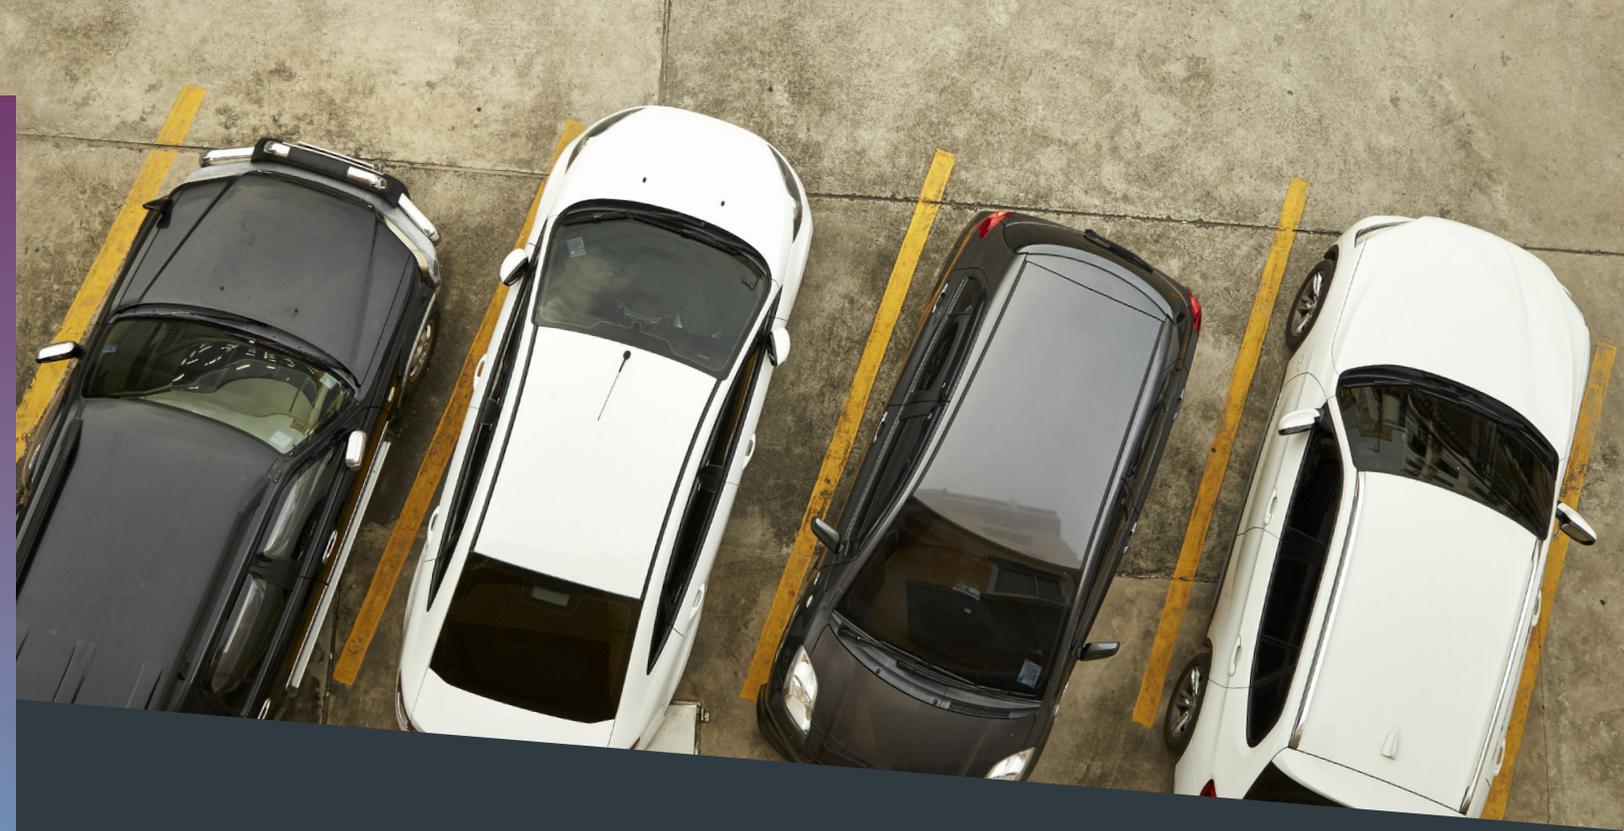


Optimizing Supply and Demand Balance Through Econometrics

By Shawn Higbee MBA PhD

RIDECCELL 

June 24, 2020



The classic problem of inventory balancing is as ubiquitous as your most basic supermarket item. Think about a single brand of cereal in the breakfast aisle.

Consider, how many boxes should be here?

The answer is certainly not two... because after just a few cereal-minded customers walk by, the store will be out of inventory and lose sales. The answer isn't 200 either. Too many boxes and inventory will end up sitting a long time, taking up space, tying up capital, and potentially expiring on the shelf.

To expand this example to high-ticket items — thus making the problem more extreme — pretend each of those boxes of cereal costs \$30,000 to stock. Now, what happens if you get the wrong number on the shelf?

