Human Agency in Last Mile Delivery

Victor Martínez-de-Albéniz[1] Celica Krigul[2]

Abstract

Despite the increasing use of route optimization algorithms, last mile delivery couriers often deviate from these solutions. Using a large dataset of 1.4 million package deliveries in Shanghai, we examine the behavioral factors influencing such deviations. We find that couriers generally favor shorter distances over time-efficient routes. Results show that, while algorithmic recommendations remain influential, proximity to the next stop plays a dominant role, particularly under complex routing conditions or during peak traffic hours. These deviations lead to inefficiencies, with routes and actual travel times being longer than the predicted optimum, especially for inexperienced drivers and on long routes. The findings highlight a need to realign routing algorithms with courier preferences to enhance efficiency in last-mile logistics.

This version: February 10, 2025

Keywords: Behavioral Operations, Logistics and Transportation.

1 Introduction

Last-mile delivery is an essential part of modern economy. In 2022 there were 161 billion parcels delivered, with the largest market being China (Eurosender 2024). The growth in online retailing, especially since the COVID-19 pandemic, has made nearly everyone experience the process of waiting for a courier bringing a package to one's home or workplace. This is not always a pleasant event, usually involving wait and requiring one to be available at the destination, at a certain time – in many cases without precise knowledge about the delivery time, despite promised time windows. Moreover, last-mile delivery is also the most costly portion of the supply chain (Boysen et al. 2021), making up about 50% out of the total delivery costs of a single package (Shaikley 2024). These costs are further increased due to failures to get it right the first time: around 10-15% of packages required re-deliveries, often due to incorrect addresses, failed delivery attempts, or damaged goods. These failures not only increased costs but also led to customer dissatisfaction (Bhattacharjee et al. 2024).

Last-mile delivery, in comparison to many other parts of the supply chain, like distribution centers or trucking, is still a very manual operation, typically involving human workers loading and driving the trucks, and giving the packages directly to the customer. New technologies are entering this space, such as the use of lockers for non-attended deliveries (Rutkowsky et al. 2024), which improves logistics economics –higher per-drop density– and removes temporal shipper-customer coordination constraints –faster transportation, e.g., at night. However, automation is still not generally possible, and great efforts are being made to make the human part of the process more efficient, reliable and cost-effective.

A common strategy for gaining efficiencies is to 'augment' the delivery worker with data-driven decision support. Specifically, delivery companies provide couriers with route recommendations that will help them minimize travel costs – usually either travel distance, travel time, fuel consumption, or a combination of these. These suggestions are usually based on the travelling salesman problem (TSP) or variations of it. The algorithm output is typically displayed to the delivery agent. In principle, this is the optimal route, but of course it may miss some real-time variations and could be subject to change. For instance, two leading last-mile players, Amazon and DHL, are seeking to incorporate courier behavior in these recommendations (Merchán et al. 2024, Arıkan et al. 2023).

In front of these algorithms, couriers retain total autonomy in their decision-making, which enables them to deliver the parcels from their vehicle's trunk in their desired order. They may also take breaks when needed. Due to this, it is common that only an approximated delivery time estimate is provided to the customer. When a short time window is promised, delivery efficiency and quality may be compromised, with an extra cost of about 20% for 2 hour windows (Markakis and Martínez-de Albéniz 2023).

It turns out that couriers do use their autonomy in choosing where to go next, and in the majority of cases they do not follow the suggestions from the algorithm. These deviations possibly lead to reduced performance, such as increased task completion times or higher fuel consumption. Note that this may not always be a detrimental phenomenon, as delivery agents might possess both knowledge and experience that the algorithm simply does not. For example, they may avoid congestion on school streets during the morning, or time visits to office buildings or stores to coincide with periods of high recipient presence for deliveries (Arıkan et al. 2023).

Given the significance of last-mile deliveries, it comes as a surprise that, to date, the underlying factors influencing workers' decisions and specifically deviating from algorithmic recommendations remain largely unexplored. Similar deviations have been documented in other settings before, such as healthcare (Ibanez et al. 2018). In this article, we extend this type of analysis to logistics and empirically study when couriers deviate from a simplified theoretical optimal suggested route.

For this purpose, we use the LaDe dataset from Cainiao, a leading last mile operator in China (Wu et al. 2023), and focus on last-mile routes in Shanghai, containing information about 1.4 million package deliveries over the course of 6 months. While we do not know which navigation support tools are available to the drivers, we compute the optimal route that the drivers should ideally follow – which would be suggested by standard route recommenders like Waze, Google, Baidu or Gade–, and compare it to the actual decisions made by them. Specifically, we first build a machine learning (ML) model to predict delivery speeds, and use the delivery times as input to solve the TSP. We then consider a choice model for couriers to decide on their next stop. The model is grounded on the multinomial logit (MNL), in which we incorporate the algorithmic recommendation as an additional input to the choice process. We calibrate the MNL with real decisions from LaDe.

We find that the TSP-recommended solution is generally chosen with a much higher probability compared to other options, but at the same time other factors have even greater importance: the physical distance to a stop seems to be of primary importance. Note that this is not the travel time, which reveals that there are behavioral preferences for nearby stops, which in principle should not directly contribute to overall, forward-looking route efficiency. This myopic preference for short distances is exacerbated by worse traffic conditions during rush hour, and by the complexity of the routing problem when there are more packages left to be delivered.

Furthermore, we study the impact of deviations and find that they lead to routes that are much longer than they could be, with + 91% predicted travel time on average, with a median of + 54% (this difference is due to heavy upper tails). There is also significant deviation on the actual vs. predicted travel time, with an average of + 98% and a median of + 51%. This suggests that couriers do not possess better information about traffic conditions, compared to what the algorithm could predict. Consistently with the literature, the cost of deviation is higher for inexperienced drivers. We also obtain novel insights: the cost of deviation increases when the routing problem is more complex, i.e., when routes are longer; in this case, however, the total time prediction is less volatile so the possibility of extremely large deviations is reduced.

Our research makes three relevant contributions. First, we document a significant discrepancy between human agency and classic algorithmic recommendations in last-mile deliveries. This descriptive statement is important to reflect about the practical use of routing algorithms, and specifically to perhaps reconsider how policies are structured and communicated. It also augments the contexts in which human deviations have been reported in the literature, as extensively discussed in §2. Second, we uncover evidence explaining the patterns of deviations: our empirical analysis reveals a preference of couriers for quick wins, avoiding long transitions even when those would save travel time in the course of the entire route. In particular, distance has the strongest influence in the likelihood of choosing the next stop, while time has a smaller impact. From an algorithmic perspective, this finding is useful in evaluating solution quality, so that routes with long stretches should be avoided, and thus in introducing improvements in routing algorithms. Third, we find that the consequences of human agency are negative, with losses of efficiency occurring both from longer routes, and from longer times on the actual routes. This insight should motivate last-mile delivery firms to invest in route recommmendations that perform well and at the same time are aligned with courier preferences. Overall, our work extends some general findings about human-AI interactions (prevalence of deviations, identification of predictors of deviation, and lower performance after deviation) to the context of logistics and routing in particular, which should prove valuable in the pursuit of better algorithms.

The rest of the paper is organized as follows. We discuss the related literature in §2. Section 3 describes the utility model and identification strategy, while §4 describes context and data. Section 5 reports our empirical findings. Section 6 analyzes the impact of agency on performance, and brings out the influence of courier experience and route length. Section 7 concludes. Analysis details and supporting tables are included in the Appendix.

2 Related Literature

In recent years, there has been a growing interest in better understanding human-algorithm interactions. There are numerous works looking at the impact of algorithms on performance and human perceptions (e.g., Bai et al. 2022, Dell'Acqua et al. 2023), so we focus our review on past works studying two main questions. First, behavioral studies have looked into the psychological forces that drive human deviations from algorithmic suggestions, as well as their organizational, operational and managerial implications. Second, some recent research has proposed how to integrate deviations into algorithm design.

2.1 Deviations from algorithm suggestions

A rich empirical literature has studied how humans use algorithmic 'advice' or 'suggestion', which could take the form of a prescription that either needs to be validated by the human or overwritten, or a recommendation that is meant to qualitatively guide the decision-making process. The format of those suggestions varies across contexts, but non-adherence generally leads to poorer performance.

- In inventory ordering decisions, Van Donselaar et al. (2010) is one of the first papers showing that managers do not always follow automated replenishment quantity recommendations. They show that deviations are driven by operational factors – missing in the algorithm – and propose improvements to integrate these into better recommendations.

- In assortment planning, Kawaguchi (2021) shows in a field experiment that vending-machine agents do not frequently follow recommendations, although not using the advice does not result in lower performance. Interestingly, incorporating their beliefs into the algorithm increases adherence.

- In pricing decisions, Kesavan and Kushwaha (2020) report that automobile industry's merchants' deviations from data-driven decision-making reduce profitability. Caro and de Tejada Cuenca (2023) make a similar observation and show that markdown pricing adherence increases when the products' inventory was perceived to be more scarce by the pricing managers. This is one of the examples that suggests that humans may be optimizing for a target variable that is different from the one used by the algorithm. Another relevant article on pricing deviations is Elmaghraby et al. (2015), who develop an econometric model for how sales agents deviate from recommended prices.

- In healthcare, Ibanez et al. (2018) study how radiologists deviate from a prescribed sequence of tasks, and found that radiologists tend to prioritise faster tasks, and that they group similar tasks together, in order to make their workload cognitively easier to handle. This results in lower productivity rates. Interestingly, more experienced doctors show higher productivity rates even though they deviate more.

- In warehouse tasks, Sun et al. (2022) study deviations in box size packing recommendations. They classify this deviation in two categories: information and complexity causes. They suggest that the algorithm, which may be overly simplistic, could therefore lack information that the human has, and consequently the algorithm's proposed solution is impossible to implement, thus making deviations inevitable. Furthermore, they find cases where the algorithm proposes an executable solution that requires more effort from the worker, and this leads to the suggestion being rejected.

- In inspection processes, Ibanez and Toffel (2020) report that food-safety inspectors report less violations at the end of their work day, and when there is a risk that this inspection will extend the duration of their workday. The inspection is also found to be stricter when it was preceded by a larger number of violations. This shows that an important factor in deviations from the inspection standard is fatigue; the latter has also been identified as a first order effect in decision making in Bavafa and Jónasson (2024) and Aouad et al. (2022).

