Investment Decisions: An Economic Tutorial

Peter F. Colwell

The ideas of *net present value* (NPV) and *internal rate of return* (IRR) are well known to real estate investment analysts. These two popular *capital budgeting* techniques are typically introduced in basic accounting and corporate finance courses in collegiate business programs. What users of these techniques often do not realize is that NPV and IRR are rooted firmly in standard microeconomics. The specific microeconomic model on which NPV and IRR analyses are based is known as *state preference theory*.

In the following pages, we use state preference theory in answering questions such as whether the NPV or IRR criterion leads to superior investment decisions, and whether an investment's mix of debt and equity financing matters. We limit the discussion to relatively simple cases, in which the investor considers only one or two possible projects. Furthermore, we allow for an investment to last only two time periods. Finally, we assume that there is a *perfect credit market*, in which interest rates do not differ with a transaction's size (large or small) or situation (lending or borrowing). These simplifications can affect the conclusions drawn from the analysis, but we include them as a means of better focusing on the underlying economic relationships.

Building Blocks of the Analysis

The basic elements of the analysis are economic concepts known as indifference curves, endowment, credit market line, cost and return of an investment, and present value (PV), along with NPV and IRR. In a nutshell, we consider someone who wants to use the wealth at his disposal in timing his consumption. If he were concerned only with today, he could spend all his money immediately on desired purchases. However, if he must be supported for more than one period, he must manage his wealth (by investing) so that he can select the best combination of consumption bundles over the relevant time span; making the best investment decisions therefore becomes crucial.

More technically, we depict the goods that someone would prefer to consume, without regard to what that person actually could afford, by a series of *indifference curves*. We define what the individual could afford (that person's *opportunities*) by analyzing the endowment of wealth, investments' costs and returns, and the credit market line. We then find the individual's optimum by considering preferences together with opportunities. Finally, we compute PV, NPV, and IRR to guide our decisions.

Of course, people do not actually make investment decisions by drawing graphs of their indifference curves and opportunity constraints. Yet there is no question that today's business practice relies heavily on the PV, NPV, and IRR tools. What we should, therefore, ask is whether these popular computations lead to decisions that are consistent with the microeconomic principles that should form the basis for *all* business analysis.

Our discussion will involve two time periods. We consider consumption in each of periods 1 and 2, with dollars spent in these periods represented symbolically as D_1 and D_2. This situation is represented graphically in Figure 1. We consider only two periods, because the central issues are easier to visualize with two periods than with more. Of course, a potential criticism of using this simplification is that a principle developed for a two-period model might not work with three or more periods. Such criticism is not valid in this case, however; note that we might conceive of the two periods as representing today and tomorrow, this year and next, or, with a little tweaking, the current period and all future periods.

Figure 1

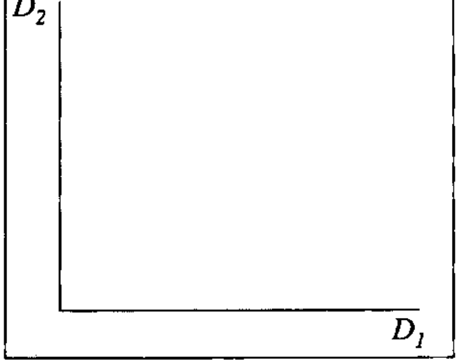

Indifference Curves

Each indifference curve identifies a group of consumption possibilities to which the subject individual is indifferent. A very simple example is that a given person might be equally happy with 11 apples and 1 orange, 5 apples and 5 oranges, 1 apple and 11 oranges, or a large number of other possible combinations. Because apples and oranges are both good things, an indifference curve on which they were represented would have to be negatively (downward) sloped, and higher curves (those farther outward from the *origin* where the vertical and horizontal axes meet) would indicate higher levels of satisfaction than lower ones. For example, a higher curve than that described above might include 17 apples and 2 oranges, 7 apples and 7 oranges, 2 apples and 17 oranges, or a range of other possibilities.

