Chapter 13

Applied Consumer Theory

In this chapter, we apply some of the analysis of the previous chapter to specific problems like the supply of labor, real estate, search, and risk.

13.1 Labor Supply

LEARNING OBJECTIVE

1. If we want people to work more, should we pay them more or will that cause them to work less?

Consider a taxi driver who owns a car, a convenience store owner, or anyone else who can set his or her own hours. Working has two effects on this consumer: more goods consumption but less leisure consumption. To model this, we let x be the goods consumption, L the amount of nonwork time or leisure, and working time $T - L$, where T is the amount of time available for activities of all kinds. The variable L includes a lot of activities that aren't necessarily fun—like trips to the dentist, haircuts, and sleeping—but for which the consumer isn't paid and which represent choices. One could argue that sleeping isn't really a choice, in the sense that one can't choose zero sleep; but this can be handled by adjusting T to represent "time available for chosen behavior" so that $T - L$ is work time and L is chosen nonwork activities. We set L to be leisure rather than labor supply because it is leisure that is the good thing, whereas most of us view working as something that we are willing to do provided we're paid for it.

Labor supply is different from other consumption because the wage enters the budget constraint twice: first as the price of leisure, and second as income from working. One way of expressing this is to write the consumer's budget constraint as $px + wL = M + wT$.

Here, M represents nonwork income, such as gifts, government transfers, and interest income. We drop the subscript on the price of X and use w as the wage. Finally, we use a capital L for leisure because a lowercase L looks like the number one. The somewhat Dickensian idea is that the consumer's maximal budget entails working the total available hours T, and any nonworked hours are purchased at the wage rate w. Alternatively, one could express the budget constraint so as to reflect that expenditures on goods px equal the total money, which is the sum of nonwork income M and work income $w(T - L)$, or $px = M + w(T - L)$.

These two formulations of the budget constraint are mathematically equivalent.

The strategy for solving the problem is also equivalent to the standard formulation, although there is some expositional clarity used by employing the budget constraint to eliminate x. That is, we write the utility $u(x, L)$ as

$$h(L) = u\left(\frac{M+w(T-L)}{p}, L\right).$$

As before, we obtain the first-order condition $0 = h'(L^*) = -u_1\left(w/p\right) + u_2$, where the partial derivatives u_1 and u_2 are evaluated at $\left(\frac{M+w(T-L^*)}{p}, L^*\right)$. Note that the first-order condition is the same as the standard two-good theory developed already. This is because the effect, so far, is merely to require two components to income: M and wT, both of which are constant. It is only when we evaluate the effect of a wage increase that we see a difference.

To evaluate the effect of a wage increase, differentiate the first-order condition to obtain

$$0 = \left(u_{11}\left(w/p\right)^2 - 2u_{12}\left(w/p\right) + u_{22}\right)\frac{dL}{dw} - \frac{u_1}{p} - \left(\frac{w}{p}\right)u_{11}\frac{T-L}{p} + u_{12}\frac{ }{ }$$

Since $u_{11}\left(w/p\right)^2 - 2u_{12}\left(w/p\right) + u_{22} < 0$ by the standard second-order condition, $\frac{dL}{dw} > 0$ if and only if $\frac{u_1}{p} + \left(\frac{w}{p}\right)u_{11}\frac{T-L}{p} - u_{12}\frac{T-L}{p} < 0$. That is, these expressions are equivalent to one another. Simplifying the latter, we obtain

$$\frac{-\left(\frac{w}{p}\right)u_{11}\frac{T-L}{p} + u_{12}\frac{T-L}{p}}{\frac{u_1}{p}} > 1, \text{ or} (T-L)\frac{-\left(\frac{w}{p}\right)u_{11} + u_{12}}{u_1} > 1, \text{ or}$$

$$\frac{\partial}{\partial L}Log(u_1) > \frac{1}{T-L} = -\frac{\partial}{\partial L}Log(T-L), \text{ or} \frac{\partial}{\partial L}Log(u_1) + \frac{\partial}{\partial L}Log(T-L) > 0, \text{ or}$$
$$\frac{\partial}{\partial L}Log(u_1(T-L)) > 0.$$

Since the logarithm is increasing, this is equivalent to $u_1(T-L)$ being an increasing function of L. That is, L rises with an increase in wages and a decrease in hours worked if the marginal utility of goods times the hours worked is an increasing function of L, holding constant everything else, but evaluated at the optimal values. The value u_1 is the marginal value of an additional good, while the value $T - L$ represents the hours worked. Thus, in particular, if goods and leisure are substitutes, so that an increase in L decreases the marginal value of goods, then an increase in wages must decrease leisure, and labor supply increases in wages. The case where the goods are complements holds a hope for a decreasing labor supply, so we consider first the extreme case of complements.

Example (Perfect complements): $u(x, L) = \min\{x, L\}$.

In this case, the consumer will make consumption and leisure equal to maximize the utility, so $\frac{M+w(T-L*)}{p} = L*$ or $L* = \frac{\frac{M+wT}{p}}{1+\frac{w}{p}} = \frac{M+wT}{p+w}$.

Thus, L is increasing in the wages if $pT > M$; that is, if M is sufficiently small so that one can't buy all of one's needs and not work at all. (This is the only reasonable case for this utility function.) With strong complements between goods and leisure, an increase in wages induces fewer hours worked.

Example (Cobb-Douglas): $h(L) = \left(\frac{M+w(T-L)}{p}\right)^{\alpha} L^{1-\alpha}$.

The first-order condition gives
$$0 = h'(L) = -\alpha\left(\frac{M+w(T-L)}{p}\right)^{\alpha-1} L^{1-\alpha} \frac{w}{p} + (1-\alpha)\left(\frac{M+w(T-L)}{p}\right)^{\alpha} L^{-\alpha} \text{ or}$$
$$\alpha L \frac{w}{p} = (1-\alpha)\frac{M+w(T-L)}{p}, \frac{w}{p} L = (1-\alpha)\frac{M+wT}{p} \text{ or } L = (1-\alpha)\left(\frac{M}{w} + T\right).$$

If M is high enough, the consumer doesn't work but takes $L = T$; otherwise, the equation gives the leisure, and labor supply is given by
$$T - L = Max\left\{0, \alpha T - (1-\alpha)\left(\frac{M}{w}\right)\right\}.$$

Labor supply increases with the wage, no matter how high the wage goes.

The wage affects not just the price of leisure but also the income level. This makes it possible for the income effect of a wage increase to dominate the substitution effect. Moreover, we saw that this is more likely when the consumption of goods takes time; that is, the goods and leisure are complements.

