Uber Freight

Can empty miles in freight be eliminated?

Uber Freight explores how the industry can build a more efficient and sustainable future

Contents

It is estimated that 20-35% of trucking miles in the US are driven empty

Mostly when a truck drives empty to its next pick up, failing to find a nearby load that meets the driver's schedule. The freight industry frequently talks about reducing this colossal inefficiency–likely closer to the upper end of this range–but, to our knowledge, no research has ever answered the question: **what is the achievable empty miles reduction and how can it be done?**

In this research piece, we use Uber Freight's industry-leading access to freight shipment data and optimization techniques to answer this question. Our findings show a major opportunity for the industry to reduce empty miles by up to 64%, representing a 23% reduction in overall miles driven by freight trucks. These gains can be achieved through optimization techniques coupled with broad visibility, and the growth of autonomous trucking technology.

Such a massive reduction in trucking miles will have ramifications well beyond the freight industry. First, a reduction in empty miles will have positive effects on road congestion and fatal accidents. The effect a single truck can have on traffic flow is the equivalent of 2-15 passenger cars, depending on terrain and traffic conditions.[1] In addition, trucks travel 5.5X more miles per vehicle than light-duty vehicles, [2] and account for 9% of all vehicles involved in fatal crashes. [3] Furthermore, freight trucking is responsible for a staggering 7% of US greenhouse gas (GHG) emissions. [4]

Our analysis relies on the experience we have developed in optimizing freight networks, and uses Uber Freight's broad and detailed dataset of shipments, backed by more than 10 million dry van shipments under its management annually. Our approach strives for an optimal assignment of trucks to loads enabled by: (1) full visibility over all the shipments taking place on the US freight network and the trucks available to haul them; (2) a centralized ability to route trucks and assign them to load itineraries; and (3) the flexibility resulting from deploying hybrid human-autonomous fleets.

Eliminating these empty miles will save 1.5% - 2.5% of US GHG emissions, equivalent to transitioning 7% - 11% of the US national power grid to renewable sources.

We find that empty miles in a freight network can be broken down into three principal categories with varying impact. While some of these can be mitigated through network optimization, others cannot.

Figure 1. Current empty miles driven by tractor-trailers in the US.

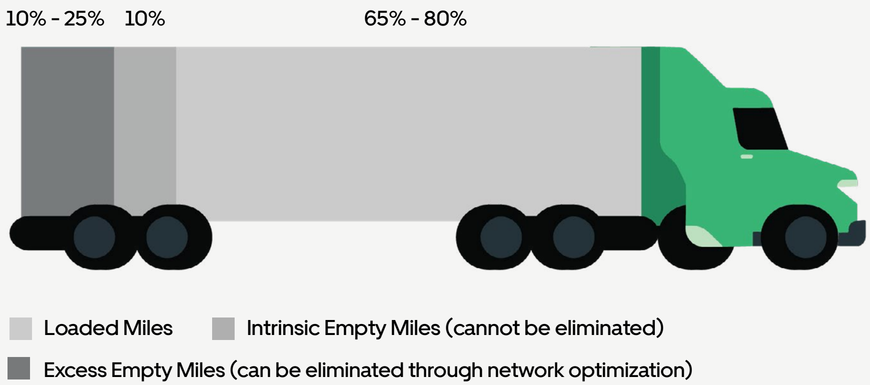

1. Fragmentation and poor planning:
In 2021, there were 708K active interstate freight carriers, [5] 91% of which had 10 or fewer trucks. This fragmentation results in lost efficiency, particularly with small carriers and private fleets, [6] due to: (a) insufficient visibility into available loads; and (b) lack of advanced optimization solutions.

2. Intrinsic network imbalance:
The US freight network is imbalanced by nature. Some large populated cities, such as Miami, have more load deliveries than pickups, while other industrial and agricultural markets might have more pickups than deliveries. On average, we expect more loaded trucks to travel into Miami than trucks coming out of Miami. This results in intrinsic empty miles in freight networks that no network optimization can eliminate.

3. Supply and demand density:
Higher truck density means that it would be easier to find a nearby truck to haul a load. Therefore, when more trucks are available to haul a load, carriers can reduce their empty miles. However, carriers do not optimize their fleet sizes to empty miles, but to the total operating costs, which include the purchase and financing of equipment. Similarly, higher load density means that a truck is more likely to find a nearby convenient load.

In the following sections, we develop these findings and discuss how the freight industry can accelerate its transition towards a more efficient and sustainable future.

What is the status quo?

There is no consensus on empty miles traveled by tractor-trailer trucks in the US. Estimates vary between 15% and 35%. [7] Figures falling on the lower end of the range are likely underestimated. First, these figures are based on surveys, which are prone to underreporting. Second, these surveys target mostly medium and large fleets, and therefore, are not representative of the fragmented trucking landscape.

Our back-of-the-envelope analysis of vehicle miles traveled (VMT) shows that empty miles likely fall on the higher end of the above range. Trucks move approximately 300 million truckloads annually in the US, according to ACT Research. [8] For-hire fleets move truckloads with an average length of haul (LOH) of about 450 miles. The average LOH is significantly shorter for private fleets, which travel on average 238 miles per load. [9] This indicates that tractor-trailers travel about 103 billion miles loaded per year. In comparison, the US Bureau of Transportation Statistics estimates the annual mileage of tractor-trailers at 175 billion miles. [10] Assuming 10% of these miles are traveled by LTL trucks, [11] we estimate that the fraction of empty miles traveled by US truckload fleets is approximately 35%, and use this figure as the empty miles baseline throughout the paper.

Figure 2. Estimates of empty miles from different studies.

