

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

Richard Faltings *

First version: July 25, 2024

This version: June 3, 2025

[Click here for the latest version.](#)

Abstract

In many markets, advance interim contracts include an explicit right to renege, granting one party the option to switch to more efficient matches that emerge later in the search process. This paper studies the formation and welfare implications of such interim contracts, leveraging novel data from a brokerage firm in the trucking industry. The broker allocates advance shipment contracts to carriers through a dynamic auction mechanism and penalizes cancellations through a reputational mechanism. I develop a theoretical model linking the carriers bidding problem to the firms cancellation penalties through a dynamic job-search problem and structurally estimate the model from rich data on bids and cancellations. In counterfactual simulations, I show that the firm is incentivized to lower cancellation penalties as the option value of the right to renege is priced into carrier bids. The results rationalize the large degree of contractual flexibility observed in the trucking industry as an efficient market outcome rather than one constrained by limited enforcement.

*The University of Texas at Austin, richard.faltings@utexas.edu. Special thanks go to Eugenio Miravete, Dan Akerberg, 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, Yingkang Xie, and Alexander McKay. This paper benefitted from comments by seminar participants at UT Austin and conference participants at WoPA 2025, and IIOC 2025. 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 (McCall 1970). 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, reducing market efficiency relative to a centralized clearing mechanism where all potential trading partners are known (Roth and Xing 1994). Intuitively, efficiency in such markets can be improved by an explicit *right to renege*, allowing at least one party to continue searching for better opportunities at a predetermined price, or penalty.

Rights to renege against a penalty are a prominent feature of many markets. Travelers are familiar with cancellation terms for hotel and flight bookings, which vary across firms and over time (NYT 2008; NYT 2015; WaPo 2024). Job searchers often face exploding offers in the labor market, making it harder for them to compare offers simultaneously (Niederle and Roth 2009). Digital platforms have made it easier than ever for a range of small businesses to charge cancellation fees (WSJ 2024). Despite the widespread prevalence and variety of explicit penalties for renegeing, there is scant empirical research on the economic forces that determine firms' optimal choices of penalty terms, or the welfare implications thereof. I seek to fill this gap.

The direct economic trade-offs of a right to renege are straightforward: the party with the option to cancel conserves some of their option value from future opportunities, while the other party sacrifices their option value for the duration of the agreement.¹ The price of this trade is the cancellation penalty, which acts as a foregone deposit on the transaction.² Unlike financial options that involve the trade of fungible assets, the right to renege on a match pertains to the right to trade with a specific counterparty. Therefore, we cannot apply standard option pricing theory based on efficient markets (Black and Scholes 1973) to understand the pricing of the right to renege. Instead, we must consider the pricing problem within the strategic context of a search and matching market. I examine the problem of a firm pricing the right to renege on its standard terms of trade, and whether the profit-maximizing penalties are also socially efficient.

1. For example, the federal regulation requiring U.S. airlines to allow free cancellation within 24 hours of booking added an exemption for flights departing within 7 days, after airlines opposed it arguing that it “takes inventory off the market for the duration of the refund period, blocking it from sale to other customers [...]” (DOT 2011)

2. Some airlines allow passengers to hold a reservation for an upfront fee. On the introduction of United's FareLock program, its president said: “It's a value to people like a stock option is a value [...] Well, this is an option on a seat” (NYT 2012).

The equilibrium efficiency of such agreements also depends on the nature of the penalty. Enforcement costs might render a monetary penalty infeasible: parties might instead use reputational penalties ([Kandori 1992](#)), which use the threat of future exclusion to penalize a reneging party. Unlike monetary penalties, which transfer utility from one party to the other, reputational penalties only punish the reneging party without compensating their counterparty.³ This creates a direct deadweight loss, and may also under-incentivize the provision of flexibility. On the other hand, a monetary penalty turns reneging into a revenue-generating opportunity for the firm, which may distort firm incentives towards extracting additional rents from consumers.⁴ In this sense, monetary penalties can be included under the umbrella of “drip pricing”, a practice which presents consumers with additional fees at various stages of the transaction process, which has attracted recent regulatory scrutiny ([FTC 2024](#)).

I study the problem of a firm’s optimal pricing of a right to renege (through reputational or monetary cancellation penalties) in an important empirical context: the U.S. trucking industry, which transports 74.8% of domestic freight by value ([USDOT 2021](#)) and accounts for 2.3% of GDP ([USDOT 2022](#)). The market is suitable for this research question for several reasons. One, match quality matters because each shipment is unique, with its own pickup and dropoff locations and schedule, which makes it more or less costly for different drivers to haul, depending on their own schedules and locations. Two, finding the ideal match is difficult because the market is fragmented on both the demand and supply side (91% of carriers operate fleets of ten or fewer trucks ([FMCSA 2023](#))) and there is limited time to search over alternatives (lead times are typically 2 weeks down to less than 24 hours). Consequently, the timing friction described above is a significant issue in the industry. This is evident in the significant flexibility built into advance contracts throughout the market.⁵

The data for the empirical analysis are provided by a large trucking brokerage firm that matches shipments with freight carriers through an online auction platform. The platform is designed to allow carriers to renege at two stages. Upon winning an auction, the carrier receives a notification and has the option to confirm or cancel the match. After a match is confirmed by the carrier, they still have the option to cancel easily through an online interface, but at the

3. [Hubbard \(2001\)](#) draws a similar distinction between reputational and contractual enforcement in the context of trucking.

4. Change and cancellation fees accounted for 3.2% of U.S. airline passenger revenue in 2009 ([WSJ 2009](#))

5. To illustrate, [Caplice and Sheffi \(2006\)](#) discuss how this need for flexibility hindered efforts to allocate shipments via combinatorial auctions: “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.” Consequently, the industry couldn’t reliably guarantee bundles to bidders, as would be expected in most combinatorial auctions.

risk of a reputational penalty on future bids on the platform.⁶ Attrition is high: over half of winning bids are reneged on, while 15% of confirmed matches are eventually canceled before pickup. Self-reported reasons for cancellation and within-platform substitution behavior suggest the cancellations are by and large opportunistic: carriers cancel their original shipment when they find a better deal. In the majority of cases, carriers that cancel do not substitute to an alternative shipment on the brokerage platform itself, but they are much more likely to cancel shipments with a low payout, suggesting that opportunistic behavior extends to off-platform alternatives. These features make the data uniquely suited to the research question, as most other sources of market data only record final realized transactions.

Guided by the empirical evidence, I build a dynamic model of carrier search and bidding behavior which also incorporates the flexibility to renege into the bidding strategy. In the analysis of the model, I disentangle the direct effects of a penalty on the propensity to cancel from the indirect effects on the bidding behavior. The exercise demonstrates that accounting for the full auction equilibrium effects of the penalty is crucial for understanding the firm incentives to offer flexibility, and the ensuing welfare implications; without the effect on bids, the firm would require full commitment from carriers. I also show that the variance of the carriers' outside offers and the type of penalty—reputational or monetary—affects firm incentives. When the variance of outside offers is low, firm profits increase with higher penalties, even at the cost of overall welfare. Conversely, when the variance is high, both firm profits and overall welfare decrease in the penalty, so that firm incentives are socially aligned. With monetary penalties, the firm can generally profit by raising the penalty level, at the cost of overall welfare, showing that monetary penalties offer an additional channel for the platform to extract rents from carriers.

To take the model to the data, I enrich the baseline model with additional dynamics, allowing for stochastic arrivals of carriers and shipments over time, and accounting for the multi-round and multi-unit nature of the auctions. I then estimate the primitives governing carrier policies, including the stochastic process of outside offers, and their subjective valuation of the platform's reputational cancellation penalty. To support the estimation strategy, I develop a simple partial identification argument based on how the penalty enters into the carrier's bidding function and a non-negativity condition on carrier marginal costs. The dynamics of carrier and shipment arrivals, as well as the search process governing which shipments carriers observe and when, are taken as exogenous, and estimated by matching moments of the arrival processes in the data.

6. Auction winners are selected according to the sum of their bids and an algorithmic quality score computed on the basis of past cancellation rates and other metrics such as on time pickup and delivery.

To find the optimal penalty design, I use estimates of the model to analyze the profit and welfare consequences of two classes of counterfactual penalties: reputational penalties (which solely decrease carrier utility) and monetary penalties (which generate revenue for the firm). Among reputational penalties, I find that the current near-zero reputational penalties are nearly optimal for both social welfare and firm profits. Completely free cancellations would modestly increase both profits and welfare, while a full commitment policy would cause a significant decline in both. A decomposition exercise reveals that higher reputational penalties decrease profits through both of fewer successful matches (despite a decline in the number of cancellations) and higher costs per successful match.

I also find that moving to monetary penalties would allow the platform to increase short-run profits, as predicted by the theoretical exercise. The resulting 48.7% increase in profits, relative to the status quo, is mainly the result of transfers from carriers, whose welfare declines by 18.5%, for a net welfare loss of 5.5%. To rationalize why the firm does not use monetary penalties, I explore alternative objective functions which take a weighted average of firm profits and carrier welfare to proxy for long-run motives of carrier acquisition and retention to the platform. I find that a 45% weight on carrier welfare is sufficient to explain the lack of monetary penalties. I also explore the effect of transaction or enforcement costs incurred by the firm for every cancellation fee it collects. Under reasonable levels of transaction costs, the firm finds it optimal to avoid monetary fees at lower weights on carrier welfare. However, transaction costs alone are not sufficient to explain the lack of monetary penalties without some weight on carrier welfare.

While the baseline analysis restricts attention to schedules of cancellation penalties that follow the same time profile as the status quo penalty (which increases sharply for cancellations within the last 48 hours before pickup), the model can be used to analyze alternative profiles. I compare the increasing time profile to a flat one, whereby the cancellation penalty is constant over time. For a given maximal penalty level, carrier welfare is lower with a flat penalty as earlier cancellations are punished more harshly, but this only increases firm profits if the penalty is a monetary fee. Thus, monetary cancellation penalties may not only incentivize harsher cancellation penalties than is socially optimal, but earlier penalties as well.

To fully illustrate the value of contractual flexibility in this market, I also investigate the welfare effects of an infinite penalty. Using the same auction format, the full commitment policy reduces total welfare by 37.4% compared to the status quo, driven by a 25.5% decline in the number of final successful matches. The average bid increases by 10.2% and the transaction price of successful matches increases by 4.3%, reflecting the increased opportunity cost that

carriers price into their bids.

A reasonable conjecture is that full commitment is less costly if carriers are matched to shipments *later*. I consider another full commitment policy in which auctions are only cleared in the last 24 hours. This indeed leads to a reduction in the transaction price of *successful* matches, but the total number of matches is even lower, with a 71.4% reduction relative to the status quo. This can be explained by a greater rate of attrition to offers outside the platform; carriers waiting on the uncertain outcome of a bid are more likely to prefer the sure payoff of an immediate outside offer, as compared to a carrier who has already been matched with a shipment on the platform. Thus, flexible cancellation policies are a more effective *unilateral* policy for mitigating timing frictions when the market cannot be easily coordinated through a centralized matching process.

Literature Review Timing frictions were first highlighted in entry-level labor markets ([Roth and Xing 1994](#); [1997](#); [Li and Rosen 1998](#); [Li and Suen 2000](#); [Kagel and Roth 2000](#)). This literature explores how the strategic timing of binding offers reduces market efficiency. In an experimental setting [Niederle and Roth \(2009\)](#) show that market norms enforcing non-binding agreements can improve efficiency, but that these do not arise in competitive equilibrium. I contribute to this literature, first by extending the theoretical setting to markets with search frictions ([McCall 1970](#); [Pissarides 2000](#)), where each market participant experiences a *different* sequence of opportunities. Second, I provide the first empirical application of this largely theoretical and experimental literature in the economically important market of trucking.

Other recent research has explored the strategic incentives to offer non-binding offers in search markets, in the form of both exploding offers of varying deadlines ([Lippman and Mamer 2012](#); [Lau et al. 2014](#); [Armstrong and Zhou 2016](#); [Zorc and Tsetlin 2020](#); [Hu and Tang 2021](#)) and cancellation penalties of varying levels ([Xie and Gerstner 2007](#); [Zhang et al. 2021](#); [Liu and Chen 2022](#); [Liu and Zhang 2023](#)). My model nests both forms of non-binding agreements: exploding offers can be modeled as an initial cancellation penalty of zero, jumping to an infinite penalty at some deadline. I replicate and jointly consider several key mechanisms from the theoretical literature, including the acceptance deterrence effect of holding an attractive offer ([Zorc and Tsetlin 2020](#)), the interaction between the penalty level and the transaction price ([Hu and Tang 2021](#)), and the rent extraction effect of a monetary cancellation penalty ([Xie and Gerstner 2007](#)). The rich model allows for the joint effect of all these mechanisms to be quantified in counterfactual simulations.

For the analogous problem of product returns, [Anderson, Hansen, and Simester \(2009\)](#) show

that similar incentives exist for firms to offer flexible return policies when consumers are sufficiently uncertain about their preferences for a product. The key differences in the case of the right to renege is that the relevant uncertainty is over *alternative matches* in the market, and that the cost of reneging is endogenously derived from the foregone offers in the market, rather than from the logistics of returning and restocking a product.

I also contribute to the literature on auction formats with options to renege or adapt bid amounts. The procurement mechanism for Medicare Durable Medical Equipment (DME) received significant interest ([Merlob, Plott, and Zhang 2012](#); [Cramton, Ellermeyer, and Katzman 2015](#); [Ji 2022](#)) for its auction format in which firms may renege on bids after the auction result and clearing price have been revealed. In a different setting, [Bajari, Houghton, and Tadelis \(2014\)](#) emphasize the need to provide flexibility in procurement auctions due to bidders' own uncertainties about construction project costs while [Haberman and Jagadeesan \(2023\)](#) provide a theoretical justification for withdrawals as part of the price-finding process. I show that in dynamic search settings—similar to [Backus and Lewis \(2024\)](#)—uncertainty over future substitutes also provides an economic rationale for reneging.

Finally, my paper contributes to the growing literature on the industrial organization of trucking, a market which has garnered significantly more interest from policy-makers since the supply chain disruptions of the COVID-19 pandemic highlighted the dependency of the U.S. economy on the trucking industry. Early contributions most notably include [Hubbard \(2000\)](#), [Hubbard \(2001\)](#), [Baker and Hubbard \(2001\)](#), and [Baker and Hubbard \(2003\)](#), focusing on issues of vertical relations and contracts under moral hazard, while more recent work by [Yang \(2022\)](#) explores spatial frictions caused by the need for drivers to return to their home locations. The most closely related work in this area is [Harris and Nguyen \(2022\)](#), who also examine the implications of non-binding agreements in trucking, but with a focus on long-term relational contracts. They find that, due to the inflexible prices in long-term contracts, an excessive number of relational contracts are reneged on when the spot market prices rise, and suggest that reneging can be avoided by indexing the price to the market. In contrast, my work demonstrates the beneficial role of contractual flexibility in alleviating search frictions in the spot market and rationalizing the widespread prevalence of reneging behavior in the industry even in the absence of pricing frictions.

The rest of this paper is structured as follows. In Section 2, I provide an overview of the trucking industry and the data used in this paper. In Section 3, I present a stylized model of the carrier search process, which I then extend to a dynamic setting in Section 4. In Section 5, I describe the estimation of the model, followed by the counterfactual analysis in Section 6. I conclude

with Section 7.

2 Empirical setting and data

The trucking industry is a vital sector of the U.S. economy, accounting for 2.3% of the GDP in 2022 ([USDOT 2022](#)), and transporting 74.8% of domestic freight 2023 ([USDOT 2021](#)). The market operates on a decentralized basis, through a combination of long-term contracts (typically several months to a year) between shippers and large carriers or brokers, and spot transactions, often negotiated through brokers. The industry is highly fragmented, with over 55% of carriers operating a single truck, and 91% operating 10 trucks or fewer ([FMCSA 2023](#)). This fragmentation creates costly search frictions in the spot market, which play a major role in explaining the prevalence of long-term contracts in the industry ([Hubbard 2001](#)).

Unlike most of the industry, which still relies on labor-intensive matching processes, the firm studied in this paper was an early pioneer of the digital brokerage business model and sought to heavily automate the matching process. The majority of the broker's upstream demand came through long-term contracts with large shippers (such as Nestlé or Starbucks), which would regularly tender shipments to the platform at the contracted rate. Once the firm accepts a shipment, it would seek to procure a carrier at the lowest possible cost, using an online auction platform. Despite the firm's large size in absolute terms, it only accounted for a small fraction of the overall market, with revenues slightly below 1% of the \$251 billion for-hire trucking market in 2021.

2.1 Platform design

The study covers the full two-year period of 2021 and 2022, during which the firm implemented its "timed auction" format. This format consists of several rounds, with the reserve prices increasing across rounds. Additionally, the platform features an option called Accept-Now, allowing carriers to immediately confirm a shipment at a posted price, fixed over time. 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.

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*), as well as place a bid, use the Accept-Now feature, or simply move on (referred to as willfully *ignoring* a shipment). 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.

When an auction reaches one of its pre-determined closing times, if all bids are above the reserve price, the auction proceeds to the next round. If there are bids below the reserve price, the platform accepts the lowest 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 confirmation, the carriers continue to enjoy some flexibility, as the platform's official policy is that cancellations only suffer a reputational penalty when made less than 48 hours before the shipment's pickup time.⁸ This penalty reduces carrier's probability of winning future bids through a quality adjustment factor which changes the bid ranking, but does not result in any direct monetary loss. The exact mechanics of the reputational penalty are not disclosed to carriers, which is why the empirical strategy will also focus on recovering carrier's subjective valuation of this penalty. While it is unclear how widespread explicit cancellation policies are among traditional brokers, for the period of this study the main rival digital platform has a similar policy, penalizing carriers' reputation for cancellations in an even shorter penalized window of 24 hours before pickup.

2.2 Data and Stylized Facts

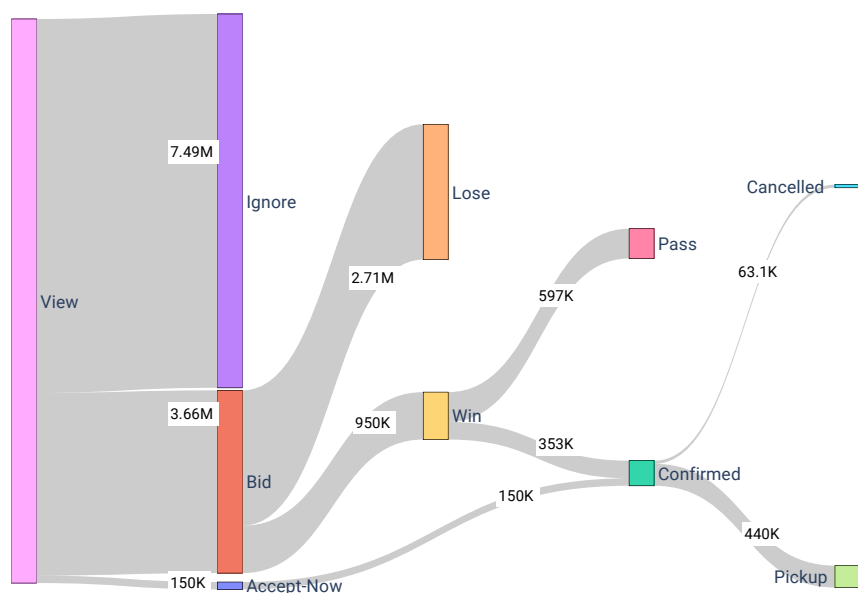
The design of the platform yields a rich dataset on the endogenous decisions by carriers at the multiple stages of the matching process, allowing me to examine the empirical determinants

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. Prior to Q1 2021, carriers had to contact a firm representative to cancel. The policy was 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."

of attrition on the platform. The process of attrition, converting views to final matches, is shown as a funnel in Figure 1. It is immediately apparent that the overwhelming majority of potential matches do not result in a match, while two-thirds of views do not even result in a bid. The attrition of carriers across this funnel is a central focus of the empirical analysis.

Figure 1: Funnel from views to matches



NOTE: Data represents the universe of non-local shipments posted to the auction platform in 2021 and 2022. Gray areas and numbers represent mass of flows between states.

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. Furthermore, the structural empirical analysis is focused on a single shipping lane—defined as a route between two metropolitan areas—specifically from Seattle, Washington, to San Francisco, California. Table 1 is a summary and comparison of the two sets of data. The sub-sample is highly representative of the full sample along the measures of attrition and bidding behavior. The main difference between subsets is the mean and variance of shipment distance, which is unsurprising given the focus on a single lane.

As my research question focuses on the short-term dynamics of the market for each pickup date, I normalize all measures of prices by the average spot market rate of a shipment's lane (defined as a metro-to-metro) at the shipment pickup time, to remove price variation that is due to aggregate market fluctuations that are caused by macroeconomic factors. This rate

was obtained from a third-party data aggregator and is used internally by the firm for various prediction and optimization tasks.

Table 1: Summary statistics

	Full data	Seattle-San Francisco
Unique carriers	87,652	4,592
Unique shipments	792,097	9,098
Number of views	14,859,358	87,887
Number of accept now	161,144	1,368
Number of bids	5,114,504	48,814
Bid acceptance rate	24.98%	28.70%
Bid confirmation rate	37.38%	36.03%
Cancellation rate	14.15%	12.83%
Average bid amount	1.16 (1.01)	1.16 (0.27)
Average shipment distance	333.08 (405.20)	757.78 (6.41)

NOTE: Full data refers to the universe of non-local shipments posted to the auction platform in 2021 and 2022. The sub-sample is restricted to the Seattle to San Francisco lane in the same period. Standard deviation of averages are in parentheses. Bid amount is normalized by the contemporaneous average spot rate of the lane at the time of shipment pickup.