Instead of trading off apples against oranges, think now of a slightly more abstract case, in which our two "goods" are consumption today (year 1) and consumption a year from now (year 2). Again, each year's consumption is a good thing, so the indifference curve characteristics described above (negative slope, higher curve denoting higher satisfaction) still apply. Figure 2 displays two indifference curves that illustrate this situation. Note that in the figure, every point on I_2 is preferred to every point on I_1, but the person whose preferences are shown is indifferent between any two points on I_2.

Because we assume that there is *diminishing marginal satisfaction* with consumption in each period, an indifference curve must be *convex* (basically

Figure 2

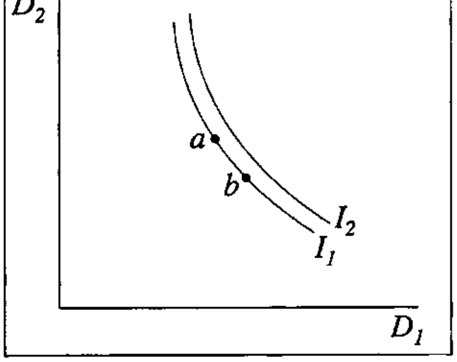

exhibiting a *c*-shape); the idea is that if someone had a lot of one consumption good but only a little of the other good, he would give up a considerable amount of the first good to get even a small additional amount of the second.

Figure 2 shows each indifference curve as convex (curving away) when viewed from below, regardless of what part of the curve is considered. The individual's willingness to trade off present for future consumption at the margin therefore changes all along an indifference curve. If year 1 consumption is much greater than expected year 2 consumption, the individual willingly gives up a fairly large amount of year 1 consumption to obtain a relatively small amount of added year 2 consumption.

Obviously, this willingness changes as the relative magnitudes of present and future consumption change. Willingness to trade present for future consumption is closely related to the individual's *discount rate* (the rate used in translating expected future dollar amounts into present value terms). The slope of an indifference curve is $-(1 + d)$, with d representing the discount rate. If the slope of the indifference curve changes continuously along the curve's length, then the discount rate must also change continuously. Specifically, it declines from northwest to southeast along the curve.[1]

For example, at point *a* in Figure 2, the individual's discount rate is about 50% (the slope is -1.50), whereas the discount rate is about 5% at point *b* (the slope is -1.05). At point *b*, the individual consumes relatively more in year 1, so it makes sense that he would require less consumption in the second year to compensate for giving up a dollar's worth of first-year consumption. Stated differently, the discount rate is lower at *b* than at *a*.

Indifference curves have three other very important features. First, it would be logically impossible for two indifference curves to intersect.[2] Second, every possible consumption combination lies on one, and only one, indifference curve; we show a few indifference curves for illustrative purposes, but it is possible to imagine any number of additional curves falling between any two that are shown (amounts of two goods that could be consumed encompass infinite possibilities). Finally, the indifference curves in this analysis should not touch the axes of the diagram; their doing so would suggest that zero year-2 consumption (possible only if the consumer died) could be equated to some level of consumption in both periods. Therefore, our indifference curves do not extend to the axes.

We specify the nature of indifference curves to see what implications may ultimately be drawn from a model with the most simple assumptions. Our only assumptions, initially, are that individuals prefer more to less, extra consumption during a given year adds less and less to satisfaction, a typical person is logically consistent, and no one would accept death tomorrow in return for consuming more today. The conclusions developed will be as general in their applicability as the assumptions are broadly valid.

Endowment

It is assumed that each person is *endowed* with the resources to consume amounts in both years. Even without engaging in investment or credit market activities, the individual could consume some amount in year 2 (perhaps through expected job earnings, or perhaps simply by hoarding some wealth not spent in year 1). While the endowment could, theoretically, be zero in either or both years, we show it as representing opportunities to consume positive magnitudes in both periods. Endowment point *e* in Figure 3 suggests that the endowment allows for greater consumption in year 1 than in year 2; *e*'s distance outward along the horizontal axis is more than its distance upward along the vertical axis. (Of course, it could be otherwise.) In numerical terms, point *e* might represent consumption of \$7,000 in year 1 and \$2,500 in year 2.