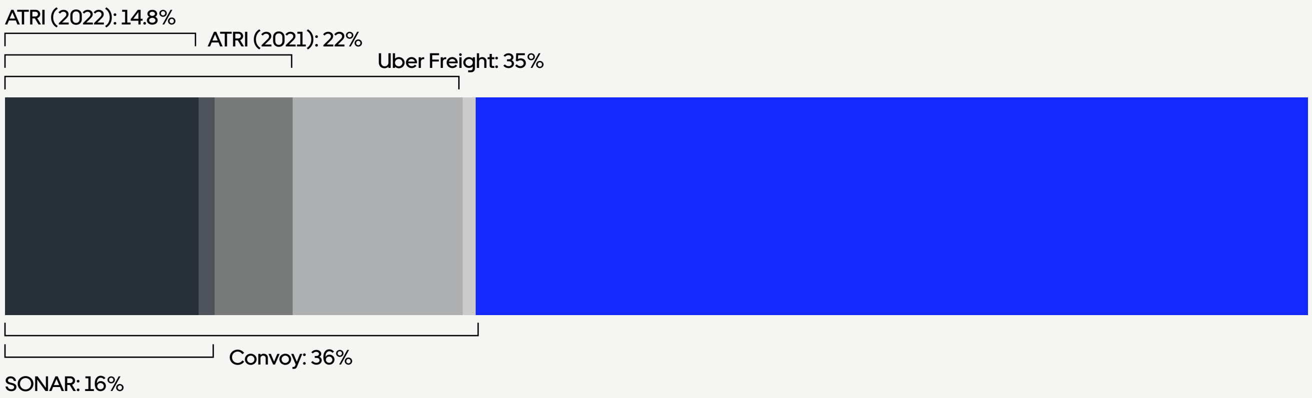

Current efforts to curb empty miles

Large fleets have been deploying network optimization to improve their efficiency. The percentage of empty miles driven by medium and large-sized fleets (surveyed by ATRI) has been decreasing in the past few years, and reached a record low level in 2021. [4] However, carriers' ability to improve their utilization remains limited by their visibility into demand and rigid contracts with shippers. For example, small carriers might not have access to the loads that will minimize their empty miles, resulting in inefficiency across the entire fleet.

Uber Freight has been offering carriers solutions to minimize empty miles, such as reloads and round trips (bundles). This makes it easier for carriers to move from point A to point B and back again as efficiently as possible. Furthermore, Uber Freight Transportation Management customers receive advanced consolidation services that reduce trucking miles by combining orders into efficient LTL and multi-stop shipments.

These fragmented efforts have had positive impacts, however, they have not yet scaled enough to tilt the needle nationally. Further, it is not clear that an increase in utilization on one truck does not come at the expense of another truck. Namely, can empty miles be reduced or just be allocated to different carriers?

A framework for optimized freight network

In order to estimate the opportunity available to the industry to cut empty miles, we assume that we have full visibility to all freight movements–approximately 6 million loads per week and the 1.4 million active trucks available to be dispatched– and the ability to assign itineraries to trucks. Equipped by this theoretical visibility and control, we set to arrange loads into itineraries that minimize the total number of empty miles driven. While this framework deviates from reality in key ways that we discuss next, its outcome, nonetheless, sheds a new, optimistic light on the opportunity ahead of the industry.

The most obvious limitation of this framework is that no entity exhibits the visibility or the control required to optimize the network, as we set out to do. However, two technologies can mature to enable these capabilities, albeit at a smaller scale. First, the past decade has seen advances in truck visibility–aided by the ELD mandate–and the emergence of digital freight platforms. Second, the emergence of hybrid human/autonomous fleets will see autonomous trucks take on itineraries guided primarily by utilization, without being constrained by Hours of Service (HOS) or the drivers' need to return home, while enabling human drivers to focus on shorter loads in the vicinity of their homes.

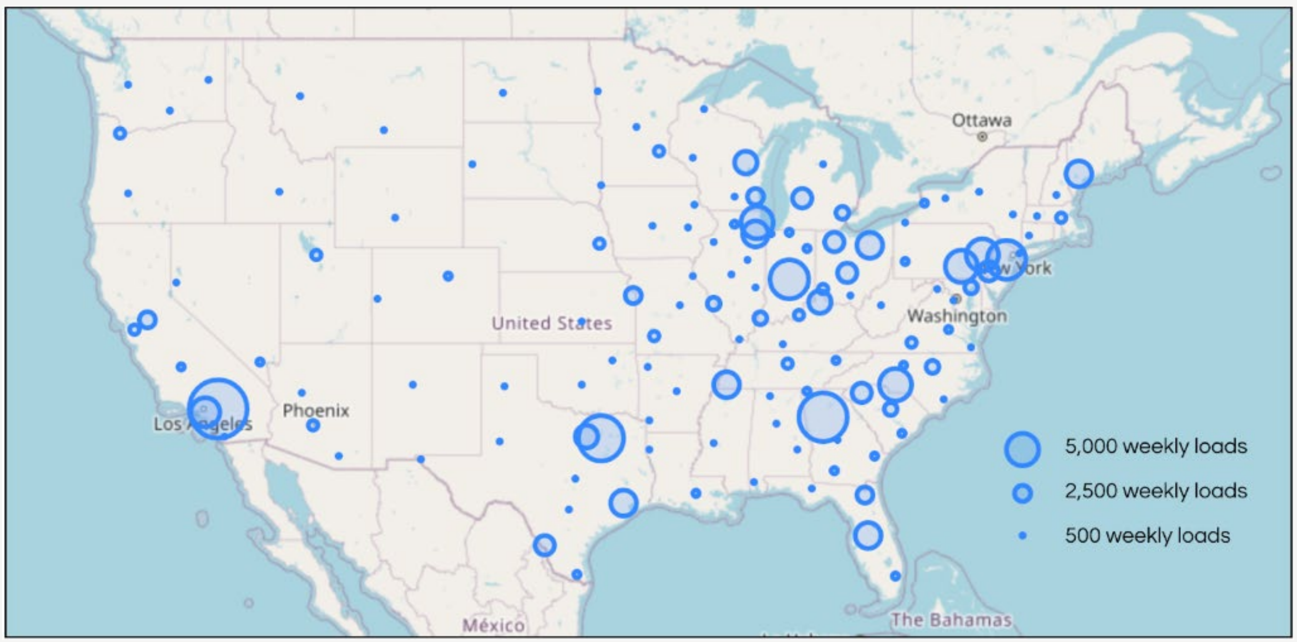

Figure 3. Uber Freight's loads in the week of Oct 22 - Oct 28, 2022 (by market).

