Supply Chain Design and the Cost of Greenhouse Gas Emissions*

Gérard P. Cachon
The Wharton School
University of Pennsylvania
Philadelphia PA, 19104
cachon@wharton.upenn.edu
opim.wharton.upenn.edu/~cachon

August 25, 2011

Abstract

The design of a supply chain not only influences a retailer's cost, it influences its customers' costs - with stores few and far between, consumers must travel a long distance to shop, whereas their shopping trips are short with a dense network of stores. This supply chain design problem is the combination of two well studied problems: the Traveling Salesmen problem and the k-median problem. A solution is provided and the following question is asked: what is the consequence of failing to charge the full externality cost of greenhouse gas emissions? A bound is derived on the "carbon penalty", the gap between realized costs and minimal costs. A key measure in this bound is the carbon efficiency of the vehicles used in the supply chain, where carbon efficiency is the ratio of the carbon emitted per kilometer to the variable operating cost per kilometer. Given current estimates of carbon efficiency, it is found that the carbon penalty is small (0.1%). Thus, the supply chain design problem is robust to misspecifications in its inputted parameters - even if the potentially costly cost of carbon emissions is ignored, the chosen supply chain design may nevertheless yield a total cost (including emission costs) that is nearly as small as possible.

Keywords: Sustainability, k-median, traveling salesman problem

*The author thanks seminar participants for their comments at The University of Maryland, The Kellogg Operations Management Workshop, Stanford University, Harvard University, New York University, NSF Sponsored Symposium on Low Carbon Supply Chains, and the University of Chicago. Thanks is also extended to the following for their helpful discussions: Dan Adelman, Saif Benjaafar, Mark Daskin, Erica Plambeck, Lawrence Robinson, Sergei Savin, Michael Toffel, Beril Toktay, and Senthil Veeraraghavan. Financial support was provided by the Fishman-Davidson Center, University of Pennsylvania.

1 Introduction

The combustion of fossil fuels is believed to contribute to climate change by adding carbon dioxide and other greenhouse gases to the atmosphere (IPCC 2007). Transportation is a major source of emissions - approximately one third of U.S. emissions (EPA 2011). Factors that influence the emissions output of a transportation network include the fuel efficiency of the vehicles used, the size of the loads they carry and the distances they travel. The latter is influenced by the design of the supply chain. We focus on the downstream supply chain - the portion that includes inbound replenishments to retail stores and the "final mile" segment between the retail stores and consumers' homes. As the number of stores increases, consumers find themselves closer to some store, so they need not travel as far to purchase what they need. However, with more stores, the retailer must travel farther to replenish its stores. It is not *a priori* clear whether it is best to have a sparse network (which reduces the distance the retailer travels) or dense retail network (which reduces the distance consumers' travel).

We consider the supply chain design that minimizes several costs: mileage related variable vehicle operating costs, fuel consumption costs and emission costs. Emission costs are the externalities associated with adding greenhouse gases to the atmosphere, thereby leading to climate change, ocean acidification, sea level rise and other environmental changes that have potential economic consequences (EPA 2009a). While there is ample scientific evidence that greenhouse gas emissions are linked to these externalities, there is no consensus as to the true cost of greenhouse gas emissions (e.g., Cline, 1992; Shelling 1992). Nevertheless, they exist and may be substantial.

A model of the downstream supply chain is developed that accounts both for the distance consumers must travel to a retail store as well as the distance the retailer must travel to replenish those stores. Consumers and the retailer use different types of vehicles that have different operating costs, fuel consumption, and load size. We ask the following question: how poorly does a supply chain design perform if it is not optimized given the true cost of emissions? In other words, compare a supply chain designed to minimize non-emission costs and ask how much worse does it perform relative to the supply chain that minimizes all costs, including emission costs. Next, under what conditions is this emissions penalty substantial? And more precisely, do we expect the emission penalty to be substantial given current vehicle technology? If failing to account for the cost of emissions leads to a poor supply chain design (in the sense that the design's true costs are substantially greater than necessary) then this provides a justification for some mechanism, such as a carbon tax, to provide incentives for proper supply chain design choices.[1]

[1] We do not consider how the mechanism should be implemented, as this is a non-trivial issue. For

2 The Supply Chain Design Problem

The supply chain design problem in this paper focuses on a retailer's downstream supply chain. The retailer's task is to choose n store locations in a single (polygon) region of area a. Consumers live uniformly throughout the region. They incur a cost t_c per unit of distance travelled per unit of product purchased, and they travel a straight line (i.e., the Euclidian norm, L_2) to shop at the nearest store from their home. Let d_c be the average round-trip distance a consumer travels to a store. Thus, the average travel cost per consumer per unit purchased is $t_c d_c$.

The retailer has a single warehouse where it receives goods from an outside supplier. The warehouse is collocated with one of the n stores. The retailer has a single truck that is used to transport goods from the warehouse to the stores - all deliveries must start at the warehouse and end at the warehouse. Furthermore, the truck travels in straight lines between stores, delivers to all n stores on each route, and completes its route instantly (i.e., fast enough such that transit time is not a major issue). The retailer incurs a cost of t_r per unit of distance the truck travels per unit of product delivered. Let d_r be the length of the truck's route, so the retailer's transportation cost per unit sold is $t_r d_r$. The cost of inbound deliveries to the retailer is not considered.

There are several components included in the transportation costs of vehicle type t, $t \in \{c, r\}$, where "c" denotes the consumer's vehicle and "r" denotes the retailer's vehicle:

$$t_t = \frac{v_t + f_t\,(p_t + c_t e)}{q_t}$$

where v_t is the non-fuel variable cost to transport the vehicle per unit of distance (e.g., \$ km^{-1}); f_t is the amount of fuel used to transport the vehicle per unit of distance (e.g., $l\ km^{-1}$); p_t is the cost of fuel per unit of fuel (e.g., \$ l^{-1}); c_t is the amount of emission released by the consumption of one unit of fuel (e.g., kg); e is the price of emissions per unit released (e.g., \$ kg^{-1}); and q_t is the average load carried by the vehicle (e.g., kg). The variable cost, v_t, includes depreciation on the vehicle, maintenance (such as tire replacement) and other costs that can be linked to the distance the vehicle is driven. Fuel usage, f_t, is measured at the average load of the vehicle, which yields and accurate measure of true fuel usage when fuel usage is linear in the vehicle's load (which is approximately true), the vehicle always carries the same amount per trip, the vehicle makes deliveries at a constant rate (e.g., the stores on the route are equally distant from each other), and the vehicle travels at the same speed during the trip. The price of emissions, e, is assumed

example, even if the externality associated with emissions is known, the appropriate tax may not merely equal the externality when other taxes are used to fund government expenditures (see Baumol, 1972; Pearce 1991; and Bovenberg and de Mooij 1994).

to be independent of the source of the emissions.[2] This price is interpreted as the true cost of emissions including all externalities (e.g., climate change, ocean acidification, sea-level rise, etc.). Furthermore, for now, it is assumed that e is fully accounted for in the cost minimization problem. The next section allows for the possibility that the costs which are minimized do not fully reflect the true cost of emissions. We refer to e as the "cost of carbon".