- In retail, Hui et al. (2009) show that shoppers do not behave according to a prescriptive efficient route within the store (even though there is no recommendation given to them). More recently, Moon (2023) and Knight et al. (2022) suggest that advice on store routes can improve picking performance.

- In last-mile delivery, Arora et al. (2024) study the reasons behind driving agents marking deliveries as failed, when in practice they never attempted to make these deliveries. They find the reasons to be previous failed attempts at this address, low chances of complaints from this customer and high probability of a successful attempt the next day.

Beyond the human reaction to guidelines or suggestions, some research has looked into how humans conduct problem-solving activities and specifically combinatorial problems, in comparison with algorithmic procedures. Adams et al. (2021) study the knapsack problem and show that people are more likely to consider additive changes over subtracting changes, to reduce the number of possibilities to consider and therefore decrease their mental workload. On the other hand, Pape et al. (2020) find that making individuals aware of their biases towards too many smaller items improves their performance. These papers all suggest that people have a preference for prioritizing faster, smaller tasks. This strategy could provide them with a sense of accomplishment, as shown by Converse et al. (2023): people find value in task completion, and show a preference towards completing sooner-to-finish tasks over higher-reward more time-consuming tasks. Similarly, we find in our study that drivers prioritize shorter distances to select the next stop in their route.

Furthermore, the adherence to algorithms has also been investigated experimentally. Kremer et al. (2011) conduct an experiment where the subjects were told to forecast demand at a retail store based on knowledge about an uninformative random walk. They find the human predictions to deviate from normative ones, and that the detrimental effects of this deviation are more visible in stable environments. Balakrishnan et al. (2022) show that human decision-makers consider both their own prediction and the algorithm's one and take the weighted average of the two, and that this results in increased prediction errors. Dietvorst et al. (2015) investigate people's aversion to algorithms, and show that reliance on the algorithm drops after finding out that it makes mistakes. Cameron (2022) observes a similar phenomenon among customer-focused Uber and Lyft ride-service providers, who report believing that the customer-driver matching algorithm works very well, even though they are unaware of the mechanisms behind this algorithm. Among efficiency-focused drivers, however, the most common belief is that the algorithm is flawed and skewed. Some suggest that the human judgment could be used as a yet another data source for the algorithm, e.g., Käki et al. (2019) or Angelopoulos et al. (2023). More generally, giving users some control on the algorithm improves adherence (Dietvorst et al. 2018).

Finally, there is some work that connects deviations to user states. Specifically, experimental studies have investigated how deviations in computer mouse cursor movements, in comparison with a straight line, are linked to both the presence of negative emotions and fraud. Hibbeln et al. (2017) find that a low emotional state of the user leads to decreased attentional control and results in deviations from the straight line. Weinmann et al. (2022) show that the deviation of mouse movements is associated with the decision to commit fraud and the magnitude of the fraud.

2.2 Design of algorithms

There is also some research on how to integrate non-adherence in the design of algorithms. Bastani et al. (2021) have suggested a way to introduce algorithm recommendations as tips instead of a ready-made decision, and shown that the implementation of these tips by the human decision-maker improves their performance in the choice-making process. In contrast, Grand-Clément and Pauphilet (2024) propose an algorithm that is in between what the human decision-maker with its own preferences would choose and what the algorithm would choose, in order to bring the human closer to system optimum.

Finally, closest to our paper is the work that seeks to build logistics routes from human decisions. Indeed, Mao et al. (2019) find that drivers with more experience and/or more local knowledge achieved lower delivery times. In the introduction to the Amazon routing challenge, Merchán et al. (2024) suggest that we can learn from effective driver routes to build routes in other contexts, based on machine learning instead of routing algorithms. Among the proposed solutions, Cook et al. (2024) suggest to augment the TSP with constraints identified from the data. Arıkan et al. (2023) describe DHL's procedures, which are also based on identifying feasible vs. impossible transitions. Dieter et al. (2023) propose a variation of the TSP with an explicit consideration of the deviation in the optimization program.

Note that machine learning has been applied to solve similar problems in the past. Firstly, Li et al. (2018) trained a graph convolutional network; when used in combination with classic heuristics, it outperforms previous deep learning approaches both in solution quality and run time. Nazari et al. (2018), Hu et al. (2021) and Zhang et al. (2020) developed models that outperform industry standards, namely Google's OR-Tools. For that, Nazari et al. (2018) used reinforcement learning to find near-optimal solutions to solve the vehicle routing problem. Hu et al. (2021) used a bidirectional graph neural network. Zhang et al. (2020) used deep reinforcement learning to solve TSP with time windows and rejection. This does not mean that industry standard tools are not competitive: Joshi et al. (2019) used deep Graph Convolutional Networks but did not see improvement.

3 Choosing Delivery Sequences

In the following section, we describe the theory behind the choice of delivery sequences. At each pair of time slot and day of the week t, a courier a is in current location i, and needs to build a route across locations within the set $J = \{1, \ldots, n\}$. The actual decision is not the entire sequence for the route, but simply the next location $j \in J$ to be visited; from j, the same problem will be solved with available locations $J/\{j\}$.

Each move from current location i to next location j is associated with a travel cost c_{ij}. In simple terms, this cost is defined as the predicted time (in minutes) taken to travel from the current location i to next location j. Note that c_{ij} may depend on the time t at which the prediction is made; in our empirical study, we allow a different time prediction for each combination of day of the week and hour of the day.

The classical, normative approach is to solve the open-ended Travelling Salesman Problem (TSP), to find the optimal route. It visits all the remaining delivery locations only once without coming back to the starting point, and in the least amount of time as possible per our predictions. The TSP can be formulated as follows (Bertsimas and Weismantel 2005):

$$\begin{aligned}
\text{minimize} \; & \sum_{i,j \in V} c_{ij} y_{ij} && (1)\\
\text{subject to} \; & \sum_{\{i|(i,j)\in A\}} y_{ij} = 1, \text{for } j \in V, \\
& \sum_{\{j|(i,j)\in A\}} y_{ij} = 1, \text{for } i \in V, \\
& \sum_{\{(i,j)\in A | i \in S, j \notin S\}} y_{ij} \geq 1, \text{for } S \subset V, S \neq \{\emptyset, V\}, \\
& y_{ij} \in \{0,1\}, \text{for } (i,j) \in A, \\
& && (2)
\end{aligned}$$

where i, j are vertices within the vertex set V, and c_{ij} is the cost of going from i to j on a feasible arc $(i, j) \in A$. y_{ij} is the binary decision to travel from i to j. The first and second constraints ensure that j only receives one incoming and outgoing arc. The third constraint ensures that the route is connected by ensuring that any partition has at least one outgoing arc. In this formulation, the TSP takes as given the cost matrix computed at time t, even though some of the transitions may take place later on. This is a reasonable assumption since traffic conditions are typically unpredictable, and current navigation decision support systems also do not usually include forecasts of future traffic conditions.

From the TSP recommended route, we let j_{TSP} be the optimal location to which we should travel to, from current location i, i.e., $y^*_{ij_{\text{TSP}}} = 1$. We can compare this decision with the actual decision of the courier j_{actual}. If $j_{\text{TSP}} = j_{\text{actual}}$ then the courier's decision coincides with that prescribed by the TSP. On the other hand, if $j_{\text{TSP}} \neq j_{\text{actual}}$, then the courier has deviated from the theoretical optimal solution. In the latter case, from the new location j_{actual}, we will need to solve TSP again. In case the context t has changed, we may need to update the cost matrix $c_{ij}(t)$.

As discussed in all past works, couriers do not follow the TSP in the majority of cases; Li and Phillips (2018) observed that 75% of deliveries did not follow the prescribed order. We thus build a choice model in which the courier has at his disposal the recommendation from the TSP, j_{TSP}, and the rest of available options in $J \setminus \{j_{\text{TSP}}\}$. Each of the options in J has a set of characteristics X_{ij}, where we explicitly mark the possible dependency of the characteristics on the courier's current location i. Furthermore, for every agent a we can think of a set of moderators Z_{at}, which may affect the importance of the features X_{ij}. These can be interpreted as effects that depend on agent a, such as experience, as well as dependencies on context t, such as traffic conditions.

Based on this information, we can build a choice model to predict the likelihood that $j_{\text{actual}} = j$, for all $j \in J$. We are thus not interested directly in whether the courier acts in line with the forward-looking optimum from the TSP, but rather on inferring which factors drive the courier choices (being the TSP solution is one of them). Specifically, we use the canonical MNL to identify the factors that explain the decisions of actual couriers. Of course, more sophisticated choice models, like nested MNL, mixed MNL or Random Utility Models such as RUMnet could be used (Aouad and Désir 2022), but given that our objective is to detect the most salient drivers of choice, we opt for the simplest specification.

To briefly introduce MNL (see Train 2009 for an excellent textbook treatment), at each stop rank, a courier has available many alternatives of where to go next. The utility that an agent obtains from each alternative can be expressed as:

$$U^{ij}_{at} = V^{ij}_{at} + \varepsilon^{ij}_{at}, \tag{3}$$

where $V^{ij}_{at} = \beta Z_{at} \cdot X_{ij}$ is the expected utility of choosing j from i, for agent a in context t;

and ε_{at}^{ij} is a random shock, which is assumed to be Gumbel distributed and independent across all decisions. Under this utility structure, the choice probability can be computed as

$$Pr[j_{\text{actual}} = j | i, a, t] = \frac{\exp(V_{at}^{ij})}{\sum_{j' \in J} \exp(V_{at}^{ij'})}. \tag{4}$$

Equation (4) is our base model. We will estimate it with actual decisions taken by couriers, as discussed next. Note that the data that we use is observational and we have no guarantee that X_{ij} are exogenous – a necessary step to interpret our results causally. However, the enormous variation in the options given to the agent, i.e., the packages that the agent has in the delivery vehicle, as well as the multitude of couriers and geographical locations to be visited suggest that our estimates should reflect the consistent biases that couriers may hold. Of course, experimental validation can be a useful next step in further strengthening the evidence that we provide.

4 Institutional Context and Data

4.1 LaDe Deliveries

LaDe is a public last-mile delivery dataset made available by Wu et al. (2023), obtained from operations of the leading logistics provider Cainiao (see Cui et al. 2020, Bray 2023, Zhan et al. 2023, Lu et al. 2023 for other articles using Chinese logistics and specifically Cainiao as an empirical context). The dataset contains both last-mile delivery and pick-up data for five cities in China, each city with different characteristics and a unique set of problems and advantages. From this dataset we choose to analyze the last-mile delivery data in Shanghai, which is a 26 million inhabitant city with significant congestion and high volume of deliveries. In Shanghai, the original dataset features 1,733 couriers, for which all their routes are available during 6 months. (Sampling is done at the courier level, so there is no censoring within a given courier). These couriers completed 1,483,864 package deliveries. They had a mean number of 40.6 working days over the 6 month sample, ranging from 1 to 184 days with a median of 12.