Furthermore, because we lack absolute visibility into the US freight movements, we have to rely on a simulation based on the partial visibility made up of Uber Freight's data. With more than 10 million dry van shipments annually under management, we believe that Uber Freight has the broadest visibility of full load level data in the US. Specifically, we use data from 200,000 dry van loads that took place during the week of Oct 22 - Oct 28, 2022. [12] We argue later that scaling this methodology to the absolute visibility case will result in even greater potential for empty mile reduction. We also consider that 80,000 trucks are available to haul these loads. We choose this number of trucks because it yields similar utilization levels to existing US carriers - about 80,000 loaded miles/truck annually. [13] Therefore, it also replicates the current supply/demand ratio of today's trucking network.

Lastly, we make some simplifications to the optimization problem in order to render the problem solvable. Those are discussed in detail in the following section with reference to whether they are unavoidable or could be relaxed in the future. Despite their limitations, these simplifications leave the results fairly robust to demand uncertainty, as we discuss in the following section.

While the scenario we depict will take decades to materialize, we believe that the findings and learnings from our analysis can be directionally generalized to today's trucking fleets.

Optimization approach

If we had unlimited computational resources and perfect knowledge of all the upcoming loads, we could use standard optimization methods to assign trucks to loads. For example, an integer program (IP) [14] would yield the optimal truck itineraries, while taking into consideration constraints on appointment windows and truck availability.

In reality, this is neither realistic nor tractable. Many of these loads are unpredictable due to short lead time. Therefore, the complete demand picture–necessary to plan itineraries–is available only a few hours ahead of a pick up window. In addition, solving a problem of this scale is difficult to impossible, even with advanced optimization tools. Even simpler problems in operations research, such as the Vehicle Routing Problem (VRP), are classified as NP-hard. This means that the size of problems that can be optimally solved using mathematical programming or combinatorial optimization is limited.

To solve this optimization efficiently, we break it down into smaller, more tractable problems. We can break it down spatially, by optimizing one geographical market at a time. In addition, we can break it down temporally, by looking at shorter time periods. This assumption is reasonable, especially because we do not have full visibility into all loads ahead of time. In our simulation, we consider all available trucks in a given market, such as the Boston metro, and all loads scheduled for pickup in that region in the next 12 hours when building the itineraries, and update these itineraries every 6 hours. Since we do not assume full knowledge of demand beyond 12 hours, our optimization results do not overfit historical demand patterns, and therefore are fairly robust to demand shocks or uncertainty.

Our optimization makes two additional simplifying assumptions, motivated by computational constraints, as well as by the anticipated proliferation of autonomous trucks. The optimization model assumes that trucks can travel without the rest dictated by the Department of Transportation's Hours of Service (HOS) regulations. Introducing HOS considerations will amount to a reduction of the available trucks that will marginally increase empty miles as discussed in the results section "Do we need that many trucks?". We further ignore any requirements by the truck driver to return to their home periodically, thus eliminating any empty miles resulting from drivers' personal considerations. While these assumptions disregard the critical considerations of safety and wellbeing of drivers, they are a good approximation of a hybrid human/autonomous freight network. Further, these assumptions simplify the optimization model dramatically and future research will seek to relax them.

Divide-and-conquer optimization model

As described above, we divide this problem spatially and temporally. We consider 133 markets [15] and 6-hour intervals. At each interval, we first assign loads to available trucks in each market with pick up time in the next 6 hours. Afterwards, we reposition idle trucks to meet the demand in each market in the following 6-hour pick up window. Periodically, trucks are dispatched across markets, to maintain a desired level of supply/demand balance. Therefore, we assume visibility of all trucks and all loads scheduled for pick up in the next 12 hours, and knowledge of the average historical demand in each market.

1. Load-to-truck assignment

In the first step, we need to assign the right truck to the right load within each market within the 6-hour pick up window to minimize empty miles. We also need to ensure that trucks pick up these loads on time. We are able to arrive at an optimal solution for a given market and time interval in a few seconds. [16]

After assigning loads to trucks in a given time interval, we update the statuses and locations of these trucks. A busy truck becomes available to pick up loads in the following intervals only after completing its delivery. This includes the time needed for (1) driving to the pickup location, (2) two hours for loading, (3) delivering the load to its final destination, and (4) two hours for unloading. [17]

2. Truck repositioning

In the second step, we need to reposition trucks proactively to ensure that enough trucks are available in each market to pick up the scheduled loads. For example, if there are 300 loads scheduled for pickup in Dallas, but only 200 trucks are available, we would dispatch 100 trucks from a nearby market, such as Oklahoma, ahead of time (in the previous time interval).

To do so, the second optimization step considers all markets simultaneously. It ensures that there will be enough trucks available in each market to satisfy the demand in the next time interval, while minimizing the total distance traveled by all repositioned trucks. The output indicates the number of trucks that are dispatched from one market to another in a given interval. As before, we also update the locations and statuses of the dispatched trucks, until they have arrived at their destination.

3. Supply - Demand rebalancing

Freight markets are imbalanced by nature. Some markets, like Miami, have more inbound freight than outbound, while others have more outbound freight than inbound. Since the first two steps above do not anticipate future demand, idle trucks will cluster in inbound markets, like Miami, and will be dispatched to outbound markets inefficiently on a load by load basis. Therefore, we find that proactively sending trucks from oversupplied to undersupplied markets based on the historical demand distribution yields a more efficient freight network. This can be done periodically, for example, once a week. [18]

The optimized freight network

About two thirds of empty miles can be eliminated through network optimization.
In our optimized network, the 80,000 trucks travel empty 12.5% of the time, significantly below our current estimate of empty miles in the US (35%). We classify empty miles into two types: intra-market empty miles driven within markets to pick up loads, and inter-market empty miles driven between markets to reposition or rebalance trucks. The results show that trucks travel on average 17 miles to pick up a load in the same market, which is only about 2.4% of the total miles traveled. However, they travel empty 10.1% of the miles–about 73 miles per load– to relocate to undersupplied markets.

Optimizing the freight network will eliminate the vast majority of empty miles in the network, carrying significant societal and environmental gain. However, it cannot eliminate all of them. Our results indicate that inter-market empty miles would constitute at least 10% of the total miles traveled across all simulation scenarios. In the following section, we show that a large fraction of our inter-market empty miles are intrinsic, and cannot be eliminated with network optimization.

Figure 4. Empty miles driven by tractor-trailers in the optimized network.