Let Z be the average cost to transport a unit from the retailer's warehouse to a consumer's home:

$$Z = t_c d_c + t_r d_r.$$

The retailer's objective is to minimize Z by choosing (i) the shape of the region a, (ii) the number of stores in the region, (iii) the location of the stores within the region, (iv) and the route the truck takes to replenish the stores. The region a is constrained to be a single polygon. Consumer costs are included in the retailer's objective function because it is assumed that the retailer must compensate consumers for their travel costs. This is consistent with the notion that consumers have an outside alternative and the retailer must present a combination of store locations and prices that leave consumers indifferent between shopping at the retailer and choosing the alternative. It follows that the net utility of consumers is independent of the supply chain design (they pay high prices to shop at stores close to their home or lower prices to shop at stores far away, leaving them equally well off in either case). Thus, aggregate demand is constant. (As discussed later, out results are also robusts to shifts in aggregate demand.) Furthermore, in this model consumers fully recognize their travel costs: e.g., if it actually costs \$0.40 to drive a mile, then consumers recognizes that driving ten miles costs \$4 and require a \$4 price reduction to drive an extra ten miles.

Parts of the supply chain design problem are familiar. Given a set of n stores, the routing subproblem is the well-known Traveling Salesman Problem (TSP) - find a route through n locations, starting and ending at the same location, and visiting each location exactly once so as to minimize the total transportation cost. There exists an extensive literature on heuristics and solutions to the TSP (see Bramel and Simchi-Levi 1997; Lawler Lenstra, Rinnooy Kan and Shmoys 1985). More generally, there is a considerable literature on vehicle routing, such as when a fleet of vehicles (as opposed to a single vehicle) must be used to make deliveries to a set of known points in a region so as to minimize travel distances (e.g., Dantzig ang Ramser 1959; Daganzo 1984; Haimovich

[2] This is accurate in our application - a kg of CO_2 has the same effect whether its source is from a truck or a passenger vehicle. This may not be true in all cases. For example, the combustion of jet fuel, because it occurs high in the atmosphere, may not have the same effect as combustion at sea level.

and Rinnooy Kan 1985). This literature is further extended by work that includes inventory management along with vehicle routing (e.g., Federgruen and Zipkin 1984; Burns, Hall, Blumenfeld, Daganzo 1985; Gallego and Simchi-Levi 1990). The key differences between this supply chain design problem and the TSP and its extensions are (i) the retailer can choose the location of the stores and (ii) the retailer accounts for the consumers' transportation costs.

Ignoring the retailer's transportation costs, the problem is analogous to the well known $k-$median problem (which is also referred to as the $p-$median problem or multi-source Weber problem or location-allocation problem). In the $k-$median problem there exists a set of demand locations. The objective is to choose k locations - call them stores - to minimize the total transportation cost from the demand locations to their nearest store. The k - median problem is generally studied in its discrete form (i.e., a finite number of possible demand and store locations) but there has also been some work on the continuous k $-$median problem, which is the retailer's supply chain design problem when the retailer's transportation costs are ignored (e.g., by assuming $t_r = 0$). Work on the $k-$median problem has focused on good solution procedures, rather than on the structure of the solution. See Daskin (1995) for an overview of the $k-$median problem. Brimberg, Hansen, Mlandinovic, Taillard (2000) study numerous solution algorithms for the discrete $k-$median problem and Fekete, Mitchell and Beurer (2005) do the same for the continuous version of the problem.

The supply chain design problem, which is the combination of the TSP and the $k-$median problem, has not been previously studied. This problem extends the traditional boundary of supply chain analysis, which typically incorporates just the firm, to include the final leg of transportation performed by consumers. The focus is on how the design of the supply chain influences emissions through variations in fuel consumption given a fixed type of transportation mode (e.g., passenger vehicles for consumers and a truck for the retailer).

There has been some work on the interaction between operational decisions and emissions. Hoen, Tan, Fransoo and van Houtum (2010) analyze the transportation mode selection problem for a single location given that modes vary in their speed (which influences inventory holdings) and their emission levels. Benjaafar, Li and Daskin (2010) analyze a single location model in which inventory management decisions (the timing and quantity of orders) influence supply chain holding costs, backorder costs and emissions. Their model includes the following emissions: (i) a fixed amount per unit held on average in inventory; (ii) a fixed amount per unit sold and (iii) a fixed amount per delivery. Section 7 includes order quantity decisions, but our model only includes emissions due to distances traveled (which Benjaafar, Li and Daskin 2010 do not include). Consequently, the amount emitted per unit in our model depends on the quantity delivered and the distance traveled per delivery. Gillerlain, Fry and Magazine (2011) also study a model in which a retailer and consumers incur transportation costs and potentially emission costs based on distances. They too evaluate the cost minimizing number of stores, but they impose a different spacial geometry - in their model, stores and consumers are located on the boundary of a circle.

3 Parameter Estimates

This section provides estimates for the parameters in the supply chain design problem.

The average mileage of passenger vehicles in the U.S. in 2009 is 21.1 $miles\ gal^{-1}$ or 8.97 $km\ l^{-1}$ (EPA 2009). Fuel consumption is then $f_c = 1/8.97 = 0.111\ l\ km^{-1}$. On the assumption that the retailer uses a truck that achieves 6 miles per gallon of diesel, the retailer's fuel consumption is $f_r = 0.392\ l\ km^{-1}$.[3] In March 2011, the national average price per gallon of gasoline and diesel, respectively, were $p_c = 0.94$ \$ l^{-1} and $p_r = 1.03$ \$ l^{-1} (DOE 2011a,b). The fuel costs to drive these vehicle are $f_c p_c = 0.111 \times 0.94 = 0.104$ \$ km^{-1} and $f_r p_r = 0.392 \times 1.03 = 0.404$ \$ km^{-1}.

From EPA (2005), 8.8 kg of CO_2 are emitted per gallon of gasoline, or $c_c = 2.325\ kg\ CO_2\ l^{-1}$.[4] From the same report 10.1 kg of CO_2 are emitted per gallon of diesel, or $c_r = 2.669\ kg\ CO_2\ l^{-1}$.

The variable cost to drive the vehicle each km is v_c. Variable costs include mileage related depreciation on the vehicle and maintenance costs (including tire replacement) that can be linked to the distance/time the vehicle is driven. It is difficult to obtain a reliable estimate of this cost. Some costs are primarily fixed and do not depend on the distance the vehicle is driven (such as licensing fees, taxes or insurance). Some costs are primarily linear in distance, such as tire usage and some maintenance repairs (e.g. a timing belt that must be replaced at 70,000 miles). Others costs can be classified as either fixed or variable, such as depreciation: a vehicle depreciates the

[3] The Bureau of Transportation Statistics (2011) Table 4-14 reports that truck fuel consumption in the United States was 5.9 in 2004 but dropped to 5.4 by 2008. Bonney (2009) reports that Walmart's fuel efficiency in 2005 was 5.9 miles per gallon, which increased to 7.1 miles per gallon by 2008. Walmart reports, in their logistics fact sheet, (http://walmartstores.com/pressroom/factsheets/), that they saved 15,000,000 gallons of diesel fuel by driving 87,000,000 million fewer miles, which works out to 5.8 miles per gallon. The Rocky Mountain Institute (http://www.rmi.org/rmi/Trucking) reports that the average long-haul truck's fuel efficiency in the United States is 6 miles per gallon.