As couriers usually operate in the same general area and metropolitan Shanghai being very large, the data does not necessarily cover the entire city; we are missing activity about some areas, which is probably due to the couriers dedicated to those areas being sampled out from the data. The distribution of deliveries is shown in Fig 1a. These areas include some more dense (e.g., Hongkou and Yangpu districts) and some less dense (e.g., Laogang) areas. There is a higher number of

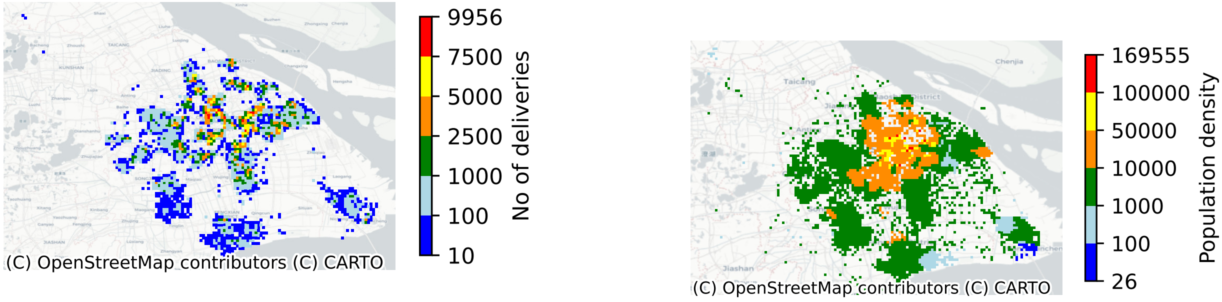

(a) Total deliveries made to approximately 1x1km squares in the full dataset of Shanghai. Only squares with at least 10 deliveries are shown.

(b) Population density in approximately 1x1km squares in Shanghai.

Figure 1: LaDe-D Shanghai dataset overview on a map.

deliveries made to the more dense central areas of the city, particularly to Jing'An.

From this raw dataset, we build a panel of sequential choices in which agents had to decide which package to deliver next, in line with the theoretical model (4). For this, we first cleaned the dataset and removed inconsistencies in geolocation (e.g., the GPS coordinates were not located in nor near Shanghai; this filter removed 80 courier-date pairs out of 70,336). We then split the full activity of a courier in one or more portions, that yielded independent routes. The criteria that we used for breaking a sequence into multiple routes was to have at least 2-hour breaks in between deliveries. In practice, this means that the courier can go back to the distribution center and reload the delivery vehicle with new packages, hence it is appropriate to handle the two portions as separate TSP problems. Within each route (the selected portions), because delivery addresses are usually in the same region and there is a consistent flow of deliveries, we no longer find long gaps (of 2 hours or more), which suggests that the tour was probably uninterrupted and moreover, the courier was not engaging in non-work-related activities that could alter optimal routing goals. Overall, we transformed 70,256 original routes into 117,096 'uninterrupted' routes. Each resulting route had a mean number of 12.65 stops, in a range of 1 to 122 stops with a median of 8. Finally, we only considered tours with a length of at least 3 stops, and at most 19 stops. Indeed, for routes consisting of only 2 stops, the TSP is not meaningful as there is only one option to choose from. Routes of 20 or more stops were computationally hard to solve, and the exact TSP solution could not be computed in reasonable time (we used the python-tsp library for this). In total we were then left with 56,916 routes, which after further cleaning and invalid responses from the API (explained below) came down to 54,390 for the final analysis. Table 1 shows an example of a 5-stop route.

Deliveries take place over the entire day (between 5:00 and midnight), according to the distri-

courier id	date	route	stop rank	current gps	next gps	time	distance
3541	14-09	2	1	31.21967, 121.68774	31.2309, 121.69242	17:07	2.1957
3541	14-09	2	2	31.2309, 121.69242	31.23253, 121.68625	17:28	0.9282
3541	14-09	2	3	31.23253, 121.68625	31.22994, 121.68844	17:38	1.1434
3541	14-09	2	4	31.22994, 121.68844	31.22022, 121.6892	17:45	1.9583
3541	14-09	2	5	31.22022, 121.6892	31.2299, 121.68968	17:54	2.0087

Table 1: Sample of observation for one courier on one route.

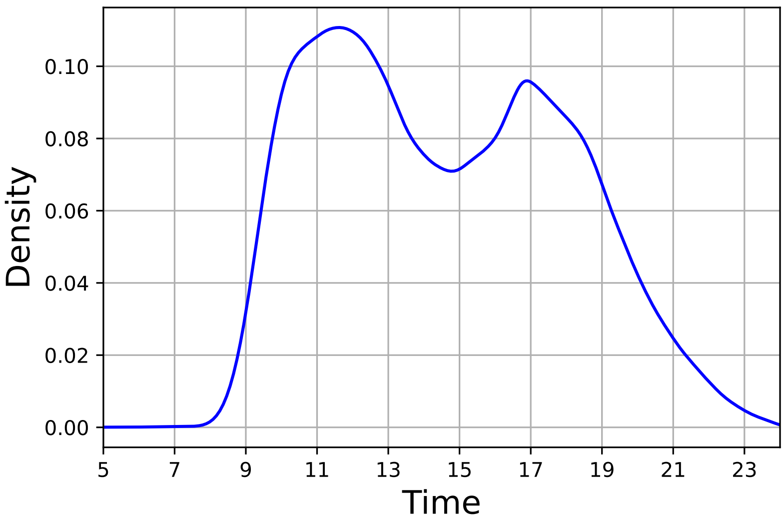

Figure 2: Distribution of delivery times over the day.

bution shown in Figure 2; most of them scheduled after the morning rush hour (between 7:30 and 9:30).

4.2 Obtaining Distance, Speed and Time

To build decision models, it is necessary to obtain travel time and distance between any two locations within a route. Of course, we need to expand this from actual transitions made to any combination of two points within a route, to be able to obtain the necessary input for the cost matrix $c_{ij}(t)$ in the TSP. Since it is prohibitively costly to query a routing recommender such as Gaode for a delivery time estimate (we can also not query past dates, since recommenders use real-time traffic conditions), we proceed in three steps.

First, we obtain the distances between delivery locations, through the publicly available OpenStreetMaps car distances, obtained by using their API (Luxen and Vetter 2011). It is important

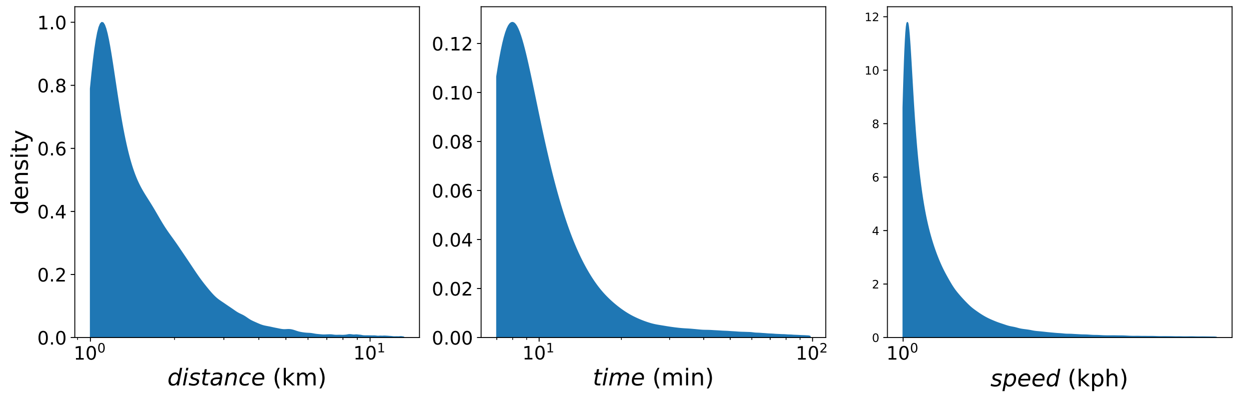

Figure 3: Density of distance (left), time (center) and speed (right) between two consecutive stops in the actual routes. The first and last 1% excluded from the data for visualization.

to note that due to this public mapping service not being perfect, and due to minor differences in OpenStreetMap and other distance-calculating services, the distance variable may be somewhat noisy. However, we check that the distance estimate has a correlation of 0.838 with bird-eye distances computed from latitude and longitude, and our estimate should in general be more reflective of actual distance compared to bird-eye distance.

Second, to obtain the expected minutes taken to travel from each current location i to each next location j at this specific t, we train a random forest regression model for the logarithm of speeds, using the following features: latitudes, longitudes, time of the day (numeric) and day of the week (categorical). This model was trained on about 750,000 observations and we tuned its hyperparameters by searching for the best combination of values during 48h, resulting in a reasonably good prediction performance for the majority of our data, as shown below on Figure 4. We show in Figure 3 the distribution of distance in km, time in minutes and speed in km per hour (kph) of any transition from A to B in the actual routes for the data used in the ML model. Applying the natural logarithm of these metrics plus one (to deal with zeros) allows us to deal with skewedness and to improve the statistical properties of these variables.

Third, we obtain a prediction of travel time by dividing actual distance with predicted speed. Note that we decide to use a model of speed rather than a model of time so that we could integrate actual distance in the prediction process (as opposed to letting the prediction directly infer distance from latitude and longitude). Specifically, the predicted travel time to each next available location j from each i at this t is obtained as:

$$time_{ij|t} = \frac{distance_{ij}}{predicted\_speed_{ij|t}} \tag{5}$$

where $predicted\_speed_{ij|t}$ is obtained from the ML model, in which we we applied a smearing correction to $predicted\_log\_speed_{ij|t}$ (Duan 1983).

	R^2	std. deviation	MAE	ME
In-sample	0.905	0.149	0.0243	-0.000302
Out-of-sample	0.596	0.0854	0.0752	0.000522

Table 2: Logspeed model overview. MAE stands for Mean Absolute Error and ME for Mean Error. Note that because we are considering logspeed, absolute errors in logspeed become percentage errors in speed.