Figure 13.1 *Hours per week*
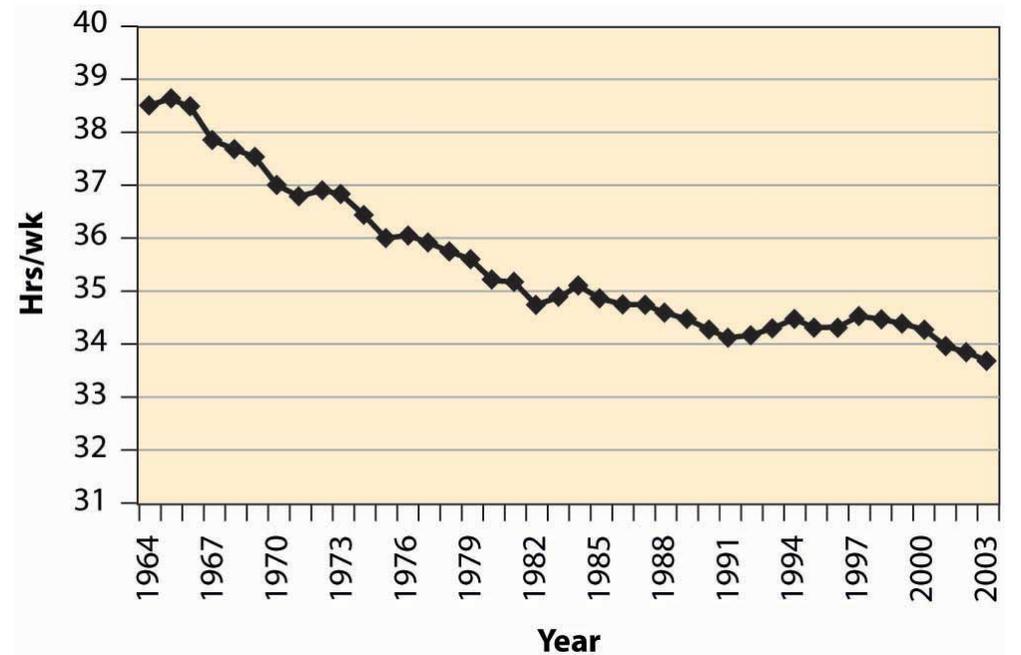
As a practical matter, for most developed nations, increases in wages are associated with fewer hours worked. The average workweek prior to 1950 was 55 hours, which fell to 40 hours by the mid-1950s. The workweek has gradually declined since then, as <u>Figure 13.1 "Hours per week"</u> illustrates.

A number of physicists have changed careers to become researchers in finance or financial economics. Research in finance pays substantially better than research in physics, and yet requires many of the same mathematical skills like stochastic calculus. Physicists who see their former colleagues driving Porsches and buying summerhouses are understandably annoyed that research in finance—which is intellectually no more difficult or challenging than physics—pays so much better. Indeed, some physicists are saying that other fields—such as finance, economics, and law—"shouldn't" pay more than physics.

The difference in income between physics' researchers and finance researchers is an example of a **compensating differential**[1]. A compensating differential is income or costs that equalize different choices. There are individuals who could become either physicists or finance researchers. At equal income, too many choose physics and too few choose finance, in the sense that there is a surplus of physicists and a shortage of finance researchers. Finance salaries must exceed physics' salaries in order to induce some of the researchers who are capable of doing either one to

1. Income or costs that equalize different choices.

switch to finance, which compensates those individuals for doing the less desirable task.

Jobs that are dangerous or unpleasant must pay more than jobs requiring similar skills but without the bad attributes. Thus, oil-field workers in Alaska's North Slope, well above the Arctic Circle, earn a premium over workers in similar jobs in Houston, Texas. The premium—or differential pay—must be such that the marginal worker is indifferent between the two choices: The extra pay compensates the worker for the adverse working conditions. This is why it is known in economics' jargon by the phrase of a compensating differential.

The high salaries earned by professional basketball players are not compensating differentials. These salaries are not created because of a need to induce tall people to choose basketball over alternative jobs like painting ceilings, but instead are payments that reflect the rarity of the skills and abilities involved. Compensating differentials are determined by alternatives, not by direct scarcity. Professional basketball players are well paid for the same reason that Picasso's paintings are expensive: There aren't very many of them relative to demand.

A compensating differential is a feature of other choices as well as career choices. For example, many people would like to live in California for its weather and scenic beauty. Given the desirability of California over, for example, Lincoln, Nebraska, or Rochester, New York, there must be a compensating differential for living in Rochester; and two significant ones are air quality and housing prices. Air quality worsens as populations rise, thus tending to create a compensating differential. In addition, the increase in housing prices also tends to compensate—housing is inexpensive in Rochester, at least compared with California.There are other compensations, besides housing, for living in Rochester—cross-country skiing and proximity to mountains and lakes, for example. Generally, employment is only a temporary factor that might compensate, because employment tends to be mobile, too, and move to the location that the workers prefer, when possible. It is not possible on Alaska's North Slope.

Housing prices also compensate for location within a city. For most people, it is more convenient—both in commuting time and for services—to be located near the central business district than in the outlying suburbs. The main compensating differentials are school quality, crime rates, and housing prices. We illustrate the ideas with a simple model of a city in the next section.

KEY TAKEAWAYS

- Leisure—time spent not working—is a good like other goods, and the utility cost of working is less leisure.
- Labor supply is different from other goods because the wage enters the budget constraint twice—first as the price of leisure, and second as income from working.
- If goods and leisure are substitutes, so that an increase in L decreases the marginal value of goods, then an increase in wages must decrease leisure, and labor supply increases in wages.
- With strong complements between goods and leisure, an increase in wages induces fewer hours worked.
- Complementarity between goods and leisure is reasonable because it takes time to consume goods.
- For most developed nations, increases in wages are associated with fewer hours worked.
- A compensating differential is income or costs that equalize different choices.
- Jobs that are dangerous or unpleasant must pay more than jobs requiring similar skills but without the bad attributes.
- The premium—or differential pay—must be such that the marginal worker is indifferent between the two choices: The extra pay compensates the worker for the adverse working conditions.
- City choice is also subject to compensating differentials, and significant differentials include air quality, crime rates, tax rates, and housing prices.

EXERCISES

1. A thought question: Does a bequest motive—the desire to give money to others—change the likelihood that goods and leisure are complements?
2. Show that an increase in the wage increases the consumption of goods; that is, x increases when the wage increases.

