CLASSROOM CAPSULES

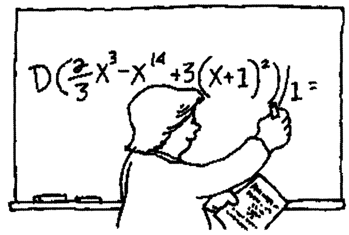

EDITOR

Thomas A. Farmer
Department of Mathematics and Statistics
Miami University
Oxford, OH 45056-1641

A Classroom Capsule is a short article that contains a new insight on a topic taught in the earlier years of undergraduate mathematics. Please submit manuscripts prepared according to the guidelines on the inside front cover to Tom Farmer.

The Rental Car Problem

Gary D. White, Northwestern State University, Natchitoches, LA 71457
Kirby Smith, Texas A&M University, College Station, TX 77845

A motorist in city A rents a car with a full tank of fuel in order to drive to city B, which is close enough that it can be reached without stopping for additional fuel. The rental company requires, however, that the tank be filled when the car is turned in at city B. Suppose the price of fuel varies along the route in a known way—say, the price per gallon x miles from A towards B is $P(x)$. What strategy should the motorist follow in purchasing fuel on the trip, to minimize the total fuel cost?

Clearly, if the price of fuel is lowest at city B, then the best strategy is to wait until reaching this destination and then fill up. But many students are surprised to discover that, in other cases, filling up at the place along the way where fuel is cheapest is usually not the optimal strategy.

This rental car problem raises a variety of interesting questions appropriate for courses from precalculus to multivariable calculus and discrete mathematics. Moreover, the total cost can be interpreted as a Riemann sum for the function $P(x)$, which then leads to an unexpected appearance of the definite integral and some surprising consequences.

We shall assume that the car's fuel consumption is a constant R miles per gallon, and that the tank is completely filled at each stop. In reality, of course, there are only finitely many places to stop for fuel, and finding the minimal cost strategy is a problem in discrete mathematics. But we will simplify the model by treating x as a continuous variable in $[0, b]$, where b is the number of miles from A to B, and we will also assume that the price P is a differentiable function of x.

Basic models. The simplest interesting case is the *two-fill problem*. Here the motorist is allowed to stop at most once along the way to fill up x miles from city A, then tops up the tank as necessary at city B. Since the cost of filling up is the number of gallons needed times the price per gallon, or

$$\frac{\text{Distance driven since last stop}}{\text{Fuel consumption rate } R} \times \text{Price per gallon},$$

the total cost C for a trip with at most two fills is

$$C(x) = \frac{x}{R}P(x) + \frac{b-x}{R}P(b) \qquad \text{for} \quad x \in [0,b].$$

Note that the fuel consumption rate R occurs only as an overall scaling factor, so the minimum of $C(x)$ and of the "scaled cost"

$$c(x) = xP(x) + (b-x)P(b) \qquad \text{for} \quad x \in [0,b] \tag{1}$$

occur at the same point. For simplicity, we will call $c(x)$ the *cost function*.

As a first example, where should you fill up in the two-fill problem if the price of fuel is increasing linearly? If

$$P(x) = mx + P_0, \tag{2}$$

where the slope m is positive, then the cost function has the quadratic form

$$c(x) = mx^2 - mbx + bP(b) \qquad \text{for} \quad x \in [0,b].$$

Since the graph of $c(x)$ is a parabola opening upward, the minimum occurs at the vertex $x = b/2$. Thus, the best strategy in the two-fill problem with linearly increasing price function is to fill up at the halfway point.

Before proceeding to other examples, observe that the scaled cost function $c(x)$ has a familiar geometric interpretation: $c(x)$ is the Riemann sum for the price function P relative to the partition $\{0, x, b\}$, using the values of P at the right-hand endpoints (see Figure 1). So we are trying to find a value of x to minimize this Riemann sum. Clearly, these observations can be generalized to the n-fill problem, in which the motorist makes at most $n-1$ stops along the way, fills the tank at each stop, and then tops up the tank at the final destination.

As Figure 2 illustrates, $c(x_1, \ldots, x_{n-1})$ is the right-hand rule Riemann sum for P using the partition $\{0, x_1, \ldots, x_{n-1}, b\}$. Minimizing c in this case (for each fixed $n > 2$) is a problem of multivariable calculus. If P is once again the increasing linear function (2) and we allow n fills, then the problem is to minimize

$$c(x_1, \ldots, x_{n-1}) = \sum_{i=1}^{n} (x_i - x_{i-1})(mx_i + P_0), \tag{3}$$

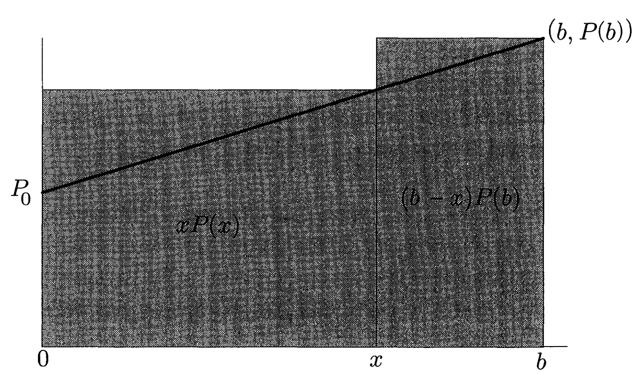

Figure 1. The cost function for a two-fill trip.