The results of the model are shown in Table 2. For training data, $predicted\_log\_speed_{ij|t}$ took a median value of 0.046 and a third quartile of 0.119, corresponding to $predicted\_speed_{ij|t}$ of 2.85 and 7.60 kph, respectively. It should be noted that this is not the speed at which the vehicle travels, but instead the average speed taking into account the time it took for the agent to travel from location i to location j and the time it took to deliver the package after already stopping the vehicle. For the test data, the distribution is nearly identical, with a median value of 0.046 and a third quartile of 0.120. On the same test data, the model predicts the exact same values with 0.046 for the median and 0.120 for the third quartile. Note that in-sample data is inherently more noisy than test data, with a higher goodness of fit (90.5% vs. 59.6%) but a larger residual error (0.149 vs. 0.085). Figure 9 in the Appendix provides additional details about model accuracy as a function of distance. We can furthermore compare prediction and actual values on the training and test data, as shown in Figure 4. We see that the model predicts without any bias for low values of logspeed and small uncertainty, and only becomes unprecise at the very high end of the distribution. This suggests that the model performs adequately for the task at hand.

These values of expected minutes are then used in the cost matrix $c_{ij}(t)$ and given as input to TSP. This matrix changes dynamically as for each decision a new matrix is created by dropping the row and column corresponding to the previous current location i, as this location has now been visited and is no longer available to be chosen. In case the time slot t has changed, the values of the cost matrix will therefore also change, as the expected minutes taken to travel from i to j are dependent on the day and the time slot, indirectly reflecting historical traffic conditions.

In addition, as the problem at hand is open-ended, that is, it is not required for the courier to return to the starting point after the end of the route, we always set the first column of the cost matrix to zeros, representing that returning from any point to the starting point results in cost zero, that is $c_{j0}(t) = 0$.

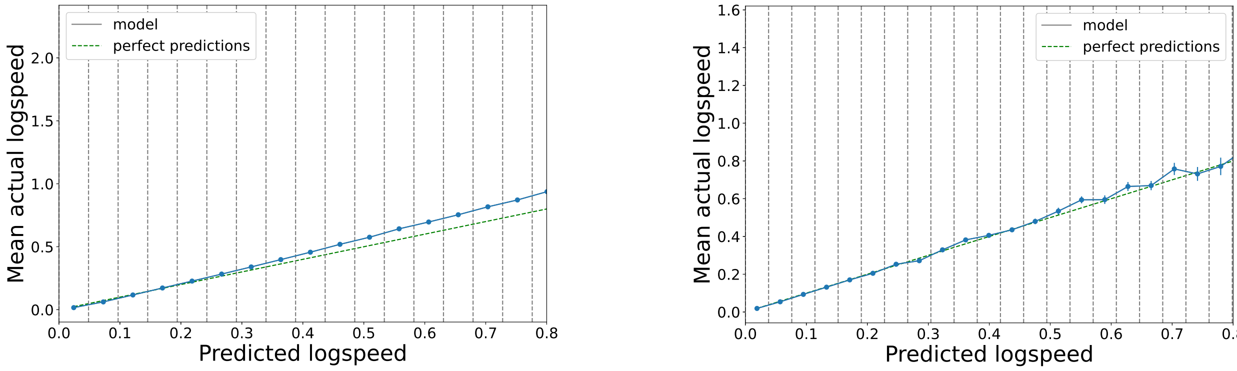

(a) Train data

(b) Test data

Figure 4: Model for $predicted\_log\_speed_{ij|t}$. In the x-axis, we group all predictions within a bin and in the y-axis, the actual average value for those predictions. 90% of the data used for training and testing falls below 0.24.

4.3 Variables

The main driver of choices for couriers should be whether a stop appears as the next stop within the TSP. We thus let is_tsp_{ij} be the binary variable capturing whether $j = j_{\text{TSP}}$, to account for the forward-looking nature of the courier.

However, couriers may deviate from the TSP. This could be due to the existence of valuable local information coming from courier's knowledge and experience (Arıkan et al. 2023, Dieter et al. 2023). To model courier's decisions of choosing one delivery location over the other, we incorporated several geometrical and non-geometrical variables into our analysis.

The first variable that we investigated is the shape of the route formed by all available delivery locations in the set J. As the minimum route length is 3, the vertices to be visited form a polygon. We are interested in looking at the convex hull corresponding to that polygon, and we are specifically interested in whether the remaining delivery locations are inside of that convex hull or on the edges. Indeed, when presented with such a convex hull, visiting first the points that are on the edges reduces the size of the remaining polygon, and removes the points that are the farthest from the others, giving a sense of accomplishment. In our model this feature is denoted by $is\_outside_{ij}$ and is defined as a binary variable.

Another geometrical variable we choose to incorporate in the analysis requires us to look at the before-mentioned convex hull and to find its center of gravity. Then we imagine that the courier is placed in the current location i and is looking towards the center of gravity of this convex hull formed by all remaining j. He then judges whether each possible next location j is to the left or to

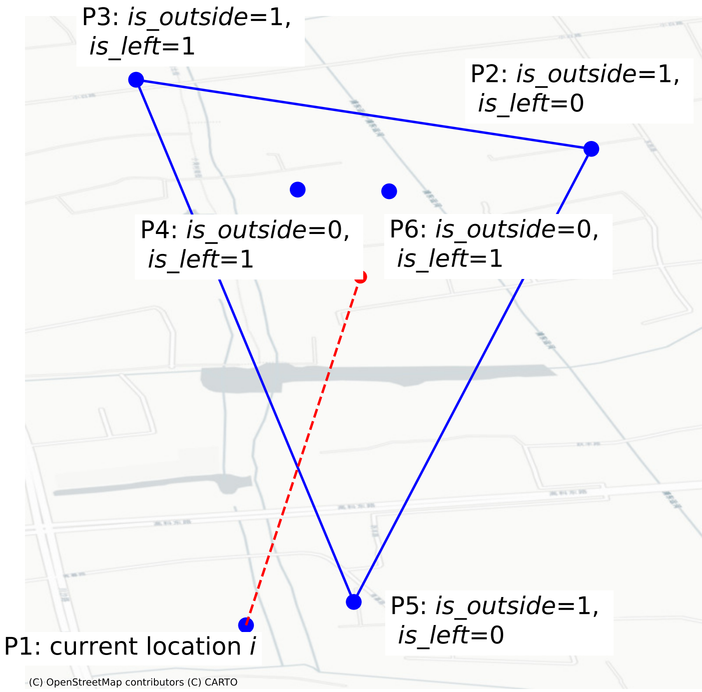

Figure 5: Choices for a 6-stop tour.

the right from the center point. This ties in with the recommendation to avoid left turns (Ludwig and Geller 2000, Rosenbush and Stevens 2015), which would lead couriers to go left first (just once) so that they can turn right later (in all remaining stops). We call this variable is_left_{ij}.

Both $is\_outside_{ij}$ and is_left_{ij} variables are illustrated on Figure 5 where we can see a convex hull presented to the courier completing a 6-stop route and currently in stop rank 1, formed by the remaining delivery locations 2, 3 and 5. These points we classify as being on the edges of the convex hull, and delivery locations 4 and 6 are inside the convex hull. Furthermore, we visualize a straight line going from the courier's current location to the center of gravity of this convex hull to classify whether each point in J is to the left or right from this line. In our example, points 2, 5 and 6 are therefore defined as being to the right, and points 3 and 4 are to the left.

In addition to the geometrical variables, we focus on two other important characteristics. We include log_time_{ij} the logged predicted time to travel from current location i to each potential next location j, to account for the preference for a quick next location, shown in Equation (5). Observe that this may introduce correlation between tsp_actual_{ij} and log_time_{ij}, but we have sufficient variability to include both variables simultaneously. We also define $log\_distance_{ij}$ the logged actual distance from i to j, as described in §4.2, to account for the preference for a near next location which may or may not be the fastest to reach.

Finally, we include three moderating factors that modulate the importance of these five features: $log\_remaining_i$ the logged remaining packages to deliver on the route, to account for the complexity of the routing problem – at the beginning of the route, the courier solves a much larger and complex problem, compared to later towards the end of the shift–; $log\_experience_i$ the logged number of days the courier has previously worked in this company, to account for experience and also familiarity with the surroundings (given that couriers are assigned to the same areas over and over); and $rush\_hour_i$ a binary indicator that the choice is made during the afternoon rush hour between 17:00 and 19:00 (recall that few deliveries take place during the morning rush hour so we do not include it in the analysis). Note that these three factors are common to all the options $j \in J$, so their direct effect cannot be estimated for lack of variation. But they will be informative in reflecting the local knowledge that agents may have, and how they shift the weight of the five main covariates.

Table 3: Summary statistics.

Variable	Type	Mean	Std. dev.	Min	Max
is_tsp	binary	0.162	0.369	0	1
$is\_outside$	binary	0.646	0.478	0	1
is_left	binary	0.502	0.5	0	1
log_time	float	2.56	0.72	0.693	10.5
$log\_distance$	float	0.84	0.591	0	4.55
$log\_remaining$	float	1.78	0.63	0.693	2.94
$log\_experience$	float	2.72	1.63	0	5.21
$rush\_hour$	binary	0.204	0.403	0	1

5 Results

5.1 Low Adherence to the TSP

Before showing the empirical results, we provide some model-free evidence of the couriers' adherence to the algorithm's suggestions. We define tsp_actual_i as the binary outcome variable that determines whether the chosen next stop was the one consistent with the TSP.