13.2 Urban Real Estate Prices

LEARNING OBJECTIVE

1. How are the prices of suburban ranch houses, downtown apartments, and rural ranches determined?

An important point to understand is that the good, in limited supply in cities, is not a physical structure like a house, but the land on which the house sits. The cost of building a house in Los Angeles is quite similar to the cost of building a house in Rochester, New York. The big difference is the price of land. A $1 million house in Los Angeles might be a $400,000 house sitting on a $600,000 parcel of land. The same house in Rochester might be $500,000—a $400,000 house on a $100,000 parcel of land.

Usually, land is what fluctuates in value rather than the price of the house that sits on the land. When a newspaper reports that house prices rose, in fact what rose were land prices, for the price of housing has changed only at a slow pace, reflecting increased wages of house builders and changes in the price of lumber and other inputs. These do change, but historically the changes have been small compared to the price of land.

We can construct a simple model of a city to illustrate the determination of land prices. Suppose the city is constructed on a flat plane. People work at the origin (0, 0). This simplifying assumption is intended to capture the fact that a relatively small, central portion of most cities involves business, with a large area given over to housing. The assumption is extreme, but not unreasonable as a description of some cities.

Suppose commuting times are proportional to distance from the origin. Let $c(t)$ be the cost to the person of a commute of time t, and let the time taken be $t = \lambda r$, where r is the distance. The function c should reflect both the transportation costs and the value of time lost. The parameter λ accounts for the inverse of the speed in commuting, with a higher λ indicating slower commuting. In addition, we assume that people occupy a constant amount of land. This assumption is clearly wrong empirically, and we will consider making house size a choice variable later.

A person choosing a house priced at $p(r)$, at distance r, thus pays $c(\lambda r) + p(r)$ for the combination of housing and transportation. People will choose the lowest cost alternative. If people have identical preferences about housing and commuting, then house prices p will depend on distance and will be determined by $c(\lambda r) + p(r)$ equal to a constant, so that people are indifferent to the distance from the city's center—decreased commute time is exactly compensated by increased house prices.

The remaining piece of the model is to figure out the constant. To do this, we need to figure out the area of the city. If the total population is N, and people occupy an area of one per person, then the city size r_{max} satisfies $N = \pi r_{max}^2$, and thus

$$r_{max} = \sqrt{\frac{N}{\pi}}.$$

At the edge of the city, the value of land is given by some other use, like agriculture. From the perspective of the determinant of the city's prices, this value is approximately constant. As the city takes more land, the change in agricultural land is a very small portion of the total land used for agriculture. Let the value of agricultural land be v per housing unit size. Then the price of housing is $p(r_{max}) = v$, because this is the value of land at the edge of the city. This lets us compute the price of all housing in the city:

$$c(\lambda r) + p(r) = c(\lambda r_{max}) + p(r_{max}) = c(\lambda r_{max}) + v = c\left(\lambda\sqrt{\frac{N}{\pi}}\right) + v$$

or

$$p(r) = c\left(\lambda\sqrt{\frac{N}{\pi}}\right) + v - c(\lambda r).$$

This equation produces housing prices like those illustrated in <u>Figure 13.2 "House price gradient"</u>, where the peak is the city's center. The height of the figure indicates the price of housing.

It is straightforward to verify that house prices increase in the population N and the commuting time parameter λ, as one would expect. To quantify the predictions, we consider a city with a population of 1,000,000; a population density of 10,000 per square mile; and an agricultural use value of $6 million per square mile. To translate these assumptions into the model's structure, first note that a population density of 10,000 per square mile creates a fictitious "unit of measure" of about 52.8 feet, which we'll call a *purlong*, so that there is one person per square purlong (2,788 square feet). Then the agricultural value of a property is v = $600 per square purlong. Note that this density requires a city of radius r_{max} equal to 564 purlongs, which is 5.64 miles.
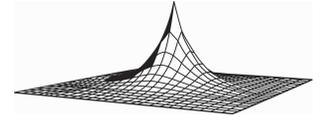
The only remaining structure left to identify in the model is the commuting cost c. To simplify the calculations, let c be linear. Suppose that the daily cost of commuting is $2 per mile (roundtrip), so that the present value of daily commuting costs in perpetuity is about $10,000 per mile. This amount is based upon 250 working days per year, for an annual cost of about $500 per mile, yielding a present value at 5% interest of $10,000. See Section 11.1 "Present Value". With a time value of $25 per hour and an average speed of 40 mph (1.5 minutes per mile), the time cost is 62.5 cents per minute. Automobile costs (such as gasoline, car depreciation, and insurance) are about 35–40 cents per mile. Thus the total is around $1 per mile, which doubles with roundtrips. This translates into a cost of commuting of $100 per purlong. Thus, we obtain

$$p(r) = c\left(\lambda\sqrt{\frac{N}{\pi}}\right) + v - c(\lambda r) = \$100\left(\sqrt{\frac{N}{\pi}} - r\right) + \$600 = \$57,000$$

Thus, the same 2,788-square-foot property at the city's edge sells for $600 versus $57,000, less than six miles away at the city's center. With reasonable parameters, this model readily creates dramatic differences in land prices, based purely on commuting time.

As constructed, a quadrupling of population approximately doubles the price of land in the central city. This probably understates the change, since a doubling of the population would likely increase road congestion, increasing λ and further increasing the price of central city real estate.

As presented, the model contains three major unrealistic assumptions. First, everyone lives on an identically sized piece of land. In fact, however, the amount of

land used tends to fall as prices rise. At \$53 per square foot, most of us buy a lot less land than at 20 cents per square foot. As a practical matter, the reduction of land per capita is accomplished both through smaller housing units and through taller buildings, which produce more housing floor space per acre of land. Second, people have distinct preferences and the disutility of commuting, as well as the value of increased space, varies with the individual. Third, congestion levels are generally endogenous—the more people who live between two points, the greater the traffic density and, consequently, the higher the level of λ. Problems arise with the first two assumptions because of the simplistic nature of consumer preferences embedded in the model, while the third assumption presents an equilibrium issue requiring consideration of transportation choices.