That's precisely the situation carsharing fleet operators face. The purchase of transportation from a car sharing fleet is a low-density event, not entirely unlike buying cereal. Lots of people do it, it's not a very high-dollar transaction, and most travel only takes up a small fraction of a day. The difference is the cost of the capital asset is tremendously high in comparison to the revenue per transaction. This discrepancy makes it critical for carsharing fleets to optimize their inventory if they're ever going to be profitable.

Measuring Missteps

There are many ways to measure the solutions to the inventory balance problem, but not all are created equal. Let's start with the definition of demand: the number of buyers who are willing and able to buy a good or service at a given price. In the modern digital economy, there's a rich set of indicators for demand — from the number of times an app is opened to the number of times a purchase is made.

Some use app-opens as a demand measurement, thinking opens translate into an intent to use a service. This is a naïve approach because app opening doesn't account for many other factors that influence demand, such as the availability of alternatives, or whether an event is simply an effort to simply learn about a service. Additionally, it doesn't include the time scale in which the purchasing decision is made. In short, demand is a VERY complicated mechanism, and we have limited insight into this complex web of factors. We can't effectively extrapolate the metric of app opens into a full view of customer behavior.

In other cases, simulation-based methods are used to hypothesize potential demand as part of a truncated or partially observed distribution. In practice these predictions seldom become reality. This is because it's difficult — sometimes even impossible — to parameterize the actual environmental factors that drive car demand. Even in the most optimistic case where we assume that someone creates a simulation of demand, and it includes the realistic contributors, it begs the question of validation and why one should believe an estimate of market size? Then if you believe it, how do you know what you should do about it?

The other obvious problem is these approaches fail to take into account the impact of competition in the market. Any estimate of demand — whether a leading or lagging indicator — needs to account for the complete picture of supply, demand, competition, convenience, and other non-price factors. Even if you have an estimate of how big demand might be, that estimate becomes inactionable in the presence of competition and how other actors participate in a local market.

The Problem Behind the Numbers

The larger issue is that focusing solely on the demand-side of the problem doesn't solve the actual economic problem fleet operators face: "I already have a fairly fixed number of cars and a considerable amount of capital tied up in them. How do I orchestrate them so I can get the highest return possible for the fleet as a whole?"

While it's attractive to speculate how much demand exists in a market — or to simulate what demand might exist — that speculation is not actionable. Worse yet, the psychology of it can cause companies to adapt their actions based on inflated numbers, and create the costly situation of over-supply.

There's a clear problem with chasing an overly optimistic estimate of demand. Fleet operators can't risk over-supplying a market and tying up even a single \$30k asset. An asset that could be earning more revenue somewhere else. This represents a tremendous inventory cost burden, and, for reasons we'll explore later, it can actually be a business risk.

Eliminating Speculation

Among the rich set of mathematical options used to estimate demand, the one that is based on the least amount of speculation is measuring market clearing at a specific time and place. This phenomenon of market clearing has a rigorous grounding in economics and enables measurement and comparison of a market, as well as optimizing the operational strategy within a market.

It's this approach that answers the actionable question:

- › How can I maximize overall revenue for my fleet by optimizing placement for the "n+1" marginal car?

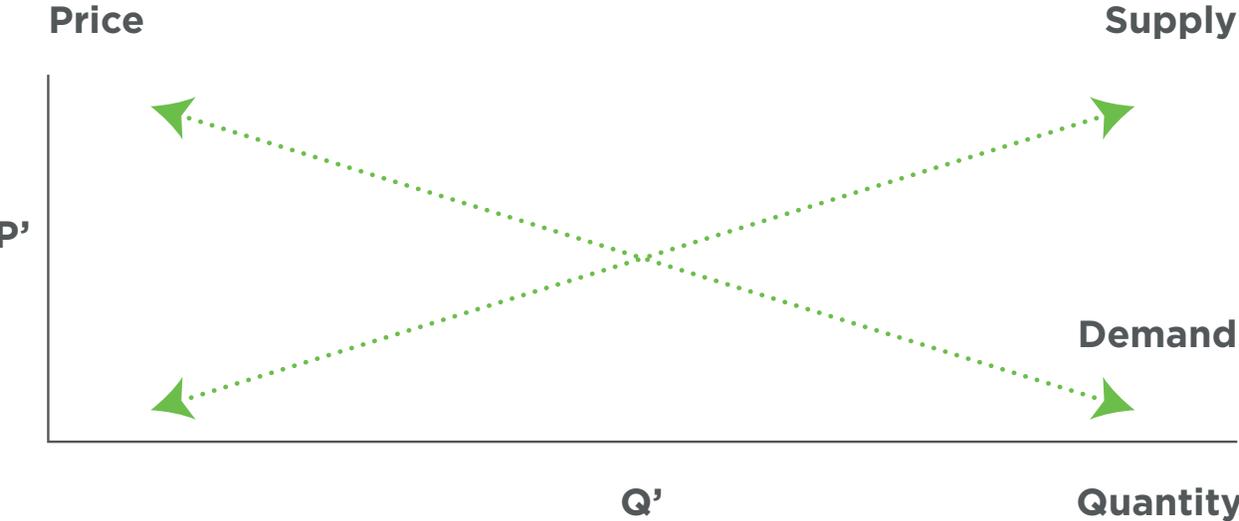
Let's Get Technical

The solve to this challenge is built into the Ridecell platform. But to help you understand the platform, we will need to introduce a few technical concepts. We'll start with a few pieces of economics and game theory that will allow us to take this problem apart, so we can arrive at an optimal solution for what a fleet operator should do to optimize their fleet strategy through econometrics.