What are intrinsic inter-market empty miles?
Economic activity varies by market: dense population areas and export hubs demand inflow of goods, while manufacturing, agricultural, and import hubs produce outflow of goods. As long as the underlying demand and supply patterns of goods are imbalanced, empty miles cannot be fully reduced. We estimate the average load imbalance in a given market to be about 18%.[19]

This means that about a fifth of all loads cannot be matched to reloads in the same market, even in an optimized network.[20]

Building on our rich dataset, we used our optimization framework to quantify the intrinsic empty miles in the network. As we increase the number of trucks available to haul the loads, both inter-market and intra-market empty miles generally decrease. However, beyond a certain threshold, inter-market empty miles plateau, because we will always need to rebalance

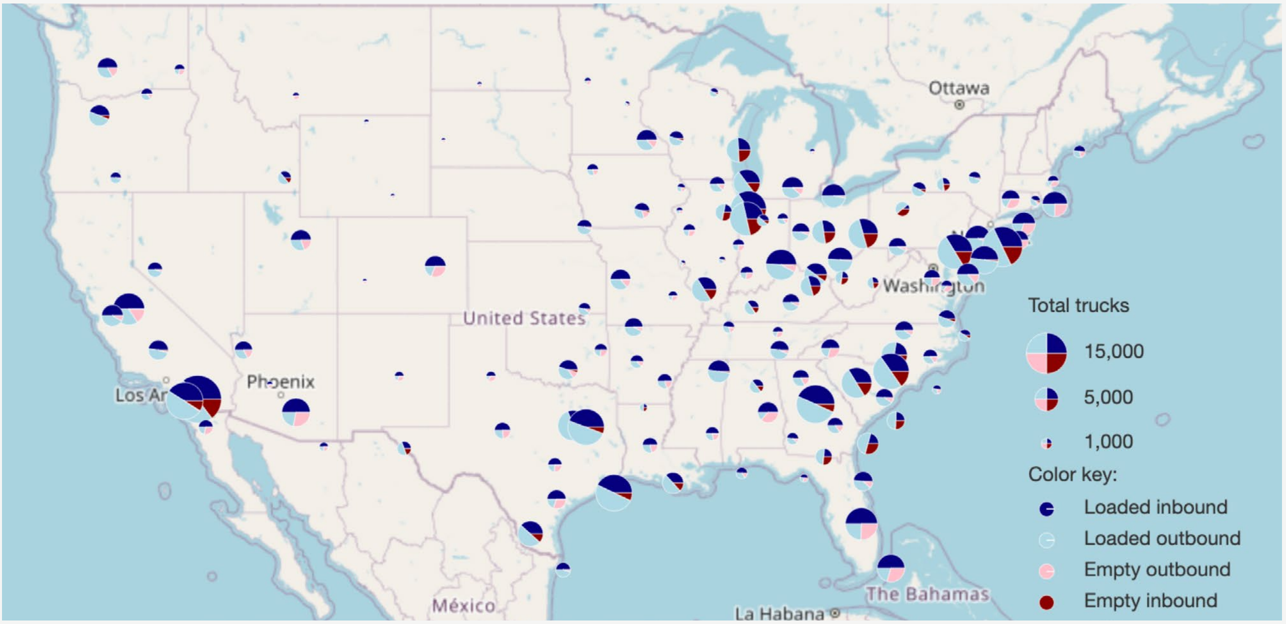

Figure 5. Distribution of loaded and empty movements across markets

trucks from inbound markets to outbound markets. Even with a very large number of trucks, these trucks will deadhead at least 10% of the total miles.

What does this tell us about survey-based estimates of empty miles?
Some survey-based studies estimate empty miles in the US at 14%-16%, which is close to what we obtain in a highly optimized network with relaxed constraints. There could be two explanations for this. First, these estimates can be severely underreported. For example, in the Vehicle Inventory and Use Survey (VIUS), several responses indicated that they drove no empty miles at all. [21] Second, the large fleets surveyed in these studies are probably sophisticated enough to construct balanced networks. This allows them to reduce their intrinsic empty miles. However, by doing so, they pass the bulk of network imbalances to smaller fleets and owner-operators, which will have to drive more empty miles, but are not captured by these surveys.

Truck utilization flows in the US

Based on our simulation results, we can analyze the distribution of empty miles across different geographies and highway corridors. Figure 5 shows the number of trucks entering and leaving each market in our simulated network. The size of each circle is proportional to the total number of inbound and outbound trucks, while the colors indicate whether these trucks are loaded or empty, and whether they are headed inbound or outbound.

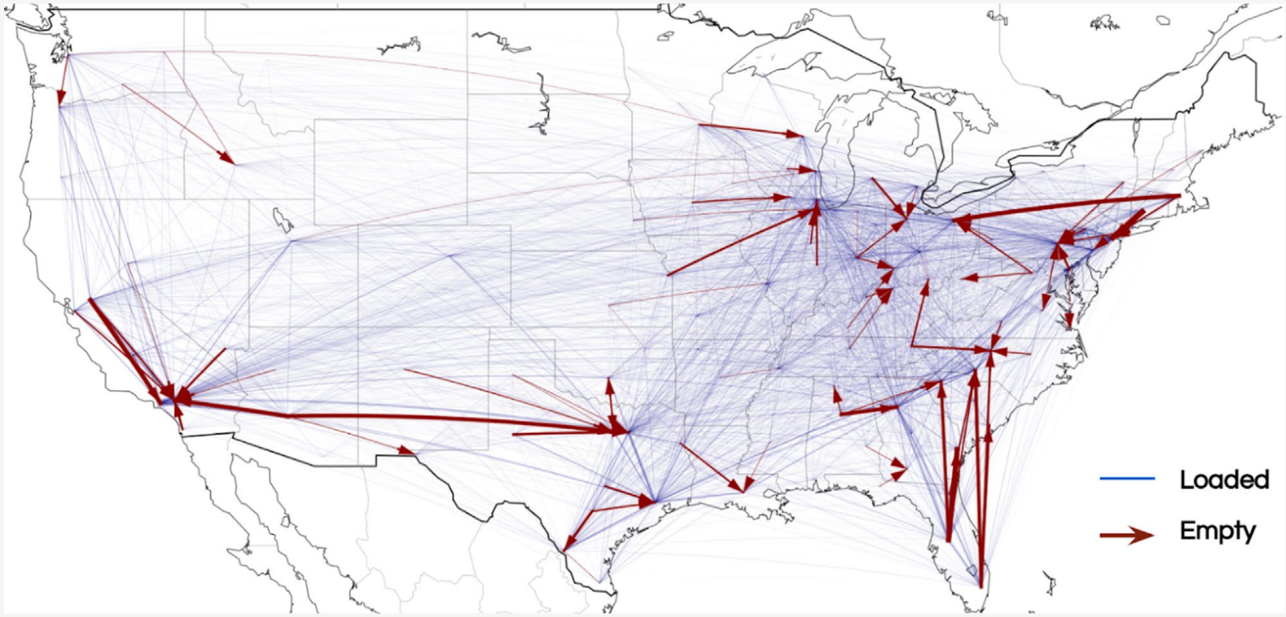

Figure 6. The resulting empty movements between markets.