2.2.1 Cancellation Behavior and Consequences

In the following, I focus on the determinants of, and consequences of carrier cancellations. Overall, 14.9% of confirmed shipments are cancelled by the carrier. Among carriers cancelling on a confirmed shipment (reneging on a bid), 27.8% (25.5%) ultimately execute a different shipment on the platform. This suggests that carriers use the flexibility to cancel and renege to substitute to other—presumably better—opportunities. The remaining cases likely represent instances of substitution to off-platform opportunities, given the small size of the platform relative to the broader market.⁹

9. As additional evidence for carrier multihoming outside the platform, I find that 95% of carriers drive no more than 10,638 miles on the platform per year (see the full distribution in Figure A2 in the Appendix). 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.

Table 2: Impact of Carrier Cancellations on Shipment Metrics

	No Cancellation	≥ 1 Cancellations
Shipment rescheduled	10.7%	18.1%
Shipment cancelled	6.9%	8.1%
Final Margin	7.7%	4.6%
Overall carrier cancellation rate	14.9%	

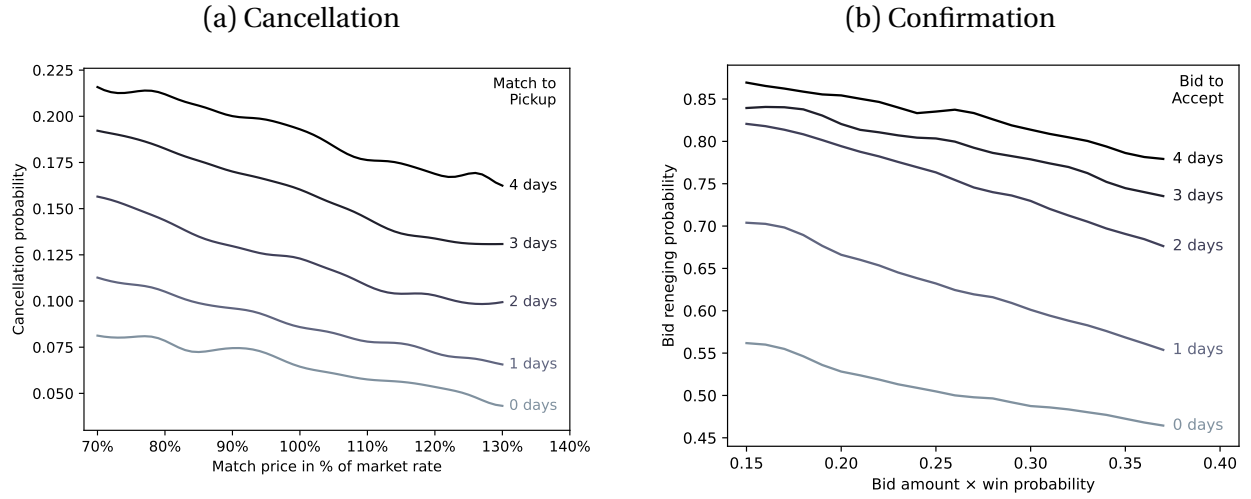
NOTE: Table compares outcomes for shipments with and without at least one carrier cancellation. Shipment cancelled refers to a cancellation from the shipper side.

In Table 2, I present the tangible consequences of cancellations for the platform. The table shows that cancellations have a large impact on the rate of shipment rescheduling, with a much more modest effect on shipment cancellations by the shipper. 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.

The data also show that in most cases, the platform is able to re-match a shipment following a carrier cancellation, but does so at a higher price, resulting in a lower margin. These financial consequences for the platform raise a natural question: what drives carriers to cancel confirmed shipments?

Figure 2 shows that both cancellations and bid reneging are affected by (i) the payoff of a shipment and (ii) the time interval available for carriers to change their minds. Panel (a) shows that the probability of a carrier cancelling a shipment is decreasing in the price of the shipment, while panel (b) shows that the probability of a carrier reneging on a bid is decreasing in a simple measure of the expected shipment payoff. Both panels also show that the probability of cancelling (reneging) is increasing in the time interval between when the carrier is matched to a shipment (makes a bid) and the pickup time (bid acceptance time). Both of these patterns are consistent with standard models of search markets ([McCall 1970](#)) where a stochastic process of outside offers composed of some rate of arrival of outside offers, and some distribution of payoffs conditional on the arrival, rationalize the carrier's decision to exit a match.

Figure 2: Confirmation/cancellation probability vs. payoff



NOTE: Figures are obtained through a two-dimensional kernel regression, conditioning both on different time horizons for attrition, and measures of payoffs. Kernel bandwidth chosen in accordance with Silverman's rule of thumb. Win probability in panel (b) is also computed with a one-dimensional kernel regression of bid acceptance on bid amount.

2.2.2 Reputational Mechanism

As described above, the platform's stated policy was to punish cancellations that occurred within 48 hours before the pickup time by penalizing future bids by some quality adjustment factor, which is simply added on to the bid dollar amount. This adjustment factor was not disclosed to carriers and its details are confidential. Nevertheless, in Table 3, I present some descriptive evidence of the effect of the reputational penalty on carriers' win probability and their bidding behavior to shed light on the severity of the penalty.

In the first regression, I estimate a linear probability model of the conditional win probability of an auction on the quality adjustment factor and the bid amount itself. The linear specification allows me to control for local market conditions through high-dimensional shipment lane and calendar month fixed effects. The results show that the quality adjustment factor's impact on the win probability is roughly two-thirds of the impact of the bid amount. This is likely due to the fact that only the raw bid determines whether a carrier beats the reserve price, while the quality adjustment factor only affects the ranking of the bids.

In the second regression, I estimate the effect of cancellations on the quality adjustment factor. The results show that the quality adjustment factor is more sensitive to recent cancellations, but generally very small—a single cancellation in the past 30 days increases the quality adjust-

Table 3: Reputational penalty effects

	(1) Win Prob (1)	(2) Quality Adjustment (2)	(3) Bid (3)	(4) Naive Win (4)
Normalized Price	-0.746*** (1.24e-03)			
Quality Adjustment	-0.490*** (0.011)			
Cancellations (30D)		3.51e-05*** (5.24e-06)	-7.79e-04*** (5.04e-05)	5.81e-03*** (1.02e-04)
Cancellations (60D)		2.24e-05*** (5.52e-06)	-6.85e-05 (5.30e-05)	6.92e-06 (1.22e-04)
Cancellations (90D)		-1.93e-05*** (5.38e-06)	-1.93e-04*** (5.17e-05)	2.24e-03*** (1.13e-04)
Observations	2244902	2244539	2244539	2249878
Lane \times Month FE	Yes	No	No	Yes
Carrier FE	No	Yes	Yes	No

NOTE: All prices and quality adjustments are normalized by each shipment's lane-specific spot rate. 30D cancellation variable is a rolling sum of cancellations in the 30 days preceding a bid, while 60D and 90D refer to the 30-60 and 60-90 day windows, respectively. Standard errors are in parentheses. * $p < 0.100$; ** $p < 0.050$; *** $p < 0.010$.

ment on a thousand dollar shipment by roughly three and a half cents. I should caution that these estimates are not strictly causal, as carriers may change their behavior in response to a cancellation to maintain their quality score, which can explain the negative coefficient on the 90 day window.

In a similar vein, the third regression shows that carriers reduce their bids in response to cancellations. A single cancellation in the past 30 days reduces the bid on a thousand dollar shipment by roughly 80 cents, an order of magnitude stronger effect than the actual impact on the quality adjustment factor. This suggests that carriers may be overestimating the reputational penalty. Ultimately, the policy's effects will largely depend on carriers' *subjective* beliefs about the penalty, which will be recovered through the estimation of the structural model.

To summarize, the empirical analysis of the platform data has revealed several stylized facts that will guide the development of the structural model. First, the platform's auction format is characterized by a high degree of attrition, with the majority of views not resulting in a bid, and the majority of bids not resulting in a match. Second, the attrition patterns are sensitive to both price and time, suggesting search behavior driven by off-platform opportunities. Third, cancellations have a large impact on the platform's bottom line, yet the reputational penalty

imposed by the platform has quantitatively small effects on carriers' win probabilities and bidding behavior. A simplified model capturing these patterns is presented in the next section.

3 Baseline Model

I have shown that cancellation and confirmation behaviors are empirically linked to the price and timing of the initial match between a carrier and a shipment. However, there is not enough variation in the data to directly measure the impact of a change in the cancellation policy.¹⁰ The dual objectives of the model are thus to *account* for the empirical patterns in the data and *predict* the counterfactual effects of different cancellation policies.

For the sake of exposition, I begin with a simplified baseline model. This captures the two essential economic mechanisms determining the effect of a change in the cancellation penalty, but abstracts away from the additional complexities of stochastic and heterogeneous arrival times of carriers, and the multi-unit nature of the auction format. The first mechanism is that, holding the match time and price fixed, a higher cancellation penalty should decrease a carrier's propensity to cancel. The second mechanism is that a carrier's forward-looking strategy in the auction depends on the cancellation penalty through their opportunity cost of continuing to search, following well-established models of dynamic auctions (Jofre-Bonet and Pesendorfer 2003; Zeithammer 2006; Hendricks and Wiseman 2022).

Since the auction format allows bidders to renege when they win, solving for the auction equilibrium is a non-trivial challenge, as it must account for the opponent's probability of reneging.

3.1 Setup

In the model, the brokerage platform runs an auction for a single shipment for every pickup date, each of which is treated as a separate market. The auction takes place in the days leading up to the departure date, with the index $d \in \{4, 3, 2, 1, 0\}$ counting down the number of days remaining until the pickup time. For each shipment, the platform earns revenue $v \sim N(\mu_v, \sigma_v)$ for a successful match; the payoff for an unsuccessful match is normalized to zero. Two carriers $i \in \{1, 2\}$ participate in each auction, with marginal costs of shipping $c_i \sim N(\mu_c, \sigma_c)$. Each of these two carriers also takes independent draws from shipment jobs outside

10. In general, it would be difficult for any firm to experimentally vary their cancellation penalties, as these are generally less salient than prices, and carriers, or customers, may take time to adapt to a new policy.

the platform, with net payoffs $\pi - c_i$, with $\pi \sim G(\cdot) \equiv N(\mu_G, \sigma_G)$. These three distributions are the primitives of the model.

The platform's only policy lever is the cancellation penalty κ . In the status quo, the cancellation penalties κ are reputational, in that they impact a carrier's utility by reducing their future profits on the platform, but do not result in any direct monetary transfer to the platform. While such reputational mechanisms may also have benefits in terms of *screening out* low-quality carriers, Appendix A.8 provides reduced-form evidence that screening has limited benefits for platform profits, as carriers who cancel more frequently also tend to match at lower prices, effectively off-setting the cost of their cancellations. Based on this evidence, I assume the platform simply receives nothing when a carrier cancels, so that κ only provides an incentive effect. In counterfactual simulations, I also consider monetary cancellation penalties, which result in a direct monetary transfer of value κ from the carrier to the platform when a cancellation occurs.

The timing of the game is as follows:

- $d = 4$: carriers submit their bid b_i
- $d = 3$: each carrier receives an outside offer $\pi_{i,3}$
 - If a carrier takes the outside offer, they exit the market and get payoff $\pi_{i,3} - c_i$
- $d = 2$: auction clears, awarding shipment to lowest bidder still in the market, subject to bidding less than v ¹¹
- $d = 1$: each carrier receives another outside offer $\pi_{i,1}$
 - If an unmatched carrier takes the outside offer, they exit the market and get payoff $\pi_{i,1} - c_i$
 - If a matched carrier takes the outside offer, they are subject to a cancellation penalty and receive payoff $\pi_{i,1} - c_i - \kappa$
- $d = 0$: if the auction winner has not cancelled, they receive a payoff $b_i - c_i$, while the platform gets payoff $v - b_i$.

Formally, I model the auction as a Bayesian game between the two carriers. I focus on the symmetric Bayesian Nash equilibrium consisting of conditional win probability beliefs, denoted $\gamma(b) \in [0, 1]$, and the carriers' optimal policies, described further below.

11. Setting the reserve price equal to the platform's valuation is not generally optimal (Riley and Samuelson 1981). However, the optimal reserve price in this model will also depend on the cancellation penalty. For the sake of tractability, the reserve price is simplified here.

3.2 Carrier policy

Carriers are searching for the most profitable shipment they can take given their costs (their type). This may be a shipment from the platform, if they win the auction, or an outside offer they receive. Their behavior can be characterized by a set of type-specific policy functions, consisting of a bidding function $b^*(c)$ and reservation wages $R^p(b, c)$, $R^m(b)$, $R^u(c)$, where the superscripts p , m and u correspond to the pending (bid), matched and unmatched states, respectively. Any outside offers π that exceed these reservation wages will be accepted.

I can solve for carrier policies through backward induction. Recall that the time index d is counting *backwards* from the pickup date. In $d = 1$, the last period before pickup, a carrier who has not left the market yet is either matched with the platform's shipment or not.

If they are matched, their reservation wage depends on the price agreed-upon with the platform, given by their own bid b_i . The carrier then sees the outside offer $\pi_{i,1}$ and decides whether to take it or not, solving $\max\{\pi_{i,1} - c_i - \kappa, b_i - c_i\}$, which results in a reservation wage of $R^m(b_i) = b_i + \kappa$. I can describe the ex-ante value function of a matched carrier in period $d = 1$ in terms of their cutoff strategy:

$$V_1(b_i, c_i) = \int_{R^m(b_i)} (\pi_{i,1} - c_i - \kappa) dG(\pi_{i,1}) + G(R^m(b_i))(b_i - c_i) \quad (1)$$

An unmatched carrier at $d = 1$ compares the outside offer to the payoff of not working at all, solving $\max\{\pi_{i,1} - c_i, 0\}$. Their reservation wage is simply $R^u(c_i) = c_i$. The ex-ante value function of an unmatched carrier in period $d = 0$ is thus:

$$U_1(c_i) = \int_{c_i} (\pi_{i,1} - c_i) dG(\pi_{i,1}) \quad (2)$$

I then move backwards to $d = 2$, the period in which the auction is cleared. The lowest bidder who clears the reserve price is awarded the shipment. I assume for now that the equilibrium probability of winning, given bid b is $\gamma(b)$, which I solve for later. The ex-ante value function for this period is thus:

$$W_2(b_i, c_i) = \gamma(b_i)V_1(b_i, c_i) + (1 - \gamma(b_i))U_1(c_i) \quad (3)$$

In $d = 3$, carriers are waiting for the outcome of the auction, while receiving another outside offer. Their reservation wage at this point is $R^p(b_i, c_i) = W_2(b_i, c_i) + c_i$. The ex-ante value

function for this period is then:

$$W_3(b_i, c_i) = \int_{R^p(b_i, c_i)} (\pi_{i,3} - c_i) dG(\pi_{i,3}) + G(R^p(b_i, c_i))(W_2(b_i, c_i)) \quad (4)$$

which is similar to Equation 1, except that the carrier incurs no penalty for taking the outside offer. Finally, the bidding problem in the first period in $d = 4$ is to maximize the continuation value in the next period:

$$\max_{b_i} W_3(b_i, c_i)$$

By applying Lemma 1 in the Appendix (which is just a special case of the Envelope theorem), I can show that the FOC is:

$$\begin{aligned} \frac{\partial W_3(b_i, c_i)}{\partial b_i} &= G(R^p(b_i, c_i)) \frac{\partial R^p(b_i, c_i)}{\partial b_i} = 0 \\ \Rightarrow \frac{\partial W_2(b_i, c_i) + c_i}{\partial b_i} &= \gamma'(b_i)(V_1(b_i, c_i) - U_1(c_i)) + \gamma(b_i) \frac{\partial V_1(b_i, c_i)}{\partial b_i} = 0 \end{aligned}$$

Using Lemma 1 again, the derivative of the matched value function is given by $\frac{\partial V_1(b_i, c_i)}{\partial b_i} = G(R_1(b_i)) = G(b_i + \kappa)$.

The optimal bid satisfies the following equation:

$$b^*(c_i) = c_i + \underbrace{\frac{(1 - G(b^*(c_i) + \kappa))\kappa + \int_{c_i}^{b^*(c_i) + \kappa} (\pi_{i,1} - c_i) dG(\pi_{i,1})}{G(b^*(c_i) + \kappa)}}_{\text{Opportunity cost}} - \underbrace{\frac{\gamma(b^*(c_i))}{\gamma'(b^*(c_i))}}_{\text{Markup}} \quad (5)$$

Thus, a carrier's optimal bid consists of their marginal cost, the opportunity cost of foregone outside offers (increasing in the cancellation penalty), and a standard markup term.¹² The bidding function is monotonic in c , so it is invertible. Thus, although the reservation wages $R^p(b, c)$ and $R^m(b)$ are written in terms of the bid, they can also be expressed solely as a function of the carrier's type.

12. Note that $\gamma'(b) < 0$ in equilibrium, so that the markup is positive.

3.3 Auction equilibrium

Carrier strategies are a function of the equilibrium conditional win probabilities $\gamma(b)$. This win probability is given by:¹³

$$\gamma(b_i) = \underbrace{P(b_i < v)}_{\text{Beat reserve price}} \left(\underbrace{P(b_i \leq b_j)}_{\text{Beat opposing bid}} + \underbrace{P(b_i > b_j) \mathbb{E}_{b_j}[1 - P(\text{confirm}|b_j)|b_j < b_i]}_{\text{Opponent wins but reneges}} \right) \quad (6)$$

where $P(\text{confirm}|b_j) = G(R^p(b_j, b^{*-1}(b_j)))$ is the probability that an opponent with bid b_j doesn't find an outside offer above their reservation wage. Thus, compared to a typical auction, each bidder has an additional chance to win if the opponent reneges. The carrier policies $b^*(c)$, $R^p(b, c)$, $R^m(b)$, $R^u(c)$ jointly influence the above conditional probability, in addition to the three primitive distributions.

3.4 Decomposing the effect of an increased penalty

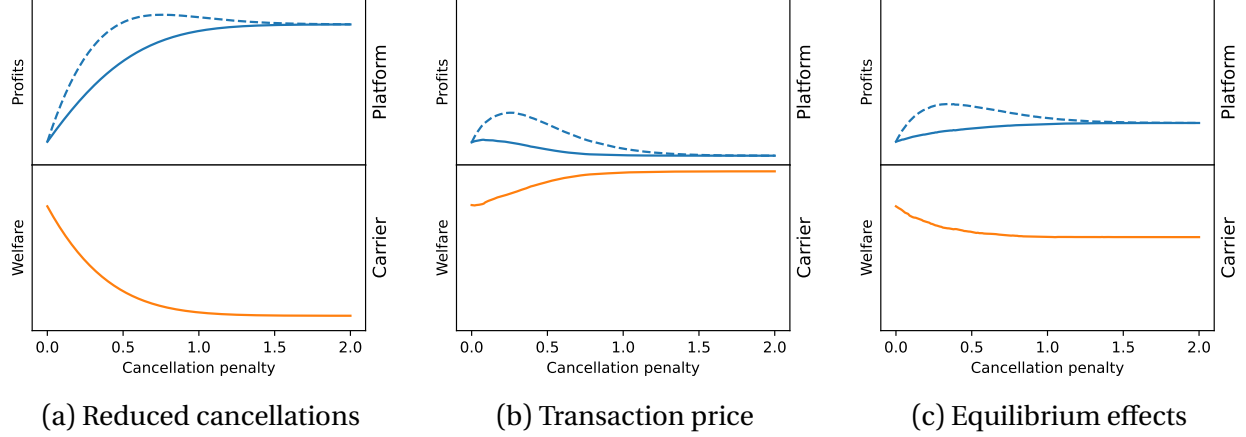
As alluded to previously, an increase in the firm's cancellation penalty will not only affect the propensity of carriers to cancel, but also their forward-looking bidding behavior, which ultimately changes the equilibrium of the auction. To understand how these effects work in conjunction, I conduct a simple numerical comparative statics exercise, progressively adding each effect. I also show the difference between reputational penalties, and monetary penalties, with the latter generating additional revenue for the firm with each cancellation. I fix $\mu_v = 1$, $\sigma_v = 0.5$, $\sigma_c = 0.65$, $\sigma_G = 0.3$, $\mu_G = 0.6$, $\sigma_G = 0.6$. As a baseline, I solve for the equilibrium of the model with the cancellation penalty κ set to 0. The results are illustrated in Figure 3.

In panel (a), I hold the carriers' bidding behavior and the auction outcomes fixed, and only adjust their reservation wages $R^m(b) = b + \kappa$. For a given bid b , the carriers' propensity to cancel, given by $1 - G(R^m(b))$, thus decreases with the penalty, resulting in a welfare loss for carriers, but a profit gain for the platform. When the penalty is reputational, the platform's profits are monotonically increasing in the penalty, but when the penalty is monetary, there is an interior optimum to the platform's profits, which suggests that the penalty is a means for the platform to extract some of the gains that carriers receive from outside offers.