[4] CO_2 represents the majority of greenhouse gas that is emitted from the burning of gasoline (95%). It is easy to include all greenhouse gases in this analysis but the results are unlikely to change.

more it is driven, but a vehicle also depreciates even if it is not driven at all (e.g., possibly because of technological advances or changes in style).

From American Automobile Association (2011), the average cost per kilometer to drive a sedan includes \$0.0278 for maintenance, \$0.006 for tires, and \$0.0283 for distance related depreciation (in \$ km^{-1})[5] for a total variable cost of v_c = \$0.0621 km^{-1}. Barnes and Langworthy (2003) provide an estimate of v_c = \$0.0681 km^{-1} in 2003 dollars. Adjusting for inflation by 3% per year, yields a 2010 estimate of v_c = \$0.0837, km^{-1} which is comparable to the measure provided by the AAA. The AAA focuses on the first five years of vehicle life, so it is expected to underestimate the variable cost of operating the existing fleet, which has an average age of approximately eight years. We use v_c = \$0.0837, km^{-1}.

For trucks, Barnes and Langworthy (2003) estimate the variable cost is \$0.169 km^{-1} (adjusted to 2010 dollars), but this does not include driver wages. If driver wages are included, at \$50,000 per year and 100,000 miles, their estimate is v_r = \$0.479 km^{-1}, which is what we use. Table 1 summarizes the parameter estimates

Table 1: Parameter estimates

	v_t	f_t	p_t	c_t
Consumers	0.0837	0.111	0.94	2.325
Retailer	0.479	0.392	1.03	2.669

The load a consumer carries can vary considerably (e.g., from a kilogram for a article of clothing, to over a hundred kilograms for home repair supplies). For the sake of an estimate, say the consumer transports $q_c = 40\ kg$ worth of goods. (Note, the relevant weight is the weight of the goods purchased, as the weight of passengers is not included.) For the retailer, a tractor-trailor's average load is about $q_r = 18,000\ kg$. Given those loads and the other parameter estimates, Table 2 provides the per unit transportation costs for various costs of carbon. In all cases, the retailer's truck is about two orders of magnitude more efficient than the consumer's vehicle (in terms of cost per kilogram of product transported one kilometer). The retailer's truck may only get 6 miles per gallon, but given its much larger load, it is considerably more efficient.

[5] Three observations are provided for the annual depreciation cost {10000,\$3471}, {15000,\$3728} and {20000,\$3924}, where the first term is the miles driven during the year and the second is the total depreciation cost. A linear regression through these observations yields a slope of \$0.0453 per mile.

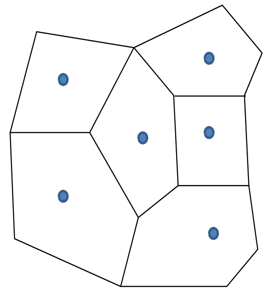

Figure 1: Voronoi diagram depicting the service regions for a six store configuration.

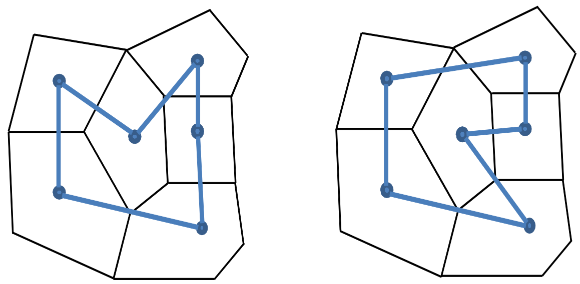

Figure 2: Two possible delivery routes through a configuration of six stores.

Table 2: Estimates of per unit transportation costs (\$ km^{-1})

	Cost of carbon, e			
	0	50	100	1000
t_c	0.004693	0.327	0.650	6.46
t_r	0.000049	0.00295	0.00586	0.0582
t_c/t_r	95.7	110.8	110.9	111.0

4 Problem analysis

Given a set of n store locations, the region of customers can be partitioned into n sub-regions that represent the stores' "service areas", i.e., all customers in a service area shop at the store in their service area because that store is the closest to them among the n available choices. This partitioning is also called a Voronoi diagram. Figure 1 displays one possible partitioning: In addition to store locations, the retailer must choose a TSP route to minimize its transportation costs. Figure 2 presents two possible routes in the Figure 1 partitioning.

The solution to the supply chain problem is unknown and may involve a complex geometry, especially given a fixed polygon of area a to serve. Therefore, we consider store configurations in which the resulting Voronoi diagram is a tesselation of a single regular polygon. Figure 3 displays a tiling with equilateral triangles. There are two other feasible tilings that consist of a single regular

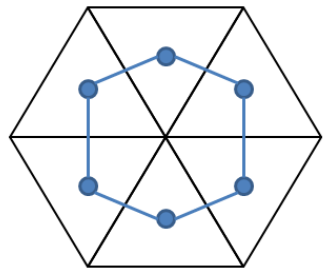

Figure 3: Triangle tesselation of stores

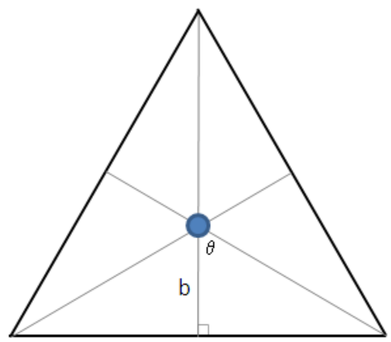

Figure 4: Sub-regions within a store's single service area given a triangle tesselation

polygon, one with squares and the other with hexagons. Therefore, for a given area a, the supply chain design problem reduces to the choice of the number of stores, n, and the tiling pattern: the stores are located in the center of each tile; and it is straightforward to find an optimal TSP route through these tiling options.

To evaluate d_c and d_r, consider one of the service regions. Within that region create a tesselation of right triangles by connecting the store to each vertex and to each face via a line perpendicular to its edge, to form sub-regions. Figure 4 demonstrates this for the triangle tesselation.

In a regular polygon with s sides, an "$s-$gon", there are $2s$ sub-regions. Consider one of the (identical) sub-regions. Let θ be the degree measure (in radians) of the acute angle formed at the vertex where the store is located and let b be the shortest distance to the edge of the sub-region (i.e., b is the height of the triangle in Figure 4). For an s-gon, $\theta = \pi/s$. The average round-trip distance of customers within a sub-region to the store is:

$$d_c(b,\theta) = 2\frac{\int_0^b dx \int_0^{x\tan\theta} \sqrt{x^2+y^2}dy}{(1/2)b^2\tan\theta}$$

where the denominator is the area of the sub-region. Using a change of variables ($t = y/x$), we obtain

$$d_c(b,\theta) = \frac{2}{3}\left(\sqrt{1+(\tan\theta)^2} + \frac{\ln\left(\tan\theta + \sqrt{1+(\tan\theta)^2}\right)}{\tan\theta}\right) b.$$

The area of the service region is a/n and the area of the sub-region is then $a/(2sn)$. Thus, $(1/2)b^2\tan\theta = a/(2sn)$, and $d_c(b,\theta)$ can be written as

$$d_c = \phi_c (a/n)^{1/2} \tag{1}$$

where ϕ_c is a constant that depends on the tesselation (i.e., the number of sides, s):

$$\phi_c = \frac{2}{3}\sqrt{\frac{1}{s\tan\theta}}\left(\sqrt{1+(\tan\theta)^2} + \frac{\ln\left(\tan\theta + \sqrt{1+(\tan\theta)^2}\right)}{\tan\theta}\right)$$