The average adherence in our dataset is quite low, at 47%. When we further separate this

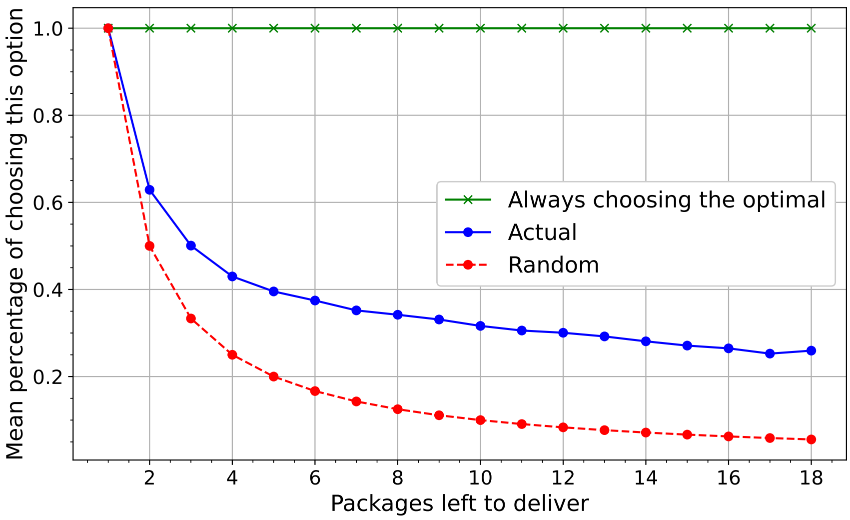

Figure 6: Adherence to the TSP solution.

adherence according to the number of packages left to deliver on this route, we see in Figure 6 adherence values decrease from 65% when there are two packages left to deliver, down to below 30% when there are 14 or more packages still to be processed. These values are higher than that of a random model (dashed line): in a random model adherence values of 50% can only be achieved when the last two packages are left to deliver. At the same time, actual adherence values are much lower than these of a perfectly obedient courier, who would comply 100% of the time, at each decision point (crossed line). Interestingly, we find that adherence depends mainly on the number of packages left to be delivered and is very stable throughout the day (see Table 9 in the Appendix) or across different levels of courier experience.

Once we have established that couriers do not adhere to TSP recommendations, we can investigate the drivers of choices for the next delivery location in §5.2. We then focus on the factors that guide adherence in §5.3.

5.2 Predicting Choices of the Next Stop

We next estimate Equation (4) using the *pylogit* package in python. In this model, each observation is uniquely identified by courier ID, date, route and current location, and faces a choice among the set of possible next stops. We first include only the five main covariates presented in Table 3, and then add the moderators to the analysis. The results are shown in Table 4.

	Dependent variable: chosen location	
	(1)	(2)
is_tsp	0.501***	0.125***
$is\_tsp \times log\_remaining$		0.193***
$is\_tsp \times log\_experience$		0.005
$is\_tsp \times rush\_hour$		-0.05***
$log\_distance$	-1.265***	-0.294***
$log\_distance \times log\_remaining$		-0.555***
$log\_distance \times log\_experience$		0.066***
$log\_distance \times rush\_hour$		-0.263***
log_time	-0.318***	0.126***
$log\_time \times log\_remaining$		-0.138***
$log\_time \times log\_experience$		-0.033***
$log\_time \times rush\_hour$		0.019
$is\_outside$	0.393***	0.358***
$is\_outside \times log\_remaining$		-0.015
$is\_outside \times log\_experience$		0.019***
$is\_outside \times rush\_hour$		-0.019
is_left	-0.009**	-0.003
$is\_left \times log\_remaining$		-0.003
$is\_left \times log\_experience$		0
$is\_left \times rush\_hour$		-0.005
Number of observations (choices)	354,190	354,190
Number of variables	5	20
AIC	1,171,772	1,166,727
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 4: MNL estimation results

Main results. The results of Model 1 in Table 4 clearly show that, even when adherence is rather low, being the TSP recommendation boosts the likelihood that an option is selected. Specifically, it is 1.65 times ($e^{0.501} = 1.65$) more likely to be chosen compared to options that are not in the TSP solution. Given the magnitude of this effect, this suggests that deviations should only occur when the rest of the covariates favor a non-TSP option.

What are these other variables? We observe that *log_distance* has the highest coefficient among all, at -1.27. This, combined with the higher standard deviation and range of this covariate (Table 3), suggests that distance has a very large effect on choices. For instance, options that are 1% further away are chosen 1.3% less often ($1 - e^{-1.27\ln(1.01)} = 0.0126$). Similarly, *log_time* is also important, with a coefficient of -0.318. These two covariates suggest that couriers prioritize *quick-win* tasks that can be completed fast, even if they impose longer routes over the entire sequence. Surprisingly, distance has a more important role than time, even though the scarce resource in last-mile deliveries should be courier available time. They turn out to have a very sizeable impact: a location is more likely to be selected than the TSP recommendation if the latter is 37% further away (both in distance and time; $e^{0.501-(1.265+0.318)\ln(1+\Delta)} = 1$ yields $\Delta = 0.372$). This effectively means that the TSP solution will not be adopted if it recommends transitions that are much further away than the nearest one.

Additionally, when it comes to geometrical variables, we see that it is 48% more likely ($e^{0.393} = 1.48$) that an option on the edges of the convex hull will be chosen, compared to the options inside the convex hull. Thus, agents have a preference for *cutting corners* so that the remaining delivery area is shrunk. Furthermore, we observe that options on the left are slightly less common than those on the right, although the effect is small. This is consistent with policies that avoid left turns (Ludwig and Geller 2000, Rosenbush and Stevens 2015), even though they will require left turns later in the sequence.

Moderation analysis. We can now look at the interaction terms to discuss how the preference for quick wins is affected by route and courier characteristics, in Model 2 of Table 4.

We observe that in the beginning of the route (when *log_remaining* is large), the choice probability is largely driven by coefficients 0.193 for *is_tsp* and -0.555 for *log_distance*, so that $e^{0.193M-(0.555+0.138)M\ln(1+\Delta)} = 1$ with a large M yields $\Delta = 0.320$. In contrast, at the end of the route (*log_remaining* small), $e^{0.125-(0.294-0.126)\ln(1+\Delta)} = 1$ yields $\Delta = 1.104$. This comparison suggests that the TSP solution is much less likely to be followed at the beginning of the route rather than at the end. This is in line with Figure 6.

Regarding the effect of agent experience, we see that as the couriers have worked longer in the company, they give less importance to distance (also slightly more to time), while they give equal importance to is_tsp. This implies that they tend to opt for the TSP solution more often. We can interpret this as learning from the part of the couriers, and that they realize that being forward-looking is a way to improve performance and finish their shift earlier.

Finally, we look at the influence of evening rush hour and see that couriers tend to follow the TSP less often during that time (coefficient of -0.005) and that shorter distances are preferred even more (coefficient of -0.263). This suggests that couriers may be aware that the traffic jams occurring during the rush hour are not inherent components of routing algorithms, and therefore smaller relocation distances are preferred to allow for the foreseen successful delivery of this package as opposed to a hypothetical uncertain delivery of the stop recommended by the TSP.

5.3 Factors Influencing the Adherence to the TSP

We have seen that courier adherence to the TSP is undermined by the presence of alternative options that are closer to the current position. To further understand what drives a courier's decision to adhere to the TSP solution, we now investigate the determinant of tsp_actual_i, which we defined in §5.1 as the binary outcome of a courier choosing the TSP option as the next location to go to. For that, we use a logit model with the five covariates of Table 3 applied to the TSP solution, as well as the three covariates relative to route and courier features. In addition, to account for other non-TSP options that could be reasonable substitutes, we include the five covariates of the two quickest potential locations that are different from the TSP recommendation. Note that adding one alternative location is equivalent to running the MNL of §5.2 with only two options – the TSP and the quickest alternative. Adding two alternative locations on the other hand results in a model different from the MNL. The results are shown in Table 5.

As we can see, the results are completely consistent with the findings of §5.2: for the TSP recommendation, shorter distance, shorter time, and being on the edge of the convex hull increase adherence to the TSP, as discussed earlier. The impact of the number of packages remaining is negative, in line with the higher number of alternatives present in the MNL. Couriers with more experience in the company are more likely to follow the TSP recommendation, given by the positive coefficient in our logit model.

The results of Table 5 also reveal that the presence of an attractive substitute (the first or second quickest alternative location) reduces adherence. The more attractive this substitute becomes – lower time or lower distance – the lower the probability of opting for the TSP solution and hence adherence drops. This is again in line with the MNL model in §5.2, but we note that the substitution

	(1)	(2)	(3)
log_distance	-0.325***	-0.650***	-0.695***
log_time	-0.252***	-0.427***	-0.420***
is_outside	0.691***	0.607***	0.591***
is_left	-0.010	-0.009	-0.009
log_experience	0.033***	0.016***	0.013***
log_remaining	-0.570***	-0.473***	-0.440***
rush_hour	0.022**	0.023**	0.024***
isnearest_log_distance		0.499***	0.386***
isnearest_log_time		0.421***	0.256***
isnearest_is_outside		-0.075***	-0.103***
isnearest_is_left		-0.012	-0.012*
is2nearest_log_distance			0.165***
is2nearest_log_time			0.248***
is2nearest_is_outside			-0.076***
is2nearest_is_left			0.013*
Observations	354,190	354,190	354,190
Pseudo R^2	0.042	0.066	0.069
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 5: Logit estimation results

effect may not work equally across options. Specifically, the nearest alternative to the TSP has a major effect on adherence, while the second nearest has only a lower impact (coefficient $0.386 > 0.165$). This suggests that couriers have a natural tendency to choose the nearest next delivery point, which is compared to the TSP recommendation. Hence, to integrate this preference for quick transitions, an ideal routing algorithm (which would no longer solve a TSP) should penalize long stretches, since those transitions will not be adopted by the courier.

These findings are consistent with the theories discussed in Section 2: couriers aim for small victories and instant gratification by prioritizing the nearest delivery locations, and prefer to make the remaining polygon smaller by choosing the points on the edges. Furthermore, we observe (again) that courier experience positively contributes to adherence, suggesting that with more experience, couriers prefer forward-looking options more, and are willing to resist the temptation of the nearest location in exchange for an optimized solution.

5.4 Robustness Checks

To check the robustness of our findings we first run the MNL and logit models again keeping only a subset of the data of where there are no transitions under 100m and then, we run a few alternative analyses which relax some of our assumptions.

First, the results for the subset are shown in Tables 11 and 12 in the Appendix. The results of the main MNL model are qualitatively similar to the model using the full dataset. We see that some feature variables become statistically insignificant in this subsample. On the other hand, the results of the logit model stay the same as for the full dataset.

Second, we consider a distance cost matrix as opposed to a time cost matrix. That is, couriers may be optimizing a different objective, for example that of minimizing fuel consumption. For this purpose, we recompute the TSP solution with a distance-based cost input, and rerun the entire analysis. The results are reported in Tables 14 and 15 in the Appendix. The results are similar to the main model and we can observe that the preference for quick wins is maintained. The only difference with the main model is that the coefficient for is_tsp becomes negative in the extended model. This may be due to the TSP capturing the 'wrong' objective, and thus becoming irrelevant for decision-making.