This model can readily be extended to incorporate different types of people, different housing sizes, and endogenous congestion. To illustrate such generalizations, consider making the housing size endogenous. Suppose preferences are represented by the utility function $u = H^\alpha - \lambda r - p(r)H$, where H is the house size that the person chooses, and r is the distance that he or she chooses. This adaptation of the model reflects two issues. First, the transport cost has been set to be linear in distance for simplicity. Second, the marginal value of housing decreases in the house size, but the value of housing doesn't depend on distance from the center. For these preferences to make sense, $\alpha < 1$ (otherwise either zero or an infinite house size emerges). A person with these preferences would optimally choose a house size of $H = \left(\frac{\alpha}{p(r)} \right)^{\frac{1}{1-\alpha}}$, resulting in utility

$$u^* = \left(\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}} \right) p(r)^{\frac{-\alpha}{1-\alpha}} - \lambda r.$$ Utility at every location is constant, so

$$\left(\frac{\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}}{u^* + \lambda r} \right)^{\frac{1-\alpha}{\alpha}} = p(r).$$

A valuable attribute of the form of the equation for p is that the general form depends on the equilibrium values only through the single number u^*. This functional form produces the same qualitative shapes as shown in <u>Figure 13.2 "House price gradient"</u>. Using the form, we can solve for the housing size H:

$$H(r) = \left(\frac{\alpha}{p(r)} \right)^{\frac{1}{1-\alpha}} = \alpha^{\frac{1}{1-\alpha}} \left(\frac{u^* + \lambda r}{\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}} \right)^{\frac{1}{\alpha}} = \left(\frac{u^* + \lambda r}{\alpha^{-1} - 1} \right)^{\frac{1}{\alpha}} = \left(\frac{\alpha}{1-\alpha} \right.$$

The space in the interval $[r, r + \Delta]$ is $\pi(2r\Delta + \Delta^2)$. In this interval, there are approximately $\frac{\pi(2r\Delta+\Delta^2)}{H(r)} = \pi(2r\Delta + \Delta^2)\left(\frac{1-\alpha}{\alpha(u*+\lambda r)}\right)^{\frac{1}{\alpha}}$ people. Thus, the number of people within r_{max} of the city's center is $\int_0^{r_{max}} 2\pi r \left(\frac{1-\alpha}{\alpha(u*+\lambda r)}\right)^{\frac{1}{\alpha}} dr = N$.

This equation, when combined with the value of land on the periphery

$$v = p(r_{max}) = \left(\frac{\alpha^{\frac{\alpha}{1-\alpha}}-\alpha^{\frac{1}{1-\alpha}}}{u*+\lambda r_{max}}\right)^{\frac{1-\alpha}{\alpha}}, \text{jointly determines } r_{max} \text{ and } u*.$$

When different people have different preferences, the people with the highest disutility of commuting will tend to live closer to the city's center. These tend to be people with the highest wages, since one of the costs of commuting is time that could have been spent working.

KEY TAKEAWAYS

- An important point to understand is that, in cities, houses are not in limited supply; but it is the land on which the houses sit that is.
- The circular city model involves people who work at a single point but live dispersed around that point. It is both the size of the city and the housing prices that are determined by consumers who are indifferent to commuting costs—lower housing prices at a greater distance balance the increased commuting costs.
- Substituting plausible parameters into the circular city model produces dramatic house price differentials, explaining much of the price differences within cities.
- A quadrupling of population approximately doubles the price of land in the central city. This likely understates the actual change since an increase in population slows traffic.

EXERCISE

1. For the case of α = ½, solve for the equilibrium values of $u*$ and r_{max}.

13.3 Dynamic Choice

The consumption of goods doesn't take place in a single instance, but over time. How does time enter into choice? We're going to simplify the problem a bit and focus only on consumption, setting aside working for the time being. Let x_1 be consumption in the first period and x_2 in the second period. Suppose the value of consumption is the same in each period, so that $u(x_1, x_2) = v(x_1) + \delta v(x_2)$, where δ is called the rate of "pure" time preference. The consumer is expected to have income M_1 in the first period and M_2 in the second. There is a market for lending and borrowing, which we assume has a common interest rate r.

The consumer's budget constraint, then, can be written $(1 + r)(M_1 - x_1) = x_2 - M_2$.

This equation says that the net savings in Period 1, plus the interest on the net savings in Period 1, equals the net expenditures in Period 2. This is because whatever is saved in Period 1 earns interest and can then be spent in Period 2; alternatively, whatever is borrowed in Period 1 must be paid back with interest in Period 2. Rewriting the constraint: $(1 + r)x_1 + x_2 = (1 + r)M_1 + M_2$.

This equation is known as the **intertemporal budget constraint**[2]—that is, the budget constraint that allows for borrowing or lending. It has two immediate consequences. First, $1 + r$ is the price of Period 2 consumption in terms of Period 1 consumption. Thus, the interest rate gives the relative prices. Second, the relevant income is the **permanent income**[3], which is the present value of the income stream. Clearly a change in income that leaves the present value of income the same should have no effect on the choice of consumption.

Once again, as with the labor supply, a change in the interest rate affects not just the price of consumption but also the budget for consumption. In other words, an increase in the interest rate represents an increase in budget for net savers but a decrease in budget for net borrowers.

2. The budget constraint that allows for borrowing or lending.

3. The present value of the income stream.

As always, we rewrite the optimization problem to eliminate one of the variables, to obtain $u = v(x_1) + \delta v\left((1 + r)(M_1 - x_1) + M_2\right)$.

Thus, the first-order conditions yield $0 = v'(x_1) - (1 + r)\delta v'(x_2)$.

This condition says that the marginal value of consumption in Period 1, $v'(x_1)$, equals the marginal value of consumption in Period 2, $\delta v'(x_2)$, times the interest factor. That is, the marginal present values are equated. Note that the consumer's private time preference, δ, need not be related to the interest rate. If the consumer values Period 1 consumption more than does the market, so that $\delta(1 + r) < 1$, then $v'(x_1) < v'(x_2)$; that is, the consumer consumes more in Period 1 than in Period 2. As usual, we are assuming that utility is concave, which in this instance means that the second derivative of v is negative, and thus the derivative of v is decreasing. In addition, to ensure an interior solution, it is useful to require what are called the *Inada* conditions: $v'(0) = \infty$ and $v'(\infty) = 0$ Similarly, if the consumer's discount of future consumption is exactly equal to the market discount, $\delta(1 + r) = 1$, the consumer will consume the same amount in both periods. Finally, if the consumer values Period 1 consumption less than the market, $\delta(1 + r) > 1$, the consumer will consume more in Period 2. In this case, the consumer is more patient than the market.

Whether the consumer is a net lender or borrower depends not just on the preference for earlier versus later consumption but also on incomes. This is illustrated in Figure 13.3 "Borrowing and lending". In this figure, the consumer's income mostly comes in the second period. As a consequence, the consumer borrows in the first period and repays in the second period.