We can draw two conclusions from this figure. First, all circles are divided equally between inbound and outbound freight. This is expected, because, in the long-term, the total number of trucks entering each market will be equal to the number of trucks leaving it. Therefore, the higher the imbalance between inbound and outbound freight, the more empty inbound or outbound miles a market will have. For example, Phoenix has 2.5X more inbound than outbound freight, resulting in a higher percentage of outbound empty miles. On the other hand, Ontario, California has more outbound freight than inbound. Because of their geographic proximity, we would dispatch empty trucks from Phoenix to Ontario.

Second, in an optimized network, empty miles from or to each market will be mostly unidirectional: for any given market, empty trucks are either mostly inbound, or mostly outbound. If we had full visibility over all freight movements, we would neither dispatch empty trucks into an already saturated market (like Miami), nor from an undersupplied market (like Ontario).[22]

Figure 6 shows that the majority of empty truck movements are concentrated along a handful of short to medium length lanes, connecting inbound and outbound markets. Cities in Florida (Miami and Lakeland) are among the largest inbound freight markets, because these are dead end regions. Markets with major ports, such as Houston, Los Angeles, Savannah, and Elizabeth (New Jersey / New York) are among the largest outbound freight markets. Industrial and agricultural markets located in the Midwest in addition to Dallas and Fort Worth are also large outbound markets.

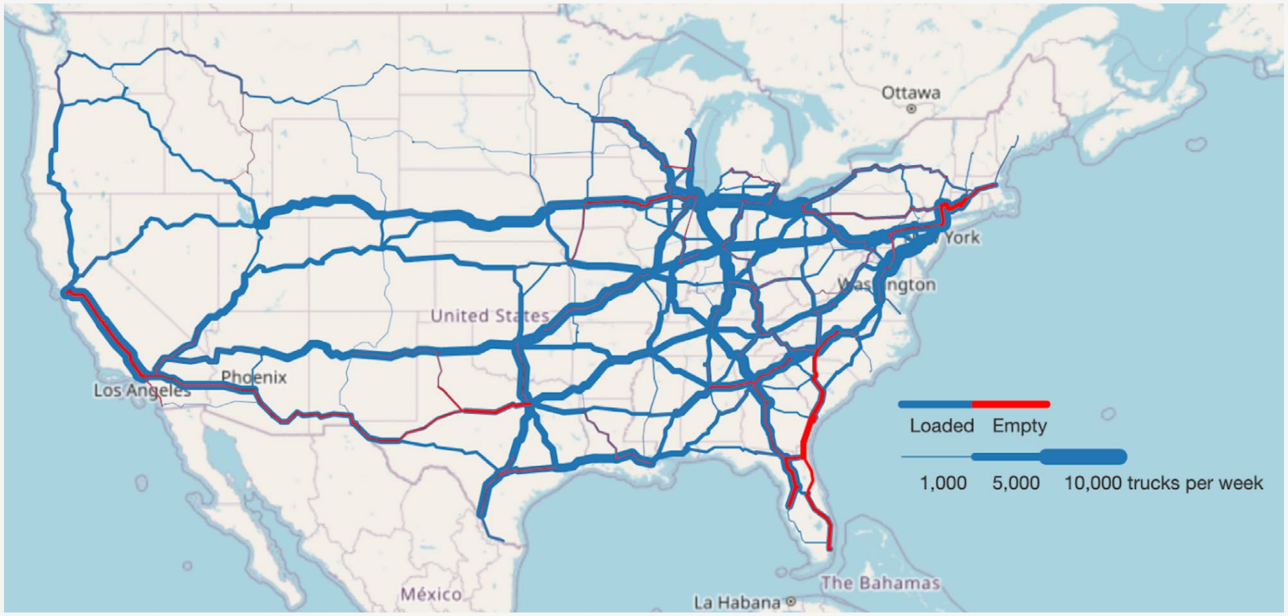

Figure 7. Distribution of loaded and empty miles across major interstate highways.

Practically, inter-market freight movements are concentrated on interstate highways, with multiple inter-market lanes overlapping on some interstate segments. To illustrate the impact of empty miles on the interstate system, we routed all freight movements in our simulated network on the US interstate system. As shown in Figure 7, in the optimized network, a few interstate corridors will carry a disproportionate amount of empty trucks. These include the southern part of the I-95 corridor between Eastern Florida and Southern Carolina, the I-5 corridor between Southern California and San Francisco, and I-10 and I-20 corridors between Dallas and Southern California.

Can intrinsic empty miles be mitigated by other means?
Because empty truck movements are concentrated on a few highway corridors, intrinsic empty miles, or at least their environmental impacts, can be mitigated by adopting efficient strategies for fleet rebalancing. For example, truck platooning, a technology that will be enabled by autonomous trucks in the future, can cut fuel consumption by 10%-17%. [23] In addition, rail cars can be used to haul tractors and empty trailers to reposition trucks at a fraction of energy emissions.

These findings also have implications for land use and transportation planning. For example, intermodal lines can be built along these corridors, either to transport freight into oversupplied markets, or transport empty containers and tractors from these markets. In addition, city planners can incentivize more freight generating activities in these markets, such as ports and industrial parks.

Do we need that many trucks?