In panel (b), I adjust the carrier bids, incorporating the increased penalty into the opportunity

13. The terms in parentheses are conditional on $b_i < v$, but since b_j and v , the expression can be simplified.

Figure 3: Cumulative effects of an increased cancellation penalty



NOTE: Figures show the cumulative effects of increasing cancellation penalty κ from a baseline equilibrium of $\kappa = 0$. Dashed lines represent platform payoffs under monetary penalties. In (a), carrier bids and auction acceptances/confirmations are held fixed, so that only cancellation behavior changes. In (b), the carrier bid is changed, while holding bidding equilibrium and auction acceptances/confirmations fixed. In (c), full equilibrium is adjusted, including equilibrium win probabilities and auction acceptances/confirmations. All quantities scaled relative to average shipment value to the platform. Carrier welfare is relative to non-existence of the platform.

cost term in Equation 5, while holding the auction outcomes and win probability beliefs $\gamma(b)$ fixed. The increase in opportunity costs is not uniform across types, since if a carrier with a high bid b is awarded the auction, they are unlikely to find a better opportunity even with no cancellation penalty. What the figure illustrates is that—in expectation—carriers increase their bids enough to more than offset the welfare loss from the increased penalty, when the auction equilibrium is held fixed. Consequently, it is less attractive for the firm to increase the penalty, though they may still collect additional revenue from doing so when the penalty is monetary.

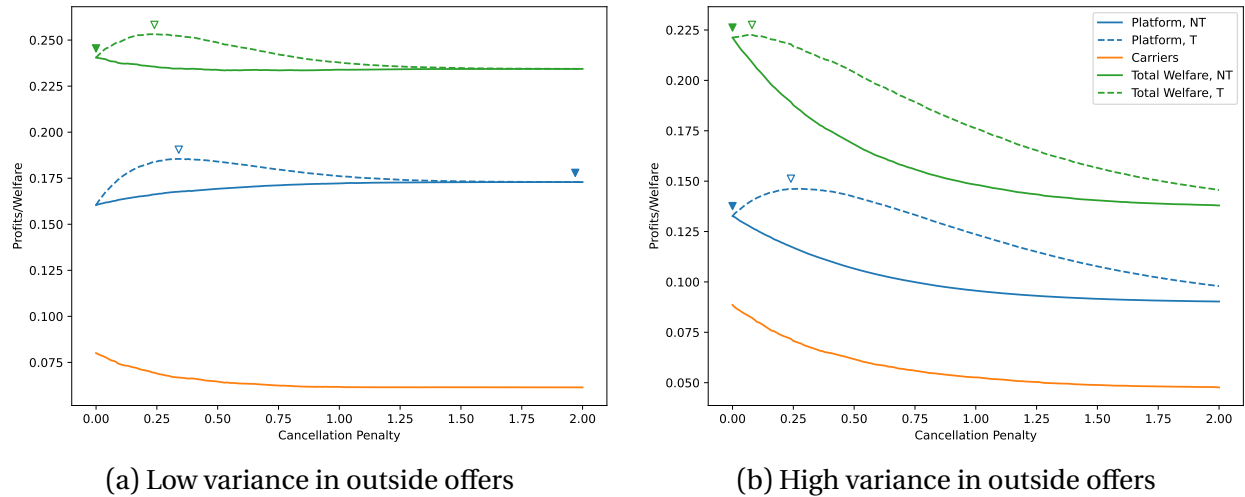
Finally, panel (c) illustrates the full equilibrium effects. This means the auction outcomes are changed, so that some of the previously winning bids now lose, as the bids have increased. It also means solving for a fixed point between carrier policies and the equilibrium win probability defined in Equation 6. As allocations change, some bids are no longer accepted, and some carriers pre-emptively choose to exit the market, reneging on their winning bids. This alters the equilibrium win probabilities, leading to a new equilibrium. Overall, this restores the pattern observed in panel (a), where the platform benefits from a higher penalty and carriers incur losses, albeit to a lesser extent. This exercise demonstrates the importance of each stage of the analysis in understanding the effects of cancellation penalties on the platform and

carriers.

3.5 Optimal Cancellation Penalties

In the next exercise I evaluate the social planner's objective and the platform's profits over a range of cancellation penalties, displayed in Figure 4. Once again, in the figure I distinguish between reputational and monetary penalties. I repeat the exercise for a scenario with higher variance in outside offers, with $\sigma_G = 1.2$. Intuitively, a higher variance in a carrier's outside option increases the value of the option to cancel.

Figure 4: Social Welfare vs. Firm Profits



NOTE: Dashed lines represent *monetary* penalties (T=Transfer), solid lines represent *reputational* penalties (NT=No Transfer). Triangles denote optimal penalty levels for respective objectives. All quantities are scaled relative to average shipment value to the platform. Carrier and total welfare is relative to non-existence of the platform. Low variance is $\sigma_G = 0.6$ and high variance is $\sigma_G = 1.2$.

I start by considering reputational penalties, given their use in the status quo. Since triggering reputational penalties does not generate any revenue, the firm would—all else equal—prefer to prevent cancellations altogether. But since carriers' bids depend on the penalty level, increasing the penalty could potentially leave the firm worse off because of the pass-through effect. I find that the net effect of these two forces depends crucially on the variance of the carriers' outside offers; when the variance is low (high) increasing the penalty monotonically increases (decreases) firm profits. In both cases, carrier welfare and overall welfare are decreasing in the penalty, although more starkly when the variance is high. The takeaway is that,

under reputational penalties, carrier and firm incentives are aligned when the welfare impact of the penalties is strong enough, thanks to the upstream effects on bidding behavior.

In the case of monetary penalties, both objective functions have interior maxima in both scenarios. However, the profit-maximizing penalty is higher than the social welfare-maximizing penalty. This aligns with the idea that penalties allow the platform to capture a portion of the rents that carriers receive from outside offers. Under monetary penalties, the platform and social planner’s objectives are generally not aligned, although worst-case divergence in optimal policies is less pronounced than under reputational penalties.

In summary, this illustrative model demonstrates that the welfare implications of different penalty regimes are theoretically ambiguous and depend heavily on the nature of the outside offer process as well as the type of penalty. The model also underscores the importance of considering the equilibrium pass-through of the penalties, rather than just the direct effects on cancellation behavior.

4 Empirical Model

The baseline model above captures the fundamental economic impacts of the cancellation penalties on carriers’ bidding strategies and reservation wages. However, it does not fully reflect all aspects of the empirical context. To enable a realistic empirical analysis, I enrich the model with additional features, including the multi-unit nature of the auctions, the stochastic arrival of carriers and shipments, the Accept-Now feature, and the multiple auction clearing rounds. With these additions the model better captures the key features of the platform described in Section 2.

Formally, this model takes place in continuous time, but as all events and decisions arrive according to Poisson processes, there is no reference to a continuous time index in what follows. The dynamics of the model are thus described by a continuous-time Markov chain, as in [Doraszelski and Judd \(2012\)](#).

The model is specified on a market-by-market basis, where one market is defined as a combination of a shipping lane (metro to metro) and a departure date. As the model parameters will be estimated on a lane-specific basis, I make no further reference to the lane below. The decision problem takes place in the days leading up to departure, with the state variable d indexing the number of full days until departure. As in the baseline model, the index counts

backwards from the pickup time, so that the period from 0 to 24 hours before departure is indexed as $d = 0$, the period from 24 to 48 hours before is indexed as $d = 1$, and so on. Carriers are indexed by i and shipments are indexed by j .

In the first half of this section I describe the individual decision problem of carriers and their optimal policies. In the second half I describe the aggregate dynamics of the platform, which are used to simulate the counterfactuals.

4.1 Carriers

As in the baseline model, the model of carrier behavior includes an optimal bidding function and a set of reservation wages. The empirical model adds a first-stage choice of whether to bid, use the Accept-Now feature, or ignore a shipment, as well as a richer cost structure to allow for match-specific costs in the multi-unit auction setting. In their interactions with a particular shipment, carriers can go through multiple states: unmatched, meaning they are certain not to take the shipment, matched, meaning they are certain to take the shipment, and pending bid, meaning they have placed and are waiting to see if they win the auction.

Carriers vary in their cost of matching with each shipment. Their shipment-specific cost c_{ij} consists of a mean cost $\bar{c}_i \sim N(\mu_d^c, \sigma_d^c)$ and an idiosyncratic term $\epsilon_{ij} \sim N(0, \sigma_\epsilon)$. The mean cost \bar{c}_i is interpreted as a combination of real marginal costs \tilde{c}_i and the opportunity cost of not taking an outside offer with payoff \bar{u}_i , so that $\bar{c}_i = \tilde{c}_i + \bar{u}_i$.¹⁴ The idiosyncratic component ϵ_{ij} captures a carrier's cost of taking on a particular shipment. Anecdotal evidence from the firm providing the data suggests that scheduling conflicts and relocation costs are the primary sources of idiosyncratic variation in costs.

In addition to shipments they view on the platform, carriers also receive outside offers off-platform. Whereas in the baseline model, the arrival times of the outside offers is deterministic, in the empirical model, they are stochastic. Each carrier is subject to a Poisson process of outside offers, with rate $\tilde{\lambda}_d$. Each offer is characterized by a payoff π_{ij} , which is distributed according to $N(\mu_d^\pi, \sigma_d^\pi)$.

Furthermore, each carrier is also subject to a stochastic attention mechanism. For each outside offer that arrives, or each auction win on the platform, carriers have a day-specific probability of paying attention α_d . It is also assumed that a carrier must be paying attention in order

14. The real marginal costs can include fuel, depreciation and other trip related costs, as well as the change in continuation values from arriving in a different location. This follows the notion of inclusive costs in [Buchholz et al. \(2020\)](#).

to confirm a bid. Ultimately, this means that the effective arrival rate of offers is $\lambda_d = \alpha_d \tilde{\lambda}_d$ and the acceptance rate of winning bids (derived in full later) is multiplied by α_d . This mechanism is reminiscent of the simple stochastic consideration sets of [Manzini and Mariotti \(2014\)](#), and is a reasonable assumption given the large proportion of carriers who are owner-operators and lack dedicated administrative staff to monitor their auctions when they are on the road or off-duty. This model feature also helps explain the high level of attrition at the pending bid stage and the low level of attrition at the matched stage.

Similarly to the baseline model, carriers' behavior can be summarized through a set of type-specific policy functions: the optimal bidding function $b_d^*(c_{ij})$, as well as the reservation wages $R_d^U(c_{ij})$, $R_d^P(b_{ij}, c_{ij})$, and $R_d^M(b_{ij}, c_{ij})$. These reservation wages determine the thresholds for accepting outside offers in the unmatched, pending bid, and matched states, respectively. The empirical model extends carrier policies with the addition of the first-stage choice probabilities $P_{d,ij}(x|c_{ij}, b_j^A)$, where x is the choice of whether to ignore, Accept Now, or bid on a shipment after viewing it. Shipment views are treated as exogenous. The model also increases the dimensionality of the carrier policies, with the time indices d denoting the current time and \bar{d}_i denoting the initial time of arrival to the market, which determines the carrier's cost distribution.

Another complication of the empirical setting is the possibility of bidding on multiple auctions simultaneously. In principle, the optimal policy of a carrier should depend on the entire *set* of shipments available to them. To reduce the dimensionality of the problem, I make a simplifying assumption about the carrier policies:

Assumption 1 *Carriers' first-stage choice probabilities $P_{d,ij}(x|c_{ij}, b_j^A)$, bidding function $b_d^*(c_{ij})$, and reservation wages $R_d^U(c_{ij})$, $R_d^P(b_{ij}, c_{ij})$, $R_d^M(b_{ij}, c_{ij})$ on a shipment j only depend on their cost c_{ij} for that shipment.*

This does not imply that carriers are myopic about other shipment opportunities on the platform. Instead, I assume that they incorporate their beliefs about future opportunities on the platform into their general belief about future outside options. This is similar to research by [Backus and Lewis \(2024\)](#), in which bidders have beliefs about future auction opportunities described by a Markov chain.

I also make a closely related assumption on carriers' knowledge of their costs:

Assumption 2 *In solving for their optimal policy, carriers only know the cost c_{ij} of the focal shipment and not their mean cost \bar{c}_i . They form Bayesian beliefs about their mean cost after*

observing c_{ij} , using $N(\mu_d^c, \sigma_d^c)$ as a prior distribution.

As will be seen further on, this assumption simplifies the computation of the likelihood of the model. If \bar{c}_i were known to carriers, but unknown to the econometrician, this would create an additional dimension of unobservables to integrate out. The strength of this assumption depends empirically on the relative variance of the mean and idiosyncratic cost components. The estimation results illustrate that the variance of the idiosyncratic component is relatively much smaller, so knowledge of the cost of one shipment is a good approximation to knowledge of the mean cost term.

However, it is a priori important to specify carriers' rational beliefs about their costs for outside offers in a manner consistent with their cost structure. For example, suppose the variance of idiosyncratic cost component ϵ_{ij} were zero, then the distribution of profits of outside offers would be fully specified by the distribution of the outside offer payoffs π_{ij} , independently of their cost for the focal shipment. On the other hand, if the variance of the idiosyncratic component is high, then a carrier that has a particularly low (high) cost on a focal shipment would expect that their costs on outside offers would be higher (lower), creating some dependence between the distribution of profits on outside offers and their cost on the focal shipment on the platform.

The Bayesian beliefs described in Assumption 2 approximate this effect. These beliefs have a straightforward analytic form under the normality assumptions on the outside offer payoff and carrier cost distributions. Formally, let the net payoff for an outside offer on alternative shipment j' be $\pi_{ij'} - c_{ij'}$ which can be rewritten as $\pi_{ij'} - (c_{ij'} - c_{ij}) - c_{ij}$. Let $\tilde{\pi}_{ij'} \equiv \pi_{ij'} - (c_{ij'} - c_{ij})$ be the net payoff of an outside offer relative to the focal shipment. Given the normality assumptions, the conditional distribution is:

$$\tilde{\pi}_{ij'} | c_{ij} \sim N\left(\mu_d^\pi + \rho(c_{ij}, \bar{d}_i), \sigma_{d, \bar{d}_i}^{\tilde{\pi}}\right) \Leftrightarrow \tilde{\pi}_{ij'} + \rho(c_{ij}, \bar{d}_i) | c_{ij} \sim N\left((\mu_d^\pi, \bar{d}_i)\right)$$

where $\rho(c_{ij}, \bar{d}_i) = \frac{\sigma_\epsilon^2(c_{ij} - \mu_{\bar{d}_i}^c)}{\sigma_\epsilon^2 + (\sigma_{\bar{d}_i}^c)^2}$, $\sigma_{d, \bar{d}_i}^{\tilde{\pi}} = \sqrt{(\sigma_d^\pi)^2 + (\sigma_{\bar{d}_i}^c)^2 + \sigma_\epsilon^2 - \frac{\sigma_\epsilon^4}{(\sigma_{\bar{d}_i}^c)^2 + \sigma_\epsilon^2}}$, and \bar{d}_i is the carrier's arrival date to the market. Thus, under Assumption 2, the beliefs about net outside offer payoffs are given by a common normal distribution (which I denote as $G_d(\cdot)$ in the following) shifted by the adjustment factor $\rho(c_{ij}, \bar{d}_i)$. This captures the fact that when a carrier is matched to a shipment at a high cost, their costs on other shipments are likely to be lower, and vice versa.

I derive the optimal carrier policies, which are composed of three parts:

1. The choice probabilities of whether to ignore, Accept Now, or bid on a shipment condi-

tional on viewing it.

2. The bidding function mapping the carrier's cost to their bid.
3. The state-dependent reservation wage to take on outside offers.

All three policies depend on the carrier's value functions in one of three states (relative to the focal shipment): unmatched, matched at price b , and having a pending bid b . These can be solved via backward induction, starting with the unmatched state.

Unmatched state The value function in the unmatched state is denoted by $U_{d,\bar{d}_i}(c_{ij})$. Whenever an unmatched carrier receives an outside offer in this state, they solve $\max\{U_{d,\bar{d}_i}(c_{ij}), \tilde{\pi}_{ij'} - c_{ij} + \rho(c_{ij}, \bar{d}_i)\}$, so that their reservation wage in this state is:

$$R_{d,\bar{d}_i}^U(c_{ij}) = U_{d,\bar{d}_i}(c_{ij}) + c_{ij} - \rho(c_{ij}, \bar{d}_i) \quad (7)$$

The unmatched value function is thus:

$$U_{d,\bar{d}_i}(c_{ij}) = \frac{1}{\eta + \lambda_d} \left[\lambda_d \left(\int_{R_{d,\bar{d}_i}^U(c_{ij})} (\tilde{\pi}_{ij'} - c_{ij} + \rho(c_{ij}, \bar{d}_i)) dG_d(\tilde{\pi}_{ij'}) \right. \right. \quad (8) \\ \left. \left. + G_d(R_{d,\bar{d}_i}^U(c_{ij})) U_{d,\bar{d}_i}(c_{ij}) \right) + \eta U_{d-1}(c_{ij}) \right]$$

with $U_{-1,\bar{d}_i}(c_{ij}) = 0$. Intuitively, the value function in the unmatched state is a weighted average between two events: the arrival of an outside offer and the arrival of a new day. When an outside offer arrives, the carrier only takes it if it exceeds their reservation wage. Otherwise, they remain in the current state.

See Appendix A.1 for the formal derivation of this value function from the continuous time setup. The other value functions that follow below can be derived in a similar manner.

Matched state The value function in the matched state is denoted by $V_{d,\bar{d}_i}(b_{ij}, c_{ij})$. The value in this state is a function of both the carrier's cost and the agreed upon payment b_{ij} , which comes from either the carrier's own bid or the Accept Now feature.

When a matched carrier receives an outside offer, they solve $\max\{V_{d,\bar{d}_i}(b_{ij}, c_{ij}), \tilde{\pi}_{ij'} - c_{ij} + \rho(c_{ij}, \bar{d}_i) - \kappa_d\}$. The platform's cancellation penalties κ_d enter into this decision in reduced form as a monetary equivalent to the reputational costs of cancelling a shipment. The carrier's

reservation wage in this state is

$$R_{d,\bar{d}_i}^M(b_{ij}, c_{ij}) = V_{d,\bar{d}_i}(b_{ij}, c_{ij}) + c_{ij} - \rho(c_{ij}, \bar{d}_i) + \kappa_d \quad (9)$$

The matched value function is thus:

$$V_{d,\bar{d}_i}(b_{ij}, c_{ij}) = \frac{1}{\eta + \lambda_d} \left[\lambda_d \left(\int_{R_{d,\bar{d}_i}^M(b_{ij}, c_{ij})} (\tilde{\pi}_{ij'} - c_{ij} - \kappa_d + \rho(c_{ij}, \bar{d}_i)) dG_d(R_{d,\bar{d}_i}^M(b_{ij}, c_{ij})) \right) + G_d(\tilde{\pi}_{ij'}) V_{d,\bar{d}_i}(b_{ij}, c_{ij}) \right] + \eta \bar{V}_{d-1}(b, c) \quad (10)$$

with $V_{-1,\bar{d}_i}(b_{ij}, c_{ij}) = b_{ij} - c_{ij}$. Similarly to the unmatched value function, the matched value function is a weighted average of the payoff from the arrival of an outside offer and the payoff from the arrival of a new day.

I can also derive the overall probability of a matched carrier cancelling a shipment, which will be used in constructing the model's likelihood. The cancellation probability of a matched carrier on day d is jointly determined by the arrival process of outside offers and the carrier's reservation wage, as follows:

$$P_{\bar{d}_i}(\text{cancel} | b_{ij}, c_{ij}) = \frac{\lambda_d (1 - G_d(R_{d,\bar{d}_i}^M(b_{ij}, c_{ij})))}{\eta + \lambda_d (1 - G_d(R_{d,\bar{d}_i}^M(b_{ij}, c_{ij})))} \quad (11)$$

Bidding A carrier with a pending bid b_{ij} has the value function $W_{d,\bar{d}_i}(b_{ij}, c_{ij})$. At any point in time, the carrier may win the auction at day-specific rate $\gamma_d(b)$. Upon winning, the carrier may accept or decline the shipment. In the model, they will only decline the shipment if they have already matched with another shipment, or if they are not paying attention.

While waiting on a pending bid, carriers continue to receive outside offers. When one arrives, they solve $\max\{W_{d,\bar{d}_i}(b_{ij}, c_{ij}), \tilde{\pi}_{ij'} - c_{ij} + \rho(c_{ij}, \bar{d}_i)\}$, so that their reservation wage in this state is

$$R_{d,\bar{d}_i}^P(b_{ij}, c_{ij}) = W_{d,\bar{d}_i}(b_{ij}, c_{ij}) + c_{ij} - \rho(c_{ij}, \bar{d}_i) \quad (12)$$

The pending bid value function is:

$$W_{d,\bar{d}_i}(b_{ij}, c_{ij}) = \frac{1}{\eta + \gamma_d(b_{ij}) + \lambda_d} \left[\lambda_d \left(\int_{R_{d,\bar{d}_i}^P(b_{ij}, c_{ij})} (\tilde{\pi}_{ij'} - c_{ij} + \rho(c_{ij}, \bar{d}_i)) dG_d(\tilde{\pi}_{ij'}) \right. \right. \quad (13) \\ \left. \left. + G_d(R_{d,\bar{d}_i}^P(b_{ij}, c_{ij})) W_{d,\bar{d}_i}(b_{ij}, c_{ij}) \right) \right. \\ \left. + \gamma_d(b_{ij}) (\alpha_d V_{d,\bar{d}_i}(b_{ij}, c_{ij}) + (1 - \alpha_d) U_{d,\bar{d}_i}(c_{ij})) + \eta W_{d-1}(b_{ij}, c_{ij}) \right]$$

with $W_{-1,\bar{d}_i}(b_{ij}, c_{ij}) = 0$. The value function is now a weighted sum of three events: the arrival of an outside offer, a new day, and the possibility of winning the auction. When a carrier wins an auction, they have the option to confirm the bid if they are paying attention, which occurs with probability α_d . If the bid is confirmed, the carrier enters the matched state. Otherwise, they remain unmatched.

