to an outside offer or a forced cancellation, so the share of these wins that carriers confirm is close to attention.

5.3 Estimates: One Penalty, Two Kinds of Outside Offers

Table 4 reports the main estimates, and Figure 5 shows the outside offers of each lane. Table C.2 reports all 21 parameters.

The perceived penalty is close to a tenth of the spot rate in both lanes: 9.8% on Seattle to San Francisco and 8.1% on Los Angeles to Phoenix. It is the one-time cost that carriers act as if

Table 4: Estimates of the carrier model

	Seattle–San Francisco	Los Angeles–Phoenix
Penalty		
Perceived penalty κ, last 48 hours	0.098	0.081
	(0.036)	(0.010)
Outside offers (one a day on average)		
Offer with a positive net payoff, $P(\pi > 0)$	0.099	0.365
	(0.025)	(0.037)
Offer that beats a typical match:		
before the last 48 hours, $P(\pi > \hat{y})$	0.034	0.054
	(0.003)	(0.003)
in the last 48 hours, $P(\pi > \hat{y} + \kappa)$	0.020	0.019
	(0.002)	(0.001)
Share of these cancellations the penalty prevents	0.397	0.639
	(0.077)	(0.032)
Attention		
On the pickup day, α_0	0.482	0.505
	(0.018)	(0.014)
On the day before pickup, α_1	0.584	0.577
	(0.016)	(0.011)
Mean cost of hauling the load		
Load viewed 4 or more days before pickup, μ^c_{4+}	0.835	0.802
	(0.011)	(0.006)
Load viewed on the pickup day, μ^c_0	0.879	0.956
	(0.036)	(0.009)
Clusters (carrier $\times$ pickup date)	27,502	116,032

NOTE: Maximum likelihood estimates, with standard errors in parentheses: sandwich, clustered by carrier and pickup date, and corrected for the inputs taken from the data; by the delta method for quantities computed from the parameters. Prices, costs and payoffs are in units of the load's spot rate. Outside offers arrive once a day on average, and $P(\pi > y)$ is the chance that an offer has a net payoff above y. The value of a typical match, $\hat{y}$, is the median value of keeping a first match won with a bid, $V_d(b,c)$, over days 2 to 4 before pickup on which the carrier was matched to the load and could cancel it. The share of these cancellations the penalty prevents is $1 - P(\pi > \hat{y} + \kappa)/P(\pi > \hat{y})$. Attention on days 2, 3 and 4 before pickup: 0.634, 0.647 and 0.603 on Seattle to San Francisco; 0.638, 0.626 and 0.596 on Los Angeles to Phoenix. All 21 parameters are in Table C.2.

they bear when they cancel inside the last 48 hours, the same for every carrier.

The two lanes have penalties of about the same size, but different outside offers. On Seattle to San Francisco, the daily chance of an offer with a positive net payoff is 9.9%, against 36.5% on Los Angeles to Phoenix. But in Figure 5, the curve for Seattle to San Francisco falls much more slowly, and the two curves cross. On Seattle to San Francisco, good offers are rare, but some are large. On Los Angeles to Phoenix, they are frequent but modest.

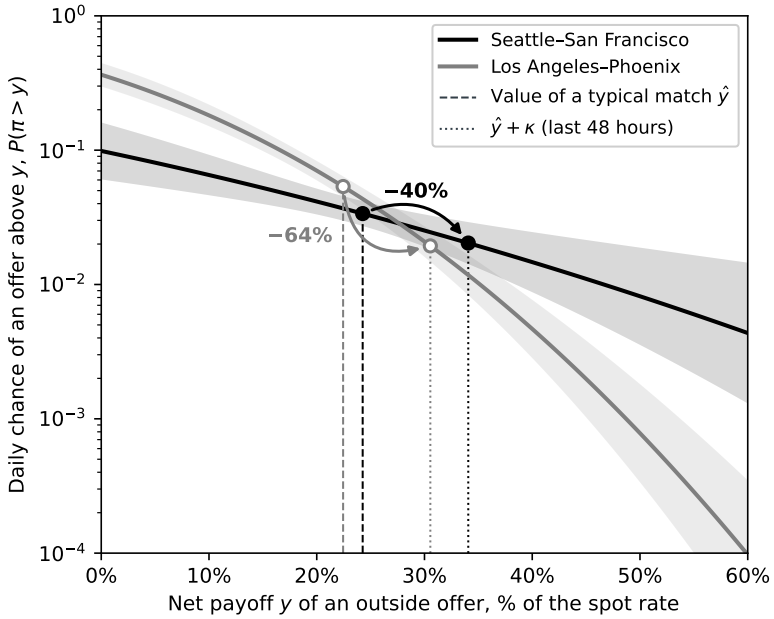
The effect of the penalty depends on these curves at the average value of being matched to a load. Take a typical load won with a bid, whose value of keeping is $\hat{y}$. Before the last 48 hours, the carrier cancels it for any offer above $\hat{y}$. The daily chance of such an offer is 3.4% on Seattle to San Francisco and 5.4% on Los Angeles to Phoenix. Inside the last 48 hours, the offer must also cover the penalty, and the chance falls to 2.0% and 1.9%. A penalty of about the same size thus prevents 40% of these cancellations on Seattle to San Francisco and 64% on Los Angeles to Phoenix. This share is one minus the ratio in equation (10), at the typical match. On Seattle to San Francisco, the cancellations that remain are for large offers, which a penalty of this size does little to stop.

This is the comparative static of the model in Section 3.1. There, a carrier cancels with probability $1 - G(r^*)$, and the penalty raises the cutoff r^*. How much this lowers the chance of cancelling depends on how fast the chance of a better offer falls above the cutoff. On Seattle to San Francisco, this chance falls slowly as the cutoff rises, so the penalty stops a smaller share of cancellations. On Los Angeles to Phoenix, it falls fast, so the penalty stops a larger share.

The difference matters for the counterfactuals. On Seattle to San Francisco, the penalty prevents a smaller share of cancellations, and each cancellation that goes ahead is worth more to the carrier. Unlike the reputational penalty, which nobody receives, a monetary penalty goes to the platform. Such a penalty would therefore bring in more money on Seattle to San Francisco for each cancellation it prevents. Section 6 examines what this means for the monetary penalty that the platform would choose.

Attention is lowest on the pickup day. A carrier notices a win on that day in time to confirm it with probability 0.482 on Seattle to San Francisco and 0.505 on Los Angeles to Phoenix, against 0.584 and 0.577 on the day before, and more on earlier days. Costs mostly rise toward pickup. The mean cost of hauling a load viewed on the pickup day is 0.879 of the spot rate on Seattle to San Francisco and 0.956 on Los Angeles to Phoenix, against 0.835 and 0.802 for a load viewed four or more days ahead. On Seattle to San Francisco, the mean cost is highest for loads viewed the day before pickup (Table C.2).

Figure 5: Outside offers in the two lanes



NOTE: The daily chance of an outside offer with a net payoff above y, $P(\pi > y) = \bar{G}(y)$, at the estimated μ^{π} and σ^{π} (Table C.2), on a log scale, with pointwise 95% confidence bands by the delta method. Dashed lines: the value of a typical match, $\hat{y}$, defined in the note to Table 4; before the last 48 hours, an offer above it ends the match. Dotted lines: $\hat{y}+\kappa$, the cutoff inside the last 48 hours. Arrows and percentages: the fall in this chance from $\hat{y}$ to $\hat{y}+\kappa$, the share of cancellations of a typical match that the penalty prevents.

The raw offer parameters of Seattle to San Francisco, the mean μ^{π} and standard deviation σ^{π} of G, are imprecise and tightly linked, because a lower mean with a wider spread fits the data almost as well (Table C.2). The chances of an offer that beats a typical match, before and inside the last 48 hours, are precise in both lanes, which is why the table reports chances rather than the raw parameters.

The simulated market also needs the platform's side: the arrivals of carriers, loads and views, the reserve and Accept-Now prices, and the cost of a load still unmatched at pickup (Appendix D).

6 Counterfactuals

Section 2.3 directly measured what a late carrier cancellation costs the platform on the specific load that was cancelled on, giving us a baseline for what a compensatory damages penalty

would look like in the conventional contract law sense. In the counterfactuals, we extend this analysis to consider the broader effects of a cancellation on the platform's other loads and on the other carriers in the market, as well as the effect of a cancellation policy on carriers' forward-looking bidding strategies, which will have implications for the number of initial matches formed.

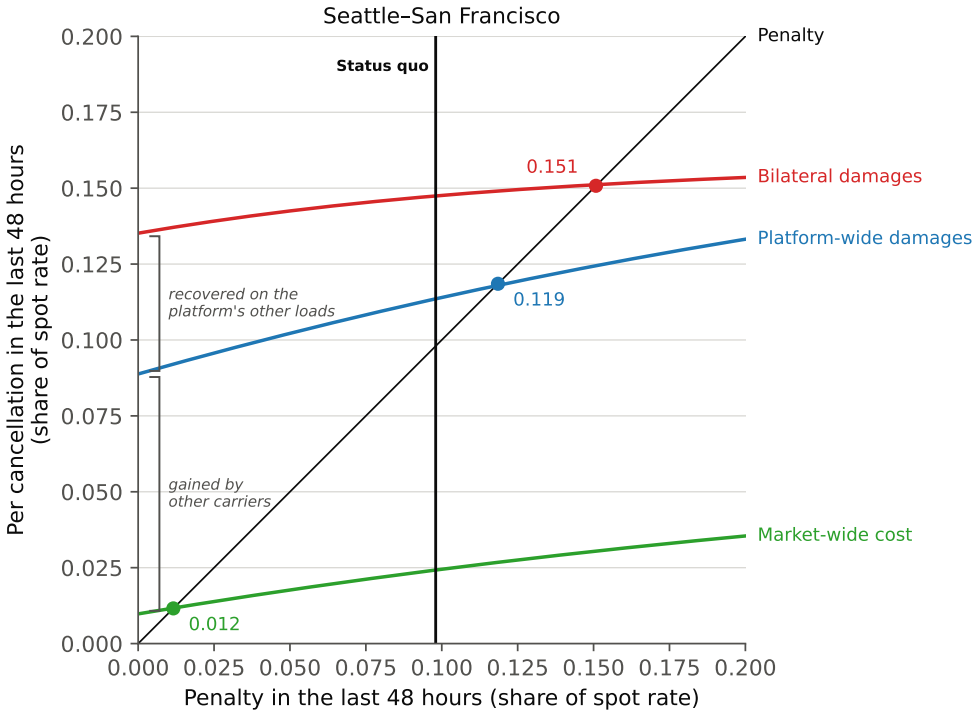
To do this, I simulate counterfactuals where I vary the level of the penalty κ charged in the last 48 hours, and resolve the market's equilibrium in each case holding the platform's reserve and Accept-Now prices fixed (Section 4.4). As a reminder (Section 3.4), the penalty can be *reputational*, in which case it purely burns carrier utility, or a *monetary* fee collected by the platform, which then acts as a transfer. This distinction has important implications both for welfare as well as the platform's incentives in penalty design. I report profit and welfare figures relative to the zero penalty scenario.

6.1 Simulating the Damages of a Cancellation

To evaluate the damage caused by a cancellation, I run an exact counterfactual copy for every simulated cancellation. A carrier cancels for an outside offer, to switch to another of the platform's loads, or because a forced event outside its control ends its participation (Section 4.1). When that occurs, the simulation state is copied exactly, including all future random draws, and only the cancellation event is negated, meaning the carrier turns down the offer, does not take the other load, or is not hit by the event. A comparison of the final counterfactual outcomes with each copy provides the exact ex-post damage caused by a cancellation. These are then averaged over all simulated cancellations occurring in the last 48 hours. The damage comparisons take place at three levels.

Bilateral damages are the platform's direct loss on the load the carrier cancelled, against the twin: a replacement's higher price, or, if no one replaces the carrier, the lost margin and eventual relational consequences. These are contract law's compensatory (expectation) damages, the loss to the party (in this case, the shipment) the carrier contracted with (Section 2.3). *Platform-wide damages* are bilateral damages less what the platform recovers on its other loads of the same lane and pickup date. A switch loses the platform one match and gains it another, so nearly all of this recovery comes from carriers who switch. Platform-wide damages are the platform's own view of what would make it whole, not a legal measure. The *market-wide cost* is platform-wide damages less what other carriers gain. It is the cost to the platform and every other carrier, leaving out the carrier who cancels, and it counts a replacement car-

Figure 6: What a cancellation costs, at every penalty



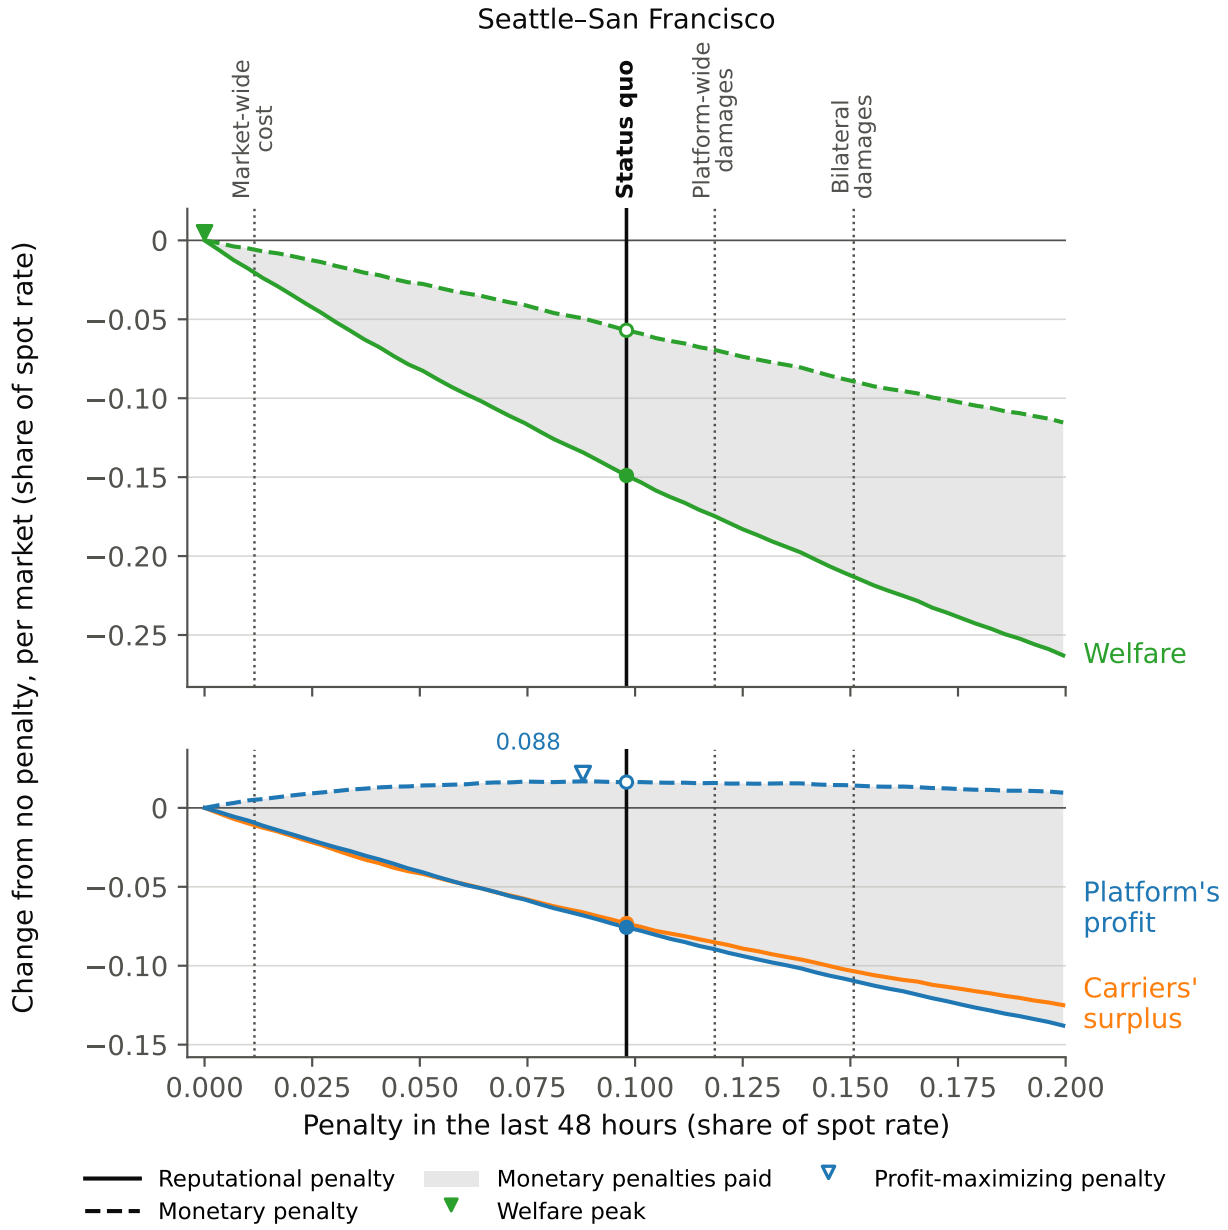
NOTE: Seattle to San Francisco, with Los Angeles to Phoenix in Figure E.1. Each curve averages one measure over the cancellations in the last 48 hours, forced and discretionary, smoothed by a spline across the grid of penalties. Dots mark the benchmarks, where each curve crosses the 45-degree line ("Penalty"), and Appendix E.3 gives their 95% intervals. The solid vertical line marks the status quo, the estimated penalty.

rier's whole surplus as a gain.