In our simulation, we considered 80,000 trucks available to haul the 200K loads, to replicate the current supply/demand ratio of the US freight network, and the current utilization levels of trucks (about 80,000 miles/year). However, in the future, autonomous trucks will operate on longer shifts, and travel more miles per week. If we were to optimize a network of autonomous trucks, do we need 80,000 of them to ship these 200K loads?

If we assume that an autonomous truck can handle 200,000 loaded miles per year, we would only need 33,000 trucks to serve this level of demand. However, reducing the available supply results in more empty miles. In total, these trucks will deadhead about 14.7% of the miles, up from the 12.5% we obtained with 80,000 trucks.

Table 1. Empty miles as a function of supply density in an optimized network.

Available supply (trucks)	80,000	33,000
Loads served per week	198,359	198,359
Loaded miles per week	124 million	124 million
Loaded miles per truck per week	1,553	3,766
Empty miles (%)	12.5%	14.7%
Truck time utilization (%)	33.4%	75.7%

Is an autonomous platform better off with a smaller fleet?

If an autonomous platform chooses to operate a smaller fleet with higher time-based utilization, [24] trucks will drive slightly more empty miles, and therefore, have higher operating costs. With a constrained supply of 33,000 trucks, we estimate that these trucks will be driving about 14.7 hours a day, significantly more than current drivers (about 6.5 hours a day [25]). Therefore, these trucks will depreciate at a faster rate, and will need to be replaced sooner.

However, by using fewer trucks, fleets can save on the opportunity cost of tied-up capital. Trucks are expensive assets, and the capital needed to purchase or finance a truck can be better utilized elsewhere. We estimate that by reducing the fleet size by 60%, fleets can save 7 cents per mile on the opportunity cost of capital alone. In addition, they can significantly cut truck parking costs, not only because there are fewer trucks, but also because these trucks are rarely idle. These two sources of cost savings are more than enough to offset the additional operating costs of running more empty miles.

This indicates that in the future, autonomous fleets will be far more efficient than human-driven ones. **A centrally-managed autonomous fleet will cut both empty miles and the required number of trucks by about 60% compared to the status quo, boosting truck utilization by about 250% compared to the current state.** But to do so, it needs sufficient demand density, state-of-the-art optimization tools, and trusted relationships with shippers, all of which Uber Freight can provide.

Previous studies have attempted to estimate the cost savings of autonomous trucks. However, they have mostly focused on reductions in the operating cost per mile resulting from the elimination of driver wages and benefits. These studies did not account for the potential reduction in empty miles, which can cut the total cost per loaded mile by up to 20%.

Are HOS limitations a barrier?
Carriers are currently constrained by hours-of-service limitations, which prevent truck drivers from driving more than 60 hours a week, or 11 hours a day. This effectively reduces the available supply hours of a truck. While our model does not account for HOS, adding this constraint will amount to reducing the number of available trucks that results in an increase in empty miles. Under this scenario, the expected fraction of empty miles would fall between the two scenarios discussed above, resulting in 12.5% and 14.7% respectively. Future research can add the HOS constraint into this methodology.

The value of density and visibility

Demand density and load visibility also play an important role in reducing empty miles. To demonstrate this, we compare our simulated Unified Freight Platform to smaller fleets (or fragmented freight platforms), each having access to only 1% of the original demand, about 2,000 loads per week. The number of trucks available to haul these loads is reduced by the same proportion, to maintain the same supply / demand ratio.

The results, shown in Figure 7, demonstrate the value of network density and load visibility. With only 1% demand density, trucks would travel empty 17.3% of the time, even if a freight planner were to optimize its network. Therefore, in a world with fragmented freight platforms, advanced technology and optimization techniques will not be enough to achieve the same level of efficiency as with a unified planner.

Following this logic, if we scaled this model to have visibility of all 6M US weekly freight movements, not only the 200k we simulated here, we would expect that empty miles will be reduced even further.

Figure 7. The value of demand density: empty miles as a function of load visibility.

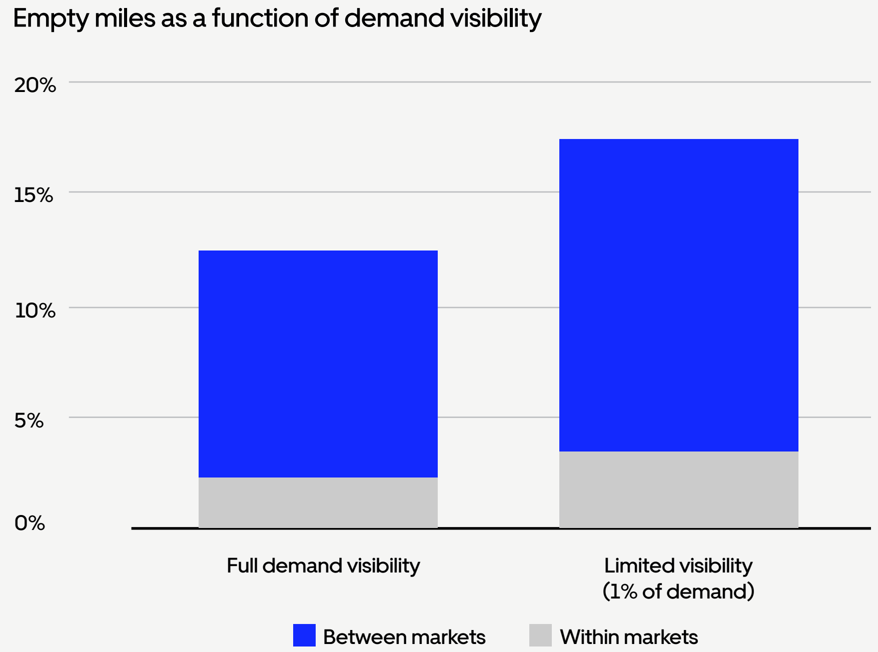

Conclusion

The trucking market is largely fragmented and inefficient. Freight imbalances between different geographies contribute to a large fraction of empty miles. However, carriers' lack of visibility and optimization expertise contribute to the majority of empty miles.