Figure 3

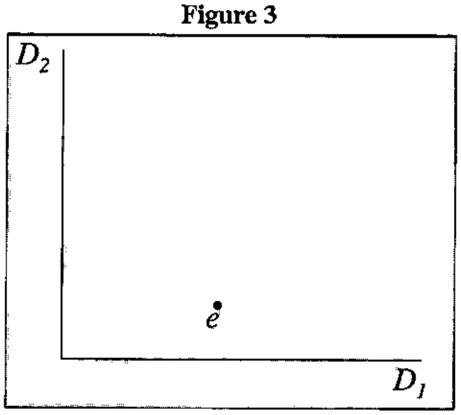

Investment and the IRR

The consumer who *invests* deviates from the consumption levels indicated by the endowment, giving up some consumption in the first year in order to consume more in the second. The reduction in year 1 consumption, called the investment's *cost*, is shown as C in Figure 4. The increase in year 2 consumption opportunities, called the dollar *return*, is labeled R in the figure. (For example, while *e* might suggest spending \$7,000 in year 1 and \$2,500 in year 2, the consumer who invested might give up consuming C = \$3,000 in year 1 and, after investing that \$3,000 for a year, get back an extra R = \$4,000 in year 2, thereby changing his consumption mix to \$4,000 in year 1 and \$6,500 in year 2.) For this one-year investment (project A), the ratio of R to C is one plus the IRR (\$4,000/\$3,000 = 1.33, such that $IRR_A = 33\%$), and the line connecting *a* to *e* has a slope equal to $-(1 + IRR_A)$, or -1.33.

The Credit Market Line

Through each investment opportunity point (representing deviation from the consumption indicated by endowment *e* because of C and R), it is possible to draw a *credit market line*. This line represents consumption opportunities the individual can attain, after reaching an investment opportunity point, by borrowing or lending. Our credit market line is shown as straight, with a constant negative slope; we implicitly assume what economists call a *perfect* credit market, in which the individual can borrow *or* lend unlimited amounts at unchanging interest rate i. This assumption obviously abstracts from the existence of financial intermediaries (such as banks), whose costs ordinarily push borrowing rates above lending

Figure 4

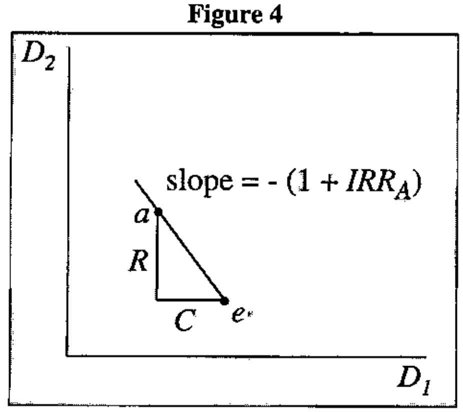

(deposit) rates. The credit market line's slope is – (1 + i); see Figure 5. For example, if i were 10%, the individual could borrow $2,000 and pay back $2,200, reaching point f with $6,000 ($4,000 + $2,000) of consumption in year 1 and $4,300 ($6,500 – $2,200) in year 2.

Consumer Optimum

Our consumer maximizes his satisfaction over the two years by reaching the highest possible indifference curve. Given an endowment, an investment opportunity, and an interest rate (a credit market line), the individual theoretically seeks the highest indifference curve that barely touches the credit market line. This *tangency* shows how much borrowing or lending is optimal. In Figure 6, someone with endowment e ($7,000 in year 1; $2,500 in year 2) would reach the highest possible indifference level by investing to reach point a ($4,000; $6,500), then borrowing to reach point f ($6,000; $4,300) on the higher attainable indifference curve shown.

Regardless of preferences (represented by the shape of the indifference curves), all consumers facing the same interest rate will have the same discount rate, because of the tangency between the credit market opportunity line and the highest reachable indifference curve. The slopes of these lines must be equal at the tangency point. In other words, $-(1 + d) = -(1 + i)$, so $d = i$, as shown in Figure 6. In this case, the interest rate is 10%, so the discount rate must be 10%.