I use this to construct the likelihood of bid confirmations in the data, using only knowledge of the day the carrier placed the bid, denoted d_{ij}^b , and the day the bid was accepted, denoted d_{ij}^a . The confirmation probability of a bid is given by:

$$P_{\bar{d}_i}(\text{confirm} | b_{ij}, c_{ij}, d_{ij}^a, d_{ij}^b) = \alpha_{d_{ij}^a} \prod_{k=d_{ij}^a}^{d_{ij}^b} \left(1 - \underbrace{\frac{\lambda_k (1 - G_k(R_{k,\bar{d}_i}^P(b_{ij}, c_{ij})))}{\eta + \gamma_k(b_{ij}) + \lambda_k (1 - G_k(R_{k,\bar{d}_i}^P(b_{ij}, c_{ij})))}}_{\text{Probability of taking an outside offer on day } k} \right) \quad (14)$$

Intuitively, the confirmation probability is simply the cumulative probability of the carrier not taking any outside offer between the bid being placed and being accepted, multiplied by the probability of paying attention.

I turn to the carrier's problem of choosing their optimal bid. I assume that bidders cannot change their bids as time goes on.¹⁵ The bidding problem at the time of first viewing the shipment can then be written as:

$$\max_{b_{ij}} W_{d_{ij}^b, \bar{d}_i}(b_{ij}, c_{ij}) \quad (15)$$

The first order condition—shown here in Equation 16—takes the form of a weighted sum of day-specific marginal utilities, each of which accounts for the benefit of a higher bid for a win

15. In reality, they can freely change their bids, but this is empirically rare, which may imply a high cognitive cost of re-optimizing the bid.

on each particular day. The full derivation is given in Appendix A.3.

$$\begin{aligned}
\frac{\partial W_{d,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} &= \sum_{k=0}^{d_{ij}^b} \left[\overbrace{\prod_{\ell=k}^{d_{ij}^b} \frac{1}{\eta + \gamma_\ell(b_{ij}) + \lambda_\ell(1 - G_\ell(R_{\ell,\bar{d}_i}^P(b_{ij}, c_{ij})))}}^{\text{Probability of bid surviving from } d \text{ to } k} \right] \eta^{-1} \\
&\quad \gamma'_d(b_{ij}) \left(\underbrace{\alpha_d V_{d,\bar{d}_i}(b_{ij}, c_{ij}) + (1 - \alpha) U_{d,\bar{d}_i}(c_{ij}) - W_{d,\bar{d}_i}(b_{ij}, c_{ij})}_{\text{Benefit of winning net of opportunity costs}} + \underbrace{\frac{\alpha_d \gamma'_d(b_{ij})}{\gamma'_d(b_{ij})} P_{d,\bar{d}_i}(\text{no cancel} | b_{ij}, c_{ij})}_{\text{Markup}} \right) \\
&= 0 \quad (16)
\end{aligned}$$

Let $b_d^*(c_{ij}, \bar{d}_i)$ denote the optimal bidding function implicitly defined by the first-order condition.

First-Stage choice When first viewing a shipment, a carrier can choose between three options: ignoring the shipments, bidding on the shipment, or using the Accept-Now feature, to immediately match with the shipment at the price set by the platform, denoted b_j^A , which is observed by the carrier. I can formulate the continuation value associated with each of these choices. Let $u_{ij}^k, k \in \{B, A, I\}$ be the continuation value of bidding, using the Accept-Now feature, and ignoring the shipment.

To increase the flexibility of the model and fit the empirical first-stage choices, I add additional non-structural shocks (which do not enter welfare calculations) to rationalize the observed behavior. I assume that carriers make their first stage choices based on a set of mean pseudo-values—combining the continuation value of a choice with a “hassle” cost—and a choice-specific Type-1 Extreme Value shock with scale parameter σ^{choice} .¹⁶ The mean pseudo-value of bidding is given by $u_{ij}^B = W_{d,\bar{d}_i}(b^*(c_{ij}), c_{ij}) - c_{bid}$. The mean pseudo-value of the Accept-Now feature is $u_{ij}^A = V_{d,\bar{d}_i}(b_j^A, c_{ij}) - c_{AN}$, with the value function of a matched shipment given by Equation 10. The mean pseudo-value of ignoring the shipment is $u_{ij}^I = U_{d,\bar{d}_i}(c_{ij})$. This formulation leads to the familiar multinomial logistic choice probabilities:

$$P_{ij}(X = k) = \frac{\exp(u_{ij}^k / \sigma^{choice})}{\sum_{\ell \in \{B, A, I\}} \exp(u_{ij}^\ell / \sigma^{choice})}, k \in \{B, A, I\} \quad (17)$$

16. Hassle costs and choice-specific shocks are omitted from welfare computations as they would mechanically increase carrier utilities for every shipment view. In addition, the estimates of these parameters are implausibly large, so their economic interpretation is unclear. However, the counterfactual results are qualitatively robust to alternative values of these parameters, so they are not the main focus of the empirical strategy.

To summarize, I have solved for the first-stage choice probabilities in Equation 17, the bidding function, defined implicitly through Equation 16, and the reservation wages in the unmatched, matched, and pending bid states in Equations 7, 9, and 12, respectively. Together, these fully specify the single-agent behavior of carriers on the platform. Next, I deal with the aggregate dynamics and equilibrium of the platform.

4.2 Platform

In the counterfactual simulations, the model of carrier behavior is combined with an aggregate dynamic model of the auction platform for each shipment pickup date. This model is used to simulate the arrival of carriers and shipments, the matching process, and the auction clearing process. As before, the model dynamics are described by a continuous-time Markov chain. At any point in time, one of the following events can occur:

1. A carrier i arrives to the market with mean cost \bar{c}_i
2. A shipment j arrives to the market with value v_j
3. An existing carrier i views an existing shipment j (drawing matching cost $\bar{c}_i + \epsilon_{ij}$), choosing to ignore, bid, or take the shipment at the Accept-Now price $p^{AN}(v_j)$
4. A carrier accepts an outside offer and exits the market, potentially cancelling their matched shipment on the platform
5. A shipment is cancelled and exits the market
6. An auction on the platform attempts to clear at reserve price $r_d(v_j)$
7. The days remaining until pickup count down from d to $d - 1$

These events occur at state-dependent *rates*, which determine the transition matrix of the Markov chain. The carrier and shipment arrivals are assumed to occur exogenously at rates $\lambda_d^{carrier}$ and $\lambda_d^{shipment}$. Each carrier-shipment pair that has not previously met can generate a view at independent rate λ_d^{view} . Following a view, carriers decide whether to ignore the shipment, bid, or take the shipment at the Accept-Now price based on the value functions derived above in section 4.1.

Carriers compute their outside offer acceptance threshold as the maximum of all thresholds on the set of pending bids or their currently matched shipment. As above, they receive outside offers at rate λ_d with payoff distribution $G(\cdot)$. When they accept an outside offer or a different offer within the platform, they take the payoff $\pi_{ij'} - c_{ij'}$ and exit the market, at the additional cost of a cancellation penalty κ_d if they were currently matched on the platform.

Shipment cancellations occur exogenously at rate $\lambda_d^{shipcancel}$. Because the focus of the paper is on the carrier side of the market, any potential welfare gains from a shipment cancellation are not modeled.

Every currently unmatched shipment reaches an auction clearing round at rate λ_d^{clear} . In a clearing round, all bidders below the reserve price $r_d(v_j)$ are offered the shipment in sequence, starting from the lowest bid, until a carrier accepts. If no bids below the reserve price confirm their bids, nothing happens. Bids that are rejected by carriers are deleted. Bids that have yet to be accepted by the platform remain pending.

As before, the rate of transition between days η is normalized to 1. Given the properties of the Poisson process, all other rates can thus be interpreted as the mean number of the corresponding event per day.

Finally, the equilibrium of the model is once again a fixed point between carrier policies and the conditional auction win probabilities. However, unlike the baseline model, these win probabilities do not have a closed form solution. They are solved for numerically through simulation of the platform dynamics.

5 Estimation

The estimation of the structural model is split into two parts. The parameters of carrier behavior, including the cost distributions and the arrival rates of offers, are estimated by maximizing the full likelihood derived from the carrier model above. The parameters governing arrivals to the platform, the distribution of shipment values, and the conditional reserve price and Accept-Now price models are estimated through the method of moments.

5.1 Carrier Model Estimation

1. **FIRST STAGE DECISIONS:** For each j , I observe the choice $x_{ij} \in \{B, A, I\}$, corresponding to the decision to bid, Accept-now, or ignore the shipment. For carrier i , let $K_i^B \subseteq K_i$ denote the set of bids, $K_i^A \subseteq K_i$ the set of Accept-Now matches, and $K_i^I \subseteq K_i$ the set of ignored shipments.
2. **BIDS:** For shipments $j \in K_i^B$, I also observe the carrier's bids $b_{ij} \in \mathbb{R}_+$.
3. **CONFIRMATIONS:** A subset of bids $\bar{K}_i^B \subseteq K_i^B$ are winning bids, giving the carrier the

opportunity to confirm their bid. I thus observe $y_{ij} \in \{1, 0\}$ for $j \in \bar{K}_i^B$, where 1 denotes a confirmation. I let $\tilde{K}_i^B \subseteq \bar{K}_i^B$ denote the set of shipments that a carrier confirms.

4. CANCELLATION DECISIONS: For shipments matched by confirming a bid or using Accept-Now ($\tilde{K}_i^B \cup K_i^A$), I also observe, for each valid day before pickup d , whether the carrier canceled on that day or not, so I have cancellation decisions $z_{ijd} \in \{1, 0\}$ for $j \in \tilde{K}_i^B \cup K_i^A$, where 1 denotes a cancellation. Let z_{ij} denote the vector of cancellation decisions.

In addition, the likelihood needs to account for the unobservable costs. This includes both the mean carrier cost \bar{c}_i and the shipment-specific costs $c_{ij} = \bar{c}_i + \epsilon_{ij}$. To start, I assume that \bar{c}_i is also observed, so that I can condition the likelihood on it, before eventually integrating it out.

I start by forming the likelihood of the observed bids. Recall that I can use the first-order condition of the carrier's bidding problem to invert the bid into their cost, so the likelihood of bids is obtained from the distribution of costs that I aim to estimate. I denote this inversion as $c_{ij} = b_{d_i}^{*-1}(b)$. Recall that $c_{ij} = \bar{c}_i + \epsilon_{ij}$, $\epsilon_{ij} \sim N(0, \sigma_\epsilon)$. I can then write the likelihood of a bid conditional on \bar{c}_i and parameters θ as:

$$\mathcal{L}(b_{ij}|\bar{c}_i, \theta) = \phi((b_{d_i}^{*-1}(b) - \bar{c}_i)/\sigma_\epsilon) \left| \frac{\partial b_{d_i}^{*-1}(b)}{\partial b_{ij}} \right| \quad (18)$$

where $\phi(\cdot)$ denotes the density function of a standard Normal distribution. The derivative term inside the absolute value is needed to properly account for the transformation of the random variable.

Below, I derive the remaining component of the likelihood by first assuming that I know the shipment specific cost c_{ij} . This is true for shipments a carrier has bid on, as I can invert the bid to obtain the cost. For the remaining shipments, I integrate out the unobservable cost.

The first stage choices are obtained straightforwardly from Equation 17. I denote their conditional likelihood as $P(x_{ij}|c_{ij}, \theta)$.

Next, the bid confirmation probabilities are obtained from Equation 14. These are only relevant for observations with bids, and can thus all be conditioned on the bid and the corresponding inverted cost $c_{ij} = b^{*-1}(b_{ij})$

$$P(y_{ij}|b_{ij}, \theta) = P_d(\text{confirmed}|b_{ij}, c_{ij})^{y_{ij}} (1 - P_d(\text{confirmed}|b_{ij}, c_{ij})^{y_{ij}})^{(1-y_{ij})} \quad (19)$$

where y_{ij} is shorthand for the decision to confirm a bid, conditional on winning. Finally, I ob-

tain the cancellation probabilities from Equation 11. These cover every day before pickup d during which a shipment was matched, which rules out multiple cancellations for one shipment.

$$P(z_{ij}|b_{ij}, c_{ij}, \theta) = \prod_{d=\underline{d}_{ij}}^{\bar{d}_{ij}} P_d(\text{cancel}|b_{ij}, c_{ij})^{z_{ij}} (1 - P_d(\text{cancel}|b_{ij}, c_{ij}))^{(1-z_{ij})} \quad (20)$$

where z_{ij} is shorthand for the vector of cancellation decisions. Computing the full likelihood for a single carrier index i requires careful consideration of which likelihoods are jointly integrated.

The panel structure of the data requires that the unobservable mean cost \bar{c}_i be jointly integrated out over the whole set of decisions of a single carrier. Formally, suppose a carrier bids on every shipment they see. Then the joint probability of their actions is:

$$\int \prod_{j \in K_i^B} [\mathcal{L}(b_{ij}|\bar{c}_i, \theta)] \underbrace{\prod_{j \in \tilde{K}_i^B} [P(y_{ij}|b_{ij}, \theta)] \prod_{j \in \tilde{K}_i^B} [P(z_{ij}|b_{ij}, c_{ij}, \theta)]}_{\text{Independent of variable of integration } \bar{c}_i} dF(\bar{c}_i)$$

Since the terms relating to confirmation and cancellation decisions are fully determined by the observed bid b_{ij} and the inverted cost c_{ij} , these can be moved outside the integral. This simplification is directly attributable to Assumption 2, which requires that carriers only make use of c_{ij} in the computation of their optimal policies. For carriers that also use the Accept-Now option or ignore some viewed shipments, additional terms are added to the integral, as described in the following.

For any shipments where a carrier used Accept-Now, no bid is observed. I first write the likelihood conditional on \bar{c}_i , and integrate out the idiosyncratic component ϵ_{ij} jointly over the first-stage choice and cancellation decisions, as follows:

$$P^A(x_{ij}, z_{ij}|\bar{c}_i, b_{ij}^A, \theta) = \int P(x_{ij} = A|\bar{c}_i + \epsilon, \theta) P(z_{ij}|b_{ij}^A, \bar{c}_i + \epsilon, \theta) dF(\epsilon) \quad (21)$$

where x_{ij} and z_{ij} are shorthand for the first-stage choice and cancellation decisions, respectively.

For shipments that were ignored by a carrier, I can simply integrate the likelihood of the first stage choice:

$$P^I(x_{ij}|\bar{c}_i, \theta) = \int P(x_{ij} = \text{Ignore}|\bar{c}_i + \epsilon, \theta) dF(\epsilon) \quad (22)$$

Now, I can gather all the shipments of carrier i , and write the likelihood of all terms to be

integrated over the unobserved \bar{c}_i :

$$\mathcal{L}_i^{int}(\theta) = \int \left[\prod_{j \in K_i^B} \mathcal{L}(b_{ij}|\bar{c}_i, \theta) \prod_{j \in K_i^A} P^A(x_{ij}, z_{ij}|\bar{c}_i, b_{ij}^A, \theta) \prod_{j \in K_i^I} P^I(x_{ij}|\bar{c}_i, \theta) \right] dF(\bar{c}_i) \quad (23)$$

And thus, the full likelihood of the model is:

$$\mathcal{L}(\theta) = \prod_{i=1}^N \left[\mathcal{L}_i^{int}(\theta) \underbrace{\prod_{j \in \bar{K}_i^B} P(y_{ij}|b_{ij}, \theta)}_{\text{Confirmations}} \underbrace{\prod_{j \in \bar{K}_i^B} P(z_{ij}|b_{ij}, b^{*-1}(b_{ij}), \theta)}_{\text{Cancellations}} \right] \quad (24)$$

While in principle it is possible to use the full likelihood to jointly estimate all parameters of the model, it is computationally expensive to do so and it is more sensitive to model misspecification. For example, if the model of first-stage choices—which is not a primary focus of this paper—is misspecified, the parameters relating to the outside offer process and cancellation penalties may be biased to create a better fit. This is a particular concern because there are many more observations of first-stage choices than of confirmations and cancellations. For these reasons, I adopt a two-step estimator. Let $\theta_1 = (\kappa, \lambda, \mu^\pi, \sigma^\pi)$ denote the parameters of the outside offer process and the cancellation penalties, and $\theta_2 = (\mu_c, \sigma_c, \epsilon_c, c_{bid}, c_{AN}, \sigma_{choice})$ denote the parameters of the cost distribution and the first-stage choices. The first step estimates θ_1 by maximizing the likelihood of the confirmation and cancellation decisions, while the second step estimates θ_2 by maximizing the full joint likelihood of the first-stage choices, taking the estimated θ_1 as given. The identification arguments I develop below focus on the first step of the estimation to justify the sufficiency of cancellation decisions to identify the outside offer process and the penalty schedule.

Implementation Details Here I discuss additional restrictions imposed on the model for estimation purposes, as well as the implementation of an economically motivated assumption of non-negative costs, which aid in identification. For details on the estimation of the asymptotic variances of the estimates, the reader is referred to Appendix A.6.

To start, while the model is derived by normalizing utilities into monetary equivalents, they are not estimated as dollar amounts. Instead, they are estimated as a proportion of the current market rate, as in Section 2. The trucking industry experiences regular seasonal fluctuations, in addition to cost variations due to macroeconomic factors (such as the fuel price shock caused by the Russian invasion of Ukraine in 2022). To remove the noise from these fluctuations and pool the data across periods, all price variables in the data are first divided

by the platform's internal estimate of the prevailing spot price for shipments along the lane, obtained from third-party data. This will be referred to as the *market rate*, and all relevant structural parameters will be expressed as a fraction of this market rate.

I then impose monotonicity and non-negativity assumptions on κ_d , so that $\kappa_{d-1} \geq \kappa_d \geq 0$, $\forall d$. If cancellation penalties were not monotonically increasing as the pickup time draws closer, it would be in the interest of a carrier to hold off on cancelling even when they already have a better offer in hand. To discipline the time path of the parameters governing the outside offer distribution $(\mu_d^\pi, \sigma_d^\pi, \lambda_d)$, I impose a second-order polynomial structure in the number of days until pickup. Furthermore, I assume that all day-specific parameters are shared between $d = 4, 5, 6, 7$ as no bids are accepted by the platform before then, so there is limited information on cancellation and bid renegeing behavior.

Before estimating the parameters of the structural model, I first estimate the conditional win probabilities of the auction at the lane level. The FOC in Equation 16 involves the day-specific win rates $\gamma_d(b)$ and their derivatives $\gamma'_d(b)$. While these can be solved for through market equilibrium conditions (and are, in the counterfactual simulations), doing so for every evaluation of the likelihood is computationally expensive, as it involves solving for a fixed point in the space of bidding strategies (the optimal bid is a function of the win probabilities and the win probabilities are a function of the optimal bidding strategies). To avoid this computational burden, I estimate the win probabilities off-line from the data directly, in the spirit of [Guerre, Perrigne, and Vuong \(2000\)](#).¹⁷ This estimation step is detailed in Appendix A.4.

Furthermore, as the model treats the confirmation and cancellation decisions of a carrier on different shipments as separate, I filter the data to only include confirmations and cancellations on the *first* bid of each carrier-day pair that is accepted by the platform. This approach sidesteps the issue of correlated confirmation and cancellation decisions, which would otherwise necessitate integrating over the arrivals of unobserved external offers in the likelihood calculation, adding a significant computational burden.

Finally, the maximum likelihood estimation is augmented with a penalty on negative inverted costs. Given the interpretation of costs in the model, there is no economic rationale for negative costs. This implies a restriction on the model parameters θ , which is difficult to characterize analytically, which is why it is enforced through a penalty term. Specifically, rather than

17. In standard auction settings as studied by [Guerre, Perrigne, and Vuong \(2000\)](#), the win probability is a straightforward function of the bid distribution (which can be efficiently nonparametrically estimated) and the number of opponents. The complex auction format studied here does not permit such a closed-form relationship, so the conditional win probability is estimated directly.

simply maximizing the log-likelihood of the model, the estimator solves:

$$\max_{\theta} \log \mathcal{L}(\theta) + \psi \left[\sum_{i=1}^N \sum_{j=1}^{k_i^B} \min\{c_{ij}, 0\} \right]$$

where the costs are obtained by the inversion $c_{ij} = b^{*-1}(b_{ij})$ and where the tuning parameter ψ is chosen to be large enough to ensure that the overwhelming majority of estimated costs are non-negative, but small enough to avoid erratic behavior in the optimization routine. This non-negativity penalty plays a key role in the identification argument that follows.

Identification I begin by discussing the main identification challenge that is unique to this paper, which concerns the distribution of outside offers. It is not directly observed, but can be inferred from the conditional cancellation probabilities in the raw data, presented in Figure 2. However, as the cancellation decision depends linearly on the cancellation penalty, the *location* of the outside offer distribution is not separately identified from the penalty level based on the conditional cancellation probability alone.

To deal with this, I develop a partial identification argument to separately identify these objects based on the non-negativity condition on inverted costs. To enable a tractable analysis, I develop the argument using the baseline model described in Section 3. Note that the following argues that the parameters are *at least* partially identified, without making use of any parametric restrictions imposed in the estimation, as discussed above. It is likely that these additional restrictions provide further identifying power.