Third, we recognize that the MNL structure may be overly restrictive (e.g., independence from irrelevant alternatives problem). We consider two alternative specifications. Instead of the logit specification for the adherence decision, we use a probit error structure in §5.3 (note that probit is not easy to manipulate for choice). The results, shown in Table 6, are qualitatively similar to those in Table 5. We then consider a random effect choice model. Results are included in Table 7. As we can see the coefficients of the key covariates turn out to be rather stable (low uncertainty), which suggests that the MNL adequately captures the choices made by the couriers. In the extended model we see a few differences with the extended MNL model. The results show that the more packages are left to deliver on this route, the less likely the couriers are to follow the TSP; and during rush hour, the couriers are more likely to follow the TSP.

	(1)	(2)	(3)
log_distance	-0.188***	-0.348***	-0.371***
log_time	-0.159***	-0.269***	-0.266***
is_outside	0.417***	0.365***	0.355***
is_left	-0.006	-0.005	-0.005
log_experience	0.020***	0.009***	0.007***
log_remaining	-0.349***	-0.288***	-0.268***
rush_hour	0.012**	0.012**	0.013**
isnearest_log_distance		0.273***	0.207***
isnearest_log_time		0.253***	0.152***
isnearest_is_outside		-0.044***	-0.061***
isnearest_is_left		-0.007	-0.007
is2nearest_log_distance			0.093***
is2nearest_log_time			0.154***
is2nearest_is_outside			-0.045***
is2nearest_is_left			0.008*
Observations	354,190	354,190	354,190
Pseudo R^2	0.042	0.065	0.068
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 6: Probit estimation results

6 The Impact of Driver Agency

Our analysis so far has focused on comparing what decisions were taken by the couriers, and in particular whether they adhered to the TSP recommendation. However, it remains to be seen that deviations actually had a negative impact on delivery efficiency. Preference for quick wins could perfectly have a rational basis, because it minimizes the risk that traffic conditions drastically change between transitions. Indeed, they are best at updating information as quickly as possible. In this section, we study the actual performance implication of courier choices.

	Dependent variable: chosen location	
	(1)	(2)
const	0.193***	0.204***
is_tsp	0.220***	0.301***
is_tsp × *log_remaining*		-0.0542***
is_tsp × *log_experience*		0.0029***
is_tsp × *rush_hour*		0.0042**
log_distance	-0.0342***	-0.105***
log_distance × *log_remaining*		0.0345***
log_distance × *log_experience*		-0.0009**
log_distance × *rush_hour*		-0.01***
log_time	-0.0390***	0.0159***
log_time × *log_remaining*		-0.0253***
log_time × *log_experience*		0.0003*
log_time × *rush_hour*		0.0016***
is_outside	0.0640***	0.213***
is_outside × *log_remaining*		-0.0786***
is_outside × *log_experience*		0.0022***
is_outside × *rush_hour*		0.0018*
is_left	-0.0015***	0.0138***
is_left × *log_remaining*		-0.0076***
is_left × *log_experience*		0.0005
is_left × *rush_hour*		0.0006
Number of observations (choices)	2,686,469	2,686,469
Number of variables	5	20
AIC	1,528,210	1,477,240
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 7: Random effects estimation results

To estimate the impact of driver agency on choosing the next delivery location, we thus proceed to compare the expected time of following the TSP route (at the time of making the initial recommendation) with what actually happened. We decompose reality into two different scenarios, in the spirit of Hui et al. (2009). We first examine the impact of sequence deviation, which can be assessed by measuring the expected time of the chosen route (vs. the TSP route). We measure this deviation as

$$\Delta R := \frac{predicted\_time\_actual\_route}{predicted\_time\_tsp\_route}. \tag{6}$$

Note that predicted time is computed at the time of each transition (so the total predicted time is the sum of predictions done at different moments of the day; for predictions lower than 1 minute, we set prediction equal to 1 to avoid unrealistic delivery times). In contrast, the total predicted time for the TSP route is computed at the start of the route, which means that because of prediction time variations over the day, it would be technically possible that $\Delta R < 1$.

We then measure the impact of time prediction errors, that is, the deviation between predicted travel times – those that enter as inputs to the TSP problem – and actual travel times. This is done for the actual route. This metric can be defined as

$$\Delta T := \frac{actual\_time\_actual\_route}{predicted\_time\_actual\_route}. \tag{7}$$

As a result, the total cost of the process to deliver all the items in the route can be expressed as $\Delta R \times \Delta T$. The larger this product, the longer the time it took to deliver the goods, in comparison with the 'ideal' time suggested by the TSP.

When looking at the distribution of ΔR values given in Table 8, we see that the routes chosen by the couriers have higher predicted completing times in comparison to the predicted time of the TSP in the beginning of the route. This suggests that going for the nearest next stop (as shown in §5) has a negative consequence on routing times, with a median increase of 54%, but with a right tail making times even more than 3 times larger. On the other hand, 6.2% of routes chosen by couriers are able to beat the TSP time predicted at the start time of the route. This shows that algorithm predictions in the morning are not perfect, and that it is possible to outperform the algorithm. However, this is an exception and deviations from the TSP are typically detrimental to performance.

Similarly, when looking at ΔT, we also see that couriers tend to take more time to complete the routes compared to the predicted time to complete the same route during this time of the day. The median deviation is 51%. Note that the distribution of delivery time is quite skewed (see

	Mean	Std. dev.	P1	P10	Median	P90	P99	% where ≤ 1
Overall								
ΔR	1.91	9.4	0.959	1.02	1.54	2.5	4.85	6.24
ΔT	1.98	2.55	0.455	0.863	1.51	3.16	9.06	17
Total routes: 54,390								
Experienced couriers								
ΔR	1.83	9	0.962	1.02	1.54	2.46	4.36	6.18
ΔT	1.89	2.35	0.463	0.858	1.48	2.99	7.94	17.6
Total routes: 44,698								
Inexperienced couriers								
ΔR	2.25	11	0.947	1.02	1.55	2.76	8.59	6.52
ΔT	2.42	3.29	0.428	0.895	1.72	4.09	13.8	14.2
Total routes: 9,692								
Short routes								
ΔR	1.86	12.8	0.925	1	1.29	2.17	5.25	11.5
ΔT	2.32	3.36	0.396	0.787	1.63	3.98	13.1	20.2
Total routes: 29,176								
Long routes								
ΔR	1.97	1.09	1.09	1.31	1.8	2.73	4.61	0.178
ΔT	1.59	0.806	0.619	0.945	1.44	2.4	3.95	13.4
Total routes: 25,214								

Table 8: Key summary statistics of route performance (P1, P10, P90 and P99 denote the 1%, 10%, 90% and 99% percentiles respectively). Experience level and route length are based on a median split.

Figure 3). It turns out that long transitions have a much smaller deviation, as shown in Table 10, while short transitions have a much larger deviation. Because of this, routes in which many short transitions appear result in a larger ΔT. In contrast, routes with no transitions of less than 100 meters (a subsample of about 30% of routes) lead to a median ΔT equal to 1.37, as shown in Table 13. Overall, we nevertheless observe that the actual route is consistently longer than the predicted one most of the time. This implies that the decision to deviate from the TSP cannot be explained by an effective reaction to adverse traffic conditions, in which case $\Delta R > 1$ would be associated with $\Delta T < 1$. In fact, we find no correlation between ΔR and ΔT.

To further compare the routes, we can separate the tours into four more categories – routes underwent by experienced and inexperienced couriers, and short vs long routes. As the median value for the previous days worked in the company is 11, we will set this as the threshold value, and define all couriers who have worked for 11 days or less as inexperienced, and all couriers who have worked for more than 11 days as experienced. We follow the same procedure for route length with 9 stops as the median value.

There is a notable difference in the distributions of ΔR and ΔT values between experienced and inexperienced drivers, with the values generally being higher for the couriers that have worked less days in the company. We observe higher values of both ΔR and ΔT (with the exception of the first percentile) for agents with less total deliveries. This suggests that experienced drivers either choose better routes or are able to complete them in less time than the inexperienced couriers.

We then proceed to separate the routes based on their length. We immediately notice that shorter routes experience lower deviation ΔR (mean 1.86 vs. 1.97 and median 1.29 vs. 1.8). There is also a considerable proportion where the route chosen is in theory at least as fast as the recommended route (11.5% of cases with $\Delta R \leq 1$), while for long routes this percentage value is insignificant (0.18%). On the other hand, ΔT may be larger for shorter routes (mean 2.32 vs. 1.59 and median 1.63 vs. 1.44), although the percentage of cases in which $\Delta T \leq 1$ is larger too (20.2% vs. 13.4%). This suggests that there is more variability in ΔT when the route is short, which is intuitive, given that we have fewer stops over which to aggregate prediction errors. Overall, the effect of ΔR dominates, and shorter routes result in a lower cost of deviation $\Delta R \times \Delta T$.

7 Conclusion

In this article, we have empirically estimated the decisions of last-mile delivery couriers and show that they systematically deviate from the TSP solution. We uncover the multiple factors that lead them to make suboptimal decisions when deciding which delivery location to visit next. We find that they opt for myopic choices that are nearby, as opposed to forward-looking sequences that reduce overall travel times. This is consistent with previous literature about human compliance to algorithm recommendations and specifically the preference for quick wins. Furthermore, we show that such deviations from the TSP are harmful with respect to performance; we observe deviations both on route sequence and actual delivery times, with negative effects on both.

In light of these findings, what are the steps that delivery companies could take to ensure better performance and adherence from their couriers to routing recommendations? Since non-compliance increases operational costs, it seems necessary to integrate our insights in their navigation suggestions. In contrast with Arıkan et al. (2023) who impose constraints on the feasible routes, our analysis suggests that the human behavior of couriers makes them prefer certain type of structures. To modify such behavior, one ideal strategy would be to selectively show, in every stage, only a subset of next options to the couriers: instead of displaying all available undelivered packages and the full route, the couriers could only be presented with two or three top choices for each stop rank. The delivery addresses of all other packages would remain hidden and not appear on the map displayed to the couriers in this stop rank. On one hand, that would enable them to choose between only the optimal and one or two sub-optimal options, therefore eliminating the possibility that an odd, calamitous decision would be made. On the other hand, displaying a few other locations and not just the optimal one, would enable the couriers to still feel that they have some control and that they can make a decision autonomously, thereby avoiding algorithm aversion and ensuring high usage of routing recommendations. Obviously, these conjectures should be validated experimentally. Such experiments would be an ideal complement to our analysis, which relies on observational data.