The Role of Debt

A credit market exists because some individuals can reach higher levels of satisfaction by borrowing to consume more today than their endowments would directly permit, while others reach higher satisfaction by consuming less today than their endowments would allow, and saving the excess to lend in the credit market. The consumer depicted in Figure 6 can achieve higher satisfaction by borrowing. Without debt, the highest attainable indifference curve (which would go through point a) is lower than the indifference curve illustrated (through point f). Since this person is better off borrowing to consume more in year 1, must a lender foregoing some current consumption be worse off? The answer is no; *every* consumer is better off as a result of the borrowing and lending opportunities the credit market provides.

This result may seem curious, but the mystery is easy to solve. Consider an economy with two individuals; their endowments and investment opportunities are the same, but their preferences differ, as shown by Figure 7's two indifference curves. (Since the curves intersect, they *must* relate to different individuals.) At the interest rate implied by the credit market line's steepness, the person with the light indifference curve wants to borrow exactly what the person with the dark curve wants to lend (segments af and am are equal). In our example, the latter person maximizes satisfaction by investing to reach a ($4,000 today; $6,500 next year) and then lending $2,000 to reach m ($4,000 – $2,000 = $2,000 today; $6,500 + $2,200 = $8,700 next year). The implied interest rate must therefore be the *equilibrium* rate. We could derive demand and supply curves by rotating the credit market line around a and noting the various credit quantities demanded and supplied at every possible interest rate. Demand and supply would intersect at the rate implied in Figure 7.

The credit market could be harmed, or even destroyed, by regulation. A *usury* law could cause the maximum interest rate to be below the equilibrium, or paying interest could be outlawed entirely. In the absence of a credit market, the best these individuals could do would be to reach their respective indifference curves that go through point a; each would be at a lower indifference level than is achievable with credit. This analysis shows that credit is an important service, which benefits both borrowers and lenders.

Present Value

We compute an investment's *present value* by discounting the dollar return, R, at an appropriate rate. While some writers have presented the choice of the correct discount rate as a controversial subject, in this discussion we provide strong support for selecting a discount rate equal to the interest rate. A present value can be shown graphically as a distance along the horizontal axis (which measures present, or year 1, dollars). Think of a right triangle with the vertical side representing the dollar return, the hypotenuse representing a segment of the credit market line, and the horizontal base connecting the return line segment to the credit market segment. The slope of the hypotenuse is the ratio of return R to the base (the length of line segment hg in Figure 8); so $R/hg = (1+ i)$. Rearranging terms yields $hg = R/(1+i)$, the familiar present value calculation. Therefore, the base of this triangle is the PV of the investment's return. In Figure 8, this PV is the distance between points h and g. Continuing with our numerical example, a dollar return of $4,000 and a discount rate of 10% produce a PV of $3,636 (the value represented by line segment hg).