Figure 13.3 *Borrowing and lending*
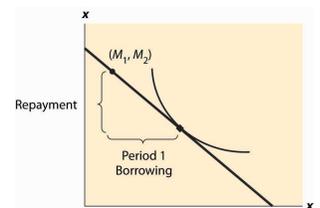
The effect of an interest rate increase is to pivot the budget constraint around the point (M_1, M_2). Note that this point is always feasible—that is, it is feasible to consume one's own endowment. The effect of an increase in the interest rate is going to depend on whether the consumer is a borrower or a lender. As Figure 13.4 "Interest rate change" illustrates, the net borrower borrows less in the first period—the price of first-period consumption has risen and the borrower's wealth has fallen. It is not clear whether the borrower consumes less in the second period because the price of second-period consumption has fallen, even though wealth has fallen, too—two conflicting effects.

An increase in interest rates is a benefit to a net lender. The lender has more income, and the price of Period 2 consumption has fallen. Thus the lender must consume more in the second period, but only consumes more in the first period (lends less) if the income effect outweighs the substitution effect. This is illustrated in Figure 13.5 "Interest rate increase on lenders".

The government, from time to time, will rebate a portion of taxes to "stimulate" the economy. An important aspect of the outcome of such a tax rebate is the effect to which consumers will spend the rebate, versus save the rebate, because the stimulative effects of spending are thought to be larger than the stimulative effects of saving.This belief shouldn't be accepted as necessarily true; it was based on a model that has since been widely rejected by the majority of economists. The general idea is that spending creates demand for goods, thus encouraging business investment in production. However, saving encourages investment by producing spendable funds, so it isn't at all obvious whether spending or saving has a larger effect. The theory suggests how people will react to a "one-time" or transitory tax rebate, compared to a permanent lowering of taxes. In particular, the budget constraint for the consumer spreads lifetime income over the lifetime. Thus, for an average consumer who might spend a present value of $750,000 over a lifetime, a $1,000 rebate is small potatoes. On the other hand, a $1,000 per year reduction is worth $20,000 or so over the lifetime, which should have 20 times the effect of the transitory change on the current expenditure.

Figure 13.4   Interest rate change

Figure 13.5   Interest rate increase on lenders

Tax rebates are not the only way that we receive one-time payments. Money can be found, or lost, and we can have unexpected costs or windfall gifts. From an intertemporal budget constraint perspective, these transitory effects have little significance; and thus the theory suggests that people shouldn't spend much of a windfall gain in the current year, or cut back significantly when they have a moderately sized, unexpected cost.

As a practical matter, most individuals can't borrow at the same rate at which they lend. Many students borrow on credit cards at very high interest rates and obtain a fraction of that in interest on savings. That is to say, borrowers and lenders face different interest rates. This situation is readily identified in Figure 13.6 "Different rates for borrowing and lending". The cost of a first-period loan is a relatively high
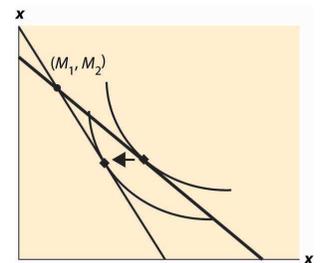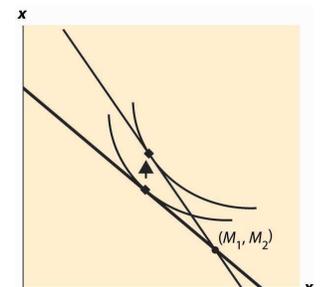
loss of x_2, and similarly the value of first-period savings is a much more modest increase in second-period consumption. Such effects tend to favor "neither a borrower nor a lender be," as Shakespeare recommends, although it is still possible for the consumer to optimally borrow in the first period (e.g., if $M_1 = 0$) or in the second period (if M_2 is small relative to M_1).

Differences in interest rates cause transitory changes in income to have much larger effects than the intertemporal budget constraint would suggest. This may go a long way toward explaining why people don't save much of a windfall gain and similarly suffer a lot temporarily, rather than a little for a long time, when they have unexpected expenses. This is illustrated in Figure 13.7 "The effect of a transitory income increase".

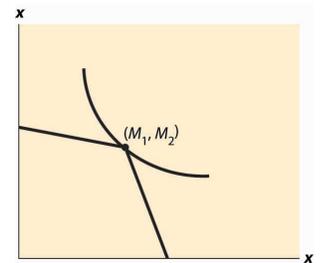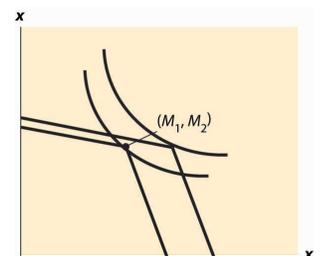*Figure 13.6   Different rates for borrowing and lending*

Figure 13.7   The effect of a transitory income increase

KEY TAKEAWAYS

- The intertemporal budget constraint takes into account the fact that savings obtain interest, producing additional money. The price of early consumption is one plus the interest rate.
- The relevant income is "permanent income" rather than "current income." A change in income that leaves the present value of income the same should have no effect on the choice of consumption.
- A change in the interest rate affects not just the price of consumption but also the budget for consumption. An increase in the interest rate represents an increase in budget for savers but a decrease in budget for borrowers.
- If the consumer values early consumption more than the market, the consumer consumes more early rather than later, and conversely.
- Whether the consumer is a lender or borrower depends not just on the preference for earlier versus later consumption but also on the timing of their incomes.
- The effect of an interest rate increase is to pivot the budget constraint around the income point. The effect of an increase in the interest rate on consumption is going to depend on whether the consumer is a borrower or a lender.
- An increase in interest rates is a benefit to a net lender. The lender must continue to lend in the present and will consume more in the future.
- People should react less to a "one-time" or transitory tax rebate than to a permanent lowering of taxes.
- Most individuals can't borrow at the same rate at which they lend. Interest rate differentials favor spending all of one's income.
- Differences in borrowing and lending interest rates cause transitory changes in income to have larger effects than the intertemporal budget constraint would suggest.

13.4 Risk Aversion

There are many risks in life, even if one doesn't add to these risks by intentionally buying lottery tickets. Gasoline prices go up and down, the demand for people trained in your major fluctuates, and house prices change. How do people value gambles? The starting point for the investigation is the von Neumann-MorgensternJohn von Neumann (1903–1957) and Oskar Morgenstern (1902–1977) are the authors of *Theory of Games and Economic Behavior* (Princeton, NJ: Princeton University Press, 1944). utility function. The idea of a **von Neumann-Morgenstern utility function**[4] for a given person is that, for each possible outcome x, there is a value $v(x)$ assigned by the person, and the average value of v is the value the person assigns to the risky outcome. In other words, the von Neumann-Morgenstern utility function is constructed in such a way that a consumer values gambles as if they were the expected utility

This is a "state-of-the-world" approach, in the sense that each of the outcomes is associated with a state of the world, and the person maximizes the expected value of the various possible states of the world. Value here doesn't mean a money value, but a psychic value or utility.