1. A DOWNWARD-SLOPING DEMAND CURVE

In microeconomics, the price/quantity relationship in a market is captured by plotting the number of units a consumer can and will buy at a given price. A downward-sloping demand curve implies that all else equal a consumer will buy fewer units of a product at a higher price than they would at a lower price. This assumes the products involved are commodities. Functionally, they're (nearly) identical, consumers have good information about the product, there's relatively low friction in purchasing one product versus another, etc. Conversely, the supply curve slopes upward, which implies that when prices are higher, suppliers of a good have an incentive to produce more units of that good.

ECON THEORY: SUPPLY/DEMAND RELATIONSHIPS



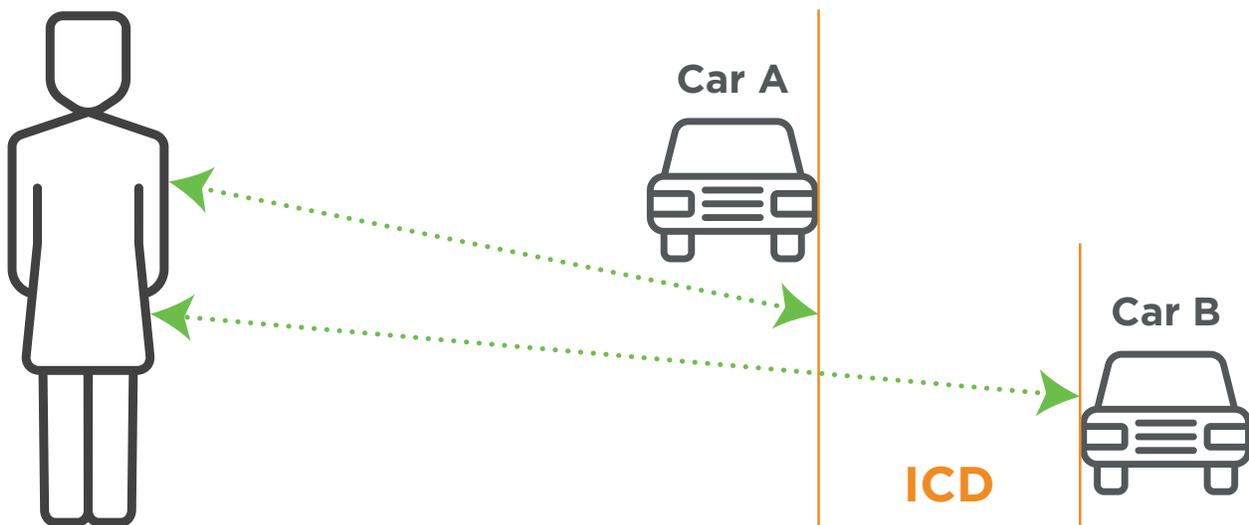
The idea is that for a market where price and quantity are allowed to adjust to economic conditions, that the market will settle on an equilibrium point where the actual price of goods sold is p' and the quantity sold at that price is q' .

All of this is generally true in fleet-managed car sharing. Operators are all trying to standardize the car-to-car experience, make renting easy, and bring transparency to the value proposition. Additionally, operators want to reach a sustainable and steady state market condition where they are making available the right amount of transportation that the market wants to consume, and at a price that serves both the consumer and the operator.

2. INVERSE-DISTANCE PREFERENCE ORDER

Because the cars are ostensibly the same — equally clean, equally fueled, equally desirable to a consumer — a potential customer's primary driver for choosing which specific car is based on location. Game theory is the discipline of predicting the actions based on the preference order of the alternatives available to them. Here game theory tells us any given consumer who sees two cars at the same price, during the same moment (the moment they need a car) and they will tend to choose the car that's closest.

INVERSE DISTANCE DEMAND MODIFICATION



ICD - “Inter-car distance” is the distance between the car that **was** rented and the next closest one that was not.

HOW TO THINK ABOUT DEMAND

On the demand side of the equation, these are called short-run demand curves where the decision to purchase or not happens in under 90 seconds. Either there's a car within the distance the consumer is willing to walk, or there isn't. The market "clears" in that 90 seconds or it doesn't and the asset had to be prepositioned in order to participate.

So let's revisit the causal chain of events that leads to a rental transaction and just like we took apart the idea of app opens, we can examine the role car placement plays in the causation of demand. Just like an app open and app downloads are prerequisites for a rental, we need to carefully think about them as signals for demand and ask why they are such weak signals of demand? If you compare the number of app downloads to the number of interactions a single parked fleet-owned car will have with the public, you'll see a huge discrepancy.

In a large city, a few thousand people will drive or walk by every single car, every day, but most operators don't have hundreds of thousands of daily app downloads. Let alone thousands of app downloads for every single car they have on the street. What this means is that it's a relatively rare event for a user to go from downloading an app to registering their account to renting a car.