Figure 5

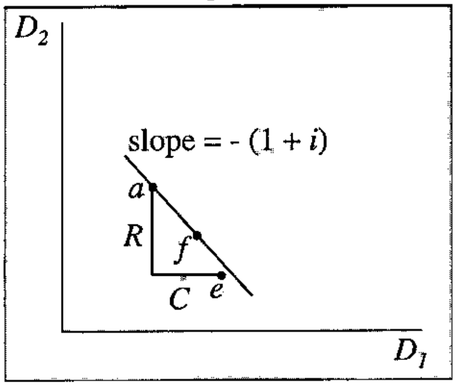

Figure 6

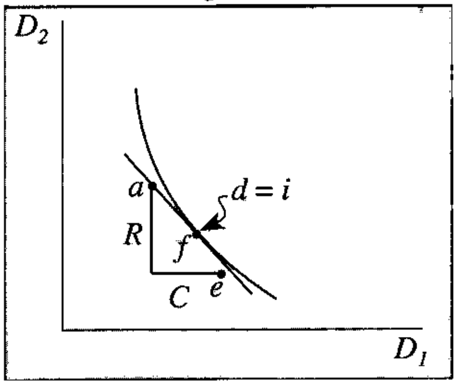

Figure 7

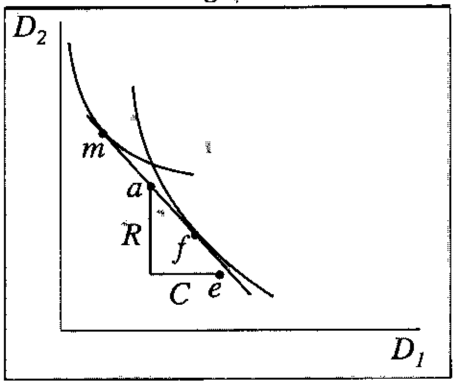

Net Present Value
Net present value is simply the present value of the dollar return, minus the investment's cost. Note that the triangle's base *hg* can be naturally divided into two line segments, one to the left and one to the right of *e*. The left segment represents the cost of the investment, so the right segment must equal NPV (in Figure 8, the distance between points *e* and *g*).

This NPV concept has an economic interpretation, although it may seem odd to those who have learned time value concepts by rote rather than through an understanding of the theory. The economic interpretation is that NPV is the minimum *subsidy* the consumer would have to receive in the present to voluntarily forego the investment opportunity. Consider that without the investment and without a subsidy, the individual could reach any of the points along a credit market line going through point *e*. But if there were a subsidy equal to the segment already identified as NPV, the individual could reach any point (such as *f*) on the credit market line going through *a* — not by investing to point *a* and then borrowing to reach *f*, but by accepting a subsidy to reach *g* and then *lending* to reach *f*.

The same consumer optimum would be available with the subsidy as with the investment opportunity, so NPV can be equated to a special kind of subsidy. If the person got a $636 subsidy, and then lent $1,636, he would be repaid $1,800 in year 2. He would, as a result, consume $6,000 ($7,000 + $636 – $1,636) in year 1 and $4,300 ($2,500 + $1,800) in year 2, shown as point *f* in Figure 8.

Ultimately, we can show that maximizing NPV is consistent with maximizing consumer satisfaction in light of the constraints regarding opportunities the individual has. Maximizing NPV is not the same thing as maximizing consumer satisfaction, because the latter requires optimal credit market activity in addition to optimal investing, while maximizing NPV is unrelated to credit market activity.

Two Mutually-Exclusive Investments
The basic capital budgeting decision rule is to select the investment, from among those available, that has the highest NPV *or* the highest IRR. In the simplest cases, the project with the highest NPV also has the highest IRR, so the two decision criteria provide equivalent answers. Yet there can be situations in which one of two projects displays a higher NPV while the other displays a higher IRR. When this inconsistency arises, which of the two should be selected? By considering the typical textbook example involving two *mutually-exclusive* (only one can be chosen) investments, we can reveal a number of important results, not the least of which is the basic superiority of NPV over IRR as an investment criterion.

Maximizing NPV is consistent with maximizing satisfaction, but maximizing IRR is not. Thus, the IRR, as it has been defined so far, is an inferior criterion for choosing investments. Because the dashed line through points *e* and *a* (corresponding to project *A*) in Figure 9 is steeper than the dashed line through *e* and *b* (corresponding to project *B*), *A*'s IRR is greater than *B*'s. But note that *B*'s NPV (line segment *eg*) is greater than *A*'s (segment *ef*). Thus, there is a conflict between these two investment criteria, which the example shown in Figure 9 has been carefully selected to illustrate. *A* has the higher IRR, but *B* has the higher NPV; which is the correct choice? Note that the consumer is better off with *B* (reaching the indifference curve that would be tangent to the line connecting *b* and *g*) than with *A*. Thus, the NPV criterion always is consistent with maximizing consumer satisfaction, whereas IRR may not be.

Returning to our computational example, assume that project *B* has a dollar return of $8,100 and a $6,000 cost. Project *A* has a $2,790 dollar return and a $1,800 cost. So project *B*'s IRR is 35%, whereas project *A*'s is 55%. On the other hand, project *B*'s NPV is $1,364, while project *A*'s is $736. There is a conflict between the two investment criteria. Yet regardless of the credit strategy employed with project *A*, the individual can achieve more year 1 *and* year 2 consumption with project *B* and a credit market transaction. Thus, the NPV criterion provides superior indications of correct investment strategy when IRR/NPV conflicts exist.