The retailer's truck must travel into and out of each region at least once (assuming there are at least two stores, $n \geq 2$). The truck's minimum distance within a sub-region is $2b$, so the length of the TSP tour is no shorter than $2bn$. Hence, a lower bound estimate for the retailer's transport distance is

$$d_r = \phi_r (an)^{1/2} \tag{2}$$

where the constant ϕ_r also only depends on the chosen tesselation:

$$\phi_r = \frac{2}{\sqrt{s\tan\theta}}.$$

The estimate (2) is exact for triangles with $n \in \{2, 6, 10, ...\}$, for squares with $n \in \{2, 4, 6, ...\}$ and for all n with hexagons.[6]

Given (1) and (2), the supply chain cost function can be written as

$$Z(n) = a^{1/2}\left(\phi_c t_c n^{-1/2} + \phi_r t_r n^{1/2}\right) \tag{3}$$

This result is consistent with several studies of probabilistic versions of the k-median and TSP problems. For example, Fisher and Hochbaum (1980) consider the k-median problem of selecting n store locations from a set of randomly chosen sites to minimize the total distance consumers must travel to the closest of the n stores. They find that the value of the optimal cost grows proportional to $\sqrt{1/n}$, as in (3). For the TSP, Beardwood, Halton and Hammersley (1959) show that the shortest distance through n randomly selected points in an area a is asymptotically proportional to $\sqrt{an}$, again, as in (3).

[6] With triangles an upper bound on the TSP distance is $2b(n+1)$, which is exact for $n \in \{4, 8, 12, ...\}$. For $n \in \{3, 5, 7, 9, ...\}$ the TSP distance is $2b(n+\sqrt{3}-1)$. With squares, the TSP distance is $2b(n+\sqrt{2}-1)$ for odd n.

Table 3 summarizes the constants in the cost function.

Table 3: Distance coefficients for different tesselations

Tesselation	ϕ_c	ϕ_r	$\phi_c\phi_r$	ϕ_c/ϕ_r
triangle	0.807293	0.877383	0.708	0.920
square	0.765196	1.000000	0.765	0.765
hexagon	0.754393	1.074570	0.811	0.702

For a given number of stores (n) and a fixed region size (a), consumers must travel the farthest with the triangle tesselation (that has the highest ϕ_c), while the firm must travel the farthest with the hexagon tesselation. In all cases, consumers on average travel a shorter distance than the retailer ($\phi_c < \phi_r$) despite the fact that some consumers must travel farther (e.g., as much as twice as far with the triangle tesselation). Consequently, if it were equally costly for consumers and the retailer to transport goods, then one would expect the optimal tesselation to have few stores, forcing consumers to drive long distances.

If we allow n to be non-integer, minimization of $Z(n)$ is straightforward, given that it is quasi-convex in n. Let n^* be the cost minimizing number of stores:

$$n^* = \frac{\phi_c t_c}{\phi_r t_r}. \tag{4}$$

From Table 3, $\phi_c/\phi_r < 1$, so $n^* > 2$ only if $t_c >> t_r$, which is likely to hold according to Table 2.

The minimum cost is

$$Z\left(n^*\right) = 2\sqrt{a\phi_c\phi_r t_c t_r}.$$

Thus, for a fixed set of transportation cost coefficients, $\{t_c, t_r\}$, according to Table 3, costs are minimized when the retailer uses a triangle tesselation (because that minimizes $\phi_c\phi_r$), an intermediate solution is obtained with a square tesselation and costs are maximized (among these three) when the firm uses a hexagon tesselation.[7] In fact, cost with the triangle tesselation is is about 7% lower than with hexagons (1-$\sqrt{0.708/0.811}$). It is also possible to show that $\phi_c\phi_r$ is decreasing for all $s \geq 3$, suggesting that the triangle tesselation may perform well even with tesselations that include more than one regular polygon.

Even though consumers travel the farthest with the triangle tesselation, the triangle tesselation is best among the regular polygons no matter the ratio of transportation cost parameters, t_c/t_r. This occurs because the relative transportation costs, t_c/t_r, interacts with the configuration parameters, $\phi_c\phi_r$, multiplicatively. Hence while the absolute cost clearly depends on the transportation parameters, the relative cost across the various tesselations does not.

[7] When exact values of the TSP are utilized, triangles are best (among the three) for $n = 2$ and $n \geq 5$, hexagons are best for $n = 3$ and squares are best for $n = 4$.

The distances d_c and d_r are measured by the Euclidian norm, L_2. Another common norm is L_1, which is sometime referred to as the "Manhattan" norm or the "city-block" norm in which the distance between points $\{x_1, y_1\}$ and $\{x_2, y_2\}$ is taken to be $|x_2 - x_1| + |y_2 - y_1|$, i.e., travel occurs along a square grid. It is possible to show that for consumers, even with the L_1 norm, their round-trip distance to the nearest store is proportional to $(a/n)^{1/2}$. With the L_1 norm, the retailer's travel distance is easiest to estimate with the square tesselation, and in that case the distance continues to be proportional to $(an)^{1/2}$. Hence, our results do not appear to be sensitive to how distance is measured.

5 Supply chain design and the cost of carbon

The optimal number of stores can be written as

$$n^* = \frac{\phi_c t_c}{\phi_r t_r} = \left(\frac{\phi_c q_r}{\phi_r q_c}\right)\left(\frac{v_c + f_c p_c}{v_r + f_r p_r}\right)\left(\frac{1+\beta_c e}{1+\beta_r e}\right)$$

where the constants β_t, $t \in \{c, r\}$, are introduced for notational convenience and represent the ratio of carbon emissions to variable operating costs:

$$\beta_t = \frac{f_t c_t}{v_t + f_t p_t}.$$

Using the estimates in Table 1, we obtain, $\beta_c = 1.333$ and $\beta_r = 1.204$.

The constants, β_t, dictate the relationship between the number of stores, n^*, and the cost of carbon, e. As carbon becomes more expensive, the number of stores increases if $\beta_c > \beta_r$, otherwise it decreases:

$$\frac{dn^*}{de} = \left(\frac{\phi_c}{\phi_r}\right)\left(\frac{q_r}{q_c}\right)\left(\frac{v_c + f_c p_c}{v_r + f_r p_r}\right)\left(\frac{\beta_c - \beta_r}{(1+\beta_r e)^2}\right).$$

When carbon becomes more expensive, stores should move closer to consumers (i.e., there should be more of them) when the ratio of the amount emitted per unit of distance traveled to (non-emission) variable costs per unit of distance traveled, β_t, is greater for consumers than for the retailer, as it appears to be ($\beta_c = 1.333$ vs. $\beta_r = 1.204$). Put another way, if the emission cost per kilometer traveled ($f_t c_t e$) relative to the total variable cost per kilometer traveled ($v_t + f_t p_t + f_t c_t e$) is higher for consumers than for the retailer, then add more stores if carbon becomes more expensive, otherwise there should be fewer stores. It is interesting to note that the cost that matters is the cost to move the vehicle and not the cost to move one unit of product (t_c and t_r), i.e., the quantities carried do not influence whether the supply chain should have more or fewer stores as the cost of carbon changes. Again, this is because the quantity ratio, q_r/q_c, appears multiplicatively in the solution to n^*. For a similar reason, nor does the particular tesselation matter: whether there should be more or fewer stores as carbon becomes more expensive does not depend on whether the tesselation is with triangles, squares, or hexagons.