Each measure averages over the cancellations in the last 48 hours that still happen at a given penalty, and each rises with the penalty (Figure 6). The damages rise modestly because a higher penalty discourages carriers from taking on loads, particularly if they would be switching from another load and would themselves be subject to a penalty. The intersections of the damage curves with the 45-degree line represents fixed points where the penalty exactly reflects the corresponding damage measure, and thus serve as natural benchmarks for setting the penalty. The estimated penalty is well above the market-wide cost; its 95% interval, with a standard error of 0.036, spans both damages benchmarks.

With no penalty, the brackets in Figure 6, with the market-wide cost below them, split the 0.135 of bilateral damages three ways: 0.046 recovered on the platform's other loads, 0.079 gained by other carriers, and 0.011 of market-wide cost. Most of the platform's loss on the load is thus not lost to the market, and much of it goes to the carrier who moves the load instead.

Figure 7: Welfare, carriers' surplus and the platform's profit at every penalty



NOTE: Seattle to San Francisco, with Los Angeles to Phoenix in Figure E.2. The levels with no penalty are in the first row of Table 5. Both panels have the same scale. Carriers' surplus is the same with a reputational penalty and with a monetary penalty. The solid vertical line marks the status quo, the estimated penalty, and the dotted vertical lines mark the benchmarks of Figure 6.

6.2 Profit and Welfare under Different Penalties

Figure 7 shows how welfare, carriers' surplus and the platform's profit on Seattle to San Francisco change as the penalty is increased from zero. Total welfare is always highest with no penalty, and lower under every benchmark penalty. The platform's payoffs, and thus incentives, depend on the type of penalty. With a reputational penalty, the platform collects nothing, and the cancellations a penalty prevents aren't enough to offset the overall loss in matches, so its profit peaks at zero penalty, the welfare-maximizing level.[15] With monetary penalties, the platform collects additional revenues, making a non-zero penalty profitable on Seattle to San Francisco, at the cost of carrier welfare. Table 5 reports welfare and profits under monetary penalties at penalty levels for both the Seattle to San Francisco and Los Angeles to Phoenix lanes.

Table 5: Welfare and the platform's profit at each benchmark, with a monetary penalty

	Seattle–San Francisco			Los Angeles–Phoenix		
	Penalty	Welfare	Platform's profit	Penalty	Welfare	Platform's profit
No penalty	0.000	2.091	−0.179	0.000	4.065	1.175
Change from no penalty						
Penalty equal to the market-wide cost	0.012	−0.006	0.005	0.002	−0.004	−0.001
Status quo (estimated penalty)	0.098	−0.057	0.016	0.081	−0.127	−0.041
Penalty equal to platform-wide damages	0.119	−0.069	0.016	0.104	−0.156	−0.052
Penalty equal to bilateral damages	0.151	−0.089	0.014	0.111	−0.166	−0.055
Profit-maximizing monetary penalty	0.088	−0.049	0.017	0.000	0.000	0.000

NOTE: Per market, in shares of the spot rate, with a monetary penalty in the last 48 hours, collected by the platform. The first row gives levels with no penalty, and the other rows give changes from no penalty, linear between grid points. The benchmarks are those of Figure 6 and, for Los Angeles to Phoenix, Figure E.1. The platform's profit counts the cost of loads still unmatched at pickup, which on Seattle to San Francisco exceeds its margins on matched loads. Standard errors of the changes, paired by market, are at most 0.003.

On both lanes, the status quo penalty level is close to the platform-wide damage level, suggesting a reasonably calibrated penalty from the classic contract law perspective. However, this penalty level is neither welfare-maximizing nor profit-maximizing for the platform. The two lanes contrast in terms of their profit-maximizing penalty levels. On the Seattle to San-Francisco lane the profit-maximizing penalty is slightly below the status quo, while on the Los Angeles to Phoenix lane it is zero. The difference between lanes can be explained by the different nature of their outside offer processes. Recall from subsection 5.3 that carriers on the

15. This is also the case on the Los Angeles to Phoenix lane, as seen in Figure E.2

Los Angeles to Phoenix lane are more likely to receive outside offers and that these outside offers have a tighter distribution. Thus, the opportunity cost of penalties is higher and scales faster with the penalty level, so the negative pass-through effect is stronger on Los Angeles to Phoenix (pass-through of 0.59 per unit of penalty collected) than on Seattle to San Francisco (pass-through of 0.30), rationalizing the lower profit-maximizing monetary penalty on the Los Angeles to Phoenix lane.

6.3 The Tax on Matching

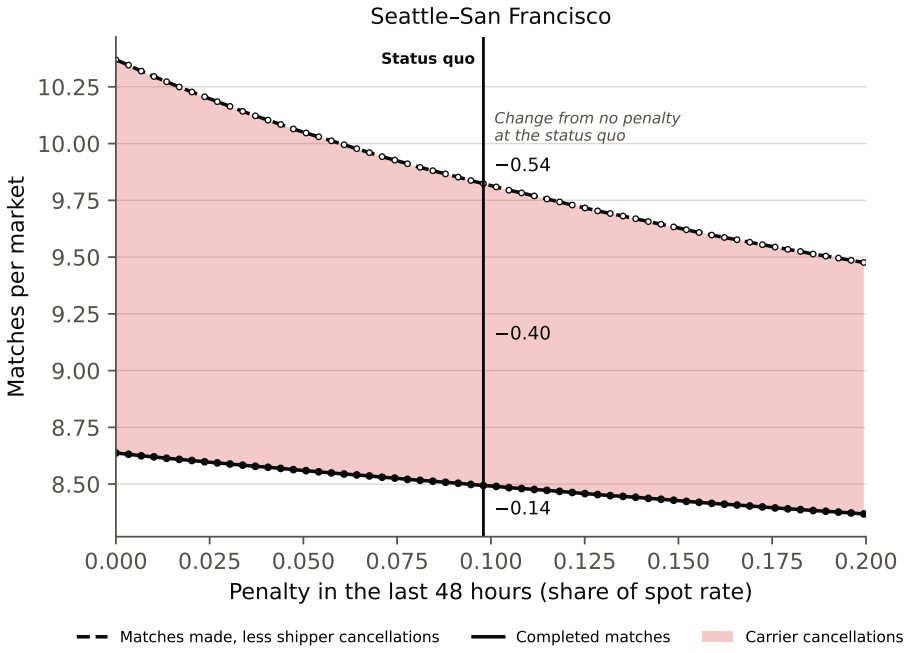
A penalty taxes two decisions: whether to cancel, and whether to match. A carrier who may face a penalty for leaving a match values the match less, so it bids higher and takes fewer loads at the Accept-Now price. In the bid equation (3) of Section 3.1, the penalty enters every bid, weighted by the chance of cancelling over the chance of completing. Figure 8 shows both decisions. At the estimated penalty, matches made, leaving out those that shippers cancel, change by -0.54 per market, and carrier cancellations by -0.40, so completed matches change by -0.14. A switch is one match made and one cancelled, so fewer switches explain part of the fall in matches made. Completed matches are the relevant measure, and they fall at every penalty even though a penalty prevents cancellations.

A match is worth more to the market than to the carrier who chooses it. Two parts of its value lie outside that choice: the platform's margin, 0.083 of the spot rate per completed match with no penalty, and the cost of not matching the load, which the match avoids, 0.159. A small fall in completed matches therefore costs a lot. Deterring a discretionary cancellation saves the platform and other carriers only 0.006. Per cancellation that a small penalty prevents, the tax on matching costs welfare 0.111, and lost completed matches alone cost more than the whole welfare loss at the estimated penalty (Appendix E.5). Welfare is therefore still rising as the penalty falls to zero, with a slope of -0.52 per unit of penalty if the penalty is monetary, since the market lacks matches rather than restraint in cancelling.

6.4 Penalty Timing

A platform that could charge a monetary penalty would have no reason to confine it to the hours in which cancellations cost the market. To see this, I move the start of the penalty window, so that the penalty applies only in the last 24 hours, or over the last 72 (Table 6). With no penalty, the market-wide cost is positive only for cancellations in the last 24 hours. Earlier

Figure 8: Matches and carrier cancellations at every penalty



NOTE: Seattle to San Francisco, with Los Angeles to Phoenix in Figure E.3. Dots mark the grid of penalties. Completed matches are those that last to pickup, and the shaded carrier cancellations are counted at any lead time. The solid vertical line marks the status quo, the estimated penalty.

ones slightly help the market, since the freed load is rematched more often and other carriers gain more than the platform loses. On Seattle to San Francisco, however, the platform's gain at the profit-maximizing monetary penalty is about the same in every window. On Los Angeles to Phoenix, the market-wide cost follows the same pattern, and the profit-maximizing penalty is zero in the 48-hour window.

The Handy settlement described in the introduction limited both when its fees applied and how large they could be. In the model, a limit on timing alone is undone by a higher monetary penalty. On Seattle to San Francisco, if the platform may charge only in the last 24 hours, the profit-maximizing penalty more than doubles (Table 6). Across the penalties within one standard error of the profit-maximizing one, welfare with a monetary penalty then changes by -0.095 to -0.048 per market, lower than with the 48-hour window. A limit on the amount is the only policy that helps. A ban keeps welfare at its peak, and a cap near the market-wide cost changes welfare by -0.006 per market, against -0.049 at the profit-maximizing penalty (Table 5).

Table 6: Market-wide cost and the profit-maximizing monetary penalty, by penalty window

	Seattle–San Francisco			Los Angeles–Phoenix		
Penalty applies in the	Market-wide cost	Profit-maximizing penalty	Platform's gain	Market-wide cost	Profit-maximizing penalty	Platform's gain
Last 24 hours	0.024	0.196	0.019	0.014	–	–
Last 48 hours	−0.014	0.088	0.017	−0.027	0.000	0.000
Last 72 hours	−0.020	0.088	0.019	–	–	–

NOTE: Shares of the spot rate, with a monetary penalty. The market-wide cost is the average per cancellation with no penalty, for the cancellations in the hours that each window adds: the last 24 hours, then 24 to 48 and 48 to 72 hours before pickup. The platform's gain is the change in its profit from no penalty at the profit-maximizing monetary penalty, per market. On Los Angeles to Phoenix, the 24- and 72-hour windows are still being computed. Appendix E.6 gives the standard errors.

6.5 Alternative Barriers to Monetary Penalties

Curbs on a monetary penalty matter only for a platform that puts little weight on carriers and collects it cheaply. A platform may value carriers' surplus because it wants to keep them. Collecting a fee also costs something, such as billing carriers and pursuing those who do not pay. Suppose the platform chooses its penalty to maximize a weighted sum of carriers' surplus and its profit, net of a cost for each payment it collects (Appendix E.7). On Seattle to San Francisco, with no cost of collection, the penalty disappears once the platform puts about a third of its weight on carriers. With no weight on carriers, it disappears once each payment costs about two percent of the spot rate to collect. Together, smaller amounts are enough (Figure E.4). On Los Angeles to Phoenix, the platform would not charge a monetary penalty even with no weight on carriers and no cost of collection. On Seattle to San Francisco, a small collection cost, however, first raises the penalty, which fewer carriers then pay, so it can lower welfare.

To sum up, bilateral damages, contract law's compensatory measure, would charge a cancellation more than ten times its market-wide cost, and the tax on matching outweighs even that cost. A reputational penalty keeps the platform in line with the market: nobody receives it, and the platform's profit, like welfare, peaks at no penalty. On Seattle to San Francisco, a monetary penalty would not, since the platform would charge one for the money. For a regulator, a limit on the amount matters more than one on timing.

7 Conclusion

Contract law prices a broken agreement at compensatory damages, the loss to the party left behind. Diamond and Maskin (1979) compare damage rules in a market where matched parties keep searching, so that a broken match also affects others. This paper measures that comparison on a freight platform, where carriers can cancel a match before pickup and the platform penalizes late cancellations through their future bids. Most of what a cancellation costs the platform on the load is not lost to the market, since the platform recovers part of it on its other loads, mostly from carriers who cancel to move another of them, and other carriers gain most of the rest. A penalty equal to compensatory damages would therefore charge a cancellation more than ten times what it costs the market, on both lanes. The penalty that carriers perceive on the platform is in the range of compensatory damages.

Even a penalty equal to the market's cost lowers welfare, because a penalty acts on two decisions: whether to cancel, and whether to match. A carrier who may face a penalty for leaving a match values the match less, so it bids higher and takes fewer loads at the posted price. Fewer matches last to pickup, and since a completed match is worth more to the market than to the carrier who chooses it, welfare is highest with no penalty. The right to renege thus lets carriers take loads before they know their other options, and the market gains more from these matches than it loses from the cancellations that follow.

The type of penalty can decide whether the platform's interest agrees with the market's. With a reputational penalty, which nobody receives, the platform's profit is also highest with no penalty. A monetary penalty would let the platform collect part of what carriers gain from their outside offers, so on Seattle to San Francisco a platform able to charge one would set it above the market's cost, and would have no reason to confine it to the last 24 hours, the only ones in which cancellations cost the market. Gig platforms that charge workers for cancelling may face the same incentive, and California's settlement with Handy limited both when its fees applied and how large they could be. In the model, a cap on the amount protects welfare, while a limit on timing alone is undone by a higher penalty.

References

Agarwal, Nikhil, Itai Ashlagi, Michael A Rees, Paulo Somaini, and Daniel Waldinger. 2021. "Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys." *Econometrica* 89 (1): 37–76.

Agarwal, Nikhil, Itai Ashlagi, Paulo Somaini, and Daniel Waldinger. 2018. "Dynamic incentives in wait list mechanisms." In *AEA Papers and Proceedings,* 108:341–347. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Airbnb. 2022. "Host Cancellation Policy," August 22, 2022. Accessed April 20, 2026. https://web.archive.org/web/20221127193530/https://www.airbnb.com/help/article/990/.

American Law Institute. 1981. *Restatement (Second) of Contracts.* St. Paul, MN: American Law Institute Publishers.

Backus, Matthew, and Greg Lewis. 2024. "Dynamic Demand Estimation in Auction Markets." *Review of Economic Studies,* rdae031.

Bajari, Patrick, Stephanie Houghton, and Steven Tadelis. 2014. "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs." *American Economic Review* 104 (4): 1288–1319.

Baker, George P, and Thomas N Hubbard. 2001. "Empirical Strategies in Contract Economics: Information and the Boundary of the Firm." *American Economic Review* 91 (2): 189–194.

———. 2003. "Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information." *American Economic Review* 93 (3): 551–572.

Caplice, Chris, and Yossi Sheffi. 2006. "Combinatorial Auctions for Truckload Transportation." Chap. 21 in *Combinatorial Auctions,* edited by Peter Cramton, Yoav Shoham, and Richard Steinberg. MIT Press.

Cramton, Peter, Sean Ellermeyer, and Brett Katzman. 2015. "Designed to Fail: The Medicare Auction for Durable Medical Equipment." *Economic Inquiry* 53 (1): 469–485.

De Groote, Olivier, Anaïs Fabre, Margaux Luflade, and Arnaud Maurel. 2025. *Sequential College Admission Mechanisms and Off-Platform Options.* Working Paper, Working Paper Series 34105. National Bureau of Economic Research, August. https://doi.org/10.3386/w34105. http://www.nber.org/papers/w34105.

Diamond, Peter A, and Eric Maskin. 1979. "An Equilibrium Analysis of Search and Breach of Contract, I: Steady States." *Bell Journal of Economics* 10 (1): 282–316.

Doraszelski, Ulrich, and Kenneth L Judd. 2012. "Avoiding the Curse of Dimensionality in Dynamic Stochastic Games." *Quantitative Economics* 3 (1): 53–93.

Federal Motor Carrier Safety Administration. 2023. *Pocket Guide to Large Truck and Bus Statistics*. Technical report. Accessed June 24, 2024. https://www.fmcsa.dot.gov/safety/data-and-statistics/commercial-motor-vehicle-facts.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong. 2000. "Optimal nonparametric estimation of first-price auctions." *Econometrica* 68 (3): 525–574.

Haberman, Alexander, and Ravi Jagadeesan. 2023. "Auctions with Withdrawal Rights: A Foundation for Uniform Price."

Harris, Adam, and Thi Mai Anh Nguyen. 2025. *Long-Term Relationships and the Spot Market: Evidence from US Trucking*.

Hubbard, Thomas N. 2000. "The Demand for Monitoring Technologies: The Case of Trucking." *The Quarterly Journal of Economics* 115 (2): 533–560.

———. 2001. "Contractual Form and Market Thickness in Trucking." *RAND Journal of Economics,* 369–386.

Huitfeldt, Ingrid, Victoria Marone, and Daniel C Waldinger. 2024. *Designing Dynamic Reassignment Mechanisms: Evidence from GP Allocation*. Working Paper, Working Paper Series 32458. National Bureau of Economic Research, May. https://doi.org/10.3386/w32458. http://www.nber.org/papers/w32458.

Jofre-Bonet, Mireia, and Martin Pesendorfer. 2003. "Estimation of a Dynamic Auction Game." *Econometrica* 71 (5): 1443–1489.

Kagel, John H, and Alvin E Roth. 2000. "The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment." *The Quarterly Journal of Economics* 115 (1): 201–235.

Kapor, Adam, Mohit Karnani, and Christopher Neilson. 2024. "Aftermarket frictions and the cost of off-platform options in centralized assignment mechanisms." *Journal of Political Economy* 132 (7): 2346–2395.