While not shown formally here, it is likely that the additional parametric restrictions on the shape of the outside offer distribution and its evolution over time provide additional constraints on the model that make it point-identified.

Using the threshold $R_1^m(b) = b + \kappa$, the cancellation probability can be written as:

$$P(\text{cancel}|b_i) = 1 - G(b + \kappa) \quad (25)$$

Let $\tilde{G}(x) = G(x + \kappa)$ be the location normalized distribution, which is directly identified from the cancellation probability above. Taking the bid b as given, the first-order condition for optimal bidding from Equation 16 can be re-written in terms of the normalized distribution as follows:

$$b - c - \frac{(1 - \tilde{G}(b))\kappa + \int_{c-\kappa}^b (\tilde{\pi} - c + \kappa) d\tilde{G}(\tilde{\pi})}{\tilde{G}(b)} + \frac{\gamma(b)}{\gamma'(b)} = 0 \quad (26)$$

The derivative of the implicit function of c with respect to κ is:

$$\frac{dc}{d\kappa} = -\frac{1 - \tilde{G}(b) + \int_{c-\kappa}^b \tilde{G}'(\tilde{\pi}) d\tilde{\pi}}{\tilde{G}(b) + \int_{c-\kappa}^b \tilde{G}'(\tilde{\pi}) d\tilde{\pi}} < 0, \quad \forall b \text{ s.t. } \tilde{G}(b) < 1 \quad (27)$$

Now, let \underline{b} be the lowest bid in the data. If $\tilde{G}(\underline{b}) < 1$, that is, the (data-derived) conditional cancellation probability at the lowest bid is less than 1, then the inverted cost $b^{*-1}(\underline{b})$ is strictly decreasing in the penalty κ . Given the lower bound $c \geq 0$, this provides an upper bound on the penalty κ . The lower bound $\kappa \geq 0$ is assumed. For any κ in the identified set, there is a corresponding G^κ . Thus, κ and the location of $G(\cdot)$ are jointly set-identified.

The Monte Carlo exercise in Appendix A.5 provides further evidence for the identification of the parameters of interest and examines the identification of the attention probability α from the addition of the confirmation probabilities.

This identification argument also supports the estimation of the outside offer distribution parameters and cancellation penalties from the confirmation and cancellation behavior alone. Conditional on these parameters, the identification of carrier's cost distribution and the parameters governing their first-stage choice probabilities are straightforward.

The cost distribution is identified from a combination of the costs obtained from the bid inversion and the first-stage choice probabilities. Carriers with very low costs for a ship will generally find it more attractive to use the Accept-Now option, as the posted price will be closer to their optimal bid, while carriers with very high costs for a shipment will generally prefer to ignore it, as they stand very low chances of winning. Thus, the first-stage choice probabilities “fill in” the truncated regions of the cost distribution. The use of the joint likelihood helps to separately identify the variance of the idiosyncratic cost term ϵ_{ij} from the variance of the mean cost term \bar{c}_i .

Finally, the parameters governing the first-stage choice probabilities are identified from the first-stage choices alone and effectively act as a residual. The variance of the idiosyncratic cost term ϵ_{ij} is identified from the elasticity of choices to the value functions associated with each choice, while the hassle costs help to fit the overall average choice probabilities.

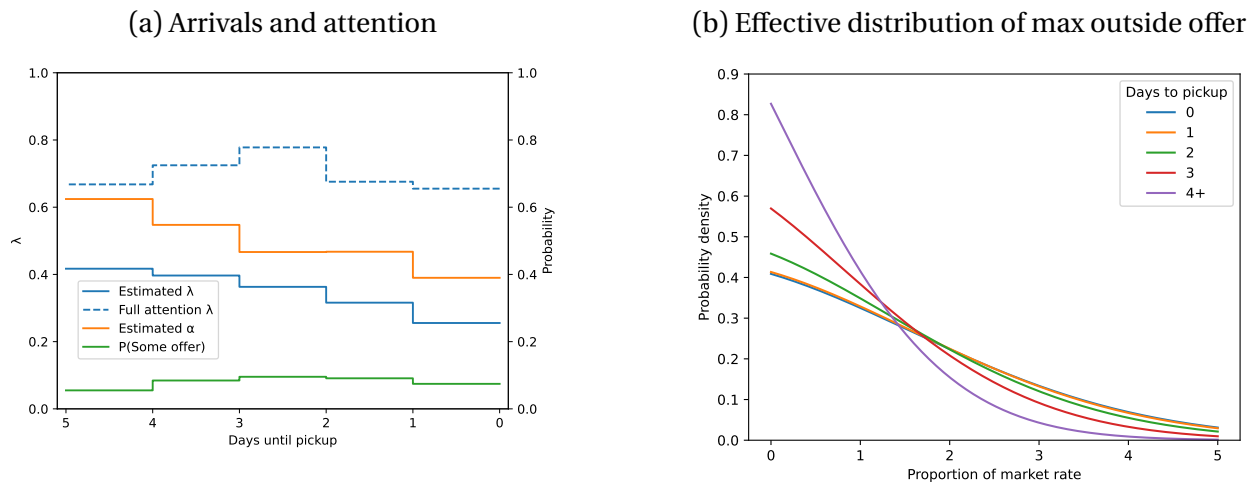
With these identification arguments in place, I now turn to the estimation results.

Carrier Model Estimates Recall that the carrier model is estimated in two steps, with the first step estimating the parameters of the outside offer process and the penalty schedule, and

the second step estimating the carrier cost distribution and the parameters governing the first-stage choice probabilities. I begin by presenting the results of the first step of the estimation. In Table A.2 in the appendix I report the raw structural parameter estimates from this step and their asymptotic standard errors, showing that all parameters are precisely estimated.

Formally, the stochastic process of outside offers consists of four elements: the arrival rates of offers λ_d , the mean and standard deviation of the distribution of outside offers μ_d^π, σ_d^π , and the attention probability α . However, what is effectively identified under the arguments in the previous section is the amalgamated distribution of the maximal outside offer received by a carrier within a day.

Figure 5: Features of the carrier search process

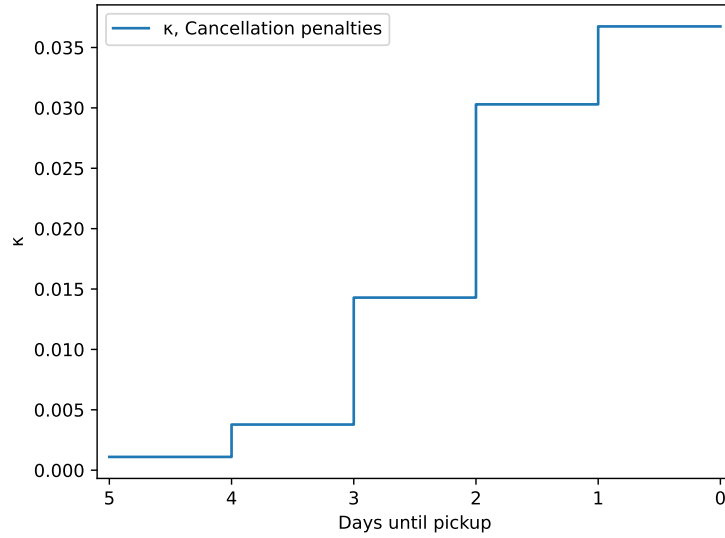


NOTE: In (a), P(Some offer) is the probability of an offer arriving and being non-negative. In (b), the effective distribution is integrated over the number of offer arrivals from the Poisson process, and is conditioned on the offer being non-negative.

The left-hand side of figure 5 presents the offer arrival rates and attention probabilities, along with the *overall probability of receiving a non-negative offer*. The estimates imply that, although the arrival rates are declining due to declining attention being paid, the overall probability of receiving an offer is slightly increasing, due to the quality of offers changing. The right-hand presents the effective distribution of the maximal outside offer received by a carrier, which shows additional mass at the right of the distribution as the pickup date approaches, which may reflect the added urgency of shippers.

The cancellation penalties are presented in Figure 6, which shows a near-zero cancellation penalty until the 48-hour pickup window, which is consistent with the platform's officially stated policy. Even within the window, the penalties are fairly modest, reaching a maximum of

Figure 6: Cancellation penalty



3.5% of the market rate, which is consistent with the notion that relational penalties are weak in the spot market.

I now turn to the estimates from the second step of the estimation, presented in Table 4, which include the parameters of the carrier cost distribution and those governing the first-stage choices. The mean carrier cost is steady over nearly all days, but is substantially lower on the very last day, which reflects that carriers, like shippers, have a greater urgency to match as the pickup date nears, though the greater variance on the last day indicates that there is substantial heterogeneity in this respect. The variance of the carrier-shipment idiosyncratic cost shocks is substantially smaller than the variance across carriers, with an average variance ratio of 21.56%.

As described in the model section, I do not ascribe a structural significance to the first-stage choice parameters. Indeed, the estimated variance of the idiosyncratic term in the first-stage choice model is quite high, which is likely due to the model's lack of unobserved heterogeneity in any dimension other than cost. Furthermore, the hassle cost of bidding is negative, while the hassle cost of using the Accept-Now option is implausibly large. These parameters are likely to be affected by the model's misspecification, but as previously discussed, this does not pose a great concern for the results of the counterfactual simulations.¹⁸

18. Counterfactual simulations with alternative parameters for the first-stage choice model (with a much lower σ^{choice}) yielded similar results.

Table 4: Second stage estimates

$d =$	0	1	2	3	4
μ_d^c	0.624	0.801	0.806	0.817	0.792
σ_d^c	0.219	0.180	0.172	0.186	0.214
<hr/>					
$\forall d$					
σ^ϵ	0.0902				
σ^{choice}	0.247				
c_{bid}	-0.0633				
c_{AN}	0.858				

NOTE: Table presents parameters estimated from the second step. Upper half reports day-specific parameters, lower half reports parameters that are common across days. μ_d^c and σ_d^c are the mean and standard deviation of the distribution of the mean cost \bar{c}_i and σ^ϵ is the standard deviation of the idiosyncratic term σ^ϵ . c_{bid} and c_{AN} are the hassle costs of bidding and using the Accept-Now option, respectively. σ^{choice} is the standard deviation of the idiosyncratic term in the first-stage choice model.

Model Fit Next, I evaluate the fit of the model to the data. Table 5 shows that the overall fit of the model in terms of confirmation and cancellation rates is high, with the model closely matching the data. Figure 7 is presented to check the fit along the temporal dimension of the model. The fit of cancellation rates over time is very close, but the confirmation rates are too low early on, and too high in the 72-24 hour window before pickup, and these cancel each other out in the aggregate.

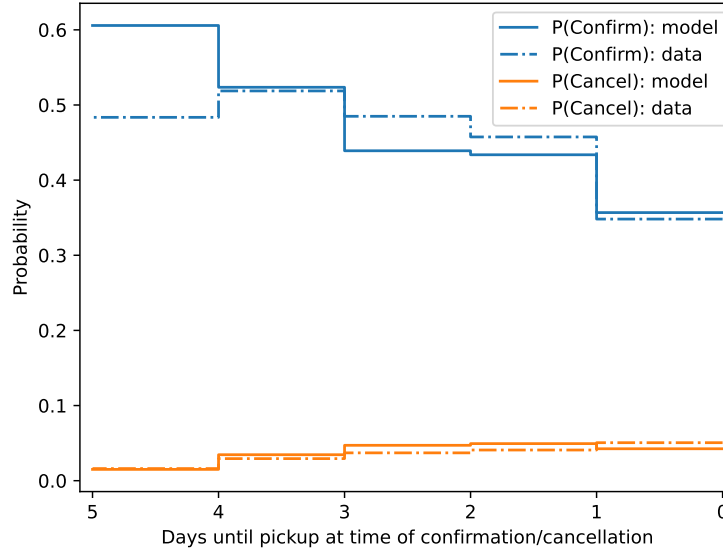
Table 5: Model fit: renegeing behavior

Type	Data	Model
Confirmation	44.29%	44.04%
Cancellations (bids)	9.08%	9.25%

5.2 Platform Parameters

The additional parameters to be estimated are the arrival rates of carriers, shipments, and views (a proxy for the search process on the platform), the distribution of shipment values to the platform, and conditional reserve price and Accept-Now price models. These parameters are all estimated directly in reduced form from their empirical analogues as they all involve observable events. More details on the estimation method, the parameter estimates and their standard errors can be found in Appendix A.7. In addition to the estimated parameters, the

Figure 7: Confirmation and cancellation probabilities in model vs. data



first-stage choice parameters of carriers are calibrated using the platform model to match the empirical first-stage choice probabilities. This calibration will eventually be replaced by the full likelihood estimation of the carrier model described above.

Figure 8 shows the estimated arrival rates of various events on the platform. The rate of view events is defined at the carrier-shipment level; for example, between 48 and 24 hours before pickup, each carrier-shipment pair that has not previously been viewed has a 15% chance of being viewed by the end of the period.

The shipment and carrier arrival rates reveal an important asymmetry: while shipments are mostly posted to the platform with at least a week of lead time,¹⁹ the arrival rate of carriers is at its highest in the last few days, before dropping substantially in the last 24 hours before pickup. This asymmetry has implications for the design of the penalty schedule, as it gives the platform many options in the last few days before pickup.

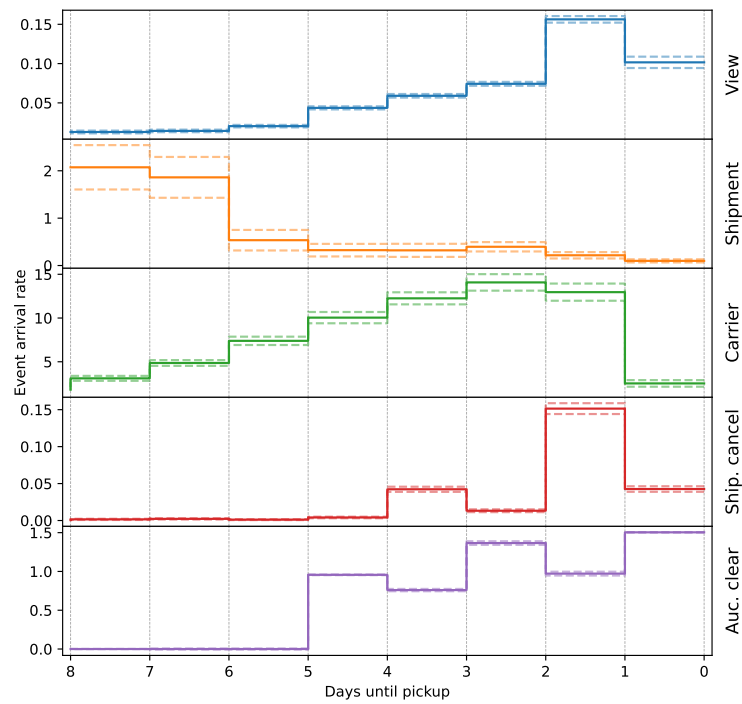
Shipment cancellations are also estimated directly from the data and are part of the counterfactual simulations. They are, however, abstracted away from the carrier's decision model, because the existing industry norm is to pay pecuniary cancellation penalties to carriers when cancelling shipments close to the pickup time.²⁰

Finally the auction clearing rounds are also modelled stochastically. In general, the platform

19. Most shipments are posted to the platform before the arrivals of the first carriers. To account for this, I draw an initial set of shipments from a Poisson distribution with rate $\lambda = 11.92$

20. Referred to as "TONU" (Truck Order Not Used) in the industry.

Figure 8: Arrival rates of events

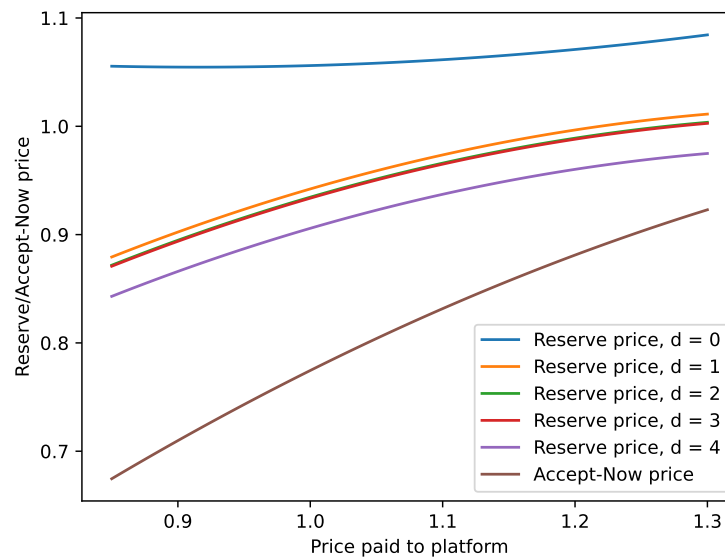


NOTE: Estimated Poisson arrival rates of carriers, shipments, views, shipper cancellations, and auction clearing rounds on the platform. Dashed lines represent the 95% confidence intervals.

would schedule at least one clearing round per auction per day, and sometimes more, beginning when there were fewer than 5 days remaining until pickup. In the last 24 hours, however, timed auction format is replaced by a continuous auction format with escalating prices and frequent manual broker intervention. The clearing rate for this phase is approximated by taking the count of unique hours in which bids were accepted on the last day.

The other feature of auctions is their reserve prices and Accept-Now prices, which I model as second-order polynomials of the rate paid *to the platform* from the shipper. As before, the estimation details of these parameters and their standard errors can be found in Appendix A.7. These are plotted in Figure 9.

Figure 9: Reserve and Accept-Now prices



NOTE: Reserve prices for $d \geq 1$ and Accept-Now prices estimated via linear regression on second-order polynomials of the shipment price paid to the platform. Reserve price regression is pooled, with a separate intercept for each d . Reserve price on $d = 0$ uses data from the fallback auction model with manual intervention, using inequality conditions on highest accepted bids and lowest unaccepted bids for each shipment.

All pricing functions are increasing in the rate paid to the platform, which is consistent with the platform's profit-maximizing behavior, while the Accept-Now price is substantially lower than the reserve prices on any day. In the last 24 hours, the true reserve price is not observed, due to the change in formats. Instead, I use the highest accepted bid and the lowest unaccepted bid as bounds to estimate the reserve price. The much higher curve may imply that the platform is more desperate to match very high-paying shipments in the last 24 hours.

Finally, the distribution of shipment values to the platform is fitted to a log-normal distribu-

tion resulting in a mean shipment value of 1.002. See Appendix A.7 for additional details. The cost of not matching a shipment is calibrated at 10% of the shipment value, which is the same heuristic used by the platform in its development of the newer auction format introduced in 2023.

The estimated structural model is then used to simulate counterfactuals, discussed in the following section.

6 Counterfactuals

The counterfactual analysis explores the effect of alternative cancellation policy designs on platform profits and overall welfare. Several features of the design are of interest. One, what is the effect of changing the overall level of the penalties? Two, how does the type—reputational or monetary—of the penalty affect outcomes and incentives?

To evaluate the optimal level and type of the penalties, I consider counterfactual policies that are linear transformations of the status quo penalty schedule. That is, given the estimated penalty schedule $\hat{\kappa}_d$ for day d , I consider a counterfactual penalty schedule $\kappa_d = s\hat{\kappa}_d$, where s is a scaling factor. This allows me to explore the effect of changing the overall level of penalties, while keeping the shape of the penalty schedule constant. I consider other variations in the shape of the penalty below.

Simulating outcomes under each counterfactual policy involves solving for the fixed point of win probabilities and carrier policies, which is why I restrict attention to a linear grid.²¹ The grid is composed of 128 linearly spaced points such that the maximal penalty varies between 0 and approximately 1.2 times the market rate for a single shipment. To reduce simulation variance, a spline is fitted to the simulation results, thereby averaging out noise across the grid. Figure A4 in the Appendix compares the raw simulation outcomes to the fitted splines.

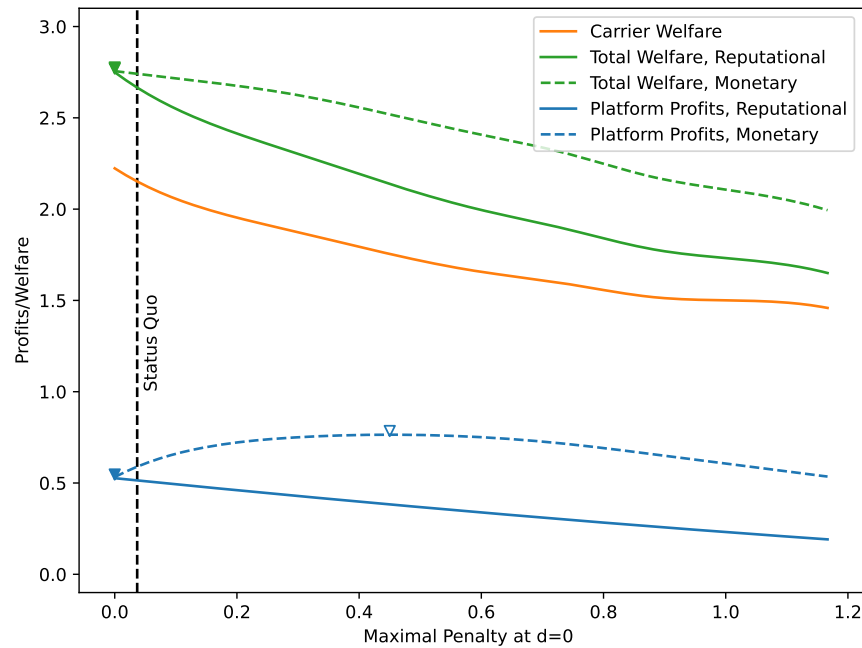
The main outcomes of interest are average platform profits, carrier welfare, and total welfare per market (defined as a lane and pickup date combination). Because a large portion of carrier welfare stems from offers *outside* the platform, I subtract the carrier’s welfare under a no-platform counterfactual to obtain welfare numbers in platform value-added terms. As in the stylized model in Section 3, I also distinguish profits and overall welfare under reputational

21. With five possible days of cancellation and the restriction that penalties are non-decreasing, exploring a grid of K values for each day would require $\binom{5+K-1}{K-1}$ simulations. Simulating a single counterfactual policy requires approximately 15 hours of computation time on a single thread.

penalties (as employed in the status quo) and monetary penalties, which involve a transfer from cancelling carriers to the platform, providing the latter with an additional source of revenue. In the following sections, I discuss the optimal reputational penalty, and then compare the trade-offs between monetary and reputational penalties.