Yet the IRR decision is not always *in*consistent with maximizing satisfaction, either. For example, if all projects are equal in cost, there will be no conflict between IRR and NPV, and choosing the project that appears superior under either criterion is consistent with maximizing consumer satisfaction. Similarly, if both projects have the same dollar returns and differ only by their cost magnitudes, then the NPV and IRR criteria both yield valid answers; there are other circumstances in which the project with the highest IRR also has the highest NPV, as well. If the rankings do not conflict, then it is perfectly acceptable to base an investment decision on the IRR criterion.

After-Debt IRR
We can attain broader validity for the IRR criterion by considering a measure that differs slightly from the traditional IRR. If IRR is measured after debt, then we may find that there is no conflict between the two investment criteria. (We do not have to modify NPV similarly, because it should make no difference whether NPV is measured before or after debt.) However, we do have to make an additional assumption about preferences in order to get after-debt IRR to work. The most modest assumption we can make is that consumption in the two time periods is *normal*, such that an individual who receives a subsidy that shifts out the credit market line will choose to consume more

Figure 8

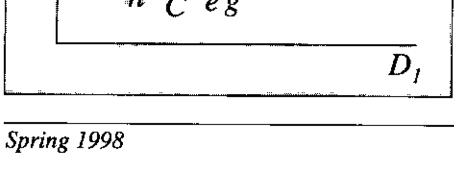

Figure 9

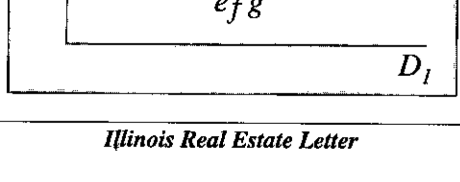

in both year 1 and year 2. In Figure 10, the steeper of the two dashed lines is associated with project *B*. We compute after-debt IRR by subtracting one from the negative of this slope. Project *B* is associated with a steeper dashed line and, therefore, a higher after-debt IRR. Project *B*'s NPV equals the length of line segment *eg*, while project *A*'s NPV equals the length of line segment *ef*. Thus, the project with the higher NPV also has the higher after-debt IRR, or IRR*.

Suppose that the individual invests in project *A* but does not borrow. Thus, he consumes $5,200 ($7,000 – $1,800) in year 1 and $5,290 ($2,500 + $2,790) in year 2. The after-debt IRR is the same as the before-debt IRR. However, if he instead invests in *B*, and then borrows $4,700 (paying back $5,170), he consumes $5,700 ($7,000 – $6,000 + $4,700) in year 1 and $5,430 ($2,500 – $5,170 + $8,100) in year 2. Note first that more is consumed in both years with project *B*, as required under our assumption regarding "normal" goods. Also note that the after-debt IRR for project *B* is 125% [($8,100 – $5,170)/($6,000 – $4,700) – 1]. So the after-debt IRR is consistent in identifying *B* as the superior investment. Of course, this outcome was already apparent from the earlier comparison of the projects' NPVs (a comparison that had not required our knowledge of consumption choices).

Negative NPV

Students are often told that positive NPV projects are good and negative NPV projects are bad. Consider a negative NPV project, such as project *A* in Figure 11. If the individual chooses not to invest, but rather to lend from point *e*, he can achieve a higher level of satisfaction (shown graphically as reaching a higher indifference curve). Thus, as intuition would suggest, *not* investing is superior to investing in a project that is expected to deliver a negative NPV.

The PV of dollar returns for project *A* in Figure 11 is equal to the length of line segment *fg*. The cost, however, is a larger magnitude, equal to the length of line segment *fe*. Therefore, NPV is a negative magnitude, equal to the length of segment *ge*. The individual is better off by lending from point *e* to achieve point *h* than by investing to point *a* and then borrowing to achieve point *k*; he attains higher satisfaction by not investing if the NPV is negative. In numerical terms, if we assume that project *A* costs $3,000 and returns $3,060, its NPV is – $218 ($3,060/1.1 – $3,000). If the individual borrows $1,500 and repays $1,650, he can consume $5,500 ($7,000 – $3,000 + $1,500) in year 1 and $3,910 ($2,500 + $3,060 – $1,650) in year 2. Alternatively, he could ignore project *A* and lend $1,400 in year 1, receiving $1,540 in year 2. This approach would allow consumption of $5,600 ($7,000 – $1,400) in year 1 and $4,040 ($2,500 + $1,540) in year 2 (more in both years).