It's even more rare for a user to go from seeing a car to downloading an app to registering their account to renting a car in a single go. We can see that this mechanism is weak through empirical observation, and that all else equal, the supply of cars on the street has very little to do with generating demand, or even signaling it. To take the example one step further, the "demand" for transportation existed even before the app was downloaded, so app downloads and app opens are certainly signals, they are definitely not "demand." While there is obviously some minimum presence required for a fleet to have name recognition and be on the top of mind for a customer, the simple demonstration above shows that the presence of the car itself is not generating demand and the observable events in the customer funnel are at best signals of the underlying need for transportation.

The causal chain that is more consistent with the data, is the process of a marketing campaign is what creates brand awareness, then the placement of cars that can be searched through an in-app experience, which inform a purchase decision amongst alternatives, then the conversion steps that lead to a purchase. Certainly there are some people that register and convert the entire acquisition funnel: from app download, registration, to rental on the curbside, but the consumer journey is typically longer. Especially for fleets with very subtle branding, it can be difficult to tell from a quick glance if the car is actually even part of a fleet, so the physical curbside presence of those cars has even less ability to generate demand.

In summary, if cars were effective at generating their own demand, fleets would immediately cease all other forms of advertising outside of online app stores, much the way that smartphone games are marketed. Because we know that the purchase of transportation has a lot more to do with other factors, we need to be careful with how actionable we think estimates or indicators of demand really are.

3. SUPPLY CONTROL

The supply side of mobility — at least until autonomous driving is ubiquitous — is generally solved on a 24/7 timescale. Operators can only make so many moves during this time. There may be parking restrictions. Or they may be legally required to maintain a minimum number of cars in or out of an area.

Operators also have to factor in the organic re-supply of assets caused by customers ending their rentals and dropping off their cars. Where these cars are dropped off has a lot to do with the journey endpoint and convenience factors such as parking, local shopping availability, etc.

Many journeys tend to be out-and-back, so there's a reasonable chance the car will be returned near its original pickup point. However, if a user does take a one-way trip — or if a different user happens to grab the car while it's parked — there's little to be done to predict or control where the car lands.

The main factor an operator can control in this complex game of supply and demand is how many cars should be in a given area to begin with.

4. PERISHABLE ASSET PRICING AS IT APPLIES TO CAR SHARING

The upside risk and reward of putting a car where the user wants it is the attractive part of the problem. But, to make fleets profitable, we have to consider the downside risk: cars are often perishable assets. Much like hotel rooms, the value a car offers is lost when it's not rented. You can never again sell the hours the car has gone unused.

Unlike hotel rooms, the downside risk of carsharing is much more impactful than lost opportunity costs. Carsharing is beholden to city rules that govern how long a car can sit in any one place. This means any sub-optimal move or over-supply of a local market could compel you to *spend* money to take an otherwise perfectly good car and move it or run the risk of being penalized.

This problem exists in *both* high-demand areas and low-demand areas. It's independent of the actual amount of demand in a market. Layer in the fact that even the highest demand areas can be oversaturated with cars, either your own or by the competition and the question changes. It isn't "is this a high-demand area?" but rather "weighing all the factors, is this the right amount and density of cars for this area?"

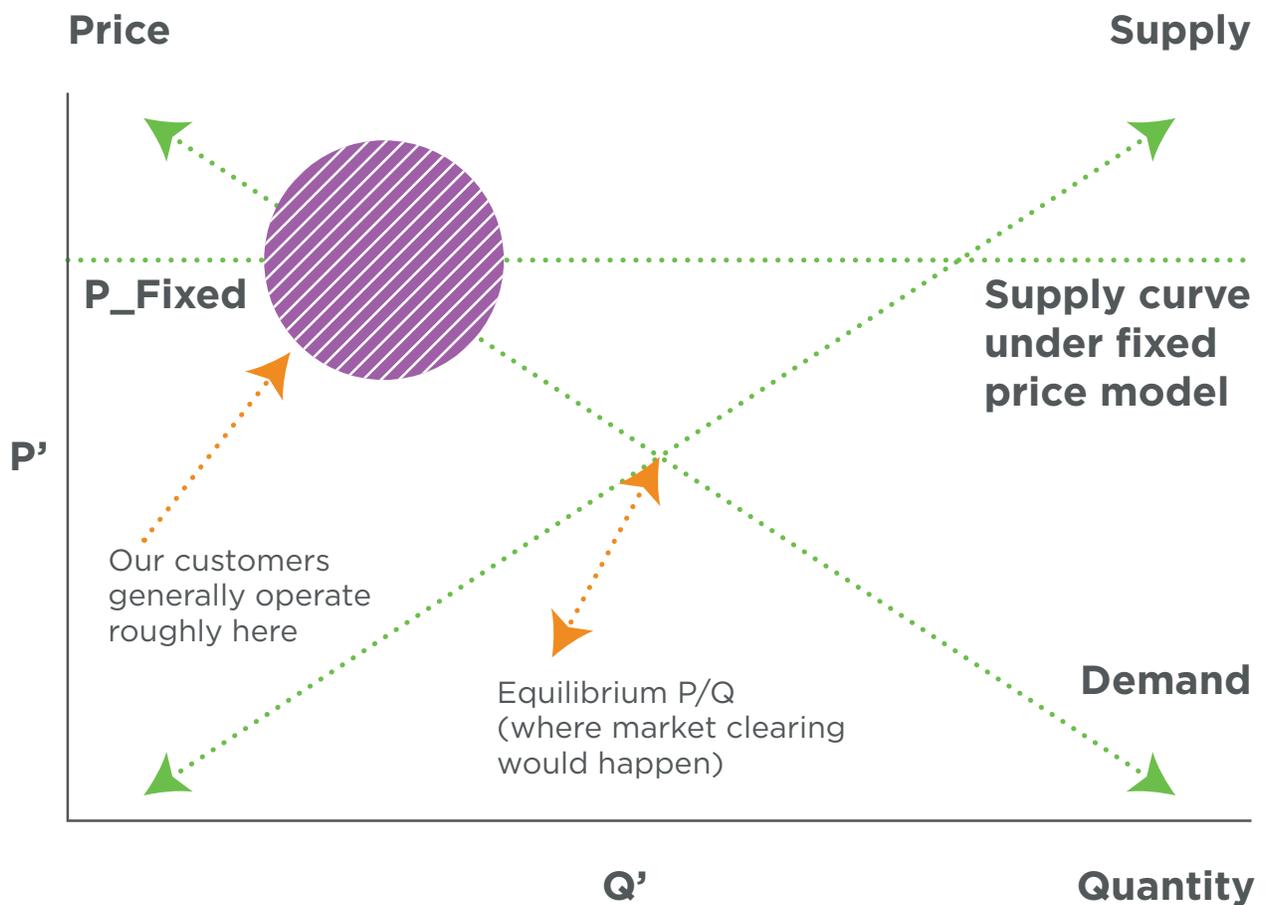
This is why it's so important for operators to have access to tools that measure the optimum quantity of cars for any given spot.