6.1 Optimal Reputational Penalty

Figure 10: Visualization of counterfactual penalty policies.

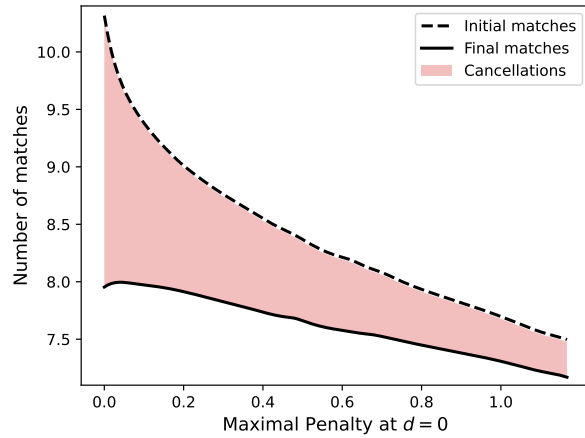


NOTE: All quantities scaled by average lane market price of a shipment. Carrier welfare is relative to platform not existing. Maxima of platform profits and total welfare are indicated by triangles in the corresponding color.

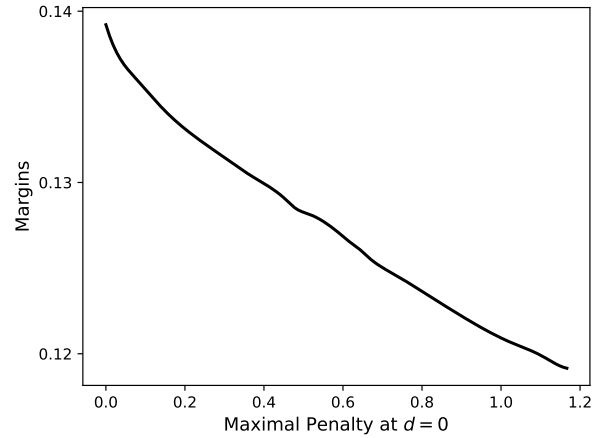
Figure 10 presents the evolution of profits and welfare over the range of counterfactual policies. It is immediately noticeable that the optimal reputational penalties for both social welfare and platform profits are at zero, so that fully flexible agreements are optimal. This is consistent with the stylized model under high variance of the outside offers, depicted in Figure 4b. The status quo reputational penalties, depicted in Figure 6, are very close to this optimum.

To further understand why the profits are decreasing in the level of the reputational penalty, I decompose the change in profits into three different margins: a change in the number of successfully completed shipments (the extensive margin), the firm's gross profit margin on each shipment (the intensive margin), and the indirect losses suffered by the firm as a consequence of shipments that failed to be matched in time.

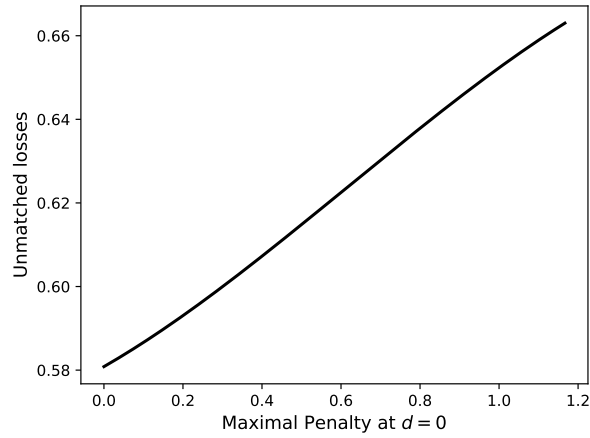
Figure 11: Decomposition of profits under reputational penalties.



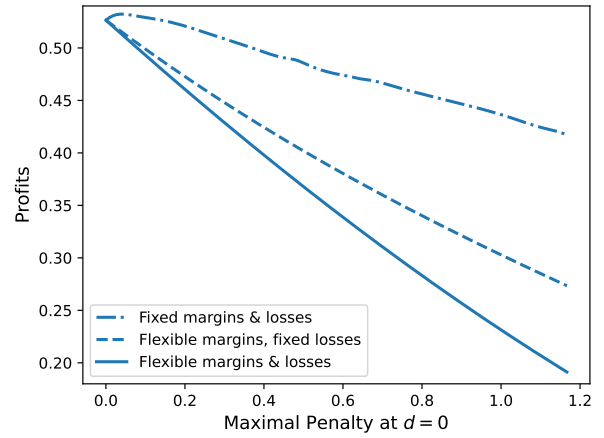
(a) Extensive margin



(b) Intensive margin



(c) Indirect losses



(d) Cumulative effects on profits

NOTE: Profits, losses, and penalties are scaled by average lane market price of a shipment. Panel (d) presents the cumulative effect of the three margins on platform profits.

The results of the decomposition exercise are presented in Figure 11. Panel (d) presents the cumulative effect of the three margins on platform profits. All three margins contribute to the decline in profits as the penalty level increases. The effect on the extensive margin can be further decomposed into an effect on the number of initial matches and an effect on the number of cancellations, which both decline as the penalty level increases. However, the decline in the number of initial matches is stronger, so that overall the extensive margin monotonically declines after a slight increase at the lowest penalty levels. The intensive margin is also declining, as the bids are increasing in response to the higher penalties, which is not enough to offset the decline in the number of matches. Finally, the indirect losses are increasing in the penalty level as a direct consequence of the decline in the number of matches. The decomposition exercise highlights the importance of accounting for the strategic response of bids to changes in the cancellation penalty, which drives both the higher transaction prices and lower matching rates.

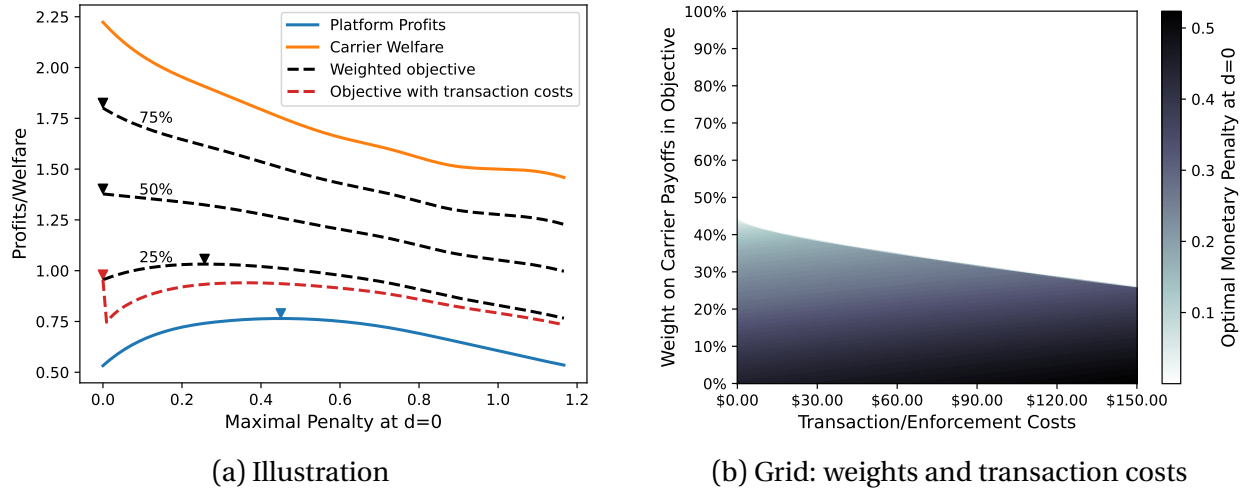
6.2 Trade-offs between Monetary and Reputational Penalties

Figure 10 also presents the results of the counterfactual analysis under monetary penalties. For a given intermediate penalty schedule, monetary penalties are more efficient than reputational penalties because they do not "burn money." Nevertheless, the counterfactual analysis shows that, under penalty schedules of the same shape as in the status-quo, the welfare-maximizing monetary penalties are still at zero. In contrast, the profit-maximizing monetary penalties are much higher. The increase in profit is accompanied by a large decline in both carrier and total welfare, suggesting that these penalties mainly transfer rents from carriers to the platform.

An immediate question that arises is why the firm does not use monetary penalties in practice. To rationalize these, I can explore two caveats to the baseline results. Firstly, the model does not account for carrier's decisions to join or leave the platform. As there may be some set-up costs and time investment to continuously monitor the platform, the number of active carriers in the market may be sensitive to the average welfare of carriers. Secondly, I have so far assumed that charging monetary fees for cancellations is costless. In practice, there may be significant transaction fees associated with collecting these fees, particularly due to enforcement costs. Below I provide some quantification of these concerns.

To approximate the effect of long-term carrier acquisition and retention on the platform, I treat the platform as maximizing a weighted sum of its own short-term profits and the welfare

Figure 12: Monetary penalties with weighted objective and transaction costs.



NOTE: Profits, losses, and penalties are scaled by average lane market price of a shipment. Panel (a) illustrates the effect of a weighted objective and transaction costs on platform profits and total welfare, with objective maximizing penalties highlighted with corresponding triangles. Panel (b) presents the grid of weights and transaction costs considered, with the corresponding optimal penalty level shown as a color (darker=higher).

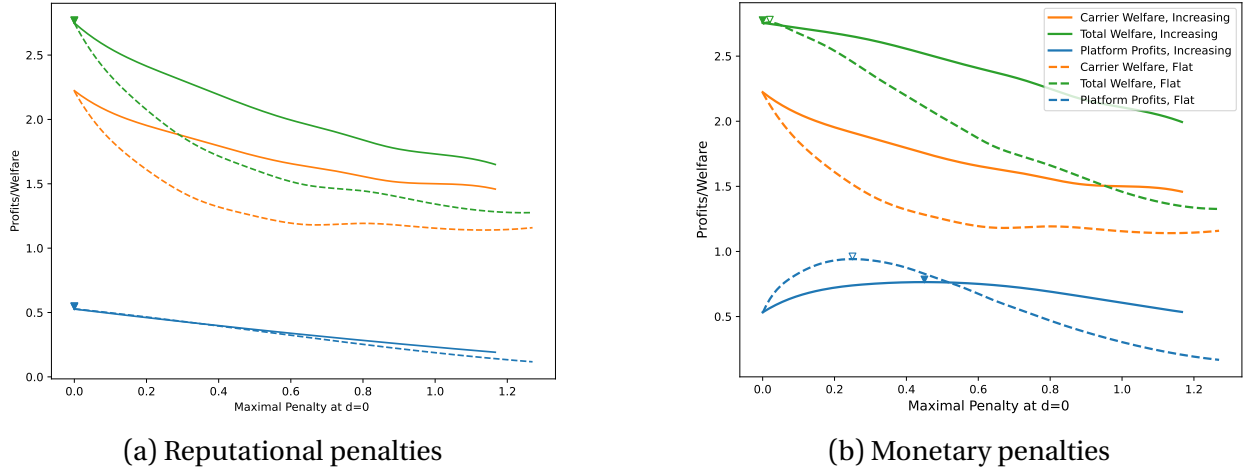
of carriers. Different weighted objectives are illustrated in Figure 12a. As the weight on carrier welfare increases, the optimal penalty level decreases. Secondly, I can suppose that the platform bears some transaction cost for each monetary cancellation fee that it collects. In Figure 12a, this transaction cost leads to a discontinuity in the objective function when moving from zero penalties to positive penalties. For a large enough transaction cost, the optimal penalty level is zero, even for a lower weight on carrier welfare.

The joint effect of the weight on carrier welfare and transaction costs is presented in Figure 12b. A weight of approximately 45% on carrier welfare is sufficient to rationalize why the firm would not use monetary penalties at all. Transaction costs on their own are not sufficient to explain the absence of monetary penalties—even for a high transaction cost of 150\$, some weight on carrier welfare is still necessary to rationalize the avoidance of monetary penalties.²² Furthermore, for a given weight on carrier welfare below the threshold, the optimal penalty level is increasing in the transaction cost, as the firm is able to pass on part of the transaction cost to the carriers.

Taken together, the results suggest that the firm avoids monetary penalties due to strong con-

22. Monetary fees for shipper cancellations, referred to as a Truck Order Not Used (TONU) fee, range from 150\$ to 300\$. The absence of such monetary fees for carrier cancellations in the industry would suggest transaction or enforcement costs of a similar order of magnitude.

Figure 13: Comparison of flat and increasing penalties.



cerns about the long-term growth of the platform, which are not captured by the short-term profit-maximizing model. The choice of reputational or monetary penalties will thus depend on the degree of lock-in that the platform has over carriers, which is likely to be low in the highly fragmented trucking industry.

6.3 Timing of the Penalty

Although the primary focus of the counterfactual analysis has been on the level and type of the penalties, the rich dynamic model also permits an analysis of different timings of the penalties—that is, not only how *much* to charge, but also *when* to charge a given amount. This can also be thought of as examining different shapes of the penalty schedule. To illustrate this, I compare the stair-shaped increasing penalty schedule of the status quo to a flat penalty schedule, where the penalty is the same for all days of cancellation.

Figure 13 presents the results of the comparison between flat and increasing penalties. For reputational penalties, the shape of the profit function is slightly steeper under reputational penalties, but welfare decreases much more dramatically as the penalty level increases. Intuitively, carriers are exposed to more outside offers in the earlier days before pickup (see Figure 5a), so being subject to higher penalties on these days is very costly in terms of opportunity costs. Monetary penalties, on the other hand, are able to extract more rents from carriers on these early days, leading to a steeper increase in platform profits. This analysis demonstrates that profit-seeking firms may reduce overall welfare not only by setting penalties too *high*, but

also by setting them too *early*.

6.4 Full commitment and late-clearing policies

Contracts with limited commitments are pervasive in the trucking industry and extend beyond cancellation policies. For example, long-term contracts between shippers and carriers or brokers typically only specify a price and a general service level, but do not contain strict guarantees on the number of shipments to be executed. Prior research on this contracting process has focused on the negative aspects of these weak commitments. [Caplice and Sheffi \(2006\)](#) highlight the difficulty of awarding long-term contracts through combinatorial auctions: potential complementarities between different lanes are not guaranteed to be exercised in practice as both carriers and shippers are free to partner with other parties at any time. [Harris and Nguyen \(2022\)](#) show how price fluctuations in the spot market can reduce the gains from long-term relationships, as carriers are incentivized to renege on long-term contracts when spot prices rise. In contrast, my model explicitly highlights the *value* of contractual flexibility in enabling market participants to act on new opportunities they encounter through the search process.

To fully illustrate the value of flexibility, consider a full-commitment policy with infinite cancellation penalties. This policy is the extreme opposite of the profit- and welfare-maximizing reputational penalty, where carriers are free to cancel at any time. Table 6 summarizes the results of multiple different counterfactual policies, including the full-commitment policy with infinite penalties. The full-commitment policy drastically reduces total welfare and the number of matches, highlighting the value of non-binding agreements in the market.

Finally, I examine the late-clearing policy, which is an alternative means of reducing the opportunity cost of transacting. We know that when the *entire market* is centralized at a later matching time, timing frictions become less significant ([Roth and Xing 1994](#)). However, when considering a *unilateral deviation* by the platform to a later matching time, there are severe negative effects on both total welfare and platform profits, with reductions of 103.1% and 216.5% respectively. The massive reduction in platform profits can be attributed to the large drop in the number of matches, resulting in large losses for the platform due to penalties for unmatched shipments. This loss occurs despite an improvement in the per-shipment margin, with a much lower average transaction price, which is likely due to most matches occurring with late arriving carriers, who have lower opportunity costs at the time of bidding.

The model suggests a mechanism to explain why the platform’s unilateral shift to a later match-

Table 6: Comparison of Counterfactual Policies

Penalty Type	Status Quo Reputational	Max Welfare Pecuniary	Max Profit Pecuniary	Zero Penalty -	Infinite Penalty ∞	Late Clearing ∞
Metric						
Total Revenues	8.10	8.14 (0.4%)	8.03 (-0.9%)	8.06 (-0.6%)	6.08 (-25.0%)	2.40 (-70.4%)
Total Welfare	2.67	2.76 (3.5%)	2.46 (-7.9%)	2.75 (3.2%)	1.67 (-37.4%)	0.19 (-93.0%)
Platform Profit	0.51	0.63 (23.4%)	0.94 (83.1%)	0.53 (2.4%)	-0.17 (-132.5%)	-0.58 (-213.6%)
Carrier Welfare	2.15	2.13 (-1.2%)	1.51 (-29.6%)	2.22 (3.3%)	1.83 (-14.7%)	0.77 (-64.1%)
Matches	7.99	8.02 (0.4%)	7.90 (-1.1%)	7.95 (-0.5%)	5.95 (-25.5%)	2.29 (-71.4%)
Match Rate	0.57	0.58 (0.3%)	0.57 (-1.1%)	0.57 (-0.4%)	0.43 (-25.6%)	0.16 (-71.5%)
Match Welfare	0.42	0.41 (-0.2%)	0.42 (1.8%)	0.42 (0.2%)	0.47 (13.2%)	0.40 (-2.8%)
Transaction Price	0.88	0.88 (-0.0%)	0.88 (0.7%)	0.87 (-0.4%)	0.91 (4.3%)	0.79 (-9.6%)
Average Bid	1.04	1.04 (-0.1%)	1.06 (1.5%)	1.04 (-0.2%)	1.15 (10.2%)	2.10 (102.1%)

NOTE: The table presents counterfactual outcomes under different penalty schedules. Change relative to status quo in parentheses. Absolute numbers are multiples of average lane market price of a shipment. Infinite penalty is approximated by a penalty of $\kappa = 10^6$. Match welfare is average over final matches on platform.

ing time proves detrimental to all parties. Carriers remain pressured to accept or reject outside offers before knowing the platform auction outcome, leading to high attrition rates. Matching strong bidders earlier allows them to reject outside offers with greater confidence (reflected in higher reservation prices), thereby reducing attrition and ultimately fostering a more efficient allocation of shipments. Consequently, in the absence of centralized matching, early matching with non-binding agreements emerges as the most effective market design.

7 Conclusion

I provide the first comprehensive empirical analysis of pricing the right to renege in matching markets, using a novel dataset from a digital brokerage platform with detailed information on reneged matches. In the stylized model, I first show the importance of considering both the downstream effect of penalties on the propensity to renege and the upstream effects on the initial matching probability and transaction price. In particular, I show that even when the penalty itself is non-pecuniary, the value of the right to renege can be priced into the final transaction price, such that the firm has a sufficiently strong incentive to offer the flexibility.

By extending the model to a dynamic setting and structurally estimating it, I can simulate counterfactual policies and evaluate the welfare implications of different penalty schedules. I find that the current near-zero reputational penalties are nearly optimal for both social welfare and firm profits. In contrast, moving to a monetary penalty, which intuitively should enable more efficient trade, instead distorts the platform's incentives by allowing it to extract more rents from carriers, at the cost of overall welfare. Barriers to the implementation of monetary penalties, whether by regulation or through transaction costs, can thus improve welfare.