6 The carbon penalty

It is possible that the true cost of carbon, e, may not be fully included in the supply chain design decision - the retailers and consumers may not be charged the full societal cost of carbon, but rather charged a portion of it. To be specific, say they are actually charged αe, $\alpha \in [0, 1)$, rather than e and decisions are based on that cost rather than the full cost of carbon. Consequently, a supply chain design is chosen that does not minimize the true total cost. Let $n^*(\alpha)$ be that choice:

$$n^*(\alpha) = \left(\frac{\phi_c}{\phi_r}\right)\left(\frac{t_c(\alpha)}{t_r(\alpha)}\right)$$

where for $t \in \{c, r\}$

$$t_t(\alpha_t) = \left(\frac{v_t + f_t p_t}{q_t}\right)(1 + \alpha\beta_t e).$$

Hence, the realized cost is $Z(n^*(\alpha))$ whereas the minimum total cost is $Z(n^*)$. The gap between $Z(n^*(\alpha))$ and $Z(n^*)$ is referred to as the "carbon penalty": it is a measure of the increase in cost that occurs due to not fully charging for the true cost of emissions. The carbon penalty can be measured by the following ratio:

$$\frac{Z(n^*(\alpha)) - Z(n^*)}{Z(n^*)} = \frac{1}{2}\left(\gamma(\alpha)^{1/2} + \frac{1}{\gamma(\alpha)^{1/2}}\right) - 1$$

where

$$\gamma(\alpha) = \frac{(1 + \beta_c e)(1 + \alpha\beta_r e)}{(1 + \beta_r e)(1 + \alpha\beta_c e)}.$$

There is no penalty when $\alpha = 1$, because then $Z(n^*(\alpha)) = Z(n^*)$. However, if carbon is not fully charged, $\alpha < 1$, then the carbon penalty is positive. Theorem 1 provides an upper bound for the cost penalty of not fully charging for carbon. As one might expect, the bound is obtained when there is no charge for carbon, $\alpha = 0$, but the true cost of carbon is very high ($e \to \infty$).

Theorem 1 *The carbon penalty, $(Z(n^*(\alpha)) - Z(n^*))/Z(n^*)$, is no greater than*

$$\frac{1}{2}\left((\beta_c/\beta_r)^{1/2} + \frac{1}{(\beta_c/\beta_r)^{1/2}}\right) - 1 \tag{5}$$

Proof. The ratio $Z(n^*(\alpha))/Z(n^*)$ is quasi-convex in γ, minimized at $\gamma = 1$ and at the minimum, $Z(n^*(\alpha))/Z(n^*) = 1$, i.e., at that point there is no penalty. If $\beta_c = \beta_r$, then $\gamma = 1$ for all α, i.e., no matter whether carbon if fully priced or not there is no penalty, $Z(n^*(\alpha))/Z(n^*) = 1$. However, there will exist a cost penalty ($Z(n^*(\alpha))/Z(n^*) > 1$) when $\beta_c <> \beta_r$. The size of that penalty depends on how far $\gamma(\alpha)$ deviates from 1:

$$\lim_{\gamma(\alpha)\to 0} = \lim_{\gamma(\alpha)\to\infty} = \infty.$$

It is straightforward to show that for all $\alpha \in [0, 1)$:

$$\begin{array}{ll} \gamma(\alpha) < 1 & \text{if } \beta_c < \beta_r \\ \gamma(\alpha) > 1 & \text{otherwise} \end{array}$$

and

$$\frac{\partial \gamma(\alpha)}{\partial \alpha} = \frac{(1+\beta_c e)}{(1+\beta_r e)} \frac{(\beta_r - \beta_c)\, e}{(1+\alpha\beta_c e)^2}.$$

Thus, for a given e, $Z\left(n^*\left(\alpha\right)\right)/Z\left(n^*\right)$ is maximized when $\alpha = 0$, i.e., when there is no charge for carbon emissions: if $\beta_c < \beta_r$, then $Z\left(n^*\left(\alpha\right)\right)/Z\left(n^*\right)$ is maximized by decreasing γ, which occurs with $\alpha = 0$ because $\partial\gamma(\alpha)/\partial\alpha > 0$; if $\beta_c > \beta_r$, then $Z\left(n^*\left(\alpha\right)\right)/Z\left(n^*\right)$ is maximized by increasing γ, which occurs with $\alpha = 0$ because $\partial\gamma(\alpha)/\partial\alpha < 0$. For $\alpha = 0$ we have

$$\gamma(0) = \frac{1+\beta_c e}{1+\beta_r e}; \qquad \frac{\partial \gamma(0)}{\partial e} = \frac{\beta_c - \beta_r}{(1+\beta_r p)^2}.$$

If $\beta_c < \beta_r$, then $\gamma(0) < 1$ and $\gamma(0)$ is minimized (which maximizes $Z\left(n^*\left(\alpha\right)\right)/Z\left(n^*\right)$) when $e \to \infty$. Similarly, if $\beta_c > \beta_r$, then $\gamma(0) > 1$ and $\gamma(0)$ is maximized (which maximizes $Z\left(n^*\left(\alpha\right)\right)/Z\left(n^*\right)$) when $e \to \infty$. In either case

$$\lim_{e\to\infty} \gamma(0) = \frac{\beta_c}{\beta_r},$$

which yields the bound in the theorem.□

Failing to charge for the total cost of carbon leads to a distortion in the supply chain design. It is also possible to consider the cost penalty when the supply chain is intentionally distorted to minimize emissions. Let n_e^* be the optimal decision if the objective is exclusively to minimize total emissions. It is straightforward to evaluate

$$n_e^* = \frac{\phi_c f_c c_c q_r}{\phi_r f_r c_r q_c} = \left(\frac{1+\beta_r e}{1+\beta_c e}\right) n^*.$$

The cost penalty for minimizing emissions (rather than total costs) is reflected in the ratio

$$\frac{Z\left(n_e^*\right) - Z\left(n^*\right)}{Z\left(n^*\right)} = \frac{1}{2}\left(\sqrt{\frac{1+\beta_r e}{1+\beta_c e}} + \sqrt{\frac{1+\beta_c e}{1+\beta_r e}}\right) - 1.$$

The upper bound of that ratio is the same as the upper bound (5):

$$\max_e \frac{Z\left(n_e^*\right) - Z\left(n^*\right)}{Z\left(n^*\right)} = \frac{1}{2}\left((\beta_r/\beta_c)^{1/2} + \frac{1}{(\beta_r/\beta_c)^{1/2}}\right) - 1$$

Given the parameter estimates in Table 1, the upper bound in Theorem 1 yields

$$\frac{1}{2}\left((\beta_c/\beta_r)^{1/2} + \frac{1}{(\beta_c/\beta_r)^{1/2}}\right) - 1 = 0.001.$$

Thus, using current estimates of the various variable costs for the two vehicle types, our main result is that failing to explicitly charge for carbon yields a relatively modest penalty in terms of supply chain design - even if carbon emissions are extremely costly (in terms of the externalities generated), and, nevertheless, there is no explicit charge for carbon, the distortions in supply chain design leads to an increase in total costs of less than 0.1%, a relatively modest amount. Similarly, relative to the current supply chain configuration, if the supply chain were configured to minimize emissions, total costs would increase by less than 0.1%. This is remarkable given the substantial differences in the vehicles' emissions per kg per km measures. (Each kg hauled by the retailer's truck emits about 100 times less carbon per km than a consumer's vehicle.) One might assume that, given the retailer's carbon advantage, a failure to fully account for the cost of carbon would lead to a substantial distortion - ignoring carbon would lead to too few stores, which are too far from consumers. However, load sizes do not matter here. (Or, more precisely, they already factor into n^* and $n^*(\alpha)$ multiplicatively whether carbon is charged properly or not.) What matters is the carbon emitted per km traveled relative to the variable costs, β_r and β_c. Given that these parameters are reasonably close to each other, the carbon penalty is small.