Li, Hao, and Sherwin Rosen. 1998. "Unraveling in Matching Markets." *American Economic Review,* 371–387.

Li, Hao, and Wing Suen. 2000. "Risk Sharing, Sorting, and Early Contracting." *Journal of Political Economy* 108 (5): 1058–1091.

McCall, John Joseph. 1970. "Economics of Information and Job Search." *The Quarterly Journal of Economics* 84 (1): 113–126.

McKinney, C. Nicholas, Muriel Niederle, and Alvin E. Roth. 2005. "The Collapse of a Medical Labor Clearinghouse (and Why Such Failures Are Rare)." *The American Economic Review* 95 (3): 878–889. ISSN: 00028282, accessed April 23, 2026. http://www.jstor.org/stable/4132746.

Newey, Whitney K, and Daniel McFadden. 1994. "Large sample estimation and hypothesis testing." *Handbook of econometrics* 4:2111–2245.

Niederle, Muriel, and Alvin E Roth. 2003. "Unraveling reduces mobility in a labor market: Gastroenterology with and without a centralized match." *Journal of political Economy* 111 (6): 1342–1352.

———. 2009. "Market Culture: How Rules Governing Exploding Offers Affect Market Performance." *American Economic Journal: Microeconomics* 1 (2): 199–219.

Owner-Operator Independent Drivers Association. 2022. *Owner-Operator Member Profile Survey.* Technical report. Accessed June 24, 2024. https://www.ooida.com/wp-content/uploads/2022/05/2022-Owner-Operator-Survey.pdf.

People v. Handy Technologies, Inc. 2023. *Stipulated Final Judgment and Permanent Injunction.* Case No. CGC-21-590442, Superior Court of California, County of San Francisco, May 18, 2023.

Roth, Alvin E, and Xiaolin Xing. 1994. "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions." *The American Economic Review,* 992–1044.

———. 1997. "Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists." *Journal of Political Economy* 105 (2): 284–329.

U.S. Department of Transportation. 2024. "Moving Goods in the United States." Bureau of Transportation Statistics. Accessed June 24, 2024. https://data.bts.gov/stories/s/Moving-Goods-in-the-United-States/bcyt-rqmu.

Uber. 2025. "Introducing the new Uber Eats Pro – better service, better status, better earnings," May 12, 2025. Accessed April 20, 2026. https://perma.cc/7CJR-L7QC.

Waldinger, Daniel. 2021. "Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation." *American Economic Review* 111 (8): 2660–2696.

Xie, Yingkang, and Huaiyu Zhu. 2022. "Platform Leakage: Disintermediation and Incentive Conflicts in Two-Sided Markets." Working paper. https://questromworld.bu.edu/platformstrategy/wp-content/uploads/sites/49/2023/06/PlatStrat2023_paper_27.pdf.

Yang, Ron. 2022. *(Don't) Take Me Home: Home Preference and the Effect of Self-Driving Trucks on Interstate Trade*. Technical report. National Bureau of Economic Research.

Zeithammer, Robert. 2006. "Forward-Looking Bidding in Online Auctions." *Journal of marketing research* 43 (3): 462–476.

Appendices

A Reduced-form evidence: construction and robustness

A.1 Screening motives of reputational penalties

In the main text, the model focuses on the incentive effects of reputational penalties, assuming they reduce the present discounted value of a carrier's future profits on the platform. The penalty affects the platform only through carriers' choices: the carrier who cancels pays it, and nobody receives it (Section 4.4). However, this assumption may not hold if reputational penalties also affect the platform's profits by *screening* carriers. A penalty makes a carrier's later bids less likely to win, so it steers the platform's loads away from carriers who cancel often. This could help the platform, as fewer of its matches would end in a cancellation. A counterpoint to this argument is that screening could also hurt the platform: carriers who cancel often may also match at lower prices, so steering loads away from them would raise what the platform pays. This subsection asks which of the two effects dominates.

To answer this question, the following reduced-form exercise, with minimal modelling assumptions, approximates the potential benefits of screening out frequent cancellers:

1. For each carrier with at least 10 matches, compute the mean cancellation rate and average margins directly from the data.
2. Rank carriers by their cancellation rate, and exclude the carriers with the highest rates, who together hold X% of the matches.
3. Evaluate the impact of this exclusion on the overall cancellation rate and margins of the remaining matches.

The sample consists of the matches (confirmed bids and accept-nows) of owner-operators on non-local loads of 300 miles or more with a pickup date in 2021–22, offered in the timed auctions (Section 2.2): 115,632 matches by 15,579 carriers. A match counts as cancelled if it ends in a carrier-fault cancellation, as in Section 3.2. The 3,102 carriers with at least 10 matches hold 80,903 of these matches (70.0%).

I focus on margins here under the assumption that the platform's remaining carriers will pick up the slack from the excluded carriers. While unrealistic, this assumption is biased in favor of screening, as it assumes that the platform can fully replace the excluded carriers with other carriers and does not suffer any other negative consequences from the reduced carrier base (in particular, reduced competition in auctions). Let b denote the match price, the price the platform agreed to pay the carrier, and let v denote the price the shipper pays for the load, its value to the platform. I focus on two types of margins:

- **Booking margin** $= \frac{v-b}{v}$
  (as if the match is kept)
- **Final margin** $= \mathbb{1}[\text{No cancel}]\frac{v-b}{v} - L\,\mathbb{1}[\text{Cancel}]$
  (accounts for the impact of cancellations)

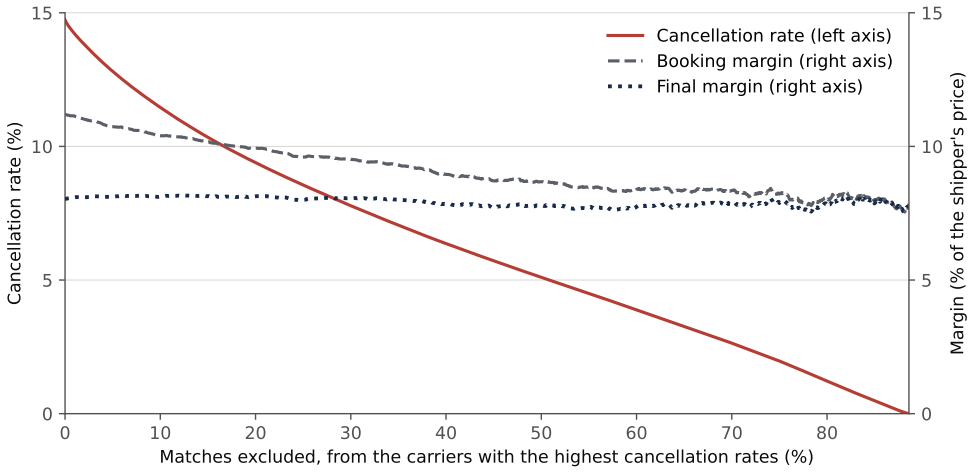
The final margin is the relevant profit margin I focus on. When the carrier cancels, the final margin charges the platform the cost of not matching the load, a share L of the shipper's price, as if the platform had to ask the shipper to reschedule or cancel the shipment. I measure L by what the platform pays to avoid this, on the 122,053 loads for which the platform accepted a bid in the last 24 hours before pickup. For each of these loads, I take the highest bid the platform accepted in those 24 hours, whether or not the carrier then confirmed it, and compare it with the shipper's price. On average, this bid is 5.3% above the shipper's price, so L is 5.3%.[16] The final margin is a more accurate measure of the platform's profitability, as it accounts for the cost of cancellations. It is again biased in favor of screening. It charges every cancellation the cost of not matching the load, even when the carrier cancels days before pickup and the platform has time to find another carrier. In the data, the load is rarely lost after a cancellation (Section 2.3).

Figure A.1 shows the impact of excluding carriers with the highest cancellation rates on the overall cancellation rate and margins. The figure shows that excluding carriers with the highest cancellation rates reduces the overall cancellation rate but also reduces the booking margin, implying that carriers who cancel more often also tend to match at lower prices. Excluding the carriers with the highest rates, who hold 10.0% of the matches, lowers the cancellation rate from 14.7% to 11.5%. It also lowers the booking margin from 11.2% to 10.4% of the shipper's price. The final margin rises only from 8.0% to 8.1%. Excluding still more carriers lowers it a little. Once only the carriers who never cancel remain, it is 7.9%, below its level with no

16. The model measures the cost of a load still unmatched at pickup in the same way, from the highest bids accepted on the pickup day, separately for each of its two lanes (Section D.1).

exclusion. The two forces nearly offset each other in terms of the final margin, suggesting that the net benefit of screening out carriers who cancel often is small.

Figure A.1: Effect of screening out frequent cancellers on the cancellation rate and margins



NOTE: Effect of excluding the carriers with the highest cancellation rates, who together hold X% of the matches (horizontal axis). Carriers are ranked by their share of matches that end in a carrier-fault cancellation. Left axis: cancellation rate of the remaining matches (solid). Right axis: margins of the remaining matches, in percent of the shipper's price. Booking margin (dashed): the shipper's price minus the match price, divided by the shipper's price. Final margin (dotted): the booking margin for matches that are kept, and minus $L = 5.3\%$ for matches the carrier cancels: the average amount by which the highest bid the platform accepted in the last 24 hours before pickup exceeds the shipper's price. Sample: the 3,102 owner-operators with at least 10 matches on non-local loads of 300 miles or more with a pickup date in 2021–22 (80,903 matches). Margins leave out the 1.4% of matches without a usable price. The curves end at 88.6%, where only the 559 carriers who never cancel remain.

In conclusion, while reputational penalties might have a screening effect by discouraging or excluding frequent cancellers, the net benefit to the platform is small under conservative assumptions. The loss of carriers who match at lower prices largely offsets the gains from fewer cancellations. Moreover, reducing the carrier base may adversely affect the platform's ability to meet demand. Therefore, the incentive effects of reputational penalties, as discussed in the main text, remain the primary channel through which these penalties influence carrier behavior and platform performance. This is why the model lets the reputational penalty affect the platform only through carriers' choices.

A.2 How the quality adjustment and the realized penalty were measured

The timed auction ranks bids by the bid plus a carrier-specific quality adjustment. Carriers saw neither the adjustment nor how it was computed. I recover it from an archive of auction rankings, which records for each bid the price used to rank it; the adjustment is that price minus the bid. I then relate the adjustment to the carrier's recent record on the platform (its recent carrier-fault cancellations and its other recent outcomes), controlling for carrier and calendar-time effects, separately for the two periods in which the platform used different versions of the rule (2021 and 2022). In both periods the adjustment rises with the share of carrier-fault cancellations in the recent record, and one cancellation moves it more when the record is short. I do not report the form or the coefficients of the reconstructed rule.

The realized penalty in Figure A.2 is computed at each later match i (105,810 matches of owner-operators on loads of 300 miles or more). Let H_i be the carrier's record just before match i, b_i its price, s_i the load's spot rate, and $q(H; b)$ the adjustment that the reconstructed rule assigns to a bid b given a record H. Then

$$\Delta_i = \frac{100}{s_i}\Big[\, q\big(H_i \oplus \text{cancellation};\, b_i\big) - q\big(H_i \oplus \text{on-time completion};\, b_i\big)\Big], \tag{A.1}$$

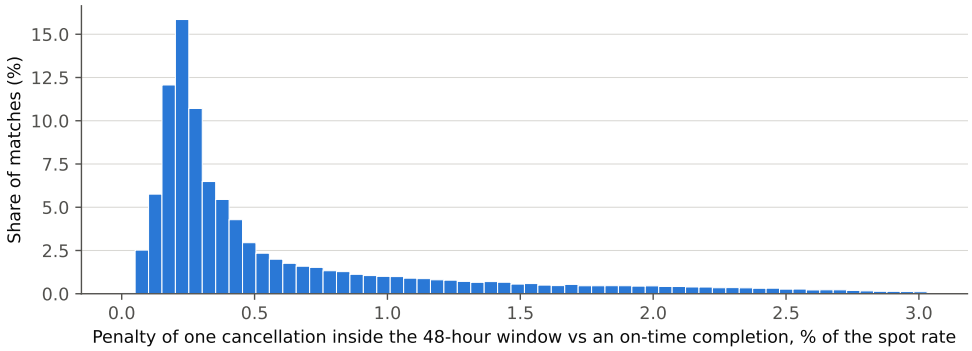
where $H \oplus e$ is the record H with one more event e: a carrier-fault cancellation inside the communicated 48-hour window, or an on-time completion. Δ_i is the percent of spot by which one such cancellation raises the carrier's adjustment at match i. The tercile analysis of Section 3.3 uses the same quantity, evaluated at the start of each match.

The code that reconstructs the adjustment from the auction archive and computes the realized penalty is part of the replication package. The platform's records are confidential and are not released.

A.3 The 48-hour hazard: specification, tests and the forced comparison series

The hazard sample consists of the matches of owner-operators on loads of 300 miles or more in the universe of Section 2.2, won with a bid or taken at the accept-now price, that were still active at some point in the last four days before pickup. Each match is split into two-hour cells of time before pickup, with the first hour after the match as a separate segment. For match h and cell c, y_{hc} is one if the match ends in a cancellation of the given group in that cell, and

Figure A.2: What one cancellation cost a carrier at each later match



NOTE: Distribution across matches of the realized penalty of one carrier-fault cancellation inside the communicated 48-hour window: the change in the quality adjustment at a later match, compared with one on-time completion, given the carrier's record and bid at that match, in percent of the load's spot rate. The mean is 0.62% of the spot rate. The histogram is drawn up to the 99th percentile. Owner-operators on loads of 300 miles or more, 2021 and 2022 rule periods (105,810 matches). The adjustment is reconstructed from archived auction rankings, as described above.

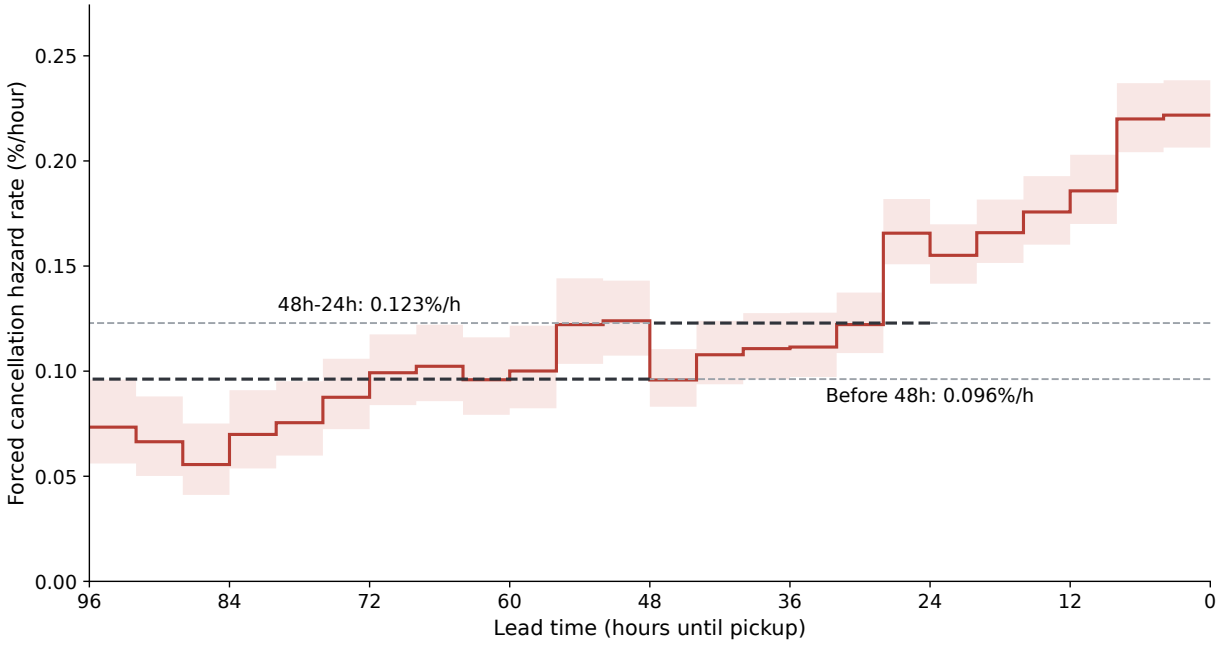
E_{hc} is the number of hours the match was at risk in the cell. The match's time at risk starts at the beginning of the window, or when the match was made if later, and ends at pickup or at a cancellation of any kind. The model is

$$\mathbb{E}[y_{hc}] = E_{hc}\,\exp\Big(\theta_{b(c)} + \omega_{k(hc)} + \phi\,F_{hc} + \psi\,A_h + \beta_{\text{bid}}\,(1-A_h)\,x_h + \beta_{\text{AN}}\,A_h\,x_h\Big), \tag{A.2}$$

where $\theta_{b(c)}$ is an effect for the four-hour bin of time before pickup, $\omega_{k(hc)}$ an effect for the two-hour clock time at the middle of the cell's exposure, F_{hc} indicates the first hour after the match, A_h indicates an accept-now match, and x_h is the match price as a share of the spot rate. The model is estimated by Poisson maximum likelihood with standard errors clustered by carrier. The two-constant version replaces $\theta_{b(c)}$ with one constant for the last two days and one for the two days before; the three-constant version splits the last two days at 24 hours before pickup, so that the drop at the deadline is measured on the first day after it. Reported hazards are evaluated at the sample's exposure-weighted mix of clock times and channels, each channel at its mean price, outside the first hour after the match.