5. PRICING:

The last piece of the economic puzzle is pricing. Economic theory for commodities — let's use wheat as an example — dictates that price is determined by the balance between supply and demand. The price gets lowered until consumers choose to buy wheat over the alternatives.

If there is wheat leftover, the price continues to fall until it sells out completely, at what is known as the market-clearing price. Consider what would happen if wheat farmers set an arbitrary price without considering the price of substitute products like oats, corn, rice or barley. Some wheat would still sell at that price, but there would likely be an unsold surplus of wheat at the end of the year.

SUPPLY/DEMAND IN FIXED PRICE CAR SHARING



The same phenomena occurs in carsharing. Some fleet operators offer demand-based pricing. The best example of this is Uber's well-known surge pricing model. This approach makes sense given the timescale Uber's supply and demand function in. When you've got a crowd of people standing on a corner in the rain after a concert gets out, demand is high, and people are willing to pay a higher price for the same ride than they would in different circumstances. This model takes advantage of the fact that in this short run, the demand curve has shifted to the left and consumers are willing to pay more for the same good under these conditions.

The carsharing market is generally different than Uber's for a few reasons. First, as mentioned earlier, the time scale of car sharing is 24/7. Many trips where a user would opt for carsharing over ridesharing have more flexibility in the timing. For these less urgent transportation needs, consumers tend to price compare, which makes it more difficult to adopt a surge pricing model.

Second, many fleet operators choose a fixed-price model. Some of this has to do with branding and simplicity. Some of it's to avoid getting involved in a price war with ridesharing competitors or other transportation alternatives.

Lastly, carsharing is different in that overall time and distance driven isn't necessarily known when the purchasing decision is made. A user simply agrees to pay per hour or per mile and is billed after the fact.

THREE KEY QUESTIONS

When you consider all these factors — from pricing constraints to supply control to regulations to how people shop for carsharing — the central questions fleet operators must address become:

1. How will I know if I have too many cars in one place?
2. For any repositioning that I do, where should I put the $n+1$ marginal car that I do move?
3. For any local market, how do I decide how far apart to space the cars? How localized is this market? How should measurement of market localization affect my marketing and promotions campaigns?

The approach Ridecell uses aims to solve this complex set of problems through a common framework.