A large carbon penalty requires that β_c and β_r deviate substantially from each other. This could occur if either the retailer or consumers adopt a different technology, thereby accentuating the gap between β_c and β_r. For example, if all consumers drove vehicles that achieved 48 miles per gallon (such as a hybrid vehicle), consumer fuel consumption would decline to $f_c = 0.049\ l\ km^{-1}$. If v_c remains unchanged, then the new, more efficient, fuel consumption leads to $\beta_c = 0.843$ and the ratio β_c/β_r decreases from 1.1 to 0.70. The upper bound on the carbon penalty would then increase, but only to 1.6%. A greater carbon penalty would occur if the retailer were to adopt more fuel efficient vehicles. For example, if the retailer were able to cut fuel consumption in half (while consumer fuel efficiency remains the same), the carbon penalty bound increases to 3.5% - still a relatively small penalty.

It is possible that errors in our parameter estimates are leading to a small carbon penalty. The greatest uncertainty with respect to the β_c and β_r estimates is probably due to the variable costs of operating the vehicles. Holding the other parameters constant, Table 4 displays the carbon penalty bound for various adjustments to v_c and v_r :

Table 4: Carbon penalty bound (as a %) for various adjustments to variable operating costs. Column and row headings are the true variable operating costs for consumers and the retailer, respectively. v_c and v_r are listed in Table 1.

	$0.10v_r$	$0.20v_r$	$0.50v_r$	$1.00v_r$	$2.00v_r$	$5.00v_r$	$10.00v_r$
$0.10v_c$	0.0	0.1	1.3	5.1	14.6	42.4	80.9
$0.20v_c$	0.1	0.0	0.8	4.0	12.5	38.7	75.4
$0.50v_c$	1.1	0.5	0.0	1.7	8.0	29.8	61.9
$1.00v_c$	3.9	2.6	0.6	0.1	3.5	20.1	46.6
$2.00v_c$	11.1	8.9	4.4	1.0	0.3	9.6	28.8
$5.00v_c$	33.8	29.6	20.4	11.4	3.3	0.5	8.4
$10.00v_c$	66.4	60.1	45.8	31.0	15.7	1.9	0.6

The largest penalty, 80.9%, occurs if the current consumer cost is overestimated by a factor of 10 and the retailer's current cost is underestimated by a factor of 10 as well - these are rather large deviations from the current estimates. Given that our current estimates has $\beta_c > \beta_r$, the largest penalties occur if the true value of β_c is even larger or if the true value of β_r is even smaller (so as to accentuate the gap between β_c and β_r).

While Table 4 provides a bound on the carbon penalty, it is possible to evaluate the exact carbon penalty for specific carbon prices. Table 5 considers $e = \$50$, a rather low estimate of the true cost of carbon:

Table 5: Actual carbon penalty (as a %) when $e = \$50$ for various adjustments to variable operating costs. Column and row headings are the true variable operating costs for consumers and the retailer, respectively. v_c and v_r are listed in Table 1.

	$0.10v_r$	$0.20v_r$	$0.50v_r$	$1.00v_r$	$2.00v_r$	$5.00v_r$	$10.00v_r$
$0.10v_c$	0.0	0.1	1.3	5.0	14.1	40.3	74.8
$0.20v_c$	0.1	0.0	0.8	3.9	12.1	36.7	69.5
$0.50v_c$	1.1	0.5	0.0	1.6	7.7	28.2	56.9
$1.00v_c$	3.8	2.6	0.5	0.1	3.4	18.9	42.5
$2.00v_c$	10.8	8.6	4.2	0.9	0.3	8.9	26.0
$5.00v_c$	32.4	28.3	19.4	10.8	3.1	0.5	7.4
$10.00v_c$	62.1	56.1	42.6	28.6	14.3	1.7	0.5

Comparison of Tables 4 and 5 suggests that the carbon penalty bound, (5), provides a reasonably good approximation of the actual carbon penalty even if the true cost of carbon is reasonably low, $e = 50$.

Overall, the main conclusion from this analysis is that the supply chain design problem is highly robust given current technologies. Remarkably, even if the true cost of carbon is very high, while failing to account for this cost does distort decisions, the distortions are unlikely to be substantial enough to make a significant difference in terms of costs. Furthermore, this result is likely to be unaffected by shifts in aggregate demand - the distortion in supply chain (i.e., the gap between n^* and $n^*(\alpha)$) neither depends on the level of aggregate demand nor the load quantities.

This robustness finding is reminiscent of results developed for the well-known EOQ problem that studies the trade-off between fixed ordering costs and inventory holding costs. In the EOQ problem, costs are the sum of two terms, one linearly increasing in the decision variable, call it x, and the other decreasing proportional to $1/x$. In the supply chain design problem with fixed quantities, costs are again the sum of two terms, but now one is linearly increasing in $\sqrt{x}$ and the other decreasing as $1/\sqrt{x}$. It has been established that the EOQ problem has a flat objective function (see Dobson 1988, Porteus 2002), but it follows that the objective function in the supply chain design problem is even flatter.

So far the carbon penalty has been measured assuming that neither the retailer nor consumers are fully charged the price of carbon. However, what if the retailer is fully charged, e.g., because the tax on diesel fuel is increased to reflect the externality it creates, but consumers are charged only αe, $\alpha \in [0, 1)$ for their carbon emissions?[8]

Following the earlier methodology, the carbon penalty is

$$(1/2)\left(\left(\frac{1+\beta_c e}{1+\alpha\beta_c e}\right)^{1/2} + \left(\frac{1+\beta_c e}{1+\alpha\beta_c e}\right)^{-1/2}\right).$$

In this situation it is possible to envision substantial carbon penalties. For example, suppose $e = \$50$, $\beta_c = 1.33$, and the retailer is charged this cost but consumers are not, i.e., $\alpha = 0$. In this case, the carbon penalty is 317%. However, the carbon penalty does decrease quickly as the consumer's portion of the cost increases as revealed in Table 6.

Table 6: Actual carbon penalty (as a %) when $e = \$50$, and the retailer is charged the full cost of carbon but consumers are charged only α percent of e given $\beta_c = 1.33$.