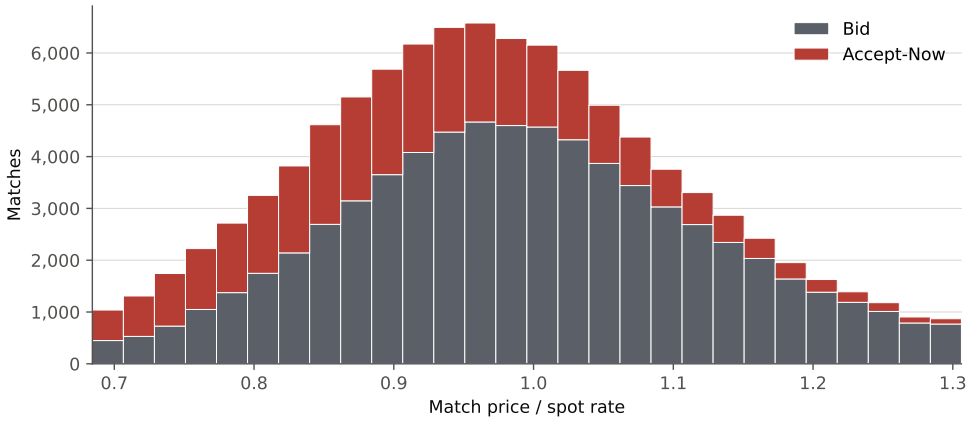
Figure A.4 shows the prices of these matches, by channel: 72,929 matches won with a bid and 32,153 taken at the posted accept-now price.

Figure A.3: Forced cancellation hazard over the last four days before pickup



NOTE: As Figure 3a, for forced cancellations (Section 3.2): estimated hazard in percent per hour by four-hour bin of time before pickup, with 95% confidence intervals clustered by carrier, the constant more than 48 hours before pickup and the constant from 48 to 24 hours before pickup (dashed; dark over the hours each is fitted to). Same matches and regression (equation (A.2)); 6,645 forced cancellations in the window.

Figure A.4: Match prices in the hazard sample, by channel



NOTE: Number of matches by match price as a share of the spot rate, stacked by channel: matches won with a bid and matches taken at the posted accept-now price. Matches of owner-operators on loads of 300 miles or more with a pickup date in 2021–22 (105,082 matches), the sample of Figure 3; drawn over the price range of Figure 3b.

A.4 Uniformity across realized-penalty terciles: details

Each match is assigned the realized penalty of equation (A.1), evaluated at the start of the match, and placed in a tercile $T(h) \in \{1,2,3\}$ of that penalty. The model replaces the four-hour bins of equation (A.2) with one constant per tercile on each side of the deadline:

$$\begin{aligned}\mathbb{E}[y_{hc}] = E_{hc}\exp\Big(&\sum_{T=1}^{3}\mathbf{1}\{T(h)=T\}\left[\theta_T^{\text{pre}}\,(1-P_c)+\theta_T^{\text{post}}\,P_c\right]\\ &+\omega_{k(hc)}+\phi\,F_{hc}+\psi\,A_h+\beta_{\text{bid}}(1-A_h)\,x_h+\beta_{\text{AN}}\,A_h\,x_h\Big),\end{aligned} \tag{A.3}$$

where P_c indicates a cell in the last two days before pickup. The deadline response of tercile T is $\rho_T = \exp\big(\theta_T^{\text{post}} - \theta_T^{\text{pre}}\big)$, and a Wald test with two degrees of freedom checks $\rho_1 = \rho_2 = \rho_3$.

A.5 Bid confirmation: sample and specification

The confirmation sample starts from the bids of owner-operators. For each carrier and pickup date in force, I find the carrier's first bid acceptance among all of its bids, on any load. I keep it if the load is in the universe of Section 2.2 and is 300 miles or more, and if the acceptance came before pickup, within the last five days before pickup (the days in which the auction clears bids, Section 2.1), and less than four days after the bid. A carrier's later acceptances for the same pickup date are left out, because by then the carrier may already hold a load for that day. This leaves 126,003 acceptances by 15,461 carriers, of which 44.7% were confirmed.

The adjusted rates in Figure 3c come from the logit

$$\Pr(\text{confirm}_i = 1) = \Lambda\Big(\sum_{w=1}^{39}\theta_w\,\mathbf{1}\{\text{wait}_i \in W_w\} + \omega_{k(i)} + \tau_{d(i)}\Big), \tag{A.4}$$

where the wait bins W_w are 15 minutes long in the first hour, one hour long up to one day, and six hours long after that; $\omega_{k(i)}$ is an effect for the two-hour clock time of the acceptance and $\tau_{d(i)}$ an effect for the days before pickup of the acceptance. Standard errors are clustered by carrier. The rate reported for bin w is the average over the sample of $\Lambda(\theta_w + \omega_{k(i)} + \tau_{d(i)})$, the predicted probability with each acceptance's wait set to bin w.

A.6 Derivation of the optimal bid

This subsection derives equation (3) of Section 3, using the notation defined there. To solve the optimal bidding problem, we can define the expected payoff of a bid b, net of the common continuation value $\delta\mathcal{V}(q^C)$, as

$$\gamma(b+q)\,\mathbb{E}\left[\max\{b-c,\pi-\kappa\}\right]+\left[1-\gamma(b+q)\right]\mathbb{E}\left[\max\{0,\pi\}\right]. \tag{A.5}$$

Raising the bid increases the payment only when the carrier completes the shipment, an event with probability $G(b-c+\kappa)$ conditional on winning. Consequently,

$$\frac{\partial}{\partial b}\mathbb{E}\max\{b-c,\pi-\kappa\}=G(b-c+\kappa).$$

Differentiating (A.5), an interior optimal bid b^* satisfies

$$\begin{aligned}\gamma'(b^*+q)\,\mathbb{E}\big[\max\{b^*-c,\pi-\kappa\}-\max\{0,\pi\}\big]\\ +\gamma(b^*+q)G(b^*-c+\kappa)=0.\end{aligned} \tag{A.6}$$

The first term captures the loss of winning probability from a higher bid, weighted by the benefit of winning rather than remaining unmatched. The second captures the additional payment when the carrier wins and completes.

Let $r^*=b^*-c+\kappa$ be the outside-offer cutoff at the optimal bid. The difference between the expected payoffs from winning and losing is

$$\begin{aligned}&\mathbb{E}\big[\max\{b^*-c,\pi-\kappa\}-\max\{0,\pi\}\big]\\ &\quad=G(r^*)(b^*-c)-\kappa[1-G(r^*)]-\int_0^{r^*}\pi\,dG(\pi).\end{aligned}$$

For $G(r^*)>0$, substitution into (A.6) yields

$$b^*=c+\frac{\kappa[1-G(r^*)]+\int_0^{r^*}\pi\,dG(\pi)}{G(r^*)}-\frac{\gamma(b^*+q)}{\gamma'(b^*+q)}, \tag{3}$$

which is the optimal bid of Section 3.

B Carrier Model Derivations

This appendix derives the carrier's values of Section 4.2, the probabilities of a cancellation and of a confirmation, and the first-order condition of the bid in Section 4.3. The notation is that of Section 4.

B.1 Values in Continuous Time

The carrier's state is its position on the load and the day. Events change the state, and each kind of event arrives at its own rate: the day ends, an outside offer arrives, a forced shock hits, or a pending bid wins. Nothing is paid between events, and I do not discount.

Consider a state s, and let $\nu_e(s)$ be the rate of event e in that state, with total rate $\bar{\nu}(s) = \sum_e \nu_e(s)$. Event e pays ϕ_e and moves the carrier to state s_e. Since nothing is paid while waiting and nothing is discounted, the time until the next event does not matter. The value $J(s)$ of the state is the expected payoff of the next event plus the value of the state it leads to. With events that arrive at independent Poisson rates, the next event is e with probability $\nu_e(s)/\bar{\nu}(s)$. So

$$J(s) = \sum_e \frac{\nu_e(s)}{\bar{\nu}(s)} \, \mathbb{E}\big[\phi_e + J(s_e)\big].$$

An outside offer brings a choice. The carrier either keeps its position, worth $J(s)$, or takes the offer, which pays its net payoff π and ends the carrier's time on the load. For an offer, $\phi_e + J(s_e)$ is the better of the two. A carrier who has left the load, after an offer or a forced shock, gets nothing more from it. Multiplying through by $\bar{\nu}(s)$ gives the form of Section 4.2: the total rate of events times the value equals the sum, over events, of each rate times what the event is worth.

A free carrier on day d faces three events. An outside offer arrives at rate λ, and the carrier takes it if π exceeds U_d. A forced shock arrives at rate λ_d^F; it ends the carrier's capacity for the pickup date and pays nothing. The day ends at rate η, and the carrier is free on day $d-1$. Hence

$$(\eta + \lambda + \lambda_d^F)\, U_d = \lambda\, \mathbb{E} \max\{U_d,\, \pi\} + \eta\, U_{d-1}. \tag{4}$$

A carrier still free at pickup earns nothing from the load, so $U_{-1} = 0$.

A matched carrier faces the same three events. Taking an offer cancels the match, so the offer

is worth $\pi - \kappa_d$, net of the penalty. A forced shock also cancels the match, and the carrier pays κ_d. At pickup, the carrier hauls the load and earns $V_{-1} = b - c$. Hence

$$(\eta + \lambda + \lambda_d^F)\, V_d = \lambda\, \mathbb{E} \max\{V_d,\, \pi - \kappa_d\} - \lambda_d^F\, \kappa_d + \eta\, V_{d-1}. \tag{5}$$

A carrier with a pending bid faces a fourth event: the bid wins, at rate $\gamma_d(b)$. The carrier then notices the win with probability α_d and becomes matched at price b on the same day, worth V_d. Otherwise it misses the win and is free again, worth U_d. Walking away from a bid costs nothing, so an offer is worth $\max\{W_d, \pi\}$. A forced shock leaves the carrier with nothing from the bid. A bid that has not won by pickup is worth nothing, so $W_{-1} = 0$. Hence

$$(\eta + \lambda + \lambda_d^F + \gamma_d)\, W_d = \lambda\, \mathbb{E} \max\{W_d,\, \pi\} + \gamma_d \big[\alpha_d V_d + (1 - \alpha_d)\, U_d\big] + \eta\, W_{d-1}. \tag{6}$$

Each value appears on both sides of its equation, because an offer the carrier turns down leaves it where it was. Taking the turned-down offers out of both sides gives the form I solve. Let $\bar{G} = 1 - G$. For any x, $\mathbb{E} \max\{x, \pi\} = G(x)\, x + \int_x^\infty \pi\, dG(\pi)$. With $\eta = 1$, as in Section 4.2, the three equations become

$$\begin{aligned}
U_d &= \frac{\lambda \int_{U_d}^\infty \pi\, dG(\pi) + U_{d-1}}{1 + \lambda \bar{G}(U_d) + \lambda_d^F}, \\
V_d &= \frac{\lambda \int_{r_d^*}^\infty (\pi - \kappa_d)\, dG(\pi) - \lambda_d^F\, \kappa_d + V_{d-1}}{1 + \lambda \bar{G}(r_d^*) + \lambda_d^F}, \\
W_d &= \frac{\lambda \int_{W_d}^\infty \pi\, dG(\pi) + \gamma_d \big[\alpha_d V_d + (1 - \alpha_d)\, U_d\big] + W_{d-1}}{1 + \gamma_d + \lambda \bar{G}(W_d) + \lambda_d^F},
\end{aligned}$$

where $r_d^* = V_d + \kappa_d$ is the cancellation cutoff of Section 4.2. Each denominator is the rate at which the carrier's position changes on day d: the day ends, the carrier takes an offer, a forced shock hits, or the bid wins. The values still appear on both sides, through the cutoffs, so I solve these equations by iterating on them until the values stop changing.

The same race between events gives the cancellation probability of Section 4.2 and the confirmation probability of Section 4.3. On day d, a matched carrier's position changes when the day ends, at rate one, when an offer above the cutoff arrives, at rate $\lambda \bar{G}(r_d^*)$, or when a forced shock hits, at rate λ_d^F. On a day without a forced cancellation, the offer comes before the end of the day with probability $\lambda \bar{G}(r_d^*)/[1 + \lambda \bar{G}(r_d^*)]$, which is equation (7). The forced rate drops out of this probability, but it still enters through V_d.

For the confirmation probability, take a bid placed on day d^b that wins on day d^a. On each day k from d^b down to d^a, the carrier leaves the bid, for an offer above W_k or after a forced shock, at rate $\lambda\bar{G}(W_k) + \lambda_k^F$. A carrier who has left does not confirm, but its bid can still win. Meanwhile, the day ends at rate one and the bid wins at rate $\gamma_k(b)$. The time until the first of these two events does not depend on which of them comes first. So on each of these days, whether or not the bid wins that day, the carrier is still there at the end of the day or at the win with probability $[1 + \gamma_k(b)]/[1 + \gamma_k(b) + \lambda\bar{G}(W_k) + \lambda_k^F]$. Multiplying over the days, and by the chance α_{d^a} that the carrier notices the win, gives equation (8).

B.2 An Envelope Lemma

The derivatives of the values with respect to the bid rest on one lemma, a special case of the envelope theorem.

Lemma 1. *Let f be a differentiable function, a a constant, and G a distribution with a continuous density G'. Then $h(x) = \mathbb{E}\max\{f(x),\, \pi - a\}$, with the expectation over $\pi \sim G$, has the derivative*

$$h'(x) = G\big(f(x) + a\big)\, f'(x).$$

Proof. Write $h(x) = G\big(f(x)+a\big)\, f(x) + \int_{f(x)+a}^{\infty} (\pi - a)\, dG(\pi)$. The derivative of the first term is $G\big(f(x)+a\big)\, f'(x) + G'\big(f(x)+a\big)\, f'(x)\, f(x)$. The derivative of the second term is $-G'\big(f(x) + a\big)\, f'(x)\, f(x)$, since the integrand equals $f(x)$ at the lower limit. The last two terms cancel. □

In words, a change in $f(x)$ changes the expected maximum only when the carrier keeps the option f, which happens with probability $G(f(x) + a)$. The change also moves the cutoff at which the carrier switches to the offer, but at the cutoff the carrier is indifferent, so this has no first-order effect. With $f = V_d$ and $a = \kappa_d$, the lemma gives the derivative of a matched carrier's expected maximum. With $f = W_d$ and $a = 0$, it gives that of a carrier with a pending bid.

B.3 The Bid First-Order Condition

The win rate of a bid depends on the day it was placed as well as on the day it wins. I write $\gamma_{g,d}(b)$ for the rate at which a bid b placed on day g wins on day $d \le g$, and $W_{g,d}(b, c)$ for the value of that bid, which solves equation (6) with $\gamma_{g,d}$ in place of γ_d. In equation (8), the win

rates and values are those of the bid's own day, $g = d^b$. Auctions clear only in the last five days, so $\gamma_{g,d} = 0$ for $d \geq 5$. In a race with the end of the day, a pending bid wins on day d with probability $\gamma_{g,d}(b)/[1+\gamma_{g,d}(b)]$. This is the win probability that Appendix C.3 estimates, separately for bids placed on each of days 0 to 4 and jointly for bids placed earlier. The values U_d and $V_d(b,c)$ do not depend on g. A carrier who views the load on day g and bids picks the bid $b^*_g(c)$ that maximizes $W_{g,g}(b,c)$. From here on, $\eta = 1$, and I leave out the arguments b and c of V and W.

Start with a match. Raising its price by one raises what the carrier earns at pickup by one, but only if the match lasts that long. I differentiate equation (5) using Lemma 1, with $f = V_d$ and $a = \kappa_d$. The forced-cancellation term drops out, because it does not depend on b. This gives

$$\frac{\partial V_d}{\partial b} = \frac{1}{1+\lambda\bar{G}(r^*_d)+\lambda^F_d}\,\frac{\partial V_{d-1}}{\partial b} = \prod_{m=0}^{d}\frac{1}{1+\lambda\bar{G}(r^*_m)+\lambda^F_m}, \tag{B.1}$$

with $\partial V_{-1}/\partial b = 1$. Each factor is the chance that day m ends before an offer above the cutoff r^*_m arrives or a forced shock hits. The product is thus the chance that a match held on day d lasts to pickup.

Next, the pending bid. Differentiating equation (6) with Lemma 1, with $f = W_{g,d}$ and $a = 0$, and using that U_d does not depend on b,

$$\begin{aligned}\Big[1+\gamma_{g,d}(b)+\lambda\bar{G}(W_{g,d})+\lambda^F_d\Big]\,\frac{\partial W_{g,d}}{\partial b} = \gamma'_{g,d}(b)\Big[\alpha_d V_d + (1-\alpha_d)\,U_d - W_{g,d}\Big] \\ + \alpha_d\,\gamma_{g,d}(b)\,\frac{\partial V_d}{\partial b} + \frac{\partial W_{g,d-1}}{\partial b},\end{aligned} \tag{B.2}$$

with $\partial W_{g,-1}/\partial b = 0$. A higher bid has two effects on day d. It lowers the win rate, since $\gamma'_{g,d} < 0$, and so loses the gain from winning: the value of a win, confirmed or missed, net of the value of keeping the bid pending. It also raises what the carrier is paid when the bid wins, the carrier confirms and the match lasts to pickup.