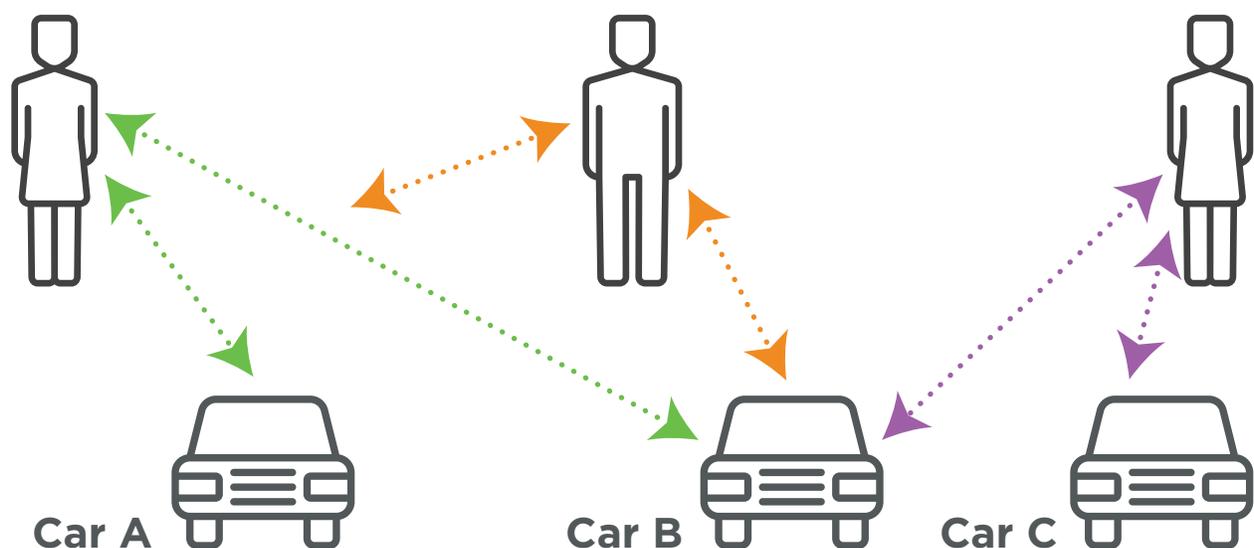
APPROACHING THE ANSWERS

With the inverse-distance preference order in mind, we start by considering what would happen if you had too many cars in one spot. This example holds even in areas that are in high demand.

Say there are 20 cars all parked around the same city block. Remember, because of the timescales involved, these cars are effectively fixed within the market's short timeframe. A user who wants to purchase transportation opens their app and rents the car closest to her. The other 19 sit there, likely at a fixed price, until enough organic demand occurs such that each car is rented one by one. Or until parking restrictions require the car to be moved (i.e., the asset perished) and the fleet operator would have to go retrieve each of the cars that went un-rented. We could empirically measure how long the 19 sat and what utilization penalty each incurred and compare it to the utilization of the first car.

Now imagine what would happen if those same 20 cars were spread farther apart at the instant the above rental happened. Now, the remaining 19 cars were positioned in different sub-markets where the consumer above still has reasonable access to a car when they want it but the other 19 would be deployed far enough away from the consumer that she would only consider the closest car. At the scale of a city, there would be many more customers like our first one in the example above and there would also be many sub-markets where they no longer compete with one another. Or at least, they have dramatically less competition and can generally be thought of as only “locally-interacting.”

INVERSE-DISTANCE PREFERENCE ORDER FOR NON-INTERACTING MARKETS



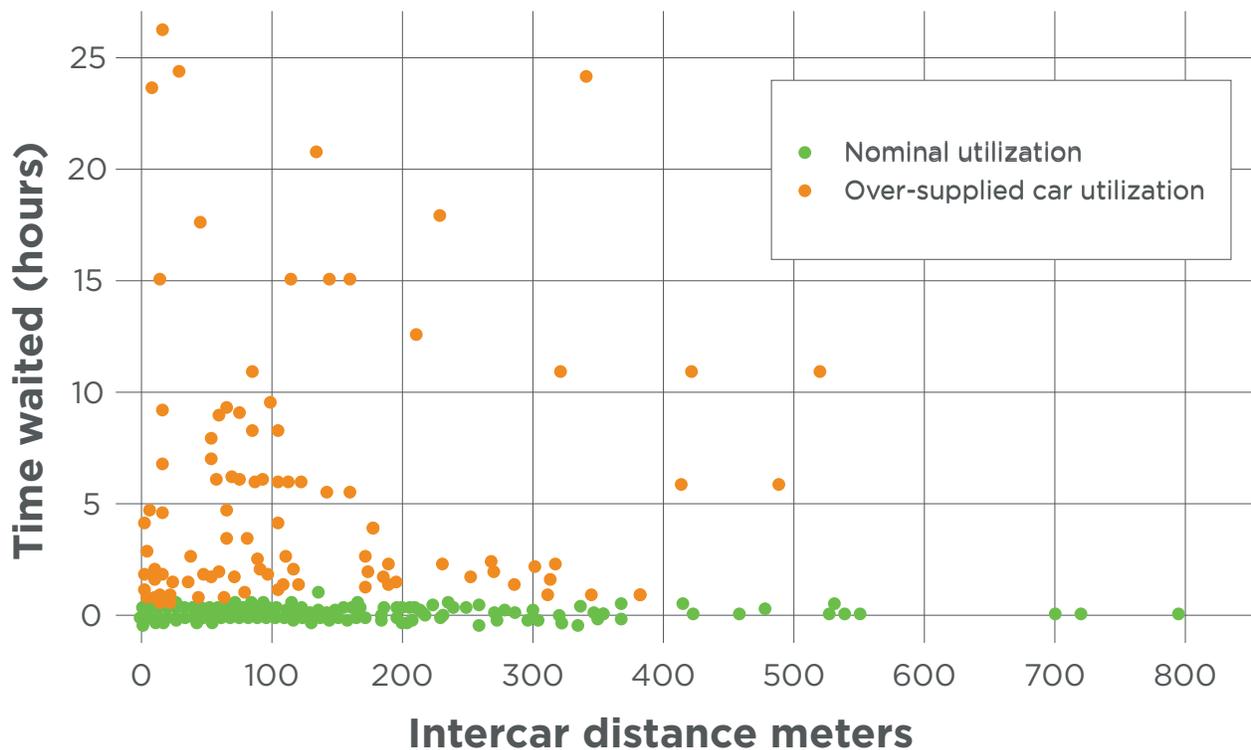
They instead participate in completely different sub-markets. Of course geography and population plays an important role here, but the point is, the density and localness of a given market is self-determined by user preference order and willingness to walk.