α	0	0.1	0.2	0.3	0.4	0.5	0.6
Carbon penalty	317%	65%	32%	18%	10%	6%	3%

The conclusion from this analysis is that a substantial carbon penalty can occur if there are substantial asymmetries in the system, such as a large tax on diesel and none on gasoline. However, the effects of such asymmetry are likely to be dampened if the retailer or consumers can change

[8] An alternative interpretation is that the retailer fully accounts for its costs, but consumers only recognize α percent of their actual cost. For example, although it may cost a consumer an additional \$5 to drive a longer distance, the consumer acts as if she only incurs \$1 in additional expense. This may be due to mental accounting on the part of consumers: the cost of driving their vehicle is not associated with the cost of procuring goods from the retailer.

their technology: if diesel becomes substantially more expensive relative to gasoline, it is likely that users of diesel will search for an alternative fuel (i.e., it may be unrealistic to assume that large disparities between the cost of the two fuels can exist). Of greater concern is an asymmetry due to decision making errors. For example, if the retailer and consumers are equally charged for carbon but consumers only account for a small fraction of that cost.

7 Delivery frequency

In the model presented in sections 2-6, consumers do not change the frequency of their store visits nor does the retailer change the frequency of deliveries, i.e., q_c and q_r are fixed. Now suppose consumers choose q_c and the retailer chooses q_r in addition to the location and number of stores, n. Consequently, each consumer visits a store every q_c/λ_c units of time, where λ_c is the consumer's consumption rate. The retailer sends a delivery to all of the stores every q_r/λ_r units of time, where λ_r is the aggregate consumption rate of consumers in the region of area a.

First, consider the q_c choice for consumer i, whose round-trip distance to the nearest store is d_i. As before, the consumer incurs cost $t_c d_i$ per unit purchased.[9] It is convenient to write $t_c = \tau_c/q_c$, where

$$\tau_c = v_c + f_c\left(p_c + c_c e\right).$$

If the consumer only considered the cost per trip, the consumer would purchase as large a quantity as possible so as to amortize the fixed cost per delivery over more units. However, purchasing large quantities also requires a large storage area, which is costly as well. In particular, say consumers incur costs at the rate h_c to maintain enough storage for one unit, so the average storage cost is $h_c q_c$ per unit of time. Consumer i's problem is then to minimize the average cost per unit, $C_c(q_c)$:

$$\min_{q_c} C_c(q_c) = \frac{h_c q_c}{\lambda_c} + \frac{\tau_c d_i}{q_c}.$$

The cost minimizing quantity is $q_c^* = \sqrt{\left(\tau_c \lambda_c / h_c\right) d_i}$: as the consumer has to travel farther to a store, the consumer purchases a larger quantity with each visit. However, doubling the distance to a store does not reduce the frequency of visits by a factor of two. Instead, it reduces the frequency of visits by only $\sqrt{2}$ because it is not optimal to double storage costs.

Given the consumer i's optimal purchase quantity, the consumer's cost per unit is

$$C_c(q_c^*) = 2\sqrt{\left(h_c \tau_c / \lambda_c\right) d_i}.$$

[9] In a different context, Hall (1992), also considers a model in which the decision maker chooses a delivery frequency to minimize costs and the cost to travel to a store is a fixed cost associated with each delivery.

In the base model, with q_c fixed, the consumer's cost per unit is proportional to d_i, whereas with an endogenous quantity the consumer's cost per unit is proportional to $\sqrt{d_i}$. Thus, if the distance to the nearest store were to double, the consumer's cost does not double. Rather, it increases by only $\sqrt{2}$ because the consumer purchases a larger quantity with each visit, thereby making fewer store visits.

Let $\hat{d}_c(b,\theta)$ be the average square root round-trip distance across all consumers in a sub-area of a store's service area in which the sub-area has height b and angle θ :

$$\hat{d}_c(b,\theta) = \sqrt{2}\frac{\int_0^b dx \int_0^{x\tan\theta} \left(x^2+y^2\right)^{1/4} dy}{(1/2)b^2\tan\theta}.$$

Using a change of variables ($t = y/x$), we obtain

$$\hat{d}_c(b,\theta) = \left(\frac{4\sqrt{2}}{5}\frac{\int_0^{\tan\theta}\left(1+t^2\right)^{1/4} dt}{\tan\theta}\right) b^{1/2}.$$

Given, $a/(2sn) = (1/2)b^2\tan\theta$, the average per-unit cost for consumers is

$$C_c(q_c^*) = \hat{\phi}_c\tau_c^{1/2}(a/n)^{1/4}$$

where

$$\hat{\phi}_c = \frac{8\sqrt{2}}{5}(\tan\theta)^{-5/4} s^{-1/4} (h_c/\lambda_c)^{1/2} \int_0^{\tan\theta} \left(1+t^2\right)^{1/4} dt$$

Now turn to the retailer's quantity decision. As with consumers, let h_r be the retailer's cost per unit of time to maintain enough space to hold one unit. Given q_r units per delivery, the retailer's cost for space per unit of time is $h_r q_r$ and the retailer's cost per unit is then:

$$\min_{q_r} C_r(q_r) = \frac{h_r q_r}{\lambda_r} + \frac{\tau_r d_r}{q_r},$$

which is minimized with $q_r^* = \sqrt{(\lambda_r/h_r)\,\tau_r d_r}$ and yields

$$C_r\left(q_r^*\right) = 2\sqrt{(h_r\tau_r/\lambda_r)\,d_r}.$$

As with consumers, the retailer's total cost is proportional to $\sqrt{d_r}$ rather than d_r. The length of the TSP tour remains $d_r = 2bn$, so

$$C_r\left(q_r^*\right) = \hat{\phi}_r\tau_r^{1/2}(an)^{1/4}$$

where

$$\hat{\phi}_r = 2\sqrt{2}\,(s\tan\theta)^{-1/4}\,(h_r/\lambda_r)^{1/2}.$$

Using numerical integration, Table 7 provides the constants $\hat{\phi}_c$ and $\hat{\phi}_r$ for the three tesselations:

Table 7: Distance coefficients for different tesselations when cost are proportional to the square root of distances

Tesselation	$\hat{\phi}_c$	$\hat{\phi}_r$
triangle	$1.7499\,(h_c/\lambda_c)^{1/2}$	$1.8734(h_r/\lambda_r)^{1/2}$
square	$1.7116\,(h_c/\lambda_c)^{1/2}$	$2.0000\,(h_r/\lambda_r)^{1/2}$
hexagon	$1.7016\,(h_c/\lambda_c)^{1/2}$	$2.0732\,(h_r/\lambda_r)^{1/2}$

As when costs are linear in the distances traveled, when costs are proportional to the square root of the traveled distance, consumers prefer hexagons and the firm prefers triangles.

For the supply chain design problem, the cost function is

$$\hat{Z}(n) = a^{1/4}\left(\hat{\phi}_c \tau_c^{1/2} n^{-1/4} + \hat{\phi}_r \tau_r^{1/2} n^{1/4}\right).$$

The cost function, $\hat{Z}$, is quasi-convex in n and minimized with $\hat{n}^*$,

$$\hat{n}^* = \frac{\hat{\phi}_c^2 \tau_c}{\hat{\phi}_r^2 \tau_r},$$

yielding minimal cost

$$\hat{Z}\left(\hat{n}^*\right) = 2a^{1/4}\hat{\phi}_c^{1/2}\hat{\phi}_r^{1/2}\tau_c^{1/4}\tau_r^{1/4}.$$

As before, the triangle tesselation is the best (about 4% better than the hexagon tesselation).