What is the "Discount Rate"?

It is sometimes said that the correct rate to use in discounting future magnitudes is the "required rate of return." A rate of return is "required" not because the investor simply wants it, but rather because his best alternative will earn that rate. Therefore the investor requires that, in order to be chosen, a project would have to earn at least this *opportunity* rate. Specifically, we interpret this prescription as meaning that we should use the IRR for the project that represents the *second* best opportunity.[3] If the two projects have the same dollar returns, but different costs, then the NPV of the superior project (the one with the lower cost), computed based on the required rate of return, is equal to the incremental NPV based on the true discount rate.

Figure 12 illustrates this situation. Note that points *a* and *b* are at the same height, reflecting the assumption that projects *A* and *B* have the same dollar returns. Note further that point *b* is to the right of point *a*, indicating that the cost of project *B* is less than that of project *A*. It is obvious that with the same return and a lower cost, project *B* is superior to project *A*. The NPV of project *B*, computed based on the true discount rate, is line segment *eh*. Similarly, the NPV of project *A* based on the true discount rate is line segment *eg*. Thus, the incremental NPV of project *B* over project *A* is line segment *gh*.

Assume hat the dollar return for each of projects *A* and *B*, as shown in Figure 12, is $3,500. However, project *A* costs $2,500, whereas project *B* costs only $1,500, so it should be obvious that *B* is preferable to *A*. But suppose that we were to attempt instead to find project *B*'s NPV by discounting at project *A*'s IRR ($3,500/ $2,500 – 1 = 40%). The resulting NPV* would be computed as $1,000 ($3,500/ 1.4 – $1,500). We can also find this $1,000 value by subtracting project *A*'s $682 NPV from *B*'s true NPV of $1,682 ($3,500/1.1 – $1,500). In economic terms, the $1,000 is the *incremental* NPV, or the *extra* addition to wealth that results from choosing the better project.

Returning to the graph, suppose that we were to compute project *B*'s NPV using the required rate of return as the discount rate. This NPV would be line segment *ef* in Figure 12. We should realize that segment *ef* is equal in length

Figure 10

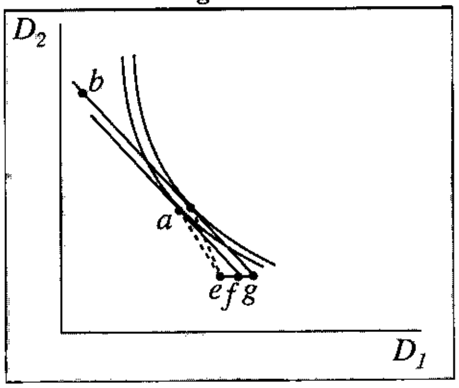

Figure 11

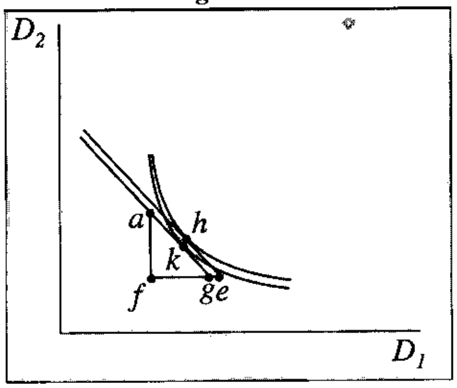

Figure 12

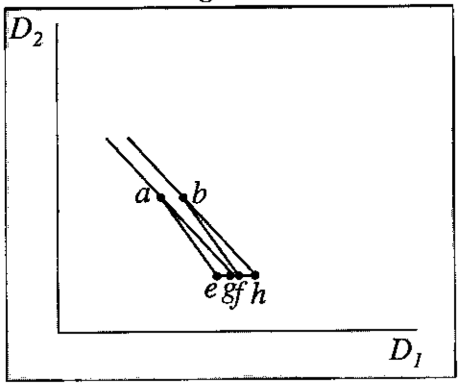

to segment *ab*. Note also that line segment *gh* is equal in length to line segment *ab*, so length *ef* is equal to length *gh*. Therefore, the NPV of the superior project, based on discounting at the required rate, is equal to that project's incremental NPV based on discounting at the true discount rate. This result is based on an assumption that the projects have equal dollar returns but unequal costs.