On days $d \geq 5$ the bid cannot win, so $\gamma_{g,d} = \gamma'_{g,d} = 0$, and equation (B.2) only multiplies $\partial W_{g,d-1}/\partial b$ by a positive factor. For a bid placed on day $g \geq 5$, the condition $\partial W_{g,g}/\partial b = 0$ is therefore the same as $\partial W_{g,4}/\partial b = 0$. Unrolling the recursion from day $\min\{g,4\}$ down to day 0, the optimal bid solves

$$\sum_{k=0}^{\min\{g,4\}} T_{g,k}\left\{\gamma'_{g,k}(b)\Big[\alpha_k V_k + (1-\alpha_k)\,U_k - W_{g,k}\Big] + \alpha_k\,\gamma_{g,k}(b)\,\frac{\partial V_k}{\partial b}\right\} = 0, \tag{B.3}$$

where

$$T_{g,k} = \prod_{\ell=k}^{\min\{g,4\}} \frac{1}{1 + \gamma_{g,\ell}(b) + \lambda \bar{G}(W_{g,\ell}) + \lambda_\ell^F}.$$

The weight $T_{g,k}$ is the expected time, in days, that the bid is pending on day k, counted from when the bid is placed, or, for a bid placed before day 4, from the start of day 4. It is the chance that the bid is still pending when day k starts, the factors for $\ell > k$, times the expected time it then stays pending, the factor for $\ell = k$. The first-order condition thus adds up the two effects of a higher bid over the days on which the bid can win, weighted by how long the bid is pending on each. Forced cancellations enter both survival terms: they shorten the time the bid is pending, in $T_{g,k}$, and lower the chance that a match lasts to pickup, $\partial V_k/\partial b$. They also lower the value of a match V_k: a shock ends the match before it pays, and inside the last 48 hours it carries the penalty. In the computations, I use the recursions (B.1) and (B.2), from day 0 up to day $\min\{g, 4\}$, rather than the sum.

Equation (B.3) is the many-day version of equation (A.6) in Appendix A.6, and it leads to a many-day version of the optimal bid (3). A match pays $b - c$ only if it lasts to pickup, so its value splits into

$$V_k = \frac{\partial V_k}{\partial b}(b - c) + O_k - K_k,$$

where O_k is the expected net payoff of the outside offer for which the carrier cancels (zero if it does not), and K_k is the expected penalty it pays, for a discretionary or a forced cancellation inside the last 48 hours. Substituting into each term of equation (B.3) and solving for b gives

$$b = \sum_{k=0}^{\min\{g,4\}} \omega_{g,k}\, b_{g,k}, \qquad b_{g,k} = c + \frac{K_k + \big[W_{g,k} - (1-\alpha_k)\, U_k\big]/\alpha_k - O_k}{\partial V_k/\partial b} - \frac{\gamma_{g,k}(b)}{\gamma'_{g,k}(b)}, \tag{B.4}$$

with weights $\omega_{g,k}$ proportional to $-T_{g,k}\,\alpha_k\,\gamma'_{g,k}(b)\,\partial V_k/\partial b$, which are positive and sum to one. Each day's bid $b_{g,k}$ has the form of equation (3). The chance that the match lasts to pickup, $\partial V_k/\partial b$, plays the role of $G(r^*)$, and dividing by it puts the other terms per completed load. The expected penalty K_k plays the role of $\kappa[1 - G(r^*)]$, and it also counts the penalties of forced cancellations. The term $[W_{g,k} - (1-\alpha_k)\, U_k]/\alpha_k$ is what the carrier gives up by winning, per confirmed win: the value of keeping its bid pending, net of what a missed win leaves it. Net of the offers O_k it still takes after cancelling, it plays the role of the forgone offers $\int_0^{r^*} \pi\, dG(\pi)$. The last term is the markup. A day gets more weight when the bid is pending longer on it, when its win rate responds more to the bid, when attention is higher, and when a match is more likely to last. As in equation (3), the condition is implicit, since the right side depends on b. For a given bid, equation (B.3) gives the cost at which that bid is optimal, which Section 5

uses to recover each bidder's cost.

C Estimation Details

This appendix backs Section 5. It defines the estimation sample, states the likelihood, describes the two inputs taken from the data before the rest of the model, sets out the units and the day groups of the parameters, and gives the standard errors with the full table of estimates.

C.1 The Estimation Sample

For each lane, the data are every carrier's decisions on the lane's loads with a pickup in 2021–22 (Section 2.2): the loads it viewed and ignored, its bids and its accept-nows. I keep all carriers. Section 5.1 explains why the estimation uses only a carrier's first decisions for each pickup date. Here I define each part of the sample, list the filters and give the counts.

I group a carrier's decisions by pickup date. Which decision comes first is measured across lanes: the carrier's bids and matches on loads of other lanes with the same pickup date count too. The sample has three parts.

The first-decision window. For each carrier and pickup date, the window holds every view of the lane's loads up to and including the carrier's first bid or accept-now for that date. A bid or accept-now on any other load with the same pickup date, on any lane, also ends the window. The window thus contains the loads the carrier ignored before its first bid or accept-now, and that first bid or accept-now if it was on the lane. If the carrier never bid or accepted now for the date, all of its views are in the window. The window is set before the filters below, so a first bid that a filter removes still ends it. The decisions in the window enter the choice part of the likelihood, and the bids among them also enter the fit of the cost distribution.

The first acceptance. For each carrier and pickup date, the first acceptance is the first of the carrier's bids, on any lane, that the platform accepts. If it is a bid on the lane, its outcome, confirmed or not, enters the confirmation part of the likelihood. It need not be the carrier's first bid.

The first match. The first match is the first load the carrier is matched to for the pickup date, on any lane: a won bid that it confirms, or an accept-now. If it is on the lane, its daily decisions to cancel or keep the load enter the cancellation part of the likelihood. An accept-now in the first-decision window comes before any other bid or match for the date, so it is always the carrier's first match. A first match won with a bid need not be the first acceptance, since the carrier may have let an earlier win pass. On the lane, I look for the first acceptance and the first match among the bids and accept-nows that pass the price filters below.

All later decisions are dropped, with one exception. A bid placed after the window closes stays in the sample when it is the first acceptance or the first match. It then enters only the confirmation or the cancellation part. Its cost is recovered from the bid like that of any other bid (Appendix C.2), but it does not enter the choice part or the fit of the cost distribution.

Filters. I drop bids below 0.7 or above 2.0 times the spot rate (589 bids on Seattle to San Francisco and 1,178 on Los Angeles to Phoenix), and decisions made when the load's Accept-Now price was above 1.2 times the spot rate (624 and 305). I also drop bids below the Accept-Now price the carrier saw (83 and 140), since accepting now would have paid more, at once. Finally, I keep the actions taken in the last eight days before pickup, which the model covers. Table C.1 gives the size of each part of the sample.

Limits. The argument of Section 5.1 has four limits. First, each carrier and pickup date contributes one first-decision window, however many loads the carrier pursues for that date. A carrier that pursues many loads for a date thus has most of its decisions left out. Second, the model gives each view its own cost draw, so the cost behind a first bid is a draw like any other. In practice, a carrier may compare several loads before it bids and send its first bid to the one it likes best. Its first bid may then come from a lower cost than its decisions on other loads would.

Third, most first acceptances and most first matches won with a bid come from bids placed after the first-decision window closed: 73% of first acceptances and 73% of first matches won with a bid on Seattle to San Francisco, and 69% and 70% on Los Angeles to Phoenix. Such a bid comes after the carrier's first bid or accept-now for the date, and the first bid may still be pending. The model's carrier has no other bid, so the first-order condition that recovers the cost from such a bid holds only approximately. These first acceptances are also confirmed less often than those from bids in the window: 48.6% against 55.8% on Seattle to San Francisco, and 40.6% against 46.6% on Los Angeles to Phoenix. Fourth, a first acceptance can come while

Table C.1: The estimation sample

	Seattle–San Francisco	Los Angeles–Phoenix
Carrier decisions on the lane's loads, pickups in 2021–22	92,590	364,174
in the first-decision window, after the filters	31,748	157,066
in the estimation sample	37,327	171,691
Accepted bids	7,930	21,443
first acceptances	6,710	19,151
First matches won with a bid	4,064	9,426
days at risk of a cancellation	8,761	20,945
discretionary cancellations	180	522
First matches accepted now	700	1,262
days at risk of a cancellation	2,168	3,456
discretionary cancellations	79	115
Forced cancellations of first matches (censored)	267	663
Loads	7,825	15,078
Clusters (carrier $\times$ pickup date)	27,502	116,032

NOTE: A carrier decision is an ignored view, a bid or an accept-now. The estimation sample is the first-decision window plus the later bids kept as first acceptances or first matches. The accepted bids are those in the estimation sample; the ones that are not first acceptances were accepted after an earlier acceptance of another of the carrier's bids for the same pickup date. A day at risk is a day on which a first match could be cancelled at the carrier's discretion: from the day the carrier takes the load down to pickup, until the match ends. Loads are the distinct loads with at least one decision in the estimation sample. A cluster holds all of a carrier's decisions for one pickup date.

the carrier is matched to a load it accepted now. Such a carrier may turn the win down because it is already matched, which the model would read as low attention. The sample keeps these first acceptances; on the lane itself, they are rare.

C.2 The Likelihood

The likelihood takes the win rates $\gamma_{g,d}(b)$ and the forced-cancellation rates λ_d^F as given, at the estimates of Appendix C.3. It estimates the 21 parameters listed in Appendix C.4: the penalty κ, attention α_d, the outside offers G, the cost distributions, the logit scale σ^{choice} and the hassle costs c_{bid} and c_{AN}. I write F_d for the cost distribution of a decision made on day d, $N(\mu_d^c, \sigma_d^c)$, and f_d for its density, with the day groups of Appendix C.4. For a view on day d of a load with Accept-Now price p^{AN}, I write $\Pr_d(k \mid c, p^{AN})$ for the probability (9) of choice k at cost c: ignoring (I), accepting now (A) or bidding (B). Finally, I write $P_t(b,c)$ for the daily cancellation probability (7) of a match at price b and cost c on day t, with $\lambda = 1$,

$$P_t(b,c) = \frac{\bar{G}\big(r_t^*(b,c)\big)}{1+\bar{G}\big(r_t^*(b,c)\big)}, \qquad r_t^*(b,c) = V_t(b,c) + \kappa_t.$$

The cost behind a bid. A bid reveals the carrier's cost. For a bid b_i placed on day g_i, the inverted cost c_i is the cost at which b_i is the optimal bid: the cost that solves the first-order condition (B.3) at $b = b_i$. For short, I write β_g for the inverse of the optimal bid b_g^* of Section 4.3, so that $c_i = \beta_{g_i}(b_i)$. The inversion depends on the penalty, attention, the outside offers and the inputs taken from the data, but not on the cost distribution or on the parameters of the choice. It is redone at every trial value of the parameters. Since the bid is a function of the cost, the density of a bid is the density of its inverted cost times the derivative of the inversion, $|\beta_g'(b)|$.

The five parts. The sample of Appendix C.1 has five kinds of observations, and each contributes one part.

Bids in the first-decision window (the set $\mathcal{B}$). Each contributes the density of its inverted cost, the derivative of the inversion, and the probability of bidding at that cost:

$$\ell_i^B = \log f_{g_i}(c_i) + \log\left|\beta_{g_i}'(b_i)\right| + \log \Pr_{g_i}\left(B \mid c_i, p_i^{AN}\right). \tag{C.1}$$

In the choice probability, the value of bidding is that of the bid placed, $W_{g_i,g_i}(b_i,c_i)$, since b_i is

optimal at the cost c_i.

Accept-nows in the first-decision window (the set $\mathcal{A}$). The cost is not revealed. Each accept-now is also the carrier's first match, so its daily decisions to cancel or keep the load follow. Let $\mathcal{T}_i$ be the days on which the match is at risk, and z_{it} equal one if the carrier cancels on day t and zero if it keeps the load. The part is the probability of accepting now times the likelihood of these decisions, integrated over the cost distribution of the day:

$$\ell_i^A = \log \int \Pr_{d_i}\left(A \mid c, p_i^{AN}\right) \prod_{t \in \mathcal{T}_i} P_t\big(p_i^{AN}, c\big)^{z_{it}} \Big[1 - P_t\big(p_i^{AN}, c\big)\Big]^{1 - z_{it}} \, dF_{d_i}(c). \tag{C.2}$$

The product sits inside the integral because the same cost drives the choice to accept now and the later decisions to keep the load. A carrier with a low cost is both more likely to accept now and less likely to cancel.

Ignored loads (the set $\mathcal{I}$). The probability of ignoring the load is integrated over the cost distribution of the day:

$$\ell_i^I = \log \int \Pr_{d_i}\left(I \mid c, p_i^{AN}\right) dF_{d_i}(c). \tag{C.3}$$

First acceptances (the set $\mathcal{C}$). Let y_j equal one if the carrier confirms the bid. The confirmation probability (8) is evaluated at the bid b_j and its inverted cost c_j, for the bid's own day d_j^b and acceptance day d_j^a, with the win rates and values of bids placed on day d_j^b (Appendix B.3):

$$\ell_j^C = y_j \log \Pr\left(\text{confirm} \mid b_j, c_j\right) + (1 - y_j) \log \Big[1 - \Pr\left(\text{confirm} \mid b_j, c_j\right)\Big]. \tag{C.4}$$

First matches won with a bid (the set $\mathcal{H}$). The daily decisions to cancel or keep the load are evaluated at the bid and its inverted cost:

$$\ell_j^H = \sum_{t \in \mathcal{T}_j} \Big\{ z_{jt} \log P_t\big(b_j, c_j\big) + (1 - z_{jt}) \log \Big[1 - P_t\big(b_j, c_j\big)\Big] \Big\}. \tag{C.5}$$

Each decision has its own cost draw, so the parts add up:

$$\ell(\theta) = \sum_{i \in \mathcal{B}} \ell_i^B + \sum_{i \in \mathcal{A}} \ell_i^A + \sum_{i \in \mathcal{I}} \ell_i^I + \sum_{j \in \mathcal{C}} \ell_j^C + \sum_{j \in \mathcal{H}} \ell_j^H. \tag{C.6}$$

A bid can appear in up to three parts: as a bid in the window, as a first acceptance and as a first match. These parts share its cost, but the inversion recovers that cost, so they need no

integral.

Days at risk and censoring. A match is at risk from the day the carrier takes the load, the day its bid is accepted or the day it accepts now, down to the pickup day. On the day of a discretionary cancellation, $z_{jt} = 1$, and the match leaves the sample. A match that ends for any other reason, a forced cancellation or a cancellation by the shipper for instance, also leaves the sample, and the day on which it ends is left out: it is censored. For forced cancellations, this follows from the model. On a day without a forced cancellation, the chance of a discretionary one is equation (7), whatever the forced rate (Appendix B.1). Censoring the days with a forced cancellation therefore gives the model's likelihood of the discretionary ones, given that no forced shock arrives. For an accept-now, whose cost is integrated out, this holds only approximately, since the chance of escaping a forced shock varies slightly with the cost. The forced rate still enters through the value of the match, and the forced cancellations themselves give the forced rates (Appendix C.3).

Integration. I compute the integrals in equations (C.2) and (C.3) as averages over 2,000 scrambled Halton draws from the standard normal distribution, the same for every decision and scaled to the cost distribution of its day, with the values that enter them computed on grids of costs and Accept-Now prices and interpolated between grid points.

C.3 Win Probabilities and Forced-Cancellation Rates

Two inputs of the likelihood come straight from the data: the win rates of bids and the daily rates of forced cancellations. I estimate them first and hold them fixed when I maximize the likelihood. Appendix C.5 adds their sampling error to the standard errors of the other parameters.

Win probabilities. The bid's first-order condition (B.3) involves the win rates $\gamma_{g,d}(b)$ and their derivatives. In the counterfactuals, the win rates follow from carriers' bidding as an equilibrium (Section 4.4). Solving for that equilibrium at every evaluation of the likelihood would be very slow. As in Guerre, Perrigne, and Vuong (2000), I instead estimate the win probabilities that bidders face directly from the bids.

The sample is every bid on the lane's loads with a pickup in 2021–22, placed in the last eight

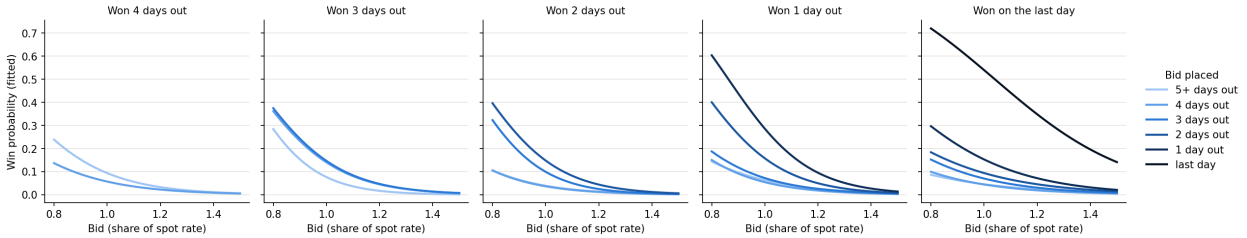
days before pickup, after the same trims of bids and Accept-Now prices as the estimation sample (Appendix C.1). Unlike the likelihood, it uses every bid, not only first decisions, since every bid competes in the auction. For the same reason, it keeps the bids below the Accept-Now price. I group bids by the day g on which they were placed: 0, 1, 2, 3, 4, and 5 or more days before pickup. A bid is at risk of winning on each day d from the day it is placed down to the pickup day, until it wins, whatever else happens to the load or to the carrier. This matches the model, in which a bid can still win after its carrier has left (Appendix B.1). For each bid day g and each day $d \leq \min\{g, 4\}$ on which the bid can win, I fit a logistic curve in the bid to the chance that a pending bid wins that day:

$$w_{g,d}(b) = \frac{1}{1 + \exp\left[(b - m_{g,d})/s_{g,d}\right]}. \tag{C.7}$$

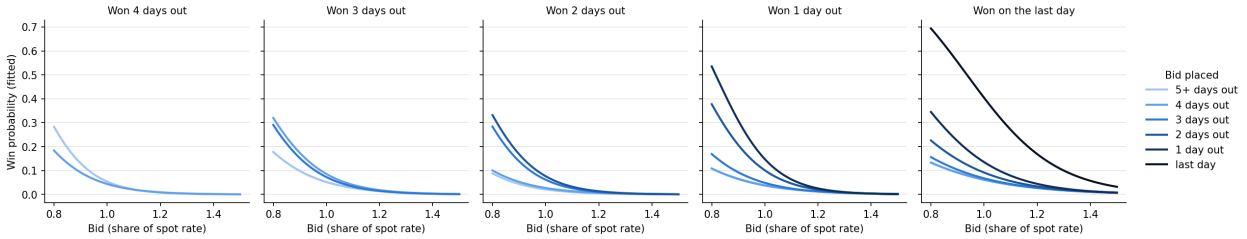
The location $m_{g,d}$ is the bid at which the chance of winning is one half, and the scale $s_{g,d}$ sets how fast the chance falls as the bid rises. I fit the 20 cells by maximum likelihood, one at a time. Figure C.1 shows the fitted curves for both lanes. A bid placed close to pickup wins far more often on the last days than an earlier bid at the same price, so the day the bid is placed matters as much as the bid.