The result is excess deadhead, unnecessary emissions, wasted time for drivers, undue road congestion, and avoidable accidents. First, empty miles cost carriers about as much as loaded ones do. A large percentage of these costs is passed to shippers, and eventually to consumers. Second, these empty miles are wasting road users' lives and livelihoods. The Bureau of Transportation Statistics (BTS) estimates that tractor-trailers travel 175 billion miles annually. This means that US truck drivers could be wasting 3.5 billion hours each year driving empty, the equivalent of 5,200 lifetimes. [26] In addition, empty miles are responsible for more than 1,000 fatalities annually, due to truck crashes. [27]

Although the results we developed here will not be achievable in the near future, there are various ways in which we can improve our network efficiency in the short term. With easy load search and personalized recommendations of loads and load bundles, Uber Freight and similar platforms help carriers book convenient loads that are not too far out in space and time. In the coming years, autonomous trucks operating alongside human drivers can boost network efficiency. These trucks can be managed by a central network with the intent of minimizing empty miles. In addition, they can complement human drivers on capabilities and preferences, by handling loads that are infeasible or unattractive to drivers. However, these fleets, and especially in their early stages, need to be more efficient than traditional fleets to remain competitive. By becoming the platform of choice for autonomous fleets, Uber Freight can benefit from its visibility, technology, and scale to allow these fleets to make the best use of their assets and resources.

Authors

Bar Ifrach has worked in data science, applied science and marketplace domains for 14 years, most recently as Head of Marketplace at Uber Freight. Prior to joining Uber, Bar spent 5 years at Airbnb's data science team, where Bar worked on key marketplace problems and built Airbnb's Homes Data Science team. Prior to his work in the tech industry, Bar pursued a PhD in Operations Management from Columbia Business School and obtained a postdoctoral fellowship position at Stanford University.

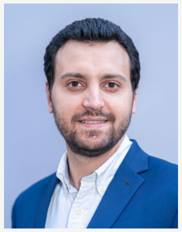

Mazen Danaf is Senior Economist and Applied Scientist at Uber Freight. Mazen's work focuses on analyzing the freight transportation landscape and producing short- and long-term forecasts based on supply and demand dynamics. He is also a research affiliate with the Intelligent Transportation Systems (ITS) Lab at MIT, where he completed his PhD in 2019. His work falls at the intersection of ITS, economic modeling, and analytics.

Appendix A: Load to truck assignment optimization

Objective

The goal of this optimization model is to determine the optimal assignment of trucks to loads in a given market and in a given time period (t), which minimizes the total distance traveled to pick up these loads. We define binary variables x_{ij} which take a value of 1 if truck i picks up load j and 0 otherwise. The objective function is specified as:

$$min \sum_{i,j} x_{ij} d_{ij} \qquad \forall\, i,\, j$$

Such that:

- Load *j*'s pickup appointment window overlaps with the 6-hour time interval *t*.
- Truck *i* is available in the same market (*m*) during time interval *t* (trucks in a market cannot pick up loads from other markets).
- A truck can be assigned to a load only if it can arrive at the pickup location before the end of the appointment window:

$$AvailTime_i + e_{ij} \leq MaxPickup_j$$

Where:

- d_{ij} and e_{ij} represent the empty driving distance and travel time between location of truck *i* and the origin of load *j* respectively,
- $AvailTime_i$ represents the time at which truck *i* becomes available (after dropping off its last load or finishing a rebalancing trip).
- $MaxPickup_j$ is the maximum allowable pickup time of load *j*.

For simplicity, we omit all subscripts of time interval *t* and market *m*, since this optimization is limited to a single time period and market.

Constraints

The constraints of this optimization are formulated as follows:

1. Each truck can pick up one load in a 6-hr interval. While this might be a conservative assumption, it is reasonable to expect that a truck will serve at most one load in a 6-hour period, considering that loading and unloading times alone can take up to 4 hours. In addition, imposing this constraint simplifies the optimization problem significantly.

$$\sum_{j} x_{ij} \leq 1 \qquad \forall\, i$$

2. Each load can be picked up only once:

$$\sum_{i} x_{ij} \leq 1 \qquad \forall\, j$$

3. If the load appointment window ends before the end of time period *t*, then this load has to be picked up during period *t*

$$x_{ij} = 1 \qquad \text{If } MaxPickup_j \text{ is before the end of the time period } t.$$

4. At least N_t loads should be served during time period *t*.

$$\sum_i x_{ij} \geq N_t$$

The last constraint is needed to guarantee that we move a minimum number of loads during each time interval. In our application, N_t is specified as the number of loads scheduled for pick up in time period t based on the beginning of their pickup window.

Post optimization

Based on the optimization results, we update the locations and availability times of trucks, by setting the location of truck *i* to the destination of load *j*.

$$Location_i = Destination_j$$

We also mark the truck as temporarily unavailable, until it has delivered the load to its destination, including the loading time, unloading time, deadhead time, and travel time:

$$TimeAvailable_i = max(TimeAvailable_i, MinPickupTime_j) + e_{ij} + LoadingTime + TravelTime_j + UnloadingTime$$

Appendix B: Repositioning & rebalancing optimizations

Objective

In this optimization, our goal is to rebalance trucks in preparation for the next period, or to achieve a desired distribution of supply and demand. By rebalancing the fleet, we send trucks, either loaded or empty, from oversupplied markets to undersupplied markets. Ideally, we would send these trucks loaded, however, this might not always be possible if loads are not available.

We define two sets of variables, x_{mn} and y_{mn} indicating the number of loaded and empty trucks respectively, to be dispatched from market m to market n. Our objective is to minimize the total empty miles traveled:

$$\min \sum_{m,n} y_{mn} d_{mn} \qquad \forall \text{ zones } m, n$$

Where d_{mn} is the average travel distance between markets m and n.

Constraints

1. We need to guarantee that enough trucks are available in each market to satisfy the demand at each time period. In each market, the number of available trucks, plus the number of inbound trucks, minus the number of outbound trucks, must meet or exceed the anticipated demand.

$$A_{mt} - \sum_{\text{all zones } n} y_{mn} - \sum_{\text{all zones } n} x_{mn} + \sum_{\text{all zones } n} y_{nm} + \sum_{\text{all zones } n} x_{nm} \geq D_{mt} \qquad \forall \text{ zones } m$$

Where A_{mt} is the number of trucks available in market m at time period t, and D_{mt} is the number of loads scheduled for pickup in market m at time period t.