Once the optimal density is reached, the utilization of the first car and the other 19 will start to look the same. By comparing the utilization of the cars that weren't rented to the utilization of the one that was, we can determine where there are too many cars. Because we observe MANY such market-clearing events, they comprise an unstructured experiment with every single rental and allow us to infer the properties of the local market by comparing the market dynamics of the rented car, relative to the other nearby alternatives that weren't.

THE INTER-CAR DISTANCE METRIC

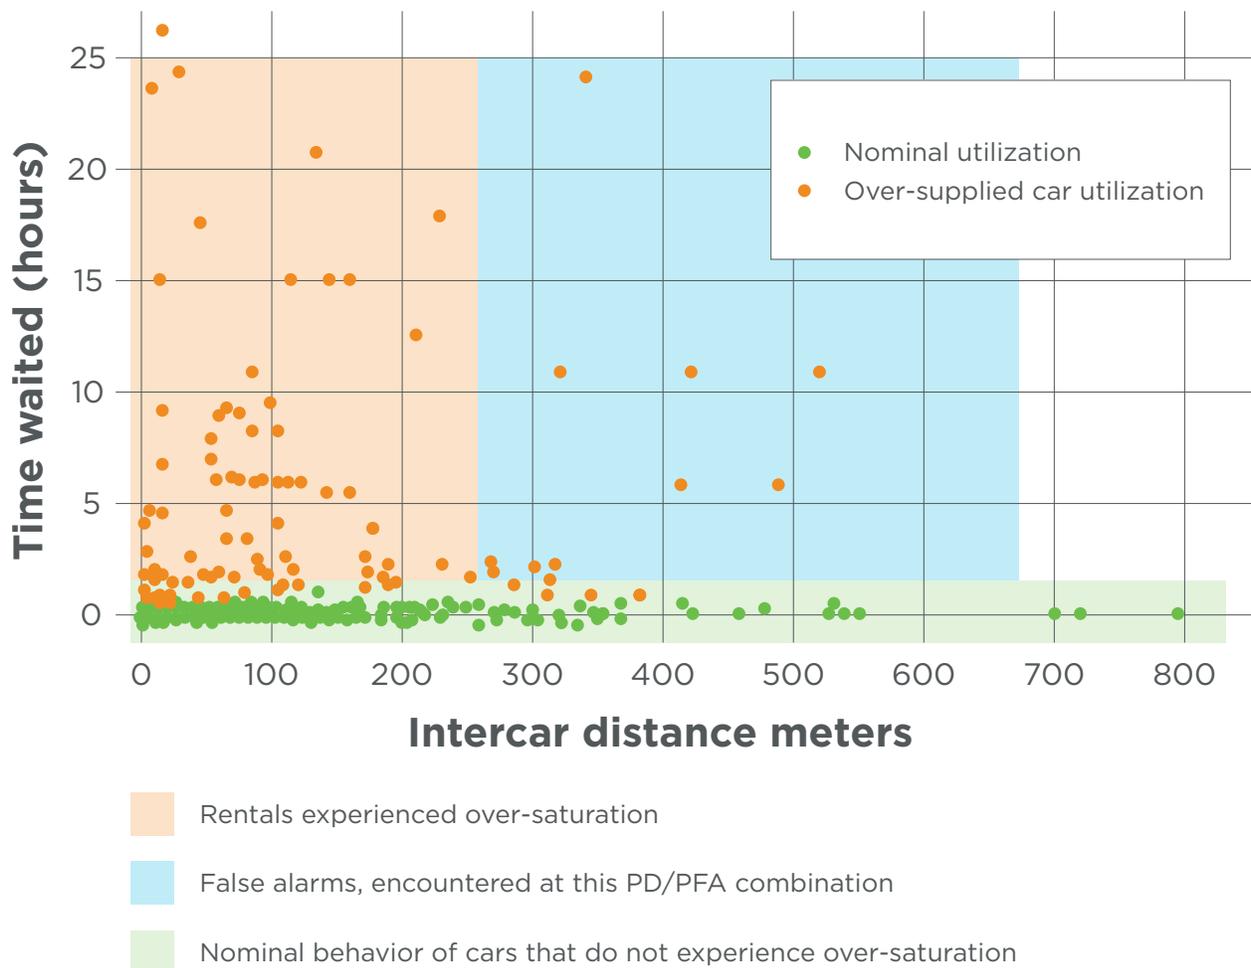
By compiling transaction data for a given spot, we can create a metric to measure this market clearing we call inter-car distance (ICD). This is defined as the distance between the car that was rented and its nearest unrented neighbor. By plotting ICD versus utilization, we can see when cars that were too close (low ICD) had longer utilization than cars with a higher ICD.