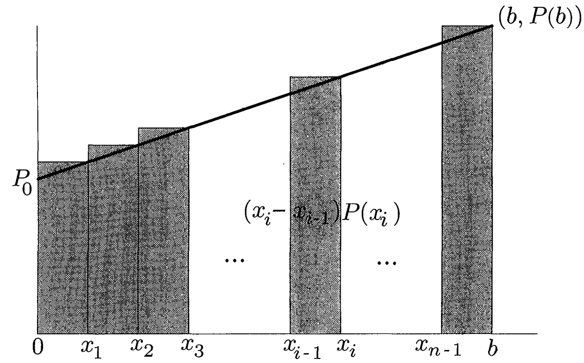

Figure 2. The cost function for an n-fill trip.

where $x_0 = 0$ and $x_n = b$, subject to the constraint

$$0 \le x_1 \le x_2 \le \cdots \le x_{n-1} \le b. \tag{4}$$

Since (3) is a right-hand Riemann sum of an *increasing* function, its minimum cannot occur at a boundary point of the region D defined by (4). A boundary point of D, where one or more of the inequalities is an equality, corresponds to a partition with fewer subintervals; but such a Riemann sum would be decreased by refining the partition. Thus, since the cost is a continuous function on the closed and bounded region D, its minimum must occur at a critical point inside D. An easy calculation shows that the only critical point of (3) is

$$x_1 = \frac{b}{n}, \quad x_2 = \frac{2b}{n}, \quad \ldots, \quad x_{n-1} = \frac{(n-1)b}{n},$$

the point in D that corresponds to the motorist's making $n-1$ evenly spaced stops.

What if we want to find the minimal cost strategy without restricting the number of stops? Our Riemann sum interpretation of the cost function allows us to make a nice observation in this situation. If the price P is increasing, then the set of right-hand rule Riemann sums has the integral of P (or the area under the graph of P) as its greatest lower bound. We can approach this limiting optimal strategy as closely as we please by making a sufficiently large number of stops to refill—it does not even matter where the stops are made, so long as the longest interval between stops (the mesh of the partition of $[0, b]$) is sufficiently small. But if the number of stops is unrestricted except that it must be finite, then the cost minimization problem has no optimal solution! Students are usually amused by the thought of someone running alongside the car with a gas hose, keeping the tank always filled, as the appropriate way to picture the integral of the price function in this setting.

Options and add-ons. After the linear price function in the two-fill problem, the next best example has parabolic price function and cubic cost function. Surprisingly, the minimal cost strategy in this case does not generally involve a stop at the place with the lowest price. In fact, if the parabola has its minimum halfway between A and B, then you should fill up two-thirds of the way in order to minimize the total cost (Figure 3). On the other hand, if the number of stops is unlimited, the minimizing prescription involves driving to the midpoint and filling up continuously from that

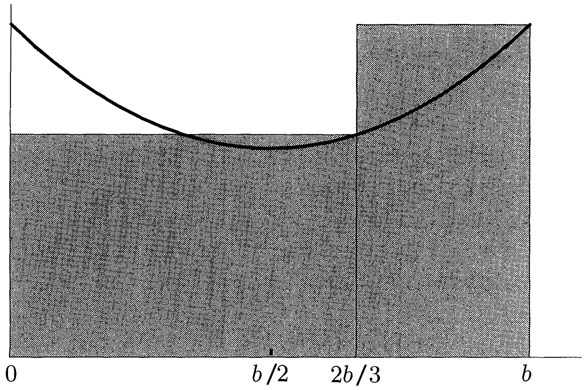

Figure 3. The parabolic price function.

point to the trip's end. This combination of the discrete and continuous is interesting both pedagogically and mathematically.

Other variations on this theme involve relaxing the tank-filling assumption, to allow partially filling the tank when stopping. In the case of one intermediate stop, two variables are needed to describe the situation: one for the distance to the stop and one for the amount of fuel purchased. This also results in a nontrivial boundary region for a multivariate optimization.