To illustrate the assumption, consider equal probabilities of winning $100 and winning $200. The expected outcome of this gamble is $150—the average of $100 and $200. However, the actual value of the outcome could be anything between the value of $100 and the value of $200. The von Neumann-Morgenstern utility is ½v($100) + ½$v$($200).

The von Neumann-Morgenstern formulation has certain advantages, including the logic that what matters is the average value of the outcome. On the other hand, in many tests, people behave in ways not consistent with the theory.For example, people tend to react more strongly to very unlikely events than is consistent with the theory. Nevertheless, the von Neumann approach is the prevailing model of behavior under risk.

4. The value of each outcome, constructed in such a way that a consumer values gambles as if they were the expected utility.

To introduce the theory, we will consider only money outcomes, and mostly the case of two money outcomes. The person has a von Neumann-Morgenstern utility function v of these outcomes. If the possible outcomes are $x_1, x_2, \ldots, xn$ and these occur with probability $\pi_1, \pi_2, \ldots, \pi n$ respectively, the consumer's utility is

$$u = \pi_1 v(x_1) + \pi_2 v(x_2) + \ldots + \pi_n v(x_n) = \sum_{i=1}^{n} \pi_i v(x_i).$$

This is the meaning of "having a von Neumann-Morgenstern utility function"—that utility can be written in this weighted sum form.
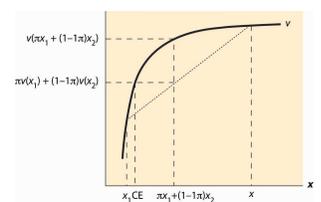
The first insight that flows from this definition is that an individual dislikes risk if v is concave. To see this, note that the definition of concavity posits that v is concave if, for all π in [0, 1] and all values x_1 and x_2,
$v(\pi x_1 + (1 - \pi)x_2) \geq \pi v(x_1) + (1 - \pi)v(x_2).$

For smoothly differentiable functions, concavity is equivalent to a second derivative that is not positive. Using induction, the definition of concavity can be generalized to show $v(\pi_1 x_1 + \pi_2 x_2 + \ldots + \pi_n x_n) \geq \pi_1 v(x_1) + \pi_2 v(x_2) + \ldots + \pi_n v(x_n).$

That is, a consumer with concave value function prefers the average outcome to the random outcome. This is illustrated in Figure 13.8 "Expected utility and certainty equivalents". There are two possible outcomes: x_1 and x_2. The value x_1 occurs with probability π, and x_2 with probability $1 - \pi$. This means that the average or expected outcome is $\pi x_1 + (1 - \pi)x_2$. The value $v(\pi x_1 + (1 - \pi)x_2)$ is the value at the expected outcome $\pi x1 + (1 - \pi)x2$, while $\pi v(x1) + (1 - \pi)v(x2)$ is the average of the value of the outcome. As is plainly visible in the figure, concavity makes the average outcome preferable to the random outcome. People with concave von Neumann-Morgenstern utility functions are known as **risk-averse**[5] people—they prefer the expected value of a gamble to the gamble itself.

Figure 13.8 Expected utility and certainty equivalents

A useful concept is the certainty equivalent of a gamble. The **certainty equivalent**[6] is an amount of money that provides equal utility to the random payoff of the gamble. The certainty equivalent is labeled CE in the figure. Note that CE is less than the expected outcome, if the person is risk averse. This is because risk-averse individuals prefer the expected outcome to the risky outcome.

5. Preferring the expected value of a gamble to the gamble.

6. The amount of money that provides equal utility to the random payoff of the gamble.

The **risk premium**[7] is defined to be the difference between the expected payoff (this is expressed as $\pi x1 + (1 - \pi)x2$ in the figure) and the certainty equivalent. This is the cost of risk—it is the amount of money an individual would be willing to pay to avoid risk. This means as well that the risk premium is the value of insurance. How does the risk premium of a given gamble change when the base wealth is increased? It can be shown that the risk premium falls as wealth increases for any gamble, if and only if $-\frac{v''(x)}{v'(x)}$ is decreasing.

The measure $\rho(x) = -\frac{v''(x)}{v'(x)}$ is known as the **Arrow-Pratt measure of risk aversion**[8]The measure was named after its discoverers, Nobel laureate Kenneth Arrow and John Pratt., and also as the measure of absolute risk aversion. It is a measure of risk aversion computed as the negative of the ratio of the second derivative of utility divided by the first derivative of utility. To get an idea about why this measure matters, consider a quadratic approximation to v. Let μ be the expected value, and let δ^2 be the expected value of $(x - \mu)^2$. Then we can approximate $v(CE)$ two different ways.

$$v(\mu) + v'(\mu)(CE - \mu) \approx v(CE) = E\big\{v(x)\big\} \approx E\big\{v(\mu) + v'(\mu)(x - \mu) + \tfrac{1}{2}v$$

Thus,

$$v(\mu) + v'(\mu)(CE - \mu) \approx E\big\{v(\mu) + v'(\mu)(x - \mu) + \tfrac{1}{2}v''(\mu)(x - \mu)^2\big\}.$$

Canceling $v(\mu)$ from both sides and noting that the average value of x is μ, so $E(x - \mu)$ = 0, we have $v'(\mu)(CE - \mu) \approx \tfrac{1}{2}v''(\mu)\sigma^2$.

Then, dividing by $v'(x)$, $\mu - CE \approx \tfrac{1}{2}\frac{v(\mu)}{v'(\mu)}\sigma^2 = \tfrac{1}{2}\rho(\mu)\sigma^2$.

That is, the risk premium—the difference between the average outcome and the certainty equivalent—is approximately equal to the Arrow-Pratt measure times half the variance, at least when the variance is small.

The translation of risk into dollars, by way of a risk premium, can be assessed even for large gambles if we are willing to make some technical assumptions. If a utility has **constant absolute risk aversion (CARA)**[9], the measure of risk aversion doesn't change with wealth; that is $\rho = -\frac{v''(x)}{v'(x)}$ is a constant. This turns out to imply, after setting the utility of zero to zero, that $v(x) = \frac{1}{\rho}(1 - e^{-\rho x})$.(This formulation is derived by setting $v(0)$ = 0, handling the case of $\rho = 0$ with appropriate limits.) Now

7. The difference between the expected payoff and the certainty equivalent.

8. A measure of risk aversion computed as the negative of the ratio of the second derivative of utility divided by the first derivative of utility.