To consider the robustness of this system, let $\tau_t(\alpha)$ be the actual operating cost of the vehicle of type $t \in \{c, r\}$, where $\alpha \in [0,1]$ is again the portion of the true cost of carbon that is explicitly charged. The design choice is $\hat{n}^*(\alpha) = \phi_c^2 \phi_r^{-2} \tau_c(\alpha) \tau_r(\alpha)^{-1}$ and the realized cost is then $\hat{Z}\left(\hat{n}^*(\alpha)\right)$. The carbon penalty is now

$$\frac{\hat{Z}\left(\hat{n}^*(\alpha)\right) - \hat{Z}\left(\hat{n}^*\right)}{Z\left(\hat{n}^*\right)} = \frac{1}{2}\left(\gamma(\alpha)^{1/4} + \frac{1}{\gamma(\alpha)^{1/4}}\right) - 1$$

which is even flatter (i.e., smaller) than the penalty with fixed delivery quantities - note that $x^y + x^{-y}$ is increasing in y, so

$$\frac{\gamma(\alpha)^{1/4} + \gamma(\alpha)^{-1/4}}{\gamma(\alpha)^{1/2} + \gamma(\alpha)^{-1/2}} \leq 1.$$

8 Conclusion

This paper's main finding is that the supply chain design problem is remarkably robust to errors in inputted costs. This conclusion is based on the following thought experiment. Suppose the true cost of carbon is very high (in the form of consequential externalities) but nevertheless, there is no explicit charge for carbon imposed on the retailer or consumers. The supply chain is designed to minimize the costs incurred by the agents rather than the true total costs. Despite this error, given current estimates regarding vehicle operating costs and fuel consumption, the total supply chain cost is unlikely to be significantly higher than its minimum cost. This finding is strengthened even further if the agents in the system can regulate their delivery frequency. Thus, even if the true cost of carbon were imposed on the retailer and consumers, possibly through a carbon tax, overall supply chain performance would not be significantly better.

Although this model does not provide a motivation for market intervention, an explicit charge for carbon may be necessary to correct other supply chain related decisions, such as where consumers choose to live (i.e., the degree of clustering they adopt), or what modes of transportation are employed by consumers and the retailer (e.g., walking, biking, or low emission vehicles). However, in terms of the total distance goods travel from a retailer to a consumer's home, total costs are insensitive to the density of the retail network. Therefore, in this context, there is minimal value for measures that are meant to correct any distortion in decisions regarding the retail supply chain structure.

References

American Automobile Association. 2011. Your driving costs: how much are you really paying to drive? http://www.aaaexchange.com/Assets/Files/201145734460.DrivingCosts2011.pdf

Baumul, W. 1972. On taxation and the control of externalities. *The American Economic Review*. 62. 307-322.

Beardwood, J., J.H. Halton, J.M. Hammersley, J. M. 1959. The shortest path through many points. *Proc. Cambridge Philos. Soc.* **55** 299-327.

Benjaafar, S., Y. Li., M. Daskin. 2010. Carbon footprint and the management of supply chains: insights from simple models. Working paper. University of Minnesota.

Bonney, J. 2009. Size, scale, innovation. Always. *The Journal of Commerce*. Oct 19.

Bramel, J., D. Simchi-Levi. 1997. *The Logic of Logistics: Theory, Algorithms, and Applications for Logistics Management*. Springer Series in Operations Research. P. Glynn. S. Robinson, eds. Springer. New York.

Brimberg, J., P. Hansen, N. Mladenovic, E. Taillard. 2000. Improvement and comparison of heuristics for solving the uncapacitated multisource Weber problem. *Operations Research*. **48**. 444-460.

Bureau of Transportation Statistics. 2011. *National Transportation Statistics*. U.S. Department of Transortation.

Burns, L.D., R.W. Hall., D.E. Blumenfeld, C.F. Daganzo. 1985. Distribution strategies that minimize transportation and inventory costs. *Operations Research*. **33**. 469-490.

Cline, W. 1992. *The Economics of Global Warming*. Institute for International Economics. Washington, D.C.

Daganzo, C. 1984. The distance traveled to visit N points with a maximum of C stops per vehicle: an analytical model and an application. *Transportation Science*. **4**. 331–350.

Dantzig, G.B., J.H. Ramser. 1959. The truck dispatching problem. *Management Science*. **6**. 80-91.

Daskin, M.S.. 1995. *Network and Discrete Location: Models, Algorithms, and Applications*. John Wiley & Sons Ltd. New York.

Department of Energy. 2011a. Weekly Retail On-Highway Diesel Prices. http://www.eia.gov/oog/info/wohdp/diesel_detail_report_combined.asp

Department of Energy. 2011b. Retail Gasoline Historical Prices. http://www.eia.gov/oil_gas/ petroleum/data_publications/wrgp/mogas_history.html

Dobson, G. 1988. Sensitivity of the EOQ to parameter estimates. *Operations Research*. **36**(4). 570-574.

Environmental Protection Agency. 2005. Emission factors: average carbon dioxide emissions resulting from gasoline and diesel fuel. Document EPA420-F-05-001.

Environment Protection Agency. 2009a. Frequently asked questions about global warming and climate change: back to basics. Document EPA-430-R08-016.

Environment Protection Agency. 2009b. Light-duty automotive technology, carbon dioxide emissions, and fuel economy trends: 1975 through 2009. Document EPA420-S-09-001.

Environmental Protection Agency. 2011. Inventory of U.S. greenhouse gas emissions and sinks: 1990-2009. Document EPC 430-R-11-005.

Federgruen, A., P. Zikin. 1984. A combined vehicle routing and inventory allocation problem. *Operations Research*. **32**. 1019-1037.

Fekete, S., J. Mitchell., K. Beurer. On the continuous Fermat-Weber problem. *Operations Research*. **53**. 61-76.

Fisher, M., D.S. Hochbaum. 1980. Probabilistic analysis of the planar k-median probkem. *Mathematics of Operations Research*. 5. 27-34.

Gallego, G., D. Simchi-Levi. 1990. On the effectiveness of direct shipping strategy for the one-warehouse multi-retailer R-system. *Management Science*. **36**. 240-243.

Haimovich, M., A.H.G. Rinnooy Kan. 1985. Bounds and heuristics for capacitated routing problems. *Mathematics of Operations Research*. **10**. 527-542.

Hall, R. 1992. A note on bounds for direct shipping cost. *Management Science*. **38**(8). 1212-1214.

Hoen, K.M.R., T. Tan, J.C. Fransso, G.J. van Houtum. Effect of carbon emission regulations on transport mode selection in supply chains. Working paper. Eindhoven University of Technology.

Intergovernmental Panel on Climate Change. 2007. Climate Change 2007: Synthesis Report.

Lawler, E.L., J.K. Lenstra, A.H.G. Rinnooy Kan, D.B. Shmoys (eds.). 1985. *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*. John Wiley & Sons Ltd. New York.

Porteus, E. 2002. Foundation of Stochastic Inventory Theory. Stanford University Press. Stanford, California.

Shelling, T. 1992. Some economics of global warming. *The American Economic Review*. **82**. 1-14.