Figure C.1: Fitted win probabilities, by the day the bid is placed and the day it wins

(a) Seattle to San Francisco



(b) Los Angeles to Phoenix



NOTE: The chance that a pending bid wins on the day named in the panel title, $w_{g,d}(b)$ of equation (C.7), against the bid in shares of the spot rate, with one line for each day g on which the bid was placed. Bids placed five or more days before pickup are pooled.

The model needs a rate rather than a probability. In a race with the end of the day, a pending bid with win rate γ wins before the day ends with probability $\gamma/(1+\gamma)$ (Appendix B.3). Setting

this equal to $w_{g,d}(b)$ gives

$$\gamma_{g,d}(b) = \frac{w_{g,d}(b)}{1 - w_{g,d}(b)} = \exp\Big(\frac{m_{g,d} - b}{s_{g,d}}\Big), \qquad \gamma'_{g,d}(b) = -\frac{\gamma_{g,d}(b)}{s_{g,d}}. \tag{C.8}$$

The win rate thus falls with the bid at a constant proportional rate. In the optimal bid (B.4), each day's markup, $-\gamma_{g,k}/\gamma'_{g,k}$, equals the scale $s_{g,k}$.

I compute the standard errors of each cell with a sandwich estimator clustered by load. A load's bids enter several cells, so I also cluster by load the joint covariance of all 40 numbers, the location and the scale of each cell. Appendix C.5 carries this covariance into the variance of the other estimates.

Forced-cancellation rates. In the model, a matched carrier's position on day d changes when the day ends, at rate one, when an offer above the cutoff arrives, at rate $\lambda\bar{G}(r^*_d)$, or when a forced shock hits, at rate λ^F_d (Appendix B.1). The shock comes first with probability $\lambda^F_d/[1 + \lambda\bar{G}(r^*_d) + \lambda^F_d]$, and the day ends first with probability $1/[1 + \lambda\bar{G}(r^*_d) + \lambda^F_d]$. The ratio of the two is λ^F_d, whatever the offers. I therefore estimate the forced rate on day d as

$$\hat{\lambda}^F_d = \frac{\text{forced cancellations of first matches on day } d}{\text{days } d \text{ on which a first match was kept}}, \tag{C.9}$$

over the first matches of the estimation sample, won with a bid or accepted now. A day counts as kept when the match lasts through it. A day with a discretionary cancellation, or on which the match ends for another reason, enters neither count. Forced cancellations are the carrier-fault cancellations other than the discretionary ones (Section 3.2). I pool all days four or more before pickup into one rate.

On Seattle to San Francisco, the estimated daily rates are 4.03% on the pickup day, 2.55% on the day before, and 1.47%, 1.16% and 1.02% two, three, and four or more days before pickup, from 267 forced cancellations. On Los Angeles to Phoenix, they are 4.82%, 2.01%, 2.07%, 1.32% and 0.35%, from 663 forced cancellations. The first matches belong to the same clusters as the likelihood, so Appendix C.5 carries each cluster's influence on these rates into the standard errors.

C.4 Units, Days and Parameters

As everywhere in the paper, prices, costs, penalties and payoffs are in units of the load's spot rate (Section 2.2). I divide every price in the data by the load's spot rate before estimation, so the estimates are shares of it. A day d before pickup is a full 24-hour period counted back from the pickup time: $d = 0$ is the last 24 hours. The model covers days 0 to 7.

I estimate 21 parameters by maximum likelihood. The penalty κ is a single number, perceived on days 0 and 1, the communicated 48 hours, and zero before. Attention takes one value for each of days 0 to 4, the days on which a bid can win: $\alpha_0, \ldots, \alpha_4$. The outside offers have one distribution, $G = N(\mu^\pi, \sigma^\pi)$, the same on every day, with one offer a day on average. The cost distribution has a mean and a standard deviation for each of days 0 to 3, and one pair for all views four or more days before pickup, written μ^c_{4+} and σ^c_{4+}. The choice among ignoring, accepting now and bidding adds the logit scale σ^{choice} and the hassle costs c_{bid} and c_{AN}. The forced-cancellation rate, taken from the data, uses the same day groups as the costs. The win probabilities have their own cells, by the day a bid is placed and the day it wins (Appendix C.3).

I estimate the logarithm of κ and, for each day, a logistic transformation of α_d, so that the penalty is positive and attention lies between zero and one. I compute the cost inversion and the values on grids of bids, costs and Accept-Now prices, interpolating between grid points, and I maximize the likelihood with the Nelder–Mead algorithm, followed by a quasi-Newton (BFGS) search from its optimum.

C.5 Standard Errors and the Full Estimates

I compute standard errors with a sandwich estimator clustered by carrier and pickup date, the unit at which the sample is built: a cluster holds all of a carrier's decisions for one pickup date. Let H be the Hessian of the log-likelihood (C.6) at the estimates and s_k the score of cluster k. The two inputs taken from the data add their own sampling error, which I carry through as in two-step estimation (Newey and McFadden 1994, Section 6). Each input moves the estimates through the derivative of the score with respect to that input: $H_{\theta\lambda}$ for the forced-cancellation rates and $H_{\theta\gamma}$ for the win probabilities. The forced rates come from the first matches, which belong to the same clusters, so I add each cluster's influence φ_k on them to its score. The win probabilities come from every bid on the lane, and I take their errors, with the load-clustered covariance V_γ of Appendix C.3, to be independent of the scores. The variance of the estimates

is then

$$\hat{V} = H^{-1}\Big[\sum_{k}\big(s_k + H_{\theta\lambda}\varphi_k\big)\big(s_k + H_{\theta\lambda}\varphi_k\big)'\Big]H^{-1} + H^{-1}H_{\theta\gamma}V_{\gamma}H'_{\theta\gamma}H^{-1}.$$

I compute all derivatives by central finite differences, with the values solved to a tight tolerance, so that the likelihood is smooth enough to differentiate twice. The standard errors of κ, of attention and of the derived quantities follow by the delta method.

Table C.2 reports all 21 parameters for both lanes, with their standard errors, and the derived quantities behind Table 4. Attention α_d is the chance that a carrier notices a win on day d, and 4+ pools all views four or more days before pickup. The derived quantities are functions of the parameters. $P(\pi > 0)$ is the daily chance of an offer with a positive net payoff, and $E\max(\pi, 0)$ the expected positive part of an offer. A free carrier takes an offer only if it beats the value of staying free, U_d, which is above zero, so fewer offers than $P(\pi > 0)$ are worth taking for it. The value of a typical match, $\hat{y}$, is the median value $V_d(b, c)$ of a first match won with a bid, at its price and its inverted cost, over the days 2 to 4 before pickup on which the match could be cancelled. $P(\pi > \hat{y})$ is the daily chance of an offer that beats this value before the last 48 hours, and $P(\pi > \hat{y} + \kappa)$ the same chance inside the last 48 hours, where cancelling also costs the penalty. The share of these cancellations the penalty prevents is one minus their ratio. The condition number is that of the Hessian H.

Table C.2: Estimates of the carrier model, all parameters

	Seattle–San Francisco		Los Angeles–Phoenix	
	Estimate	SE	Estimate	SE
Penalty and attention				
Perceived penalty, last 48 hours, κ	0.098	0.036	0.081	0.010
Attention on the pickup day, α_0	0.482	0.018	0.505	0.014
Attention 1 day before pickup, α_1	0.584	0.016	0.577	0.011
Attention 2 days before pickup, α_2	0.634	0.019	0.638	0.013
Attention 3 days before pickup, α_3	0.647	0.018	0.626	0.013
Attention 4 days before pickup, α_4	0.603	0.028	0.596	0.016
Outside offers, net payoff $\pi \sim N(\mu^\pi, \sigma^\pi)$				
Mean, μ^π	−0.580	0.207	−0.061	0.022
Standard deviation, σ^π	0.450	0.114	0.178	0.013
Cost of hauling the load, $c \sim N(\mu_d^c, \sigma_d^c)$*, by day of the view*				
Mean, on the pickup day, μ_0^c	0.879	0.036	0.956	0.009
Mean, 1 day before pickup, μ_1^c	0.940	0.011	0.915	0.005
Mean, 2 days before pickup, μ_2^c	0.897	0.011	0.871	0.006
Mean, 3 days before pickup, μ_3^c	0.887	0.011	0.847	0.007
Mean, 4+ days before pickup, μ_{4+}^c	0.835	0.011	0.802	0.006
Standard deviation, on the pickup day, σ_0^c	0.315	0.009	0.254	0.004
Standard deviation, 1 day before pickup, σ_1^c	0.238	0.005	0.200	0.003
Standard deviation, 2 days before pickup, σ_2^c	0.220	0.006	0.190	0.003
Standard deviation, 3 days before pickup, σ_3^c	0.223	0.006	0.188	0.003
Standard deviation, 4+ days before pickup, σ_{4+}^c	0.208	0.005	0.177	0.002
Choice at the view				
Logit scale, σ^{choice}	0.189	0.036	0.091	0.005
Hassle cost of bidding, c_{bid}	0.225	0.038	0.159	0.008
Hassle cost of accepting now, c_{AN}	0.692	0.116	0.462	0.020
Derived quantities				
Hassle cost of bidding / scale, c_{bid}/σ^{choice}	1.191	0.028	1.759	0.010
Hassle cost of accepting now / scale, c_{AN}/σ^{choice}	3.659	0.097	5.102	0.072
Offer with a positive net payoff, $P(\pi > 0)$	0.099	0.025	0.365	0.037
Expected positive part of an offer, $E\max(\pi, 0)$	0.021	0.002	0.044	0.004
Value of a typical match, $\hat{y}$	0.243	0.004	0.224	0.003
Offer that beats a typical match:				
before the last 48 hours, $P(\pi > \hat{y})$	0.034	0.003	0.054	0.003
in the last 48 hours, $P(\pi > \hat{y} + \kappa)$	0.020	0.002	0.019	0.001
Share of these cancellations the penalty prevents	0.397	0.077	0.639	0.032
Log-likelihood	−24,742.5		−83,829.8	
Clusters (carrier × pickup date)	27,502		116,032	
Condition number of the Hessian	4.4×10^4		1.1×10^5	

NOTE: Maximum likelihood estimates on each lane's estimation sample (Appendix C.1), in units of the load's spot rate. Standard errors are clustered by carrier and pickup date and include the sampling error of the win probabilities and the forced-cancellation rates; those of the derived quantities are by the delta method. The text above defines the derived quantities.

D Platform Parameters

D.1 Platform Parameters Estimation

This section describes the estimation of the remaining platform parameters, namely:

- The arrival rates of carriers, shipments, views, shipper cancellations, and auction clearing rounds.
- The distribution of shipment values to the platform.
- The conditional reserve and Accept-Now prices.

For the purposes of this section, let m denote each separate *market*, defined as a unique combination of origin and destination, and pickup date.

Arrival Rates We begin by estimating the exponential rates of the three event types which can be simply counted: carriers, shipments, and auction clearing rounds. For each of these events $e \in \{carrier, shipment, clear\}$, let $N^e_{m,d}$ denote the number of events of type e in market $m \in 1, \ldots, M$ on day d from departure. As the events are assumed to follow a Poisson process, the count of events follow a Poisson distribution, so that:

$$\mathbb{E}[N^e_{m,d}] = \lambda^e_d.$$

Thus, a simple method-of-moments estimator for each of these arrival rates is the sample mean of the counts:

$$\hat{\lambda}^e_d = \frac{1}{M} \sum_{m=1}^{M} N^e_{m,d}$$

with the asymptotic variance given by the sample variance of the counts divided by M.

Shipment cancellations and views have rates that are specified on a per-shipment and per-pair basis, and can only occur once, respectively. Let $Y^{shipcancel}_{j,m,d} \in \{0,1\}$ denote the event that shipment j in market m on day d is cancelled, and $Y^{view}_{ij,m,d} \in \{0,1\}$ denote the event that shipment j in market m on day d is viewed by carrier i. Given the properties of the Poisson

process, the probabilities of these events are

$$P(Y_{j,m,d}^{shipcancel} = 1) = \frac{\lambda_d^{shipcancel}}{\eta + \lambda_d^{shipcancel}} \equiv p_d^{shipcancel}$$
$$P(Y_{ij,m,d}^{view} = 1) = \frac{\lambda_d^{view}}{\eta + \lambda_d^{view}} \equiv p_d^{view}$$

where η is the rate of transition between days, normalized to 1. The sample analogues of these probabilities are the sample proportions of cancellations and views for each market and day, among the set of *uncancelled* shipments and *unmatched* pairs, respectively. These are used to estimate the arrival rates of cancellations and views:

$$\hat{\lambda}_d^{shipcancel} = \frac{\hat{p}_d^{shipcancel}}{1 - \hat{p}_d^{shipcancel}}$$
$$\hat{\lambda}_d^{view} = \frac{\hat{p}_d^{view}}{1 - \hat{p}_d^{view}}$$

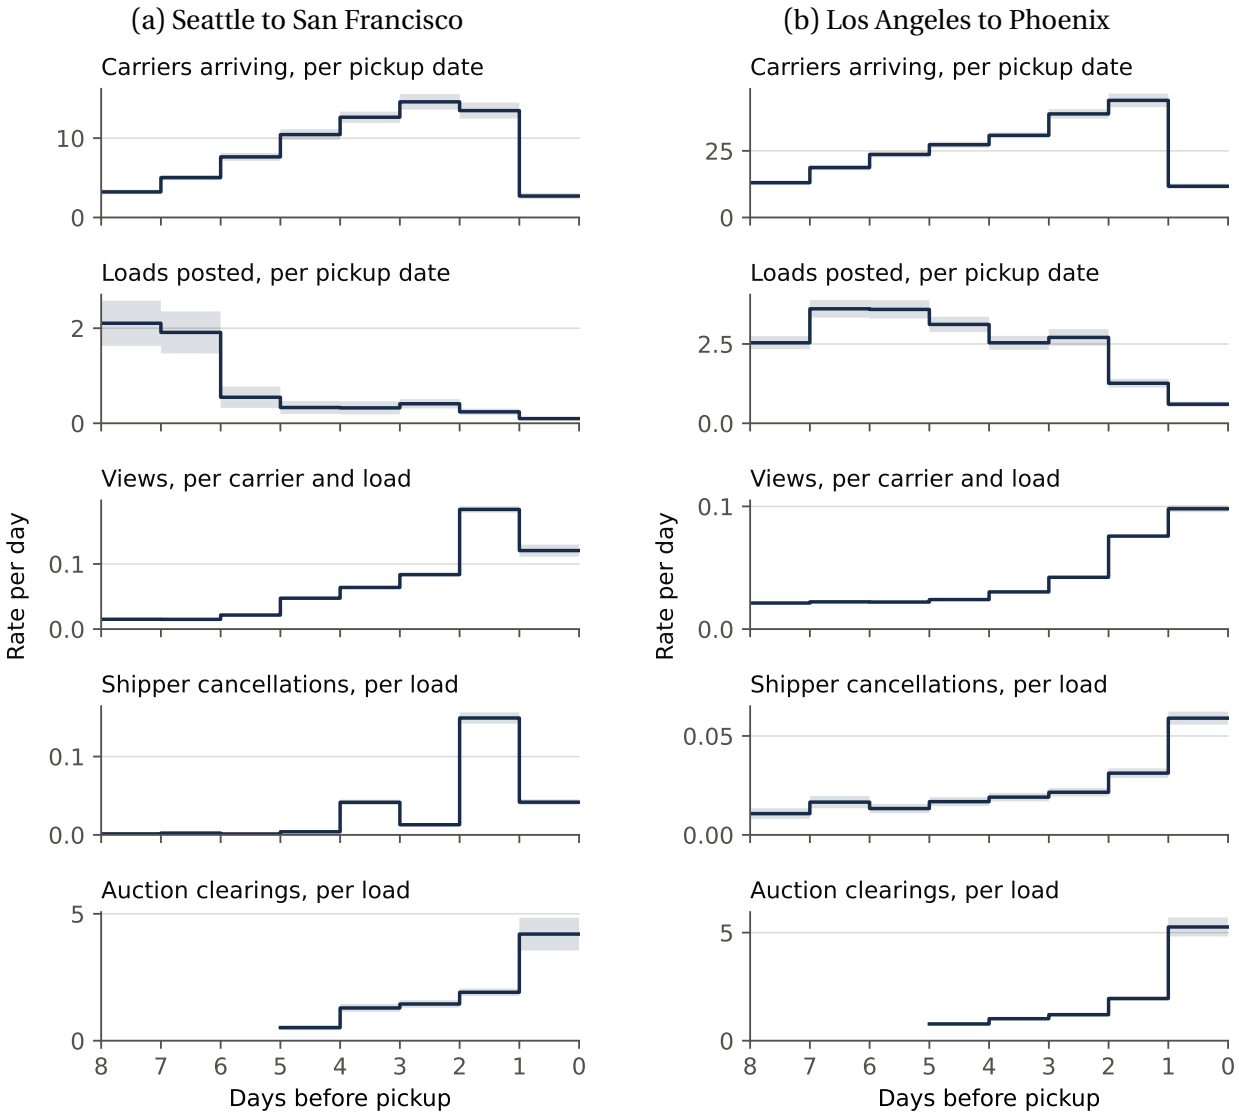
with the asymptotic variances of the estimates obtained by multiplying the variances of the sample proportions by the square of the derivative of the transformation function.