9. Situation in which the measure of risk aversion doesn't change with wealth.

also assume that the gamble x is normally distributed with mean μ and variance δ^2.

Then the expected value of $v(x)$ is $Ev(x) = \frac{1}{\rho}\left(1 - e^{-\rho\left(\mu - \frac{\rho}{2}\sigma^2\right)}\right)$.

It is an immediate result from this formula that the certainty equivalent, with CARA preferences and normal risks, is $\mu - \frac{\rho}{2}\sigma^2$. Hence, the risk premium of a normal distribution for a CARA individual is $\frac{\rho}{2}\sigma^2$. This formulation will appear when we consider agency theory and the challenges of motivating a risk averse employee when outcomes have a substantial random component.

An important aspect of CARA with normally distributed risks is that the preferences of the consumer are linear in the mean of the gamble and the variance. In fact, given a choice of gambles, the consumer selects the one with the highest value of $\mu - \frac{\rho}{2}\sigma^2$. Such preferences are often called **mean variance preferences**[10], and they describe people who value risk linearly with the expected return. Such preferences comprise the foundation of modern finance theory.

10. Preference that describe people who value risk linearly with the expected return.

KEY TAKEAWAYS

- The von Neumann-Morgenstern utility function for a given person is a value $v(x)$ for each possible outcome x, and the average value of v is the value the person assigns to the risky outcome. Under this theory, people value risk at the expected utility of the risk.
- The von Neumann approach is the prevailing model of behavior under risk, although there are numerous experiment-based criticisms of the theory.
- An individual dislikes risk if v is concave.
- For smoothly differentiable functions, concavity is equivalent to a second derivative that is not positive.
- People with concave von Neumann-Morgenstern utility functions are known as risk-averse people.
- The certainty equivalent of a gamble is an amount of money that provides equal utility to the random payoff of the gamble. The certainty equivalent is less than the expected outcome if the person is risk averse.
- The risk premium is defined to be the difference between the expected payoff and the certainty equivalent.
- The risk premium falls as wealth increases for any gamble, if and only if $-\frac{v''(x)}{v'(x)}$ is decreasing.
- The measure $\rho(x) = -\frac{v''(x)}{v'(x)}$ is known as the Arrow-Pratt measure of risk aversion, and also as the measure of absolute risk aversion.
- The risk premium is approximately equal to the Arrow-Pratt measure times half the variance when the variance is small.
- Constant absolute risk aversion provides a basis for "mean variance preferences," the foundation of modern finance theory.

EXERCISES

1. Use a quadratic approximation on both sides of the equation to sharpen the estimate of the risk premium. First, note that
$$v(\mu) + v'(\mu)(CE - \mu) + \tfrac{1}{2}v''(\mu)(CE - \mu)^2 \approx v(CE)$$
$$= E\{v(x)\} \approx E\{v(\mu) + v'(\mu)(x - \mu) + \tfrac{1}{2}v''(\mu)(x - \mu)^2\}.$$

2. Conclude that $\mu - CE \approx \frac{1}{\rho}\left(\sqrt{1 + \rho^2\sigma^2} - 1\right)$. This approximation is exact to the second order.

3. Suppose that $u(x) = x^{0.95}$ for a consumer with a wealth level of $50,000. Consider a gamble, with equal probability of winning $100 and losing $100, and compute the risk premium associated with the gamble.

4. Suppose that $u(x) = x^{0.99}$ for a consumer with a wealth level of $100,000. A lottery ticket costs $1 and pays $5,000,000 with the probability $\frac{1}{10,000,000}$. Compute the certainty equivalent of the lottery ticket.

5. The return on U.S. government treasury investments is approximately 3%. Thus, a $1 investment returns $1.03 after one year. Treat this return as risk-free. The stock market (S&P 500) returns 7% on average and has a variance that is around 16% (the variance of return on a $1 investment is $0.16). Compute the value of ρ for a CARA individual. What is the risk premium associated with equal probabilities of a $100 gain or loss given the value of ρ?

6. A consumer has utility $u(x) = x^{7/8}$ and a base wealth of $100,000. She is about to take part in a gamble that will give her $10,000 (bringing her to $110,000) if a fair die rolls less than 3 (probability 1/3), but will cost her $5,000 (leaving her with $95,000) otherwise.

    a. What is the certainty equivalent of participating in this gamble?

    b. How much would she be willing to pay in order to avoid this gamble?

13.5 Search

<div style="background:#e8e4d8">

LEARNING OBJECTIVE

1. How should a consumer go about finding the lowest price when available prices are random?

</div>

In most communities, grocery stores advertise sale prices every Wednesday in a newspaper insert, and these prices can vary from week to week and from store to store. The price of gasoline can vary as much as 15 cents per gallon in a one-mile radius. Should you decide that you want to buy a specific Sony television, you may see distinct prices at Best Buy and other electronics retailers. For many goods and services, there is substantial variation in prices, which implies that there are gains for buyers to search for the best price.

The theory of consumer search behavior is just a little bit arcane, but the basic insight will be intuitive enough. The general idea is that, from the perspective of a buyer, the price that is offered is random and has a probability density function $f(p)$. If a consumer faces a cost of search (e.g., if you have to visit a store—in person, telephonically, or virtually—the cost includes your time and any other costs necessary to obtain a price quote), the consumer will set a **reservation price**[11], which is a maximum price that he or she will pay without visiting another store. That is, if a store offers a price below p^*, the consumer will buy; otherwise, he or she will visit another store, hoping for a better price.

Call the reservation price p^*, and suppose that the cost of search is c. Let $J(p^*)$ represent the expected total cost of purchase (including search costs). Then J must equal $J(p^*) = c + \int_0^{p^*} pf(p)dp + \int_{p^*}^{\infty} J(p^*)f(p)dp.$

This equation arises because the current draw (which costs c) could either result in a price less than p^*, in which case observed price, with density f, will determine the price paid p; or the price will be too high, in which case the consumer is going to take another draw, at cost c, and on average get the average price $J(p^*)$. It is useful to introduce the cumulative distribution function F, with $F(x) = \int_0^x f(p)dp$. Note that something has to happen, so $F(\infty) = 1$.

---

11. The maximum acceptable price a consumer will pay without visiting another supplier.

We can solve the equality for $J(p^*)$, $J(p^*) = \dfrac{\int_0^{p^*} pf(p)dp+c}{F(p^*)}$.