Alternatively, we can assume that there are only two gas stations at known locations with known fuel prices, P_1 and P_2. Then the cost function can be expressed in terms of the amounts of fuel purchased, t_1 and t_2, as

$$c(t_1, t_2) = \left[t_1 P_1 + t_2 P_2 + \left(\frac{b}{R} - t_1 - t_2 \right) P(b) \right].$$

The restrictions that we must impose on t_1 and t_2 so as not to overfill the tank again produce nontrivial boundary regions, but this time the cost function is linear. Under these assumptions, the rental car problem is transformed into a linear programming problem!

A student project. While versions of the rental car problem can be designed for several mathematics courses, here is how we have used it in a beginning calculus course late in the semester. The assignment came in the context of learning about the exponential function and was also intended as an exercise in the use of *Maple*:

Very often rental car agencies require you to return the car with a full tank of gas—or else they will fill it up for you at an exorbitant price. Imagine that you are in city A with a full tank of gas in a car that gets 25 *miles per gallon and has a* 20*-gallon tank. You are supposed to drive to city B, which is* 500 *miles away. Furthermore, assume that gas prices in city A are a dollar a gallon while in city B they are* \$1.50 *per gallon.*

1. If the gas price P increases linearly from A to B, where should you stop for gas so that you minimize your total cost? (Remember that you have to fill your tank at B paying \$1.50 per gallon. Assume you stop only once between A and B.)
2. If the price increases in a nonlinear fashion according to $P(x) = p + e^{kx}$, where p and $k > 0$ are constants and x is the distance from city A, where should you stop? (First get p and k, then minimize using *Maple*.)

3. Graph the function $(1/25)P$, where P is the price function from part 1. Since the car gets 25 miles per gallon, $(1/25)P(x)$ is your cost per mile during the first x miles if you fill the tank x miles from city A. Find a geometric interpretation of the total cost of gasoline as the sum of areas of rectangles overlaid on this graph. Repeat for the price function of part 2.
4. Now repeat parts 1 and 2 allowing two stops between A and B. Is this cheaper? What is the least possible cost if there is no limit on the number of stops? Show the geometric interpretation as areas for the case of unlimited stops.

After the problem was assigned, we spent a few minutes of the following class meetings discussing the students' progress and giving hints. It was surprising to see how many students had trouble writing down the linear expression for the price function and finding the values of p and k without help. Although only about half of the students got the point in part 4 about unlimited stops and integration, many of those who did confessed to being fascinated by it. Some also appreciated using *Maple* in the context of a larger "real" problem rather than just another drill. In the end-of-semester critiques, many students mentioned that they like doing "real-world" problems such as the rental car problem. One even claimed to have used these ideas to save money on the trip home during a break.

———o———

Complex Eigenvalues and Rotations: Are Your Students Going in Circles?

James Duemmel, Western Washington University, Bellingham, WA 98225

In an elementary linear algebra class, when you encounter a real matrix with complex eigenvalues, what do you say? Do you comment that it represents essentially a rotation in an unusual coordinate system? This approach is well explained in D. Lay's *Linear Algebra and Its Applications* (Addison Wesley, Reading, MA, 1994), where one finds this theorem.

Theorem. *Let A be a real 2×2 matrix with a complex eigenvalue $\lambda = a - bi$ $(b \neq 0)$ and associated eigenvector $\mathbf{w}$ in $\mathbb{C}^2$. Then $A = PCP^{-1}$, where P is the 2×2 real matrix $[\mathrm{Re}(\mathbf{w})\ \ \mathrm{Im}(\mathbf{w})]$ and $C = \begin{bmatrix} a & -b \\ b & a \end{bmatrix}$.*

If $|\lambda| = 1$, then $C = \begin{bmatrix} a & -b \\ b & a \end{bmatrix} = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$ is the standard matrix of a rotation by the angle θ; otherwise C represents the composition of this rotation with a scaling by the factor $|\lambda|$. Since the matrix P is the standard matrix of the linear transformation τ that sends the standard basis $\{\mathbf{e}_1, \mathbf{e}_2\}$ to the basis $\{\mathbf{u}, \mathbf{v}\} = \{\mathrm{Re}(\mathbf{w}), \mathrm{Im}(\mathbf{w})\}$, it follows that A is the standard matrix of the composition of τ^{-1}, then the rotation through the angle θ, then the scaling by the factor $|\lambda|$, and finally τ. Without further information about the basis $\{\mathbf{u}, \mathbf{v}\}$, however, it is hard to picture the geometric effect of τ and consequently of A.

A 3-D perspective. Perhaps we can do better by considering the 3×3 matrix $\begin{bmatrix} A & 0 \\ 0 & 1 \end{bmatrix}$ and the associated transformation $\alpha : \mathbb{R}^3 \to \mathbb{R}^3$. In fact, we will show that