There are many arrival rates, so Figure D.1 reports them graphically, with their confidence intervals, for both lanes.

Shipment Values The shipment values to the platform are directly observed. I fit a log-normal distribution to the data by simply taking the sample mean $\hat{\mu}$ and sample variance $\hat{\sigma}^2$ of the log-transformed shipment values, reported in Table D.1. In addition, a load still unmatched at pickup costs the platform a share L of its value, since the platform then has to ask the shipper to reschedule or cancel the shipment, with consequences for its relationship with the shipper. I measure L as the average amount by which the pickup-day reserve price exceeds the shipment value, where the pickup-day reserve price is the typical highest bid the platform accepts on that day. Since the platform pays up to this price on the pickup day rather than leave the load unmatched, I take this premium as the cost of not matching it. L is 0.158 on Seattle to San Francisco and 0.040 on Los Angeles to Phoenix, with standard errors of 0.008 and 0.005.

Reserve and Accept-Now Prices The conditional Accept-Now price function $p^{AN}(v_j)$ and conditional reserve price functions $r_d(v_j)$ and are fitted through second-order polynomials.

Figure D.1: Estimated arrival rates of the platform's events



NOTE: Estimated rate per day of each event, by days before pickup, with pointwise 95% confidence intervals (shaded). Each rate is constant within a day, so day d is drawn from $d+1$ to d days before pickup. Carriers arriving: carriers per pickup date whose first view of the lane for that date falls on the day. Loads posted: loads per pickup date posted on the day. Views: the rate at which a carrier views a given load it has not yet viewed. Shipper cancellations: the rate at which a load is cancelled other than by its carrier. Auction clearings: the rate at which the auction clears a load with a bid at or below the reserve price, on the last five days before pickup only.

Table D.1: Shipment value distribution parameters

Parameter	Estimate	Standard Error
μ	-0.008	0.002
σ^2	0.021	0.0003

The Accept-Now price function is estimated through a simple linear regression:

$$p^{AN}(v_j) = \theta_0^{AN} + \theta_1^{AN} v_j + \theta_2^{AN} v_j^2 + \epsilon_j$$

Reserve prices up until the last 24 hours before pickup, namely $r_d(v_j)$ for $d \in \{1, 2, 3, 4\}$, are estimated through a similar regression:

$$r_d(v_j) = \theta_{0d}^{R} + \theta_1^{R} v_j + \theta_2^{R} v_j^2 + \epsilon_j$$

where the intercept term θ_{0d}^R is specific to the number of remaining days until departure d. This allows for the reserve price to increase as the departure date approaches, and avoids any potential "crossing" of the reserve price functions across different days.

The effective reserve price for the last 24 hours at $d = 0$, $r_0(v_j)$, is less straightforward to estimate as there is no unique observed reserve price during this period. We utilize a moment-inequality approach. Let $\mathcal{J}$ denote the set of shipments unmatched at the start of the last 24 hours, and $\mathcal{J}^U \subseteq \mathcal{J}$ those still unmatched at the end. For a shipment j, the theoretical reserve price must be bounded below by the highest accepted bid (if any) and above by the lowest rejected bid (if the shipment remained unmatched). We define these bounds as:

$$\underline{b}_j \equiv \max_{i \in \mathcal{I}_j^A} b_{ij} \quad \text{and} \quad \overline{b}_j \equiv \min_{i \in \mathcal{I}_j \setminus \mathcal{I}_j^A} b_{ij},$$

where $\underline{b}_j = -\infty$ if no bid was accepted and $\overline{b}_j = +\infty$ if no rejected bid exists to form the upper bound. The model implies the moment inequalities $r_0(v_j) \geq \underline{b}_j$ for all $j \in \mathcal{J}$ and $r_0(v_j) \leq \overline{b}_j$ for all $j \in \mathcal{J}^U$.

To estimate θ^{R_0}, we minimize a positive-part quadratic objective function that penalizes violations of these bounds:

$$\mathcal{Q}_N^{\text{ineq}}(\theta^{R_0}) = \frac{1}{N} \sum_{j \in \mathcal{J}} \left[r_0(v_j, \theta^{R_0}) - \underline{b}_j \right]_-^2 + \frac{1}{N} \sum_{j \in \mathcal{J}^U} \left[\overline{b}_j - r_0(v_j, \theta^{R_0}) \right]_-^2$$

where $[x]_- \equiv \min\{x, 0\}$. Standard errors are computed by treating the estimator as first-order equivalent to a just-identified GMM estimator formed by the set of binding inequalities. The variance-covariance matrix is estimated using a cluster-robust plug-in estimator, allowing for correlations within market-day clusters.

The results of the OLS regressions ($d \geq 1$) and the inequality estimator ($d = 0$) are reported

together in Table D.2.

E Counterfactual Methods and Further Results

This appendix backs Section 6. It describes how I solve each counterfactual, defines welfare under the two accountings and the three measures of what a cancellation costs, and reports the exhibits for Los Angeles to Phoenix, the details of the tax on matching, the timing of the penalty and the curbs on a monetary penalty.

E.1 Simulation and Equilibrium

In the main counterfactuals, the penalty is $\kappa_d = \kappa$ on days $d \in \{0, 1\}$, the last 48 hours, and zero on earlier days. The grid has 72 values of κ, evenly spaced from 0 to 0.24. The penalty windows of Section 6.4 set $\kappa_d = \kappa$ on day 0 only (24 hours) or on days 0 to 2 (72 hours).

Carriers' choices depend on the win rates $\gamma_d(b)$, and the win rates depend on carriers' choices. At each penalty, I solve for win rates that reproduce themselves, by iteration:

1. Given the current win rates, solve the carrier's problem of Section 4: the values, the optimal bids and the choice probabilities.
2. Simulate markets with this policy. In each market, place artificial bids at a grid of prices on every load that is not matched, and keep each at risk until it would win or the load reaches pickup.
3. Refit the win rates to the artificial bids, by the day the bid is placed and the day it wins, as in the data (Appendix C.3).
4. Move the current win rates part of the way toward the refit ones.

The number of markets grows over the first iterations and then stays fixed, so later refits carry no new simulation noise. The iteration stops when no daily win probability, over the usual range of bids, changes by more than 0.0005. A final simulation of 36,200 markets under the last policy gives the outcomes. Every grid point converged, on both lanes.

Every penalty simulates the same markets with the same random numbers. Each market draws its choices, costs, load values and events from separate streams. Events arrive at rates set by the number of carriers and loads that ever entered the market, which does not depend on the penalty, and an event that falls on a carrier or load that has left does nothing. So out-

Table D.2: Reserve and Accept-Now price estimates

	Seattle–San Francisco		Los Angeles–Phoenix	
	Reserve	Accept-Now	Reserve	Accept-Now
Day 0	1.212	0.975	1.086	0.899
	(0.008)	(0.007)	(0.004)	(0.002)
Day 1	1.065	0.858	0.997	0.829
	(0.003)	(0.002)	(0.002)	(0.002)
Day 2	1.015	0.853	0.971	0.821
	(0.003)	(0.001)	(0.001)	(0.001)
Day 3	1.000	0.858	0.964	0.819
	(0.002)	(0.002)	(0.001)	(0.001)
Day 4	0.985	0.857	0.945	0.818
	(0.003)	(0.001)	(0.002)	(0.001)
Day 5		0.862		0.822
		(0.002)		(0.001)
Day 6		0.871		0.827
		(0.002)		(0.002)
Day 7		0.868		0.834
		(0.003)		(0.002)
Unmatched load, L	0.158		0.040	
	(0.008)		(0.005)	
Log forecast / spot, mean	−0.067		−0.057	
standard deviation	0.116		0.121	
Listing-days	23,296	26,239	35,390	51,704
with an acceptance	5,215		11,604	
Listings, forecast	9,333		16,490	

NOTE: Prices as multiples of the platform's price forecast for the load, by day before pickup (day 0 is the last 24 hours), with bootstrap standard errors over listings in parentheses. Reserve: the median of the highest bid the platform accepted that day, over the listings on the market at the start of the day that accepted one; the auction clears bids only on days 0 to 4. Accept-Now: the median Accept-Now price on carriers' views that day. L: the cost of a load still unmatched at pickup, as a share of its value. Log forecast / spot: the log of the forecast over the spot rate across the lane's listings. Listing-days: listings on the market at the start of a day (reserve) or viewed that day (Accept-Now).

comes at two penalties differ only because of the penalty, and I compare them market by market. In particular, the same loads reach pickup at every penalty. Carriers' payoffs without the platform come from a run with the same carriers and events but no loads.

E.2 The Two Accountings and Welfare

At pickup, the platform earns on load j

$$\pi_j = \begin{cases} v_j - p_j & \text{if the load is matched at price } p_j, \\ -L\, v_j & \text{if it is unmatched,} \\ 0 & \text{if the shipper has cancelled it,} \end{cases}$$

where v_j is what the shipper pays and L is the cost of not matching it, as a share of v_j. Let $\Pi(\kappa) = \sum_j \pi_j$ be the platform's profit per market before any monetary penalties, $F(\kappa)$ the penalties carriers bear, $U(\kappa)$ carriers' payoffs net of those penalties, and U^0 carriers' payoffs without the platform. Carriers' surplus is $U(\kappa) - U^0$. With a reputational penalty, nobody receives $F(\kappa)$; with a monetary penalty, the platform does. Welfare and the platform's profit are

$$\begin{aligned} W^{\text{rep}}(\kappa) &= \big[U(\kappa) - U^0\big] + \Pi(\kappa), && \text{profit } \Pi(\kappa), && \text{with a reputational penalty,} \\ W^{\text{mon}}(\kappa) &= W^{\text{rep}}(\kappa) + F(\kappa), && \text{profit } \Pi(\kappa) + F(\kappa), && \text{with a monetary penalty.} \end{aligned}$$

Carriers behave the same in both, so U, Π and F are the same. Changes from $\kappa = 0$ have standard errors paired by market: the standard deviation across markets of each market's change, over the square root of the number of markets. The profit-maximizing monetary penalty is the grid point where $\Pi + F$ is highest. The penalties within one standard error of it are the grid points whose profit falls short of the best by less than the paired standard error of that shortfall.

Over the whole grid, up to 0.24 of the spot rate, welfare falls at every grid step under both accountings, on both lanes. The platform's profit with a reputational penalty peaks at no penalty on both lanes.

E.3 Twin Markets and the Three Measures

Take a cancellation e in the last 48 hours, by carrier i, of its match on load j. Its twin market undoes only the event behind it:

- an outside offer: the carrier turns the offer down and gets nothing from it;
- a switch to another load taken at the Accept-Now price: the carrier ignores that load, which stays open to others;
- a switch to another load won in its auction: the carrier does not take it, its bid is lost, and the load goes to the next bid;
- a forced cancellation: the forced event does not happen.

From then on, both markets use the same random numbers, and every carrier follows its policy in each. Later cancellations happen in both markets but get no twins of their own. A superscript "tw" marks outcomes in the twin; unmarked outcomes are in the market with the cancellation. Let f_e be the cancellation's own monetary penalty, F the penalties over the whole market, and U_h carrier h's payoff net of the penalties it pays. The three measures of what the cancellation costs are

$$\begin{aligned} D_e^B &= \pi_j^{\text{tw}} - \pi_j && \text{bilateral damages,} \\ D_e^P &= \left(\Pi^{\text{tw}} + F^{\text{tw}}\right) - \left(\Pi + F - f_e\right) && \text{platform-wide damages,} \\ C_e^M &= D_e^P + \textstyle\sum_{h \neq i} \left(U_h^{\text{tw}} - U_h\right) && \text{market-wide cost.} \end{aligned}$$

A positive value means the cancellation lowers the payoff. Platform-wide damages leave out the cancellation's own penalty, so they measure what that penalty would have to cover. In the market-wide cost, the penalties other carriers pay appear in both F and U_h, so they count as transfers. Bilateral damages count the platform's profit on the load left without the penalties that later cancellations of the same load may bring in; with no penalty, no payments flow, and this makes no difference. Let $g_e = U_i + f_e - U_i^{\text{tw}}$ be the cancelling carrier's gain before its own penalty. Then, with a monetary penalty, the change in welfare that the cancellation causes is $g_e - C_e^M$, for every cancellation, up to rounding.

Bilateral damages split into three parts,

$$D_e^B = \underbrace{\left(D_e^B - D_e^P\right)}_{\text{recovered on the platform's other loads}} + \underbrace{\textstyle\sum_{h \neq i} \left(U_h - U_h^{\text{tw}}\right)}_{\text{gained by other carriers}} + \underbrace{C_e^M}_{\text{market-wide cost}} .$$

With no penalty, the first part is what the platform gains on its other loads of the same lane

and pickup date. The two brackets in Figure 6 and Figure E.1 show the first two parts; the third is the market-wide cost curve at no penalty.

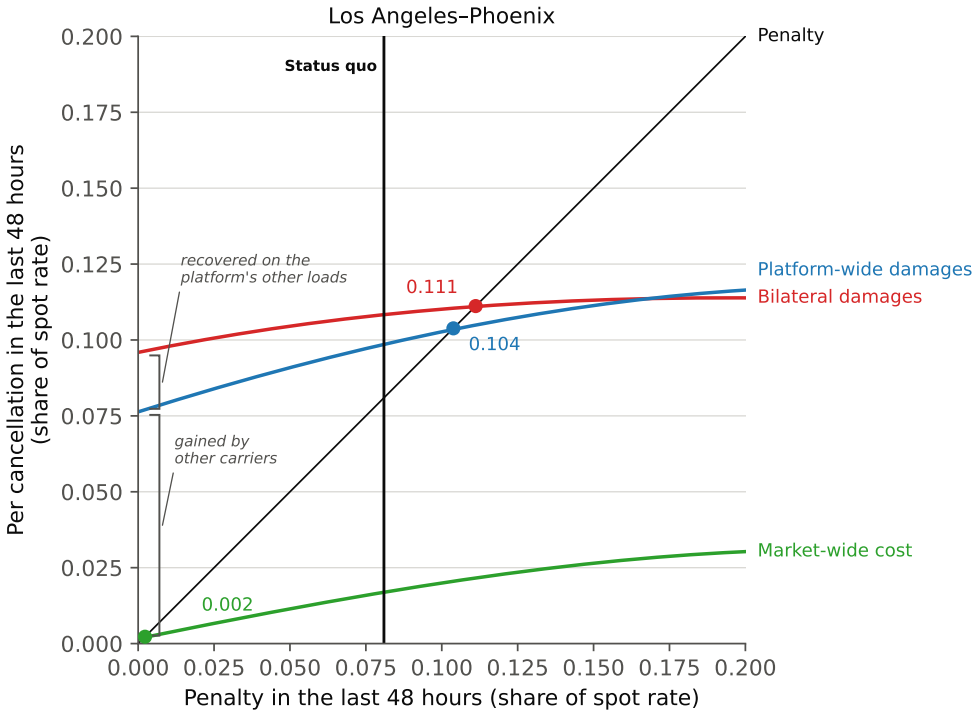
Let $E(\kappa)$ be the cancellations in the last 48 hours at penalty κ, forced and discretionary. The average $\bar{X}(\kappa)$ of each measure $X \in \{D^B, D^P, C^M\}$ weighs every cancellation in $E(\kappa)$ equally. Cancellations in the same market are separate experiments, so I average them and never add them up. The benchmark is the first κ^* at which $\bar{X}(\kappa^*) = \kappa^*$, by linear interpolation between grid points. For 95% intervals, I resample whole markets 1,000 times, with the same resamples at every grid point and for every measure, and take each resample's first crossing. On Seattle to San Francisco, the benchmarks are 0.151 [0.148, 0.153] for bilateral damages, 0.119 [0.115, 0.122] for platform-wide damages and 0.012 [0.008, 0.016] for the market-wide cost. On Los Angeles to Phoenix, they are 0.111 [0.109, 0.114], 0.104 [0.101, 0.107] and 0.002 [0.000, 0.005].