2. The number of trucks we can dispatch from a given market (loaded or empty) should be less than the number of trucks available in that market.

$$\sum_{\text{all zones } n} y_{mn} + \sum_{\text{all zones } n} x_{mn} \leq A_{mt} \qquad \forall \text{ zones } m$$

3. The number of loaded trucks we can dispatch from market m to market n should be less than the number of loads available for pickup between these two markets (L_{mnt}) in time interval t.

$$x_{mn} \leq L_{mnt} \qquad \forall \text{ zones } m, n$$

Post optimization

Based on the optimization results, we update the locations and availability times of trucks, depending on whether they are loaded or empty. We update the locations and availability times of loaded trucks as before (see Appendix A).

Empty trucks are dispatched to load locations in the destination market (n), using an optimization formulation similar to that in Appendix A. Their arrival time is updated as follows:

$$TimeAvailable_i = TimeAvailable_i + e_{mn}$$

Where e_{mn} is the expected travel time between zones m and n (or more accurately, the travel time between the original truck location and its assigned load in the destination market).

Supply / Demand Rebalancing

This optimization problem is similar to truck repositioning. However, the number of trucks dispatched to/from each market should not only account for that market's demand in the next time interval. Instead, it should match a target distribution of trucks across different geographies. For example, this target distribution can be proportional to the number of load pickups in each market. Our simulation results show that by introducing this rebalancing step, we can achieve higher efficiency in the long-term. [28]

This can be achieved by adding the following constraint to the above formulation:

$$S_{mt} - \sum_{all\ zones\ n} y_{mn} - \sum_{all\ zones\ n} x_{mn} + \sum_{all\ zones\ n} y_{nm} + \sum_{all\ zones\ n} x_{nm} \geq T_m \qquad \forall\ zones\ m$$

Where S_{mt} is the number of trucks either in market m or on their way to market m at time period *t*, and T_m is the desired number of trucks in market m based on historical demand patterns.

Appendix C: Endnotes

1. US Federal Highway Administration: https://www.fhwa.dot.gov/reports/tswstudy/Vol3-Chapter9.pdf.
2. US Federal Highway Administration: https://www.fhwa.dot.gov/policyinformation/statistics/2018/pdf/vm1.pdf
3. National Security Council: https://injuryfacts.nsc.org/motor-vehicle/road-users/large-trucks/
4. According to EIA: https://www.epa.gov/greenvehicles/fast-facts-transportation-greenhouse-gas-emissions
5. Federal Motor Carrier Safety Administration (FMCSA): https://www.fmcsa.dot.gov/safety/data-and-statistics/commercial-motor-vehicle-facts
6. Private fleets make up 27% of the total number of carriers registered with FMCSA. They are owned by a shipper and typically only haul their owner's freight. In addition, 9% of carriers are classified as both private and for-hire carriers.
7. On the lower end of the spectrum, the American Transportation Research Institute (ATRI) estimates fleets' empty miles at 14.8%, and FreightWaves SONAR puts them at 16%. However, these estimates are mostly based on medium and large fleet surveys, which are more efficient than smaller fleets and owner-operators. On the other extreme, some studies argue that fleets significantly underreport their empty miles, which could be in excess of 35%.
8. Americas Commercial Transportation Research Company: Monthly Freight Forecasts (2023): https://www.actresearch.net/reports-data/forecasts/freight-rate-forecast.
9. National Private Truck Council, Benchmarking Survey Report (2023): https://www.nptc.org/benchmarking/benchmarking-report/
10. U.S. Department of Transportation, Federal Highway Administration, Highway Statistics (Washington, DC: annual issues), table VM-1, available at
11. This is estimated using the respective market shares of for-hire fleets (~$300 Bn), private fleets (~$300 Bn), and LTL ($70 Bn), according to ACT Research.
12. The freight activity of some of Uber Freight's shippers is concentrated in certain geographic areas, and on certain lanes. To account for such biases, we used a stratified sample of loads, based on shippers and shipper-lanes.
13. Class 8 trucks currently travel about 79,808 miles per year, according to ATRI. Similarly, ACT Research estimates loaded miles per tractor to be between 80,000 and 90,000 miles annually.
14. An integer program is a mathematical optimization program in which some or all of the variables are restricted to be integers. In this case, they can be variables representing whether or not a truck is assigned to a load.
15. Based on DAT's markets: https://www.dat.com/load-boards/market-condition-index.
16. We ran this optimization using the PuLP package in Python: https://coin-or.github.io/pulp/.
17. We assume that loads are picked up at the earliest time possible. If a truck is available at the start of the pickup window of a load, it will pick up this load at the beginning of the time window. Otherwise, it will pick up the load as soon as it becomes available.
18. We have analyzed different frequencies of the supply-demand rebalancing optimization (every 6 hours, every day, every week, etc.) and concluded that a weekly frequency produces the best results in terms of minimizing empty miles.
19. This is estimated as the difference between inbound and outbound loads (in absolute value), as a fraction the maximum of inbound and outbound loads multiplied by 2.
20. Our analysis does not consider intrinsic intra-market empty miles; as we add more supply and use finer market definitions, those empty miles would approach zero.
21. US DOT: Vehicle Inventory and Use Survey (VIUS, 2002): https://www.bts.gov/vius.
22. Unless a market's imbalance between inbound and outbound freight varies with time.

23. National Renewable Energy Laboratory (NREL), Transportation and Mobility Research. https://www.nrel.gov/transportation/fleettest-platooning.html
24. The time-based utilization rate refers to the percentage of time a truck is hauling freight. This is different from the distance-based utilization rate, which is the percentage of miles a truck is moving loaded. For example, reducing the available supply from 80,000 trucks to 33,000 will increase time-based utilization by 250%, while having a small effect on distance-based utilization.
25. According to research conducted by David Correll MIT: https://www.trucknews.com/transportation/truck-driver-delays-are-the-issue-not-a-driver-shortage-mit-researcher/1003155276/
26. Assuming 35% empty miles, and an average truck speed of 50 mph.
27. According to FMCSA, there were 1.64 fatalities in large truck crashes per 100 million miles driven: https://www.fmcsa.dot.gov/safety/data-and-statistics/commercial-motor-vehicle-facts
28. Our simulation results show that this proactive rebalancing can cut empty miles by more than 5% compared to only reactive truck repositioning.