This expression has a simple interpretation. The expected price $J(p^*)$ is composed of two terms. The first is the expected price, which is $\int_0^{p^*} p\,\dfrac{f(p)}{F(p^*)}\,dp$.This has the interpretation of the average price conditional on that price being less than p^*. This is because $\dfrac{f(p)}{F(p^*)}$ is, in fact, the density of the random variable which is the price given that the price is less than p^*. The second term is $\dfrac{c}{F(p^*)}$ .This is the expected search cost, and it arises because $\dfrac{1}{F(p^*)}$ is the expected number of searches. This arises because the odds of getting a price low enough to be acceptable is $F(p^*)$. There is a general statistical property underlying the number of searches. Consider a basketball player who successfully shoots a free throw with probability y. How many balls, on average, must he throw to sink one basket? The answer is $1/y$. To see this, note that the probability that exactly n throws are required is $(1-y)^{n-1}y$. This is because n are required means that $n-1$ must fail (probability $(1-y)^{n-1}$) and then the remaining one goes in, with probability y. Thus, the expected number of throws is

$$y \; + \; 2(1-y)y \; + \; 3(1-y)^2y \; + \; 4\,(1-y)^3y + \dots$$

$$= \mathrm{y}(1 \; + \; 2(1\text{-y}) \; + \; 3(1\text{-y})^2 \; + \; 4\,(1\text{-y})^3 + \dots)$$

$$= y\big((1 \; + \; (1-y) \; + \; (1-y)^2 \; + \; (1-y)^3 + \dots) + (1-y)(1 \; + \; (1-y)$$

$$= (1-y)^2(1 \; + \; (1-y) \; + \; (1-y)^2 \; + \; (1-y)^3 + \dots) + (1-y)^3(1 \; + \;$$

$$= y\left(\frac{1}{y} + (1-y)\frac{1}{y} + (1-y)^2\frac{1}{y} + (1-y)^3\frac{1}{y} + \dots\right) = \frac{1}{y}\,.$$

Our problem has the same logic—where a successful basketball throw corresponds to finding a price less than p^*.

The expected total cost of purchase, given a reservation price p^*, is given by

$$J(p^*) = \dfrac{\int_0^{p^*} pf(p)dp+c}{F(p^*)}\,.$$

But what value of p^* minimizes cost? Let's start by differentiating

$$J'(p*) = p * \frac{f(p*)}{F(p*)} - \frac{f(p*) \int_0^{p*} pf(p)dp + c}{F(p*)^2}$$

$$= \frac{f(p*)}{F(p*)} \left(p * - \frac{\int_0^{p*} pf(p)dp + c}{F(p*)} \right) = \frac{f(p*)}{F(p*)} \left(p * -J(p*) \right).$$

Thus, if $p* < J(p*)$, J is decreasing, and it lowers cost to increase $p*$. Similarly, if $p* > J(p*)$, J is increasing in $p*$, and it reduces cost to decrease $p*$. Thus, minimization occurs at a point where $p* = J(p*)$.

Moreover, there is only one such solution to the equation $p* = J(p*)$ in the range where f is positive. To see this, note that at any solution to the equation $p* = J(p*)$, $J'(p*) = 0$ and

$$J''(p*) = \frac{d}{dp *} \left(\frac{f(p*)}{F(p*)} \left(p * -J(p*) \right) \right)$$

$$= \left(\frac{d}{dp *} \frac{f(p*)}{F(p*)} \right) (p * -J(p*)) + \frac{f(p*)}{F(p*)} (1 - J'(p*)) = \frac{f(p*)}{F(p*)} > 0$$

This means that J takes a minimum at this value, since its first derivative is zero and its second derivative is positive, and that is true about any solution to $p* = J(p*)$. Were there to be two such solutions, J' would have to be both positive and negative on the interval between them, since J is increasing to the right of the first (lower) one, and decreasing to the left of the second (higher) one. Consequently, the equation $p* = J(p*)$ has a unique solution that minimizes the cost of purchase.

Consumer search to minimize cost dictates setting a reservation price equal to the expected total cost of purchasing the good, and purchasing whenever the price offered is lower than that level. That is, it is not sensible to "hold out" for a price lower than what you expect to pay on average, although this might be well useful in a bargaining context rather than in a store searching context.

Example (Uniform): Suppose that prices are uniformly distributed on the interval $[a, b]$. For $p*$ in this interval,

$$J(p^*) = \frac{\int_0^{p^*} pf(p)dp + c}{F(p^*)} = \frac{\int_a^{p^*} p\,\dfrac{dp}{b-a} + c}{\dfrac{p^*-a}{b-a}}$$

$$= \frac{\frac{1}{2}(p^{*2}-a^2) + c(b-a)}{p^*-a} = \frac{1}{2}(p^*+a) + \frac{c(b-a)}{p^*-a}.$$

Thus, the first-order condition for minimizing cost is $0 = J'(p^*) = \frac{1}{2} - \frac{c(b-a)}{(p^*-a)^2}$, implying $p^* = a + \sqrt{2c(b-a)}$.

There are a couple of interesting observations about this solution. First, not surprisingly, as $c \to 0$, $p^* \to a$; that is, as the search costs go to zero, one holds out for the lowest possible price. This is sensible in the context of the model, but in real search situations delay may also have a cost that isn't modeled here. Second, $p^* < b$, the maximum price, if $2c < (b - a)$. In other words, if the *most* you can save by a search is twice the search cost, then don't search, because the expected gains from the search will be half the maximum gains (thanks to the uniform distribution) and the search will be unprofitable.

The third observation, which is much more general than the specific uniform example, is that the expected price is a concave function of the cost of search (second derivative negative). That is, in fact, true for any distribution. To see this, define a function $H(c) = \min_{p^*} J(p^*) = \min_{p^*} \frac{\int_0^{p^*} pf(p)dp+c}{F(p^*)}$ Since $J'(p^*) = 0$, $H'(c) = \frac{\partial}{\partial c} J(p^*) = \frac{1}{F(p^*)}$.

From here it requires only a modest effort to show that p^* is increasing in c, from which it follows that H is concave. This means that there are increasing returns to decreasing search costs in that the expected total price of search is decreasing at an increasing rate as the cost of search decreases.

KEY TAKEAWAYS

- For many goods, prices vary across location and across time. In response to price variation, consumers will often search for low prices.
- In many circumstances the best strategy is a reservation price strategy, where the consumer buys whenever offered a price below the reservation price.
- Consumer search to minimize cost dictates setting a reservation price equal to the expected total cost of purchasing the good, and purchasing whenever the price offered is lower than that level.

EXERCISE

1. Suppose that there are two possible prices, one and two, and that the probability of the lower price one is x. Compute the consumer's reservation price, which is the expected cost of searching, as a function of x and the cost of search c. For what values of x and c should the consumer accept two on the first search or continue searching until the lower price one is found?