The measures rise with the penalty mainly because the load left is rematched less often, while in the twin it reaches pickup more often. In the twin, the carrier who keeps the load is less likely to cancel it later. At the estimated penalty on Seattle to San Francisco, the share of cancelled loads matched at pickup changes by -1.2 percentage points [-1.6, -0.8] from its share with no penalty, and the twin's share by 1.2 [1.0, 1.5]. On Los Angeles to Phoenix, the changes are -2.1 [-2.4, -1.7] and 2.2 [1.9, 2.5]. Part of these changes could come from which cancellations still happen. To hold that mix fixed, I group the cancellations by cause (the four of the list above) and by day before pickup (the last 24 hours, or the 24 hours before), and weigh each group's share at κ by the group's share of cancellations at $\kappa = 0$. With the mix fixed, the changes are -1.4 and 0.6 points on Seattle to San Francisco, and -1.5 and 1.1 on Los Angeles to Phoenix. So at least half of both changes does not come from the mix. The intervals are 95% intervals of the change from no penalty, resampling whole markets.

E.4 Los Angeles to Phoenix

Figure E.1, Figure E.2 and Figure E.3 repeat Figures 6 to 8 for Los Angeles to Phoenix, and tell the same story. With no penalty, bilateral damages of 0.096 split into 0.019 recovered on the platform's other loads, 0.075 gained by other carriers, and a market-wide cost of 0.002. Welfare peaks at no penalty under both accountings, and so does the platform's profit with a reputational penalty. The profit-maximizing monetary penalty is zero. At the estimated penalty, 0.081 of the spot rate, the changes per market add up. Matches made, leaving out those that shippers cancel, change by -1.12. Carrier cancellations change by -0.74. So completed matches change by -0.38.

Figure E.1: What a cancellation costs, at every penalty, Los Angeles to Phoenix



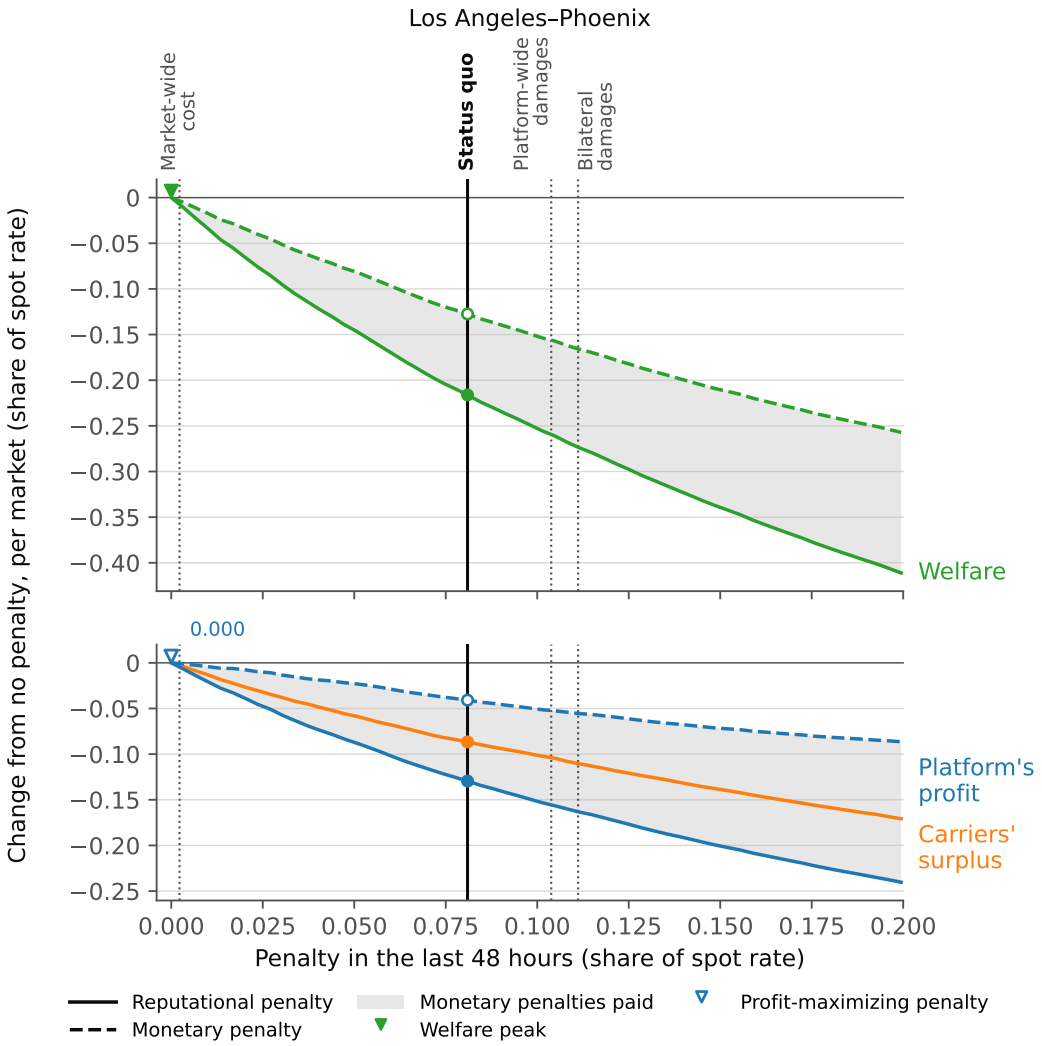
NOTE: As Figure 6, for Los Angeles to Phoenix. Appendix E.3 gives the 95% intervals of the benchmarks.

E.5 The Tax on Matching

Section 6.3 compares what a small penalty saves by deterring cancellations with what it costs by taxing matching. Both use the first grid step, from $\kappa = 0$ to κ_1. Let $N(\kappa)$ be the discretionary cancellations in the last 48 hours per market, and $\bar{C}^M_{\text{disc}}$ their average market-wide cost with no penalty: 0.006 [0.001, 0.010] on Seattle to San Francisco and −0.009 [−0.012, −0.005] on Los Angeles to Phoenix. The cancellations prevented per unit of penalty are $n = \big[N(0) - N(\kappa_1)\big]/\kappa_1$, and welfare's slope with a monetary penalty is $s = \big[W^{\text{mon}}(\kappa_1) - W^{\text{mon}}(0)\big]/\kappa_1$. The tax on matching per prevented cancellation is the part of the slope that deterrence does not explain, $-\big(s - n\,\bar{C}^M_{\text{disc}}\big)/n$: 0.111 [0.063, 0.161] on Seattle to San Francisco and 0.121 [0.089, 0.154] on Los Angeles to Phoenix. The average market-wide cost of all discretionary cancellations stands in for that of the prevented ones. Intervals resample whole markets.

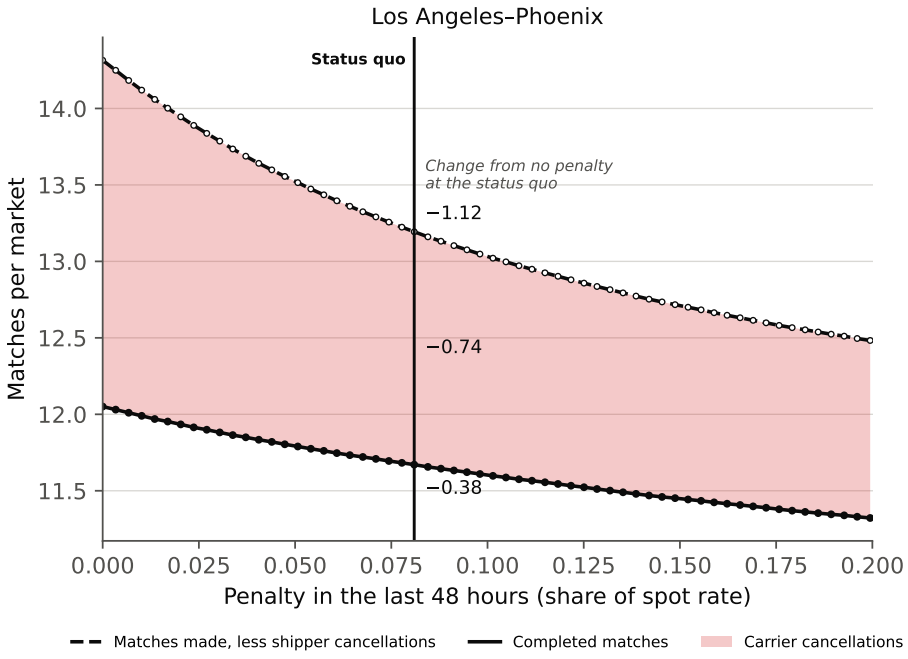
All of the welfare loss comes from lost completed matches. I split the change in W^{mon} from $\kappa = 0$ to $\hat{\kappa}$ into lost completed matches and the rest. Each lost completed match loses the load's value less the carrier's cost, at their averages per completed match with no penalty, and leaves one more load unmatched at pickup, at cost L. The rest is the change in carriers' costs

Figure E.2: Welfare, carriers' surplus and the platform's profit at every penalty, Los Angeles to Phoenix



NOTE: As Figure 7, for Los Angeles to Phoenix, with the benchmarks of Figure E.1.

Figure E.3: Matches and carrier cancellations at every penalty, Los Angeles to Phoenix



NOTE: As Figure 8, for Los Angeles to Phoenix.

on the remaining matches and in the outside offers they take. Lost completed matches make up 129% of the welfare loss at the estimated penalty on Seattle to San Francisco and 133% on Los Angeles to Phoenix. Cheaper carriers on the remaining matches more than make up for fewer outside offers taken.

E.6 When the Penalty Applies

The 24- and 72-hour windows of Section 6.4 are separate sets of counterfactuals, on the same grid of κ as the main ones. In every window run, welfare is highest with no penalty under both accountings, and so is the platform's profit with a reputational penalty.

On Seattle to San Francisco, the platform's gain at the profit-maximizing monetary penalty, per market, is 0.019, 0.017 and 0.019 in the 24-, 48- and 72-hour windows (Table 6), with standard errors of 0.0016, 0.0013 and 0.0014. Paired by market, the gain changes by -0.0018 (standard error 0.0010) from the 24-hour window to the 48-hour one, and by 0.0026 (0.0005) from the 48-hour window to the 72-hour one. Across the penalties within one standard error of the profit-maximizing one, welfare with a monetary penalty changes by -0.056 to -0.043 per market in

the 72-hour window. Averaged over those penalties, carriers' surplus changes by -0.021 per market from a 72-hour window to a 24-hour one.

The market-wide cost per cancellation with no penalty, in Table 6, has these 95% intervals, resampling whole markets: [0.019, 0.029] in the last 24 hours, [-0.019, -0.008] 24 to 48 hours before pickup and [-0.028, -0.013] 48 to 72 hours before. Each lane's longest window gives these numbers, since its twin markets cover the most cancellations.

On Los Angeles to Phoenix, the market-wide cost per cancellation with no penalty has these 95% intervals: [0.011, 0.017] in the last 24 hours and [-0.031, -0.022] 24 to 48 hours before pickup. The 24- and 72-hour windows are still being computed.

E.7 The Platform's Weight on Carriers and the Cost of Collecting a Fee

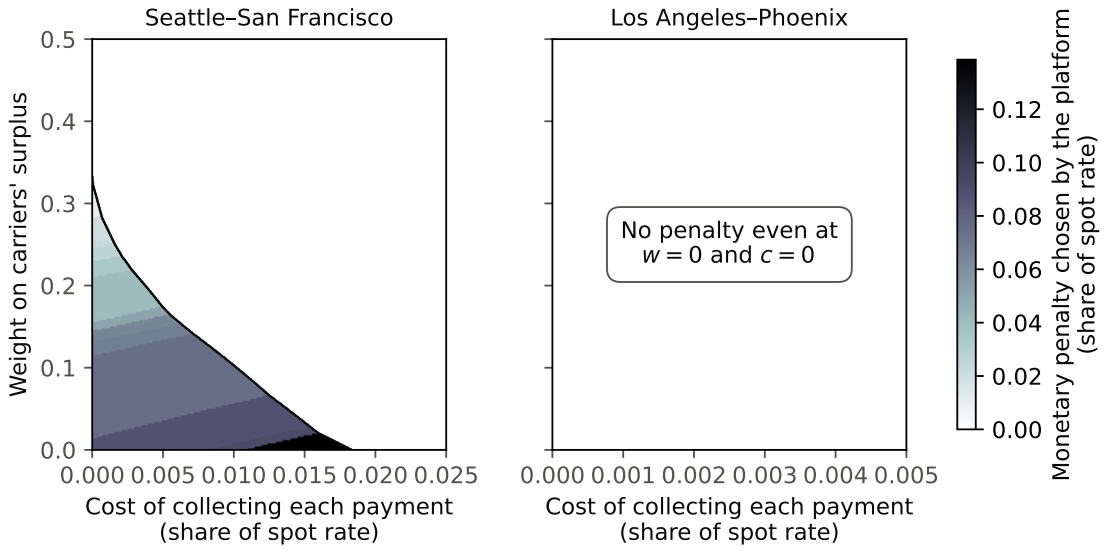
In Section 6.5, the platform puts a weight w on carriers' surplus and pays a cost c, in shares of the spot rate, for each payment it collects. It chooses the monetary penalty κ to maximize

$$O_{w,c}(\kappa) = w\left[U(\kappa) - U^0\right] + (1-w)\left[\Pi(\kappa) + F(\kappa) - c\,F(\kappa)/\kappa\right],$$

where $F(\kappa)/\kappa$ is the number of penalties paid per market (every cancellation in the last 48 hours pays κ), and no penalty is paid at $\kappa = 0$. With no cost and $w = 1/2$, the objective is half of welfare with a monetary penalty, which peaks at $\kappa = 0$. So the weight at which the penalty disappears is at most one half. I compute the objective on the grid's outcomes, for weights from 0 to 0.50 and costs from 0 to 0.025 (0.005 on Los Angeles to Phoenix).

Figure E.4 shows the monetary penalty that maximizes the platform's objective, over the weight and the cost. The weight and the cost act together: a smaller weight and a smaller cost than either needs alone remove the penalty when combined. On Seattle to San Francisco, along a row of low weights, a rising cost first raises this penalty, since a higher penalty is paid by fewer carriers, and then removes it. So a cost below the level that removes the penalty can lower welfare.

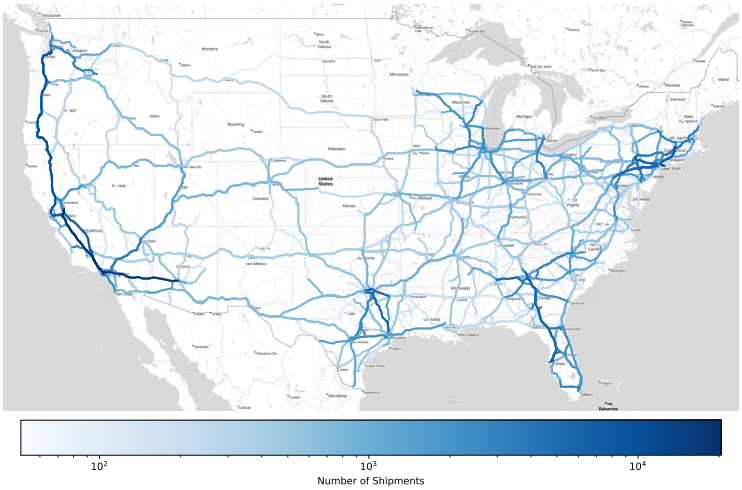
Figure E.4: The monetary penalty chosen by the platform, by its weight on carriers and its cost of collection



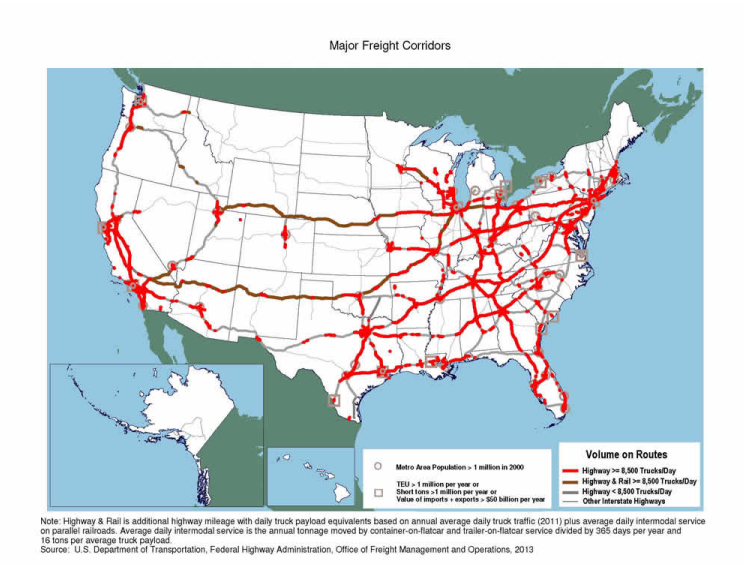
NOTE: The monetary penalty κ in the last 48 hours that maximizes $O_{w,c}(\kappa)$, in shares of the spot rate, on the grid of κ. Vertical axis: the weight w on carriers' surplus, their payoffs net of what they earn without the platform. Horizontal axis: the cost c of collecting each payment. Outlined: the region in which a positive monetary penalty is best; outside it, no penalty is best. On Los Angeles to Phoenix, no penalty is best even at $w = 0$ and $c = 0$, as its panel states.

F Additional Figures

Figure F.1: Comparison of lane frequencies on the platform and across the U.S.



(a) Platform Lane Frequencies



(b) Lane Frequencies Across the U.S.

Figure F.2: Distribution of carrier miles driven in a year