Our study also opens new research questions. First, we have investigated drivers of routing decisions focusing on courier and route characteristics. It would be interesting to additionally understand better the role of customer or package features in routing choices, which we do not observe in our data. Are there special items (e.g., more valuable or heavier to transport) or special consumer segments (e.g., VIP customers) that would increase choice probabilities? Or are choices purely driven by geography? Second, can we increase adherence by nudging couriers with appealing interfaces, like the above suggestion of partial information sharing with the courier? We have found that shorter physical distances are preferred. Could the algorithm modify the perception of distance, akin to psychological distance (Trope and Liberman 2010), so that couriers are more inclined to opt for the TSP solution? Finally, the role of unobservable information – and specifically traffic conditions that are observed by the courier but not necessarily by the algorithm at the time of the computation of the TSP – seems to be one major source of deviations. It would be useful to experimentally test how private information affects adherence, and how to embed it into algorithm design.

References

Adams GS, Converse BA, Hales AH, Klotz LE (2021) People systematically overlook subtractive changes. *Nature* 592(7853):258–261.

Angelopoulos S, Bendoly E, Fransoo JC, Hoberg K, Ou C, Tenhiälä A (2023) Digital transformation in operations management: Fundamental change through agency reversal. *Journal of Operations Management* 69(6):876–889.

Aouad A, Deshmane A, Martínez-de Albéniz V (2022) Designing layouts for sequential experiences: Application to cultural institutions. *Management science* Forthcoming.

Aouad A, Désir A (2022) Representing random utility choice models with neural networks, arXiv preprint arXiv:2207.12877.

Arıkan U, Kranz T, Sal BC, Schmitt S, Witt J (2023) Human-centric parcel delivery at deutsche post with operations research and machine learning. *INFORMS Journal on Applied Analytics* 53(5):359–371.

Arora S, Choudhary V, Kireyev P (2024) Don't fake it if you can't make it: Driver misconduct in last-mile delivery. *Management Science* Forthcoming.

Bai B, Dai H, Zhang DJ, Zhang F, Hu H (2022) The impacts of algorithmic work assignment on fairness perceptions and productivity: Evidence from field experiments. *Manufacturing & Service Operations Management* 24(6):3060–3078.

Balakrishnan M, Ferreira K, Tong J (2022) Improving human-algorithm collaboration: Causes and mitigation of over-and under-adherence. *Available at SSRN 4298669* .

Bastani H, Bastani O, Sinchaisri WP (2021) Improving human decision-making with machine learning, arXiv preprint arXiv:2108.08454.

Bavafa H, Jónasson JO (2024) The distributional impact of fatigue on performance. *Management Science* 70(5):3319–3337.

Bertsimas D, Weismantel R (2005) *Optimization over integers* (Belmont, MA: Dynamic Ideas).

Bhattacharjee D, Kamil A, Lukasiewicz M, Melnikov L (2024) Digitizing mid- and last-mile logistics handovers to reduce waste. `https://www.mckinsey.com/industries/travel-logistics-and-infrastructure/our-insights/digitizing-` Last accessed on 2024-11-26.

Boysen N, Fedtke S, Schwerdfeger S (2021) Last-mile delivery concepts: a survey from an operational research perspective. *Or Spectrum* 43(1):1–58.

Bray RL (2023) Operational transparency: Showing when work gets done. *Manufacturing & Service Operations Management* 25(3):812–826.

Cameron LD (2022) "Making out" while driving: Relational and efficiency games in the gig economy. *Organization Science* 33(1):231–252.

Caro F, de Tejada Cuenca AS (2023) Believing in analytics: Managers' adherence to price recommendations from a dss. *Manufacturing & Service Operations Management* 25(2):524–542.

Converse BA, Tsang S, Hennecke M (2023) The value of mere completion. *Journal of Experimental Psychology: General* .

Cook W, Held S, Helsgaun K (2024) Constrained local search for last-mile routing. *Transportation Science* 58(1):12–26.

Cui R, Li M, Li Q (2020) Value of high-quality logistics: Evidence from a clash between sf express and alibaba. *Management Science* 66(9):3879–3902.

Dell'Acqua F, McFowland III E, Mollick ER, Lifshitz-Assaf H, Kellogg K, Rajendran S, Krayer L, Candelon F, Lakhani KR (2023) Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality, harvard Business School Technology & Operations Mgt. Unit Working Paper 24-013.

Dieter P, Caron M, Schryen G (2023) Integrating driver behavior into last-mile delivery routing: Combining machine learning and optimization in a hybrid decision support framework. *European Journal of Operational Research* 311(1):283–300.

Dietvorst BJ, Simmons JP, Massey C (2015) Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of experimental psychology: General* 144(1):114.

Dietvorst BJ, Simmons JP, Massey C (2018) Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management science* 64(3):1155–1170.

Duan N (1983) Smearing estimate: a nonparametric retransformation method. *Journal of the American Statistical Association* 78(383):605–610.

Elmaghraby W, Jank W, Zhang S, Karaesmen IZ (2015) Sales force behavior, pricing information, and pricing decisions. *Manufacturing & Service Operations Management* 17(4):495–510.

Eurosender (2024) Parcel delivery statistics in the eu (2024). `https://www.eurosender.com/en/resources/courier-par` Last accessed on 2024-11-26.

Grand-Clément J, Pauphilet J (2024) The best decisions are not the best advice: Making adherence-aware recommendations. *Management Science* Forthcoming.

Hibbeln M, Jenkins JL, Schneider C, Valacich JS, Weinmann M (2017) How is your user feeling? inferring emotion through human–computer interaction devices. *MIS Quarterly* 41(1):1–22.

Hu Y, Zhang Z, Yao Y, Huyan X, Zhou X, Lee WS (2021) A bidirectional graph neural network for traveling salesman problems on arbitrary symmetric graphs. *Engineering Applications of Artificial Intelligence* 97:104061.

Hui SK, Fader PS, Bradlow ET (2009) Research note – the traveling salesman goes shopping: The systematic deviations of grocery paths from tsp optimality. *Marketing science* 28(3):566–572.

Ibanez MR, Clark JR, Huckman RS, Staats BR (2018) Discretionary task ordering: Queue management in radiological services. *Management Science* 64(9):4389–4407.

Ibanez MR, Toffel MW (2020) How scheduling can bias quality assessment: Evidence from food-safety inspections. *Management science* 66(6):2396–2416.

Joshi CK, Laurent T, Bresson X (2019) An efficient graph convolutional network technique for the travelling salesman problem, arXiv preprint arXiv:1906.01227.

Käki A, Kemppainen K, Liesiö J (2019) What to do when decision-makers deviate from model recommendations? empirical evidence from hydropower industry. *European Journal of Operational Research* 278(3):869–882.

Kawaguchi K (2021) When will workers follow an algorithm? a field experiment with a retail business. *Management Science* 67(3):1670–1695.

Kesavan S, Kushwaha T (2020) Field experiment on the profit implications of merchants' discretionary power to override data-driven decision-making tools. *Management Science* 66(11):5182–5190.

Knight B, Mitrofanov D, Netessine S (2022) The impact of ai technology on the productivity of gig economy workers, available at SSRN 4372368.

Kremer M, Moritz B, Siemsen E (2011) Demand forecasting behavior: System neglect and change detection. *Management Science* 57(10):1827–1843.

Li Y, Phillips W (2018) Learning from route plan deviation in last-mile delivery, technical report, MIT.

Li Z, Chen Q, Koltun V (2018) Combinatorial optimization with graph convolutional networks and guided tree search. *Advances in neural information processing systems* 31.

Lu Z, Cui R, Sun T, Wu L (2023) The value of last-mile delivery in online retail, available at SSRN.

Ludwig TD, Geller ES (2000) Intervening to improve the safety of delivery drivers: A systematic behavioral approach. *Journal of Organizational Behavior Management* 19(4):1–124.

Luxen D, Vetter C (2011) Real-time routing with openstreetmap data. *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 513–516, GIS '11 (New York, NY, USA: ACM), ISBN 978-1-4503-1031-4, URL `http://dx.doi.org/10.1145/2093973.2094062`.

Mao W, Ming L, Rong Y, Tang CS, Zheng H (2019) Faster deliveries and smarter order assignments for an on-demand meal delivery platform, available at SSRN 3469015.

Markakis M, Martínez-de Albéniz V (2023) Improving last-mile productivity at paack. IESE Business School case study OIT-17-E.

Merchán D, Arora J, Pachon J, Konduri K, Winkenbach M, Parks S, Noszek J (2024) 2021 amazon last mile routing research challenge: Data set. *Transportation Science* 58(1):8–11.

Moon K (2023) Strategic path selection in service networks: Leveraging machine learning to estimate combinatorially complex preferences for consumption and waiting, available at SSRN 3819117.

Nazari M, Oroojlooy A, Snyder L, Takác M (2018) Reinforcement learning for solving the vehicle routing problem. *Advances in neural information processing systems* 31.

Pape T, Kavadias S, Sommer SC (2020) Decision bias in project selection: Experimental evidence from the knapsack problem, available at SSRN 3448676.

Rosenbush S, Stevens L (2015) At ups, the algorithm is the driver. *The Wall Street Journal* 16.

Rutkowsky S, Morwind R, Mueller T (2024) Looking up the last mile. `https://www.es.kearney.com/en/transportation-travel/article/-/insights/locking-up-the-last-mile`, Last accessed on 2024-11-21.

Shaikley L (2024) Everything is logistics, especially last-mile delivery. https://www.forbes.com/councils/forbestechcouncil/2024/03/22/everything-is-logistics-especially- Last accessed on 2024-11-21.

Sun J, Zhang DJ, Hu H, Van Mieghem JA (2022) Predicting human discretion to adjust algorithmic prescription: A large-scale field experiment in warehouse operations. *Management Science* 68(2):846–865.

Train KE (2009) *Discrete choice methods with simulation* (Cambridge university press).

Trope Y, Liberman N (2010) Construal-level theory of psychological distance. *Psychological review* 117(2):440.

Van Donselaar KH, Gaur V, Van Woensel T, Broekmeulen RA, Fransoo JC (2010) Ordering behavior in retail stores and implications for automated replenishment. *Management Science* 56(5):766–784.

Weinmann M, Valacich J, Schneider C, Jenkins JL, Hibbeln MT (2022) The path of the righteous: Using trace data to understand fraud decisions in real time. *MIS Quarterly* 46:2317–2336.

Wu L, Wen H, Hu H, Mao X, Xia Y, Shan E, Zhen J, Lou J, Liang Y, Yang L, et al. (2023) Lade: The first comprehensive last-mile delivery dataset from industry, arXiv preprint arXiv:2306.10675.