References

- Anderson, Eric T, Karsten Hansen, and Duncan Simester. 2009. "The option value of returns: Theory and empirical evidence." *Marketing Science* 28 (3): 405–423.
- Armstrong, Mark, and Jidong Zhou. 2016. "Search Deterrence." *The Review of Economic Studies* 83 (1): 26–57.
- 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.
- Black, Fischer, and Myron Scholes. 1973. "The pricing of options and corporate liabilities." *Journal of political economy* 81 (3): 637–654.
- Buchholz, Nicholas, Laura Doval, Jakub Kastl, Filip Matjka, and Tobias Salz. 2020. *Personalized Pricing and the Value of Time: Evidence from Auctioned Cab Rides*. Working Paper, Working Paper Series 27087. National Bureau of Economic Research, May. <https://doi.org/10.3386/w27087>. <http://www.nber.org/papers/w27087>.
- 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.
- 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>.
- Federal Trade Commission. 2024. *Trade Regulation Rule on Unfair or Deceptive Fees*. Final Rule 60. Accessed: 2024-06-18. Federal Register, March. <https://www.regulations.gov/document/FTC-2023-0064-0001>.
- Greene, W.H. 2018. *Econometric Analysis*. Econometric Analysis. Pearson. ISBN: 9780134461366.
- 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. 2022. *Long-Term Relationships and the Spot Market: Evidence from US Trucking*.
- Hendricks, Kenneth, and Thomas Wiseman. 2022. "How to Sell in a Sequential Auction Market." *Theoretical Economics* 17 (4): 1451–1471.
- Higgins, Michelle. 2008. "This Season, Procrastination May Pay Off." Accessed: 2024-06-18, *The New York Times* (November). <https://www.nytimes.com/2008/11/09/travel/09prac.html>.
- Hu, Zhenyu, and Wenjie Tang. 2021. "Size Matters, So Does Duration: The Interplay Between Offer Size and Offer Deadline." *Management Science* 67 (8): 4935–4960.
- 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.
- Ji, Yunan. 2022. "Can Competitive Bidding Work in Health Care? Evidence from Medicare Durable Medical Equipment."
- 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.
- Kandori, Michihiro. 1992. "Social Norms and Community Enforcement." *The Review of Economic Studies* 59 (1): 63–80.
- Lau, Nelson, Yakov Bart, J Neil Bearden, and Ilia Tsetlin. 2014. "Exploding Offers Can Blow Up in More Than One Way." *Decision Analysis* 11 (3): 171–188.
- 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.
- Lippman, Steven A, and John W Mamer. 2012. "Exploding Offers." *Decision Analysis* 9 (1): 6–21.
- Liu, Yan, and Ningyuan Chen. 2022. "Dynamic Pricing with Money-Back Guarantees." *Production and Operations Management* 31 (3): 941–962.
- Liu, Yan, and Dan Zhang. 2023. "Intraconsumer Price Discrimination with Credit Refund Policies." *Management Science*.
- Manzini, Paola, and Marco Mariotti. 2014. "Stochastic Choice and Consideration Sets." *Econometrica* 82 (3): 1153–1176.
- McCall, John Joseph. 1970. "Economics of Information and Job Search." *The Quarterly Journal of Economics* 84 (1): 113–126.
- McCartney, Scott. 2009. "Your Bad Luck Is a Windfall For Airlines." Accessed: 2024-06-18, *The Wall Street Journal*, <https://www.wsj.com/articles/SB10001424052970204563304574318212311819146>.
- Merlob, Brian, Charles R Plott, and Yuanjun Zhang. 2012. "The CMS Auction: Experimental Studies of a Median-Bid Procurement Auction with Nonbinding Bids." *The Quarterly Journal of Economics* 127 (2): 793–827.
- Moise, Imani. 2024. "The Era of No-Show Fees Is Hereand Its Going to Cost You." Accessed: 2024-06-18, *The Wall Street Journal* (May). <https://www.wsj.com/personal-finance/appointment-fees-square-no-show-12163052>.

- Negrone, Christine. 2012. “Cashing In Before Taking Off.” Accessed: 2024-06-18, *The New York Times* (February). <https://www.nytimes.com/2012/02/28/business/new-airline-revenue-goes-beyond-baggage-fees.html?>.
- 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. 2009. “Market Culture: How Rules Governing Exploding Offers Affect Market Performance.” *American Economic Journal: Microeconomics* 1 (2): 199–219.
- Office of the Secretary, Department of Transportation. 2011. *Enhancing Airline Passenger Protections*. Final Rule Docket No. DOTOST20100140. Accessed: 2024-06-18. Federal Register, April. <https://www.govinfo.gov/content/pkg/FR-2011-04-25/pdf/2011-9736.pdf>.
- 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>.
- Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory*. MIT Press.
- Riley, John G, and William F Samuelson. 1981. “Optimal auctions.” *The American Economic Review* 71 (3): 381–392.
- 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.
- Sampson, Hannah. 2024. “Budget airlines are dropping their notorious fees. What changed?” Accessed: 2024-06-18, *The Washington Post* (June). <https://www.washingtonpost.com/travel/2024/05/23/spirit-airlines-frontier-airlines-cheap-flights-cancellation-fees/>.
- U.S. Department of Transportation. 2021. “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/bcvt-rqmu>.
- . 2022. “Contribution of Transportation to the Economy: Contribution of Transportation Services to the Economy and the Transportation Satellite Accounts.” Bureau of Transportation Statistics. Accessed June 24, 2024. <https://data.bts.gov/stories/s/smrm-36nv/>.

- Xie, Jinhong, and Eitan Gerstner. 2007. "Service Escape: Profiting from Customer Cancellations." *Marketing Science* 26 (1): 18–30.
- 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.
- Zhang, Zelin, Weishi Lim, Haitao Cui, and Ze Wang. 2021. "Partial Refunds as a Strategic Price Commitment Device in Advance Selling in a Service Industry." *European Journal of Operational Research* 291 (3): 1062–1074.
- Zipkin, Amy. 2015. "Hotel Reservations Are Becoming More Costly, and Complicated, to Cancel." Accessed: 2024-06-18, *The New York Times* (December). <https://www.nytimes.com/2015/12/22/business/hotel-reservations-are-becoming-more-complicated-and-costly-to-cancel.html>.
- Zorc, Saa, and Ilia Tsetlin. 2020. "Deadlines, Offer Timing, and the Search for Alternatives." *Operations Research* 68 (3): 927–948.

Appendices

A Appendix

A.1 Value Functions in Continuous Time Markov Models

As described in Section 4, the model takes place in continuous time, but all payoffs and decisions occur through Poisson arrival processes. This gives rise to a continuous-time Markov chain, where states proceed from one to another in a discrete way. Furthermore, there is no discounting over time in this paper, due to the short time horizons involved, which also simplifies the derivation of the value functions.

We start with a general derivation. Suppose we have states $s \in \mathcal{S}$, and $\lambda_{s'|s}$ be the Poisson arrival rates of a state s' given that the current state is s . Because Poisson arrival rates are distributed exponentially, we can also describe the transition process as a combined arrival rate $\bar{\lambda}_s = \sum_{s'} \lambda_{s'|s}$, and a probability distribution over the states $P(s'|s) = \frac{\lambda_{s'|s}}{\bar{\lambda}_s}$ which gives the probability of transitioning from s to s' *conditional* on transitioning at all. Furthermore, suppose we receive an instantaneous flow payoff $\phi(s'|s)$ when transitioning from s to s' . If we assume $\phi(s|s) = 0$, then we can ignore transitions from a state to itself without loss of generality, and let $\lambda_{s|s} = 0$.

Now, because there is no discounting nor any other time-varying component, the value function of a state s is simply given by the probabilities of all successor states s' , weighted by the flow payoff and value function of that state:

$$U_s = \sum_{s'} P(s'|s) [\phi(s'|s) + U_{s'}] \quad (\text{A.1})$$

We can now apply this to Equation 8. The implicit state space here is $(d, Unmatched)$, where the latter variable indicates that carrier i is currently unmatched with shipment j . Two events can shift a carrier out of this state. Firstly, they may receive an outside offer at rate λ_d , which they will take if it is attractive enough, thereby exiting the market, or refuse if they would rather keep searching. This event thus gives payoff

$$\int \max\{U_{d,\bar{d}_i}(c_{ij}), \tilde{\pi}_{ij'} - c_{ij} + \rho(c_{ij})\} dG_d(\tilde{\pi}_{ij'}) \quad (\text{A.2})$$

Secondly, they may transition to the next day $d - 1$, which occurs at exogenous rate η , which gives payoff $\bar{U}_{d-1}(c_{ij})$. Thus, the value function of a matched carrier is given by:

$$U_{d,\bar{d}_i}(c_{ij}) = \frac{1}{\eta + \lambda_d} \left[\lambda_d \left(\int \max\{U_{d,\bar{d}_i}(c_{ij}), \tilde{\pi}_{ij'} - c_{ij} + \nu_d + \rho(c_{ij})\} dG_d(\tilde{\pi}_{ij'}) \right) + \eta \bar{U}_{d-1}(c_{ij}) \right] \quad (\text{A.3})$$

Eventually, if the carrier hasn't found any worthwhile shipment, they will end up unmatched in the absorbing state with payoff $U_{-1}(c_{ij}) = 0$, which makes the value function finite.

Notice that this value function involves some probability of remaining in the current state, which occurs when an outside offer is not attractive enough to induce a cancellation. This occurs at the cutoff $R_{d,\bar{d}_i}^U(c_{ij}) = U_{d,\bar{d}_i}(c_{ij}) + c_{ij} - \nu_d - \rho(c_{ij})$. As stated earlier, we can ignore any transition from a state to itself, so we can rewrite the value function as follows:

$$U_{d,\bar{d}_i}(c_{ij}) = \frac{1}{\eta + \lambda_{di}} \left[\lambda_{di} \left(\int_{R_{d,\bar{d}_i}^U(c_{ij})}^{\infty} \tilde{\pi}_{ij'} - c_{ij} + \nu_d + \rho(c_{ij}) dG_d(\tilde{\pi}_{ij'}) \right) + G(R_{d,\bar{d}_i}^U(c_{ij})) U_{d,\bar{d}_i}(c_{ij}) \right] + \eta \bar{U}_{d-1}(c_{ij})$$

Which is the form used in the main text. For computational purposes we can further re-arrange this equation to:

$$U_{d,\bar{d}_i}(c_{ij}) = \frac{1}{\eta + \lambda_{di}(1 - G(R_{d,\bar{d}_i}^U(c_{ij})))} \left[\lambda_{di} \left(\int_{R_{d,\bar{d}_i}^U(c_{ij})}^{\infty} \tilde{\pi}_{ij'} - c_{ij} + \nu_d + \rho(c_{ij}) dG_d(\tilde{\pi}_{ij'}) \right) + \eta \bar{U}_{d-1}(c_{ij}) \right]$$

Benchmarks of the value function iteration have shown this re-arranged form to converge faster (when inserting the current iteration into the LHS to obtain the next iteration from the RHS). Intuitively, the original puts greater weight on the previous iteration of the value, and thus dampens the iteration, which generally reduces speed but may be more robust to oscillations in the iteration. Note that in both forms, the value function we are solving for shows up in both sides of the equation (implicitly in $R_{d,\bar{d}_i}^U(c_{ij})$), so solving the problem still requires a fixed point iteration even though we have a finite and discrete time index d .

We can similarly re-arrange the value function of a carrier with a matched shipment in Equation 10. We first define $R_{d,\bar{d}_i}^M(b_{ij}, c_{ij}) = V_{d,\bar{d}_i}(b_{ij}, c_{ij}) + c_{ij} + \kappa_d - \rho(c_{ij})$, so that the value of a pending bid is now:

$$V_{d,\bar{d}_i}(b_{ij}, c_{ij}) = \frac{1}{\eta + \lambda_d(1 - G_d(R_{d,\bar{d}_i}^M(b_{ij}, c_{ij})))} \left[\lambda_d \left(\int_{R_{d,\bar{d}_i}^M(b_{ij}, c_{ij})}^{\infty} (\tilde{\pi}_{ij'} - c_{ij} - \kappa_d + \rho(c_{ij})) dG_d(\tilde{\pi}_{ij'}) \right) + \eta \bar{V}_{d-1}(b_{ij}, c_{ij}) \right]$$

Finally, we can do the same with the value function of a carrier with a pending bid in Equation 13. We first define $R_{d,\bar{d}_i}^P(b_{ij}, c_{ij}) = W_{d,\bar{d}_i}(b_{ij}, c_{ij}) + c_{ij} - \nu_d - \rho(c_{ij})$, so that the value of a pending bid is now:

$$W_{d,\bar{d}_i}(b_{ij}, c_{ij}) = \frac{1}{\eta + \gamma_d(b_{ij}) + \lambda_d(1 - G_d(R_{d,\bar{d}_i}^P(b_{ij}, c_{ij})))} \left[\lambda_d \left(\int_{R_{d,\bar{d}_i}^P(b_{ij}, c_{ij})}^{\infty} (\tilde{\pi}_{ij'} - c_{ij} + \nu_d + \rho(c_{ij})) dG_d(\tilde{\pi}_{ij'}) \right) + \gamma_d(b_{ij})(V_{d,\bar{d}_i}(b_{ij}, c_{ij}) + \nu_d) + \eta \bar{W}_{d-1}(b_{ij}, c_{ij}) \right]$$

A.2 Special case of the envelope theorem

Lemma 1 *Given a continuous and differentiable function $f(x)$ and continuous and differentiable distribution $G(x)$, define:*

$$h(x) = \int \max\{f(x), y + g(x) + c\} dG(y) = G(f(x) - g(x) - c)f(x) + \int_{f(x)-g(x)-c}^{\infty} y dG(y)$$

Then the derivative of $h(x)$ is given by:

$$\begin{aligned} \frac{\partial h(x)}{\partial x} &= \frac{\partial}{\partial x} G(f(x) - g(x) - c)f(x) + \frac{\partial}{\partial x} \int_{f(x)-g(x)-c}^{\infty} y dG(y) = \\ &G(f(x) - g(x) - c)f'(x) + \textcolor{red}{G'(f(x) - g(x) - c)(f'(x) - g'(x))f(x)} \\ &\textcolor{red}{- G'(f(x) - g(x) - c)(f'(x) - g'(x))f(x)} + (1 - G(f(x) - g(x) - c))g'(x) \\ &= G(f(x) - g(x) - c)f'(x) + (1 - G(f(x) - g(x) - c))g'(x) \end{aligned}$$

Note that Lemma 1 is really just a special case of the envelope theorem in a form more readily applicable to the model being studied in this paper. It says that the derivative with respect to the expectation over the maximum of a binary choice with a random payoff is equivalent to the derivative with respect to the fixed utility component of each choice, weighted by the respective choice probabilities.

A.3 Derivative of bidding first order condition

To solve the bidding problem, we solve for the first order condition.

Similarly to the value functions themselves, the derivatives can be represented in recursive

form.

We begin with the derivative of $W_{d,\bar{d}_i}(b_{ij}, c_{ij})$ with respect to b_{ij} . Making use of Lemma 1, we have

$$\begin{aligned}
(\eta + \gamma_d(b_{ij}) + \lambda_d) \frac{\partial W_{d,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} + \gamma'_d(b_{ij}) W_{d,\bar{d}_i}(b_{ij}, c_{ij}) = \\
\left[\lambda_d G_d(R_{d,\bar{d}_i}^P(b_{ij}, c_{ij})) \frac{\partial W_{d,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} + \gamma'_d(b_{ij}) (\alpha_d V_{d,\bar{d}_i}(b_{ij}, c_{ij}) + (1 - \alpha) U_{d,\bar{d}_i}(c_{ij})) \right. \\
\left. + \alpha_d \gamma_d(b_{ij}) \frac{\partial V_{d,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} + \eta \frac{\partial W_{d-1,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} \right]
\end{aligned}$$

Re-arranging yields the following recursive expression:

$$\begin{aligned}
\Rightarrow \frac{\partial W_{d,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} &= \frac{1}{\eta + \gamma_d(b_{ij}) + \lambda_d (1 - G_d(R_{d,\bar{d}_i}^P(b_{ij}, c_{ij})))} \\
&\left[\gamma'_d(b_{ij}) (\alpha_d V_{d,\bar{d}_i}(b_{ij}, c_{ij}) + (1 - \alpha) U_{d,\bar{d}_i}(c_{ij}) - W_{d,\bar{d}_i}(b_{ij}, c_{ij})) + \frac{\alpha_d \gamma_d(b_{ij})}{\gamma'_d(b_{ij})} \frac{\partial V_{d,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} \right. \\
&\quad \left. + \eta \frac{\partial W_{d-1,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} \right]
\end{aligned}$$

We then need the derivative for the matched value (V_{d,\bar{d}_i}):

$$\begin{aligned}
\frac{\partial V_{d,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} &= \frac{\eta}{\eta + \lambda(1 - G_d(R_{d,\bar{d}_i}^M(b_{ij}, c_{ij})))} \frac{\partial V_{d-1,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} \\
\frac{\partial V_{-1,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} &= 1 \\
\Rightarrow \frac{\partial V_{d,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} &= \prod_{k=0}^d \frac{\eta}{\eta + \lambda(1 - G_k(R_{k,\bar{d}_i}^M(b_{ij}, c_{ij})))} \equiv P_{d,\bar{d}_i}(\text{no cancel} | b_{ij}, c_{ij})
\end{aligned}$$

We can then write the FOC as a weighted sum:

$$\begin{aligned}
\frac{\partial W_{d,\bar{d}_i}(b_{ij}, c_{ij})}{\partial b_{ij}} &= \sum_{k=0}^{d_{ij}^b} \left[\overbrace{\prod_{\ell=k}^{d_{ij}^b} \frac{1}{\eta + \gamma_\ell(b_{ij}) + \lambda_\ell(1 - G_\ell(R_{\ell,\bar{d}_i}^P(b_{ij}, c_{ij})))}}^{\text{Probability of bid surviving from } d \text{ to } k} \right] \eta^{-1} \\
&\quad \gamma'_d(b_{ij}) \left(\underbrace{\alpha_d V_{d,\bar{d}_i}(b_{ij}, c_{ij}) + (1 - \alpha) U_{d,\bar{d}_i}(c_{ij}) - W_{d,\bar{d}_i}(b_{ij}, c_{ij})}_{\text{Benefit of winning net of opportunity costs}} + \underbrace{\frac{\alpha_d \gamma'_d(b_{ij})}{\gamma'_d(b_{ij})} P_{d,\bar{d}_i}(\text{no cancel} | b_{ij}, c_{ij})}_{\text{Markup}} \right) \\
&= 0
\end{aligned}$$

While superficially complicated, this FOC is essentially a weighted sum of day-specific FOCs, each of which captures the marginal benefit of increasing the bid on the given day. The terms highlighted in braces (a day-specific derivative) take the familiar form observed in the standard theory of first-price auctions. These day-specific derivatives are then weighted according to the probability of surviving to that day.

A.4 Offline estimation of win rates

The FOC in Equation 16 involves the day-specific win rates $\gamma_d(b)$ and their derivatives $\gamma'_d(b)$. While in principle these can be derived through market equilibrium conditions, doing so is computationally expensive, as it involves solving for a fixed point in the space of bidding strategies (the optimal bid is a function of the win probabilities and the win probabilities are a function of the optimal bidding strategies). To avoid this computational burden, I simply estimate the win probabilities off-line from the data directly.

As with all other price variables in the data, I normalize all bids by the contemporaneous lane spot market rate for the respective shipment's pickup date, in order to pool data over the two-year period.

I fit a modified logistic function to the hazard rate, i.e., the probability that an auction is won on day d conditional on the bid still existing on day d :

$$\frac{\gamma(b; \theta_\gamma)}{\eta + \gamma(b; \theta_\gamma)} = \frac{\theta_{\gamma 1}}{1 + \exp((\log(b) - \theta_{\gamma 2})/\theta_{\gamma 3})} \quad (\text{A.4})$$

This yields the following expressions for γ and γ' :

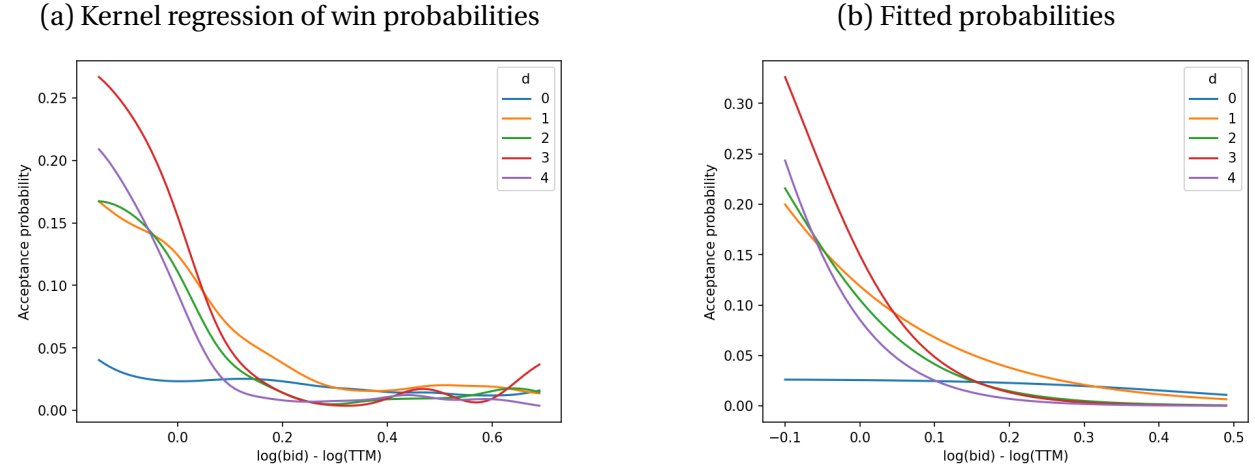
$$\gamma(b, \theta_\gamma) = \frac{\eta\theta_{\gamma 1}}{1 + \exp((\log(b) - \theta_{\gamma 2})/\theta_{\gamma 3}) - \theta_{\gamma 1}} \quad (\text{A.5})$$

$$\gamma'(b, \theta_\gamma) = -\frac{\eta\theta_{\gamma 1}\exp((\log(b) - \theta_{\gamma 2})/\theta_{\gamma 3})}{b\theta_{\gamma 3}(1 + \exp((\log(b) - \theta_{\gamma 2})/\theta_{\gamma 3}) - \theta_{\gamma 1})^2} \quad (\text{A.6})$$

The form of this function is chosen for its reasonable fit. In addition, with $\theta_{\gamma 1} \geq 0$ and $\theta_{\gamma 3} \geq 0$, the function is guaranteed to be monotonically decreasing.

The estimated win probability function is denoted by $\hat{\gamma}(b)$ and is fitted via maximum likelihood on the full bid data. The fit of the function is presented below in Figure A.1

Figure A.1: Non-parametric vs. fitted estimates of win probabilities (Seattle-San Francisco)



A.5 Monte Carlo Simulation

For the Monte-Carlo exercise, I specify $G(\cdot)$ as a Normal $N(\mu_G, \sigma_G)$, while the win probability $\gamma(b)$ is specified as a logistic distribution with location 0.7 and scale 0.2. This version of the model also adds an attention probability α to be estimated, as well as an additional component of the likelihood from the confirmation probability.