ICD VERSUS UTILIZATION



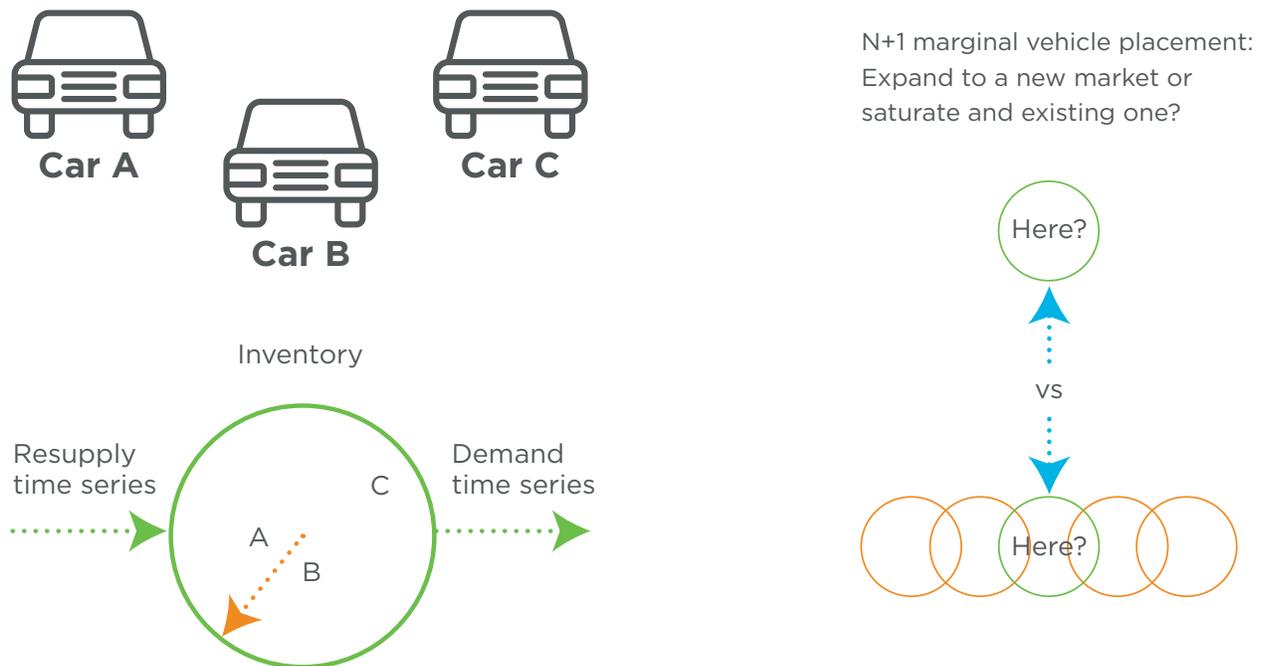
Through this method, we can simultaneously solve all three of the key questions fleet operators ask themselves every day by measuring how these localized markets function and measuring the utilization penalty for when there are too many cars. By setting thresholds in this plot, we can empirically apply our desired classification accuracy rate and balance the probability of detection with the probability of false alarm. For this notional example graph with 1000 points, for a 1% false alarm rate, we could set the threshold at the 99th percentile in the distribution, and thus arrive at, for this particular location, the optimum spacing should be on the order of 280 meters. This radius is important because it tells us the localization scale of this market, and by measuring the distance over which the market interacts, we can create tailored predictions at just the right spatial scale for every spot in their service region.

ICD VERSUS UTILIZATION WITH THRESHOLD



UNLOCKING THE POWER OF THE METRIC

By combining ICD data with Ridecell's high-volume machine learning framework, Daedalus, we can do powerful things like make time series predictions for the number of cars flowing in and out of an area, and accomplish this at scale.



Optimal Placement at scale is achieved through solving for the time-series balance for every locally-interacting market to ensure proper inventory balance and inform the decision to expand to a new market or re-supply an existing one

By combining the above pieces, we're able to automate:

- › Over-supply warnings operators can respond to by enabling price incentives in an effort to reclaim some revenue and avoid parking penalties.
- › Market opportunity recommendations that help an operator determine when they are able to move a few cars from one sub-market into another market – without cannibalizing their revenue and losing sales.
- › Price and promotional adjustments in a unified framework of supply and demand forces and measurement of market localization.

Finally, because we can actually measure the willingness a customer is willing to walk to a shared car, we can ensure that operators maintain a high quality experience for their customers while minimizing the risk of having too many in any one sub-market. So, as you can see, Ridecell's platform not only answers operators' most pressing questions but also helps them turn the answers into actions that are tied to revenue optimization. Through this ecosystem of applied econometrics and business automation, we're able to bring greater profitability to shared fleets on a daily basis.

About the Author

Shawn Higbee MBA PhD leads a data science team at Ridecell, solving a wide array of data science in the shared mobility space. He holds degrees in Industrial Engineering, Applied Math, and Imaging Science as well as an MBA from UC Berkeley. Prior to joining Ridecell he led a design team at Lawrence Livermore National Laboratory focused on putting machine learning systems onto satellites, as well as designing and building data science pipelines for imagery analysis. Earlier in his career he spent 12 years as an officer in the US Air Force including tours as a technical product manager at the National Reconnaissance Office and algorithm designer at the Air Force Research Lab.

About Ridecell

Ridecell helps companies build and operate profitable mobility businesses. With our High-yield Mobility™ SaaS platform, business services, and ecosystem partners, Ridecell customers are able to maximize three key profit drivers: customer experience; fleet utilization; and operational efficiency.

Founded in 2009, today Ridecell powers some of the most successful shared mobility services in cities across Europe and North America. These services include ZITY from Ferrovial and Groupe Renault, and Gig Car Share from AAA.

Ridecell is headquartered in San Francisco, California, and has more than 150 employees in offices across the globe.