Zhan Y, Zhang Z, Ge J (2023) Value of exclusive doorstep delivery in the last-100-meter distribution, available at SSRN 4352455.

Zhang R, Prokhorchuk A, Dauwels J (2020) Deep reinforcement learning for traveling salesman problem with time windows and rejections. *2020 International Joint Conference on Neural Networks (IJCNN)*, 1–8 (IEEE).

Appendix

Figure 7: Average compliance based on the number of packages left to deliver between 5am and 10am, and 10am and 7pm.

Figure 8: Average compliance based on the number of packages left to deliver between 10am and 7pm separated based on courier experience. Experience level are given as 'low' and 'high'. The median value is incorporated to the high category.

packages left	time	adherence	observations	packages left	time	adherence	observations
7	9	38.8	2078	4	9	42.6	1565
	10	37.4	4392		10	43	3600
	11	37.7	4449		11	43.3	5596
	12	34.3	2725		12	44.3	4630
	13	33.6	1629		13	41.9	2786
	14	31.5	1364		14	41.7	2047
	15	32.8	2101		15	42.6	2852
	16	34.1	4073		16	41.4	4536
	17	33	4102		17	42.6	5573
	18	35	2536		18	44.5	4795
					19	43.4	2743
					20	42.2	1526
6	10	39.4	4204	3	9	51.9	1408
	11	36.9	4869		10	48.7	3415
	12	36.5	3364		11	50.3	5528
	13	37	1945		12	49	5355
	14	33.5	1546		13	48.5	3320
	15	33.6	2367		14	51.3	2445
	16	38	4276		15	50.2	3187
	17	37	4690		16	48.9	4593
	18	38.1	3148		17	50.4	5837
					18	51.8	5783
					19	51.6	3582
					20	49.5	2039
					21	47.9	1082
5	11	39.8	5353	2	9	59.9	1239
	12	40.5	3924		10	59	3238
	13	38.9	2370		11	61.3	5586
	14	37.2	1765		12	61.8	6018
	15	37.1	2625		13	63	4094
	16	37.1	4371		14	62.7	2958
	17	40.2	5135		15	62.7	3507
	18	40.3	3915		16	61.5	4605
	19	40	2174		17	63.3	6098
	20	39.1	1091		18	66.3	6960
					19	65.6	4554
					20	64.6	2762
					21	62.7	1605

Table 9: Adherence based on time. Rows with at least 1000 observations.

Distances	Mean	Std. dev	P1	P10	Median	P90	P99	% where ≤ 1	Mean time (min)
all	8.76×10^8	1.66×10^{10}	0.223	0.528	1.34	8.81	713	36	18.5
< 0.1	4.87×10^9	3.88×10^{10}	0.608	1.36	7.34	156	1.1×10^{11}	4.67	12.4
$0.1 \leq x < 0.2$	3.47	6.34	0.459	0.807	1.81	6.91	28.8	19.3	12.1
$0.2 \leq x < 0.5$	2.27	3.76	0.338	0.634	1.29	4.43	16.1	33.5	14.5
$0.5 \leq x < 1$	1.77	2.36	0.248	0.501	1.08	3.66	10.9	45.2	17.7
$1 \leq x < 2$	1.6	1.85	0.203	0.442	1.01	3.4	8.9	49.7	21.6
$2 \leq x < 5$	1.6	1.85	0.203	0.442	1.01	3.4	8.9	49.7	21.6
$5 \leq x < 10$	1.2	1.08	0.107	0.3	0.881	2.48	5.34	57.1	30.2
≥ 10	1.14	0.997	0.0836	0.267	0.889	2.25	5.07	57.6	33.2

Table 10: ΔT on a transition-level depending on the distance. Note that due to the presence of zero values, predictions of small times have been converted to 1 minute.

	Dependent variable: chosen location	
	(1)	(2)
is_tsp	0.394***	-0.075
$is\_tsp \times log\_remaining$		0.201***
$is\_tsp \times log\_experience$		0.036***
$is\_tsp \times rush\_hour$		-0.071**
$log\_distance$	-1.562***	-0.326**
$log\_distance \times log\_remaining$		-0.819***
$log\_distance \times log\_experience$		0.106***
$log\_distance \times rush\_hour$		-0.372***
log_time	-0.207***	0.276**
$log\_time \times log\_remaining$		-0.098**
$log\_time \times log\_experience$		-0.073***
$log\_time \times rush\_hour$		-0.043
$is\_outside$	0.506***	0.275**
$is\_outside \times log\_remaining$		0.108**
$is\_outside \times log\_experience$		0.001
$is\_outside \times rush\_hour$		-0.013
is_left	-0.036***	0.108
$is\_left \times log\_remaining$		-0.027
$is\_left \times log\_experience$		-0.025**
$is\_left \times rush\_hour$		0.004
Number of observations (choices)	42,082	42,082
Number of variables	5	20
AIC	118,322	117,651
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 11: MNL estimation results for routes with no actual transitions with distances $<$ 100m.

	(1)	(2)	(3)
log_distance	-0.345***	-0.888***	-0.987***
log_time	-0.199***	-0.386***	-0.354***
is_outside	0.842***	0.801***	0.783***
is_left	-0.038*	-0.037*	-0.032
log_experience	0.042***	0.043***	0.041***
log_remaining	-0.446***	-0.396***	-0.359***
rush_hour	0.025	0.040	0.035
isnearest_log_distance		0.829***	0.643***
isnearest_log_time		0.462***	0.272***
isnearest_is_outside		-0.194***	-0.213***
isnearest_is_left		-0.002	-0.001
is2nearest_log_distance			0.338***
is2nearest_log_time			0.222***
is2nearest_is_outside			-0.111***
is2nearest_is_left			0.039*
Observations	42,082	42,082	42,082
Pseudo R^2	0.033	0.073	0.078
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 12: Logit estimation results for routes with no actual transitions with distances $<$ 100m.

	Mean	Std. dev.	P1	P10	Median	P90	P99	% where ≤ 1
Overall								
ΔR	1.37	0.481	0.926	1	1.24	1.9	3.02	16.7
ΔT	1.71	1.23	0.374	0.707	1.37	3.07	6.51	27.3
Total routes: 14,936								
Experienced couriers								
ΔR	1.37	0.476	0.93	1	1.24	1.9	2.99	16.3
ΔT	1.67	1.18	0.38	0.708	1.34	2.97	6.2	28
Total routes: 12,658								
Inexperienced couriers								
ΔR	1.37	0.511	0.917	1	1.21	1.95	3.24	18.6
ΔT	1.94	1.44	0.336	0.706	1.55	3.6	7.48	23.2
Total routes: 2,278								
Short routes								
ΔR	1.33	0.469	0.92	1	1.19	1.84	2.94	19
ΔT	1.76	1.29	0.362	0.685	1.4	3.21	6.75	28
Total routes: 13,059								
Long routes								
ΔR	1.64	0.479	1.03	1.19	1.52	2.18	3.47	0.533
ΔT	1.39	0.534	0.608	0.84	1.28	2.09	3.28	22.8
Total routes: 1,877								

Table 13: Key summary statistics of route performance for a subset containing no actual transitions $<$ 100m in distance. (P1, P10, P90 and P99 denote the 1%, 10%, 90% and 99% percentiles respectively). Experience level and route length are based on a median split.

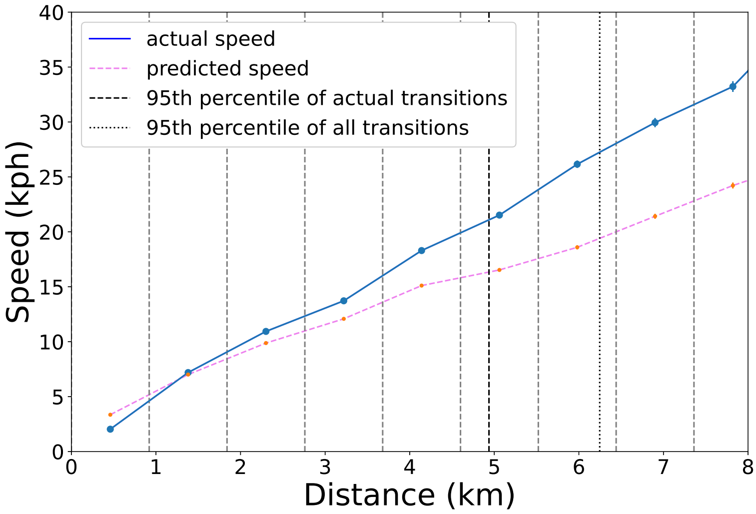

Figure 9: The dependence of predicted speed on the distance of this transition.

	Dependent variable: chosen location	
	(1)	(2)
is_tsp	0.662***	-0.056**
$is\_tsp \times log\_remaining$		0.326***
$is\_tsp \times log\_experience$		0.024***
$is\_tsp \times rush\_hour$		-0.011
$log\_distance$	-1.474***	-0.131***
$log\_distance \times log\_remaining$		-0.720***
$log\_distance \times log\_experience$		0.072***
$log\_distance \times rush\_hour$		-0.209***
$is\_outside$	0.359***	0.323***
$is\_outside \times log\_remaining$		-0.019
$is\_outside \times log\_experience$		0.020***
$is\_outside \times rush\_hour$		-0.035***
is_left	-0.011***	-0.020
$is\_left \times log\_remaining$		-0.002
$is\_left \times log\_experience$		0.004
$is\_left \times rush\_hour$		0.001
Number of observations (choices)	319,770	319,770
Number of variables	4	16
AIC	1,059,958	1,053,801
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 14: MNL estimation results for a distance cost matrix.

	(1)	(2)	(3)
log_distance	-0.588***	-1.651***	-1.668***
is_outside	0.655***	0.573***	0.566***
is_left	-0.011	-0.006	-0.005
log_experience	0.029***	0.018***	0.016***
log_remaining	-0.516***	-0.413***	-0.405***
rush_hour	0.090***	0.078***	0.072***
isnearest_log_distance		1.331***	1.052***
isnearest_is_outside		-0.093***	-0.117***
isnearest_is_left		-0.010	-0.011
is2nearest_log_distance			0.322***
is2nearest_is_outside			-0.062***
is2nearest_is_left			0.010
Observations	319,770	319,770	319,770
Pseudo R^2	0.034	0.061	0.062
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 15: Logit estimation results for distance-based cost matrix.