Table A.1 shows the results of the Monte-Carlo simulation. The mean estimates for the offer distribution and the attention probability are very close to their true values. However, estimated cancellation penalty is significantly lower than the true value, which may be due to the set identification of the penalty.

Table A.1: Results of Monte Carlo Simulation with 500 runs of 10,000 bidders each

	True	MC Mean	MC S.D.	S.D. of Mean
μ_G	1.000	0.966	0.080	0.004
σ_G	0.500	0.516	0.107	0.005
κ	0.100	0.063	0.085	0.004
α	0.500	0.497	0.032	0.001

NOTE: S.D. of Mean is simply the parameter standard deviation divided by $\sqrt{500}$. On average, 5,672 winning bids, 1,615 confirmed bids, 269 cancellations.

A.6 Raw Structural Parameters and Asymptotic Variance

In the following, I provide more detail on the raw structural parameters of the model and their asymptotic variances. Because of the penalty term on the inverted costs, the estimator is not a standard maximum likelihood estimator (MLE). For inference purposes, I draw on Theorems 7.1 and 7.3 of [Newey and McFadden \(1994\)](#) for extremum estimators with nonsmooth objective functions to derive the asymptotic variance of the estimator. Using their notation, we have the following objects:

- The sample objective function $\hat{Q}_n(\theta) = n^{-1} \sum_{i=1}^n q_i(\theta)$
- The derivative of the objective function $\hat{D}_n(\theta) = \nabla_{\theta} \hat{Q}_n(\theta)$, with $\sqrt{n} \hat{D} \xrightarrow{d} N(0, \Omega)$
- The Hessian of the true objective function at the true parameters $H = \nabla_{\theta\theta} Q_0(\theta_0)$

Then, under additional regularity conditions which I do not verify here, the asymptotic variance of the estimator is given by

$$\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, H^{-1} \Omega H^{-1}).$$

To obtain standard errors for inference, we then need an estimate of the Hessian \hat{H} and of the asymptotic variance of the derivative $\hat{\Omega}$. The former can be obtained by the sample Hessian of the objective function at the estimated parameters, as shown by [Newey and McFadden \(1994\)](#) in Theorem 7.3. Given the potential nonsmoothness of the objective function and (in this case) the lack of closed-form expressions for the Hessian, I use numerical differentiation to obtain the sample Hessian, as suggested in the reference text.

To obtain the asymptotic variance of the derivative, I use the sample variance of the derivative

of the objective function, which is

$$\hat{\Omega} = \frac{1}{n} \sum_{i=1}^n d_i(\theta) d_i(\theta)^T$$

where $d_i(\theta) = \nabla_{\theta} q_i(\theta)$ is the gradient of the objective function evaluated at the estimated parameters for the i -th observation. This approach is also used in Chapter 14.8 of [Greene \(2018\)](#) to obtain a robust covariance matrix for pseudo-maximum likelihood estimators.

Table A.2: Raw structural parameter estimates

Parameter	Estimate	Std. Error
$\log(\kappa_0 - \kappa_1)$	-5.043	0.002
$\log(\kappa_1 - \kappa_2)$	-4.135	0.008
$\log(\kappa_2 - \kappa_3)$	-4.555	0.010
$\log(\kappa_3 - \kappa_4)$	-5.922	0.003
$\log(\kappa_4)$	-6.813	0.005
$\log((1 - \alpha_0)/\alpha_0)$	0.448	0.000
$\log((1 - \alpha_1)/\alpha_1)$	0.130	0.000
$\log((1 - \alpha_2)/\alpha_2)$	0.133	0.003
$\log((1 - \alpha_3)/\alpha_3)$	-0.190	0.002
$\log((1 - \alpha_4)/\alpha_4)$	-0.508	0.005
μ^{π^0}	-1.402	0.002
μ^{π^1}	0.080	0.002
μ^{π^2}	-0.062	0.002
σ^{π^0}	2.713	0.002
σ^{π^1}	0.003	0.007
σ^{π^2}	-0.052	0.013
$\log(\lambda^0)$	-1.365	0.004
$\log(\lambda^1)$	-2.700	0.005
$\log(\lambda^2)$	-5.006	0.002

Table A.2 reports the raw structural parameter estimates and their asymptotic standard errors. The estimates are obtained by maximizing the objective function with the penalty term, and the standard errors are obtained by the method described above. As explained in the main text in Section 5, additional structure is placed on the parameters. Firstly, the cancellation penalty schedule is constrained to be non-decreasing. I enforce this by estimating the log of the difference in the penalty from one day to the next. The attention probabilities are strictly bounded between 0 and 1 by estimating the logistic transformation of the raw attention probabilities. The parameters of the outside offer distribution are constrained to be second order polynomials of the number of days until pickup. Thus, the mean of the outside offer distribution for

day d before pickup is $\mu_d^\pi = \mu^{\pi^0} + \mu^{\pi^1}d + \mu^{\pi^2}d^2$, and similarly for the standard deviation. The arrival rates are further constrained to take on a concave form by estimating the log of the rate coefficients and setting $\lambda_d = \lambda^0 + \lambda^1d - \lambda^2d^2$.

Although the parameters lack direct interpretation in the form displayed in the table, it is still possible to see that the coefficients are precisely estimated. The reader is referred back to the main text for the interpretation of the parameters in the context of the model.

A.7 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 $e \in \{\text{carrier}, \text{shipment}, \text{clear}\}$, let $N_{m,d}^e$ denote the number of events of type e in market $m \in 1, \dots, M$ 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:

$$\mathbb{E}[N_{m,d}^e] = \lambda_d^e.$$

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

$$\hat{\lambda}_d^e = \frac{1}{M} \sum_{m=1}^M N_{m,d}^e$$

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 $Y_{j,m,d}^{shipcancel} \in \{0, 1\}$ denote the event

that shipment j in market m on day d is cancelled, and $Y_{ij,m,d}^{view} \in \{0, 1\}$ 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 derivative of the transformation function.

Due to the large number of arrival rate parameters, the estimates and their confidence intervals are reported graphically in Figure 8 in the main text.

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 A.3. Finally, the cost of not matching a shipment is calibrated at 10% of the shipment value, which is the same heuristic used by the platform in its development of the newer auction format introduced in 2023.

Table A.3: Shipment value distribution parameters

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

Table A.4: Reserve and Accept-Now price estimates

	Reserve Price Model	Accept Now Price Model
v_j	1.3732*** (0.1856)	1.1939*** (0.0823)
v_j^2	-0.3820*** (0.0878)	-0.4189*** (0.0388)
$d = 1$		0.1672*** (0.0432)
$d = 2$		0.1596*** (0.0432)
$d = 3$		0.1586*** (0.0431)
$d = 4$		0.1308*** (0.0432)
Constant	-0.2167** (0.0974)	
R-squared	0.3544	0.1554
R-squared Adj.	0.3542	0.1552
N	6841	18250
R2	0.3544	0.1554

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

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:

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

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 results of the two regressions are reported in Table A.4.

The effective reserve price for the last 24 hours at $d = 0$ is less straightforward to estimate, as there is no unique reserve price during this period, and human brokers sometimes manually intervene to accept a carrier bid. I thus use a moment-inequality approach, based on the simple intuition that the theoretical reserve price must be greater than any bid that was accepted, and less than any bid that was rejected when no alternatives remained (that is, when the shipment remains unmatched).

To formalize this notion, we focus on the set of shipments that have yet to be matched by the start of the last 24 hours before pickup, which we will call \mathcal{J} . Let \mathcal{I}_j denote the set of bidders on shipment j and \mathcal{I}_j^A denote the set of bidders whose bid was accepted. The lower inequality is:

$$r_0(v_j) \geq \max_{i \in \mathcal{I}_j^A} b_{ij}, \forall j \in \mathcal{J}$$

Now, let $\mathcal{J}^U \subseteq \mathcal{J}$ denote the set of shipments that were *unmatched* by the end of the last 24 hours before pickup. The upper inequality is:

$$r_0(v_j) \leq \min_{i \in \mathcal{I}_j \setminus \mathcal{I}_j^M} b_{ij}, \forall j \in \mathcal{J}^M$$

Parametrizing the function as $r_0(v_j, \theta^{R_0})$, I construct the following objective function to be minimized:

$$\mathcal{Q}(\theta) = \sum_{j \in \mathcal{J}} \left[\max_{i \in \mathcal{I}_j^A} b_{ij} - r_0(v_j, \theta) \right] + \sum_{j \in \mathcal{J}^U} \left[r_0(v_j, \theta) - \min_{i \in \mathcal{I}_j \setminus \mathcal{I}_j^M} b_{ij} \right]$$

The derivation of standard errors for this estimator is currently in progress.

A.8 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 without directly affecting the firm's profits. However, this assumption may not hold if reputational penalties also help the firm by *screening* out low-quality carriers. Specifically, if reputational penalties lead to the exclusion of carriers with high cancellation rates, the firm might improve its overall performance by reducing cancellations. A counterpoint to this argument is that excluding carriers with high cancellation rates may also reduce the platform's margins, as these

carriers may also be offering lower bids (recall the negative relationship between cancellation rates and bid amounts in Figure 2).

To evaluate this possibility, 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 bookings, compute the mean cancellation rate and average margins directly from the data.
2. Exclude the top $X\%$ of carriers ranked by their cancellation rate.
3. Evaluate the impact of this exclusion on the overall cancellation rate and margins.

We 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 booking price of a shipment booking and let v denote the value of the shipment to the platform. We focus on two types of margins:

- **Booking Margin** $= \frac{v-b}{v}$
(assumes the booking is kept)
- **Final Margin** $= \mathbb{1}[\text{No cancel}] \frac{v-b}{v} - \frac{1}{10} \mathbb{1}[\text{Cancel}]$
(accounts for the impact of cancellations)

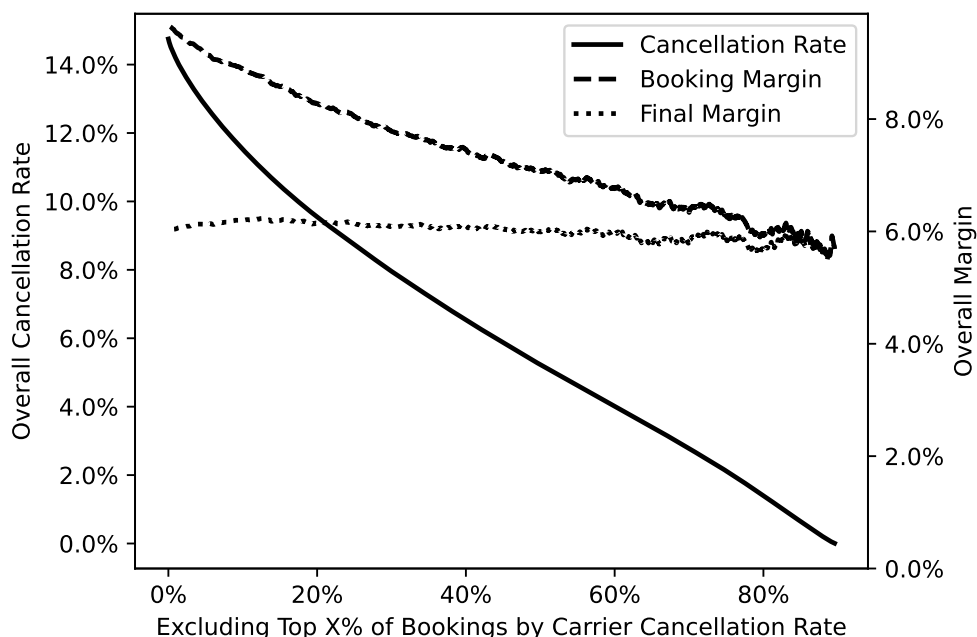
The final margin is the relevant profit margin we will focus on. It assumes the worst-case scenario in case of a cancellation: that a cancelled booking will fail to be matched in time and will incur the indirect penalty of 10% of the shipment value.²³ The final margin is a more accurate measure of the platform's profitability, as it accounts for the impact of cancellations on the platform's revenue, but is again biased in favor of screening, as it assumes a worse impact of cancellations than the platform actually experiences.

Figure A.2 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

23. Recall from Section A.7 that this heuristic was used internally by the firm as a proxy for the impact of failed shipments on their clients' contract renewal probabilities.

cancellation rates reduces the overall cancellation rate but also reduces the booking margins, implying that more frequently cancelling carriers also tend to bid lower. The overall impact of both forces almost exactly cancel each other out in terms of the final margin, suggesting that the net benefit of screening out high-cancellation-rate carriers is minimal.

Figure A.2: Effect of Screening Frequent Cancellers on Cancellation Rates and Margins

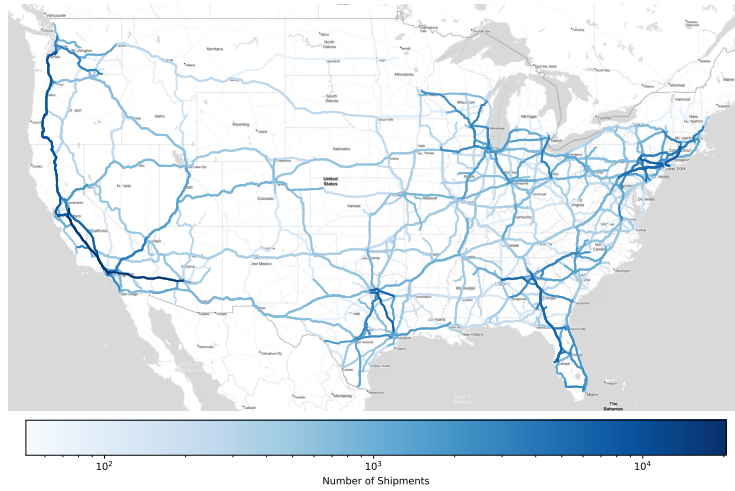


NOTE: Figure plots the effect of excluding the top X % (as a fraction of bookings) of carriers based on their average cancellation rate across all confirmed bookings. Sample restricted to carriers with at least 10 bookings. Left axis represents effect on cancellation rate and right axis represents effect on margins. Booking margin is defined as the difference between the shipment value and the booking price, divided by the shipment value. Final margin is defined as the booking margin, adjusted for the impact of cancellations (assuming a cancellation incurs a 10% penalty for a failed shipment).

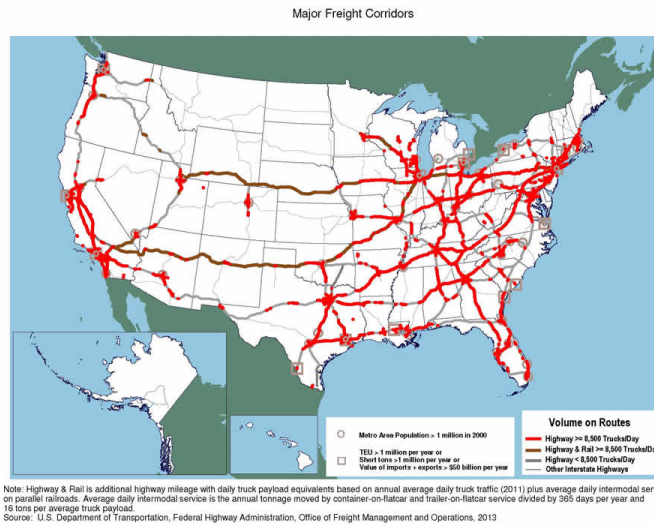
In conclusion, while reputational penalties might have a screening effect by discouraging or excluding frequent cancellers, the net benefit to the firm is minimal under conservative assumptions. The loss of low-cost carriers offsets the gains from reduced cancellations. Moreover, reducing the carrier 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.

B Additional Figures

Figure A1: Comparison of lane frequencies on the platform and across the U.S.



(a) Platform Lane Frequencies



(b) Lane Frequencies Across the U.S.

Figure A2: Distribution of carrier miles driven in a year

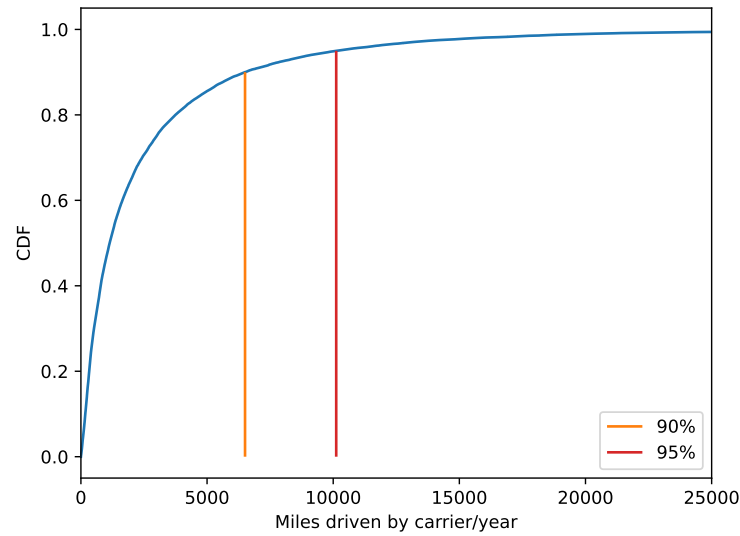
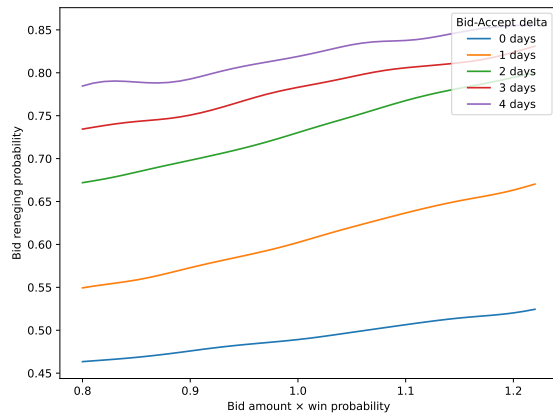
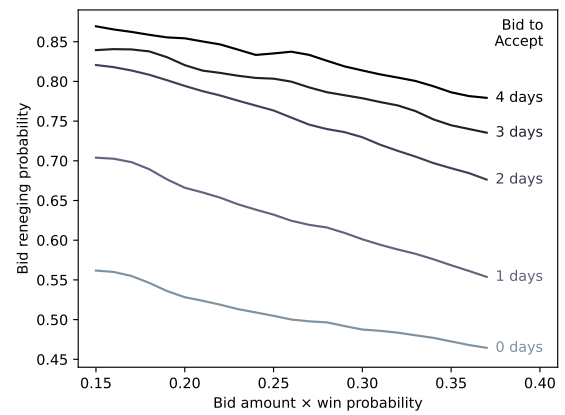


Figure A3: Confirmation probability vs. bid or expected value

(a) Bid amount

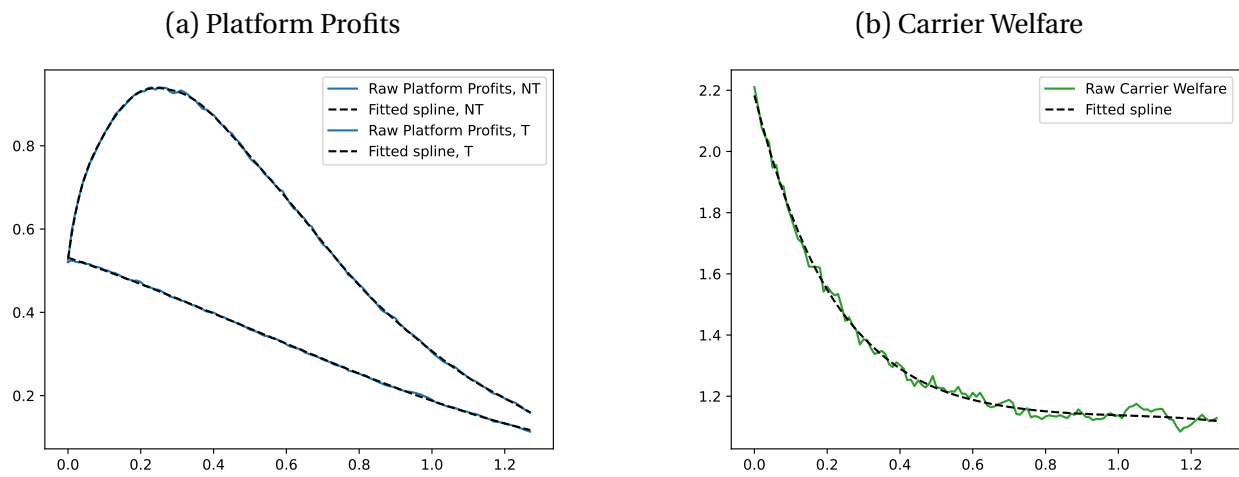


(b) Expected Payoff



NOTE: Figures are obtained through a two-dimensional kernel regression, conditioning both on different time horizons for attrition, and measures of payoffs. Kernel bandwidth chosen in accordance with Silverman's rule of thumb. Win probability in panel (b) is also computed with a one-dimensional kernel regression of bid acceptance on bid amount.

Figure A4: Raw counterfactual outcomes vs. fitted splines



NOTE: Figures compare the raw counterfactual outcomes to the fitted splines, which are used to average out simulation variance. Number of spline knots chosen to achieve a mean squared error less than a threshold s . For platform profits, the threshold is $s = 0.001$, and for carrier welfare, $s = 0.04$.