Why would we want to calculate the incremental NPV? As noted above, the incremental NPV is a reasonable indicator of whether it is worthwhile to choose one investment over the other. Viewed differently, incremental NPV represents the compensation that would leave the individual no better off or worse off *if* the superior investment project were taken away but the inferior project remained available. So we can see that the *ad hoc* use of the required rate of return might be appealing, even though its proper use would be rather restricted.

The downside is that the required rate of return or some other *ad hoc* rate might be used as the discount rate outside of the narrow cases in which doing so is reasonable. For example, the computation does not work if the two projects have different dollar returns. Another example is that someone who thinks the required rate of return is a generally applicable discount rate may "discover" some false principles. Among the most troubling of these false principles is the notion that NPV is affected by the level of debt even in the context of a perfect credit market.

Does Debt Matter?

Debt does appear to affect NPV when the chosen discount rate differs from the true discount rate. In Figure 13, we examine project *B*'s NPV, with and without borrowing and lending opportunities, using the required rate (assumed, as before, to be project *A's* IRR). NPV based on the required rate of return, NPV*, for project *B* equals line segment *ef* if there is no debt, while it equals *eg* if the individual borrows to point *h*. Note that NPV* is greater if there is debt than if there is no debt. In fact, NPV* rises in this example with each additional dollar of debt!

Figure 13

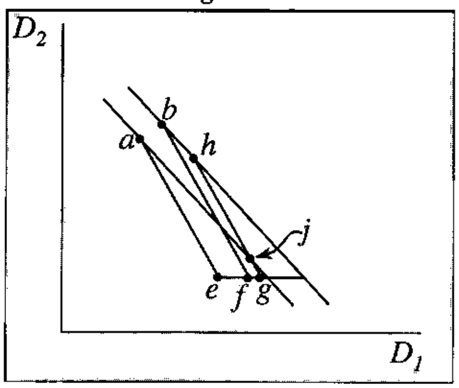

Figure 13 shows two projects, with *B* ranked higher than *A* under both the NPV and IRR criteria. Assume that project *A* costs $3,500 and returns $6,125; while project *B* costs $2,500 and returns $6,750. Project *A*'s IRR is 75%, while project *B*'s is 170%; *A*'s NPV is $2,068, while *B*'s is $3,636. Suppose that after investing in *B*, the individual were to borrow $1,500, paying $1,650 back in year 2. We compute the NPV of that position, using the required rate of return as the discount rate, as $1,914 [($6,750 – $1,650)/1.75 – ($2,500 – $1,500)]. There is no ready economic interpretation for this result.

There is, of course, something wrong with this picture. The rate we have used in discounting future returns does not equal the interest rate *or* the true discount rate (derived from the slope of the indifference curve at its tangency with a credit opportunity line). A resulting problem is that we can no longer interpret NPV as the minimum subsidy needed to get the individual to voluntarily forego an investment opportunity. In fact, the situation is even more serious. Figure 13 reveals that there is no way to identify *B*'s superiority over *A*, if we assume that the same discount rate is used. That is, there is always a debt position associated with *A* that has the same NPV* as some debt position associated with *B*; for example, points *h* and *j* have the same NPV*.

The possibility of borrowing does *not* affect NPV if the true discount rate is used. Of course, the amount of debt does matter with respect to the ultimate consumer optimum, but that issue is beyond the scope of the current discussion. It is enough to say that maximizing NPV is *consistent* with maximizing consumer satisfaction, but does not guarantee it.

Summary and Conclusions

In this discussion, we demonstrate the superiority of the NPV rule over its IRR counterpart without the usual duet of boogeymen typically found in textbook coverages of capital budgeting: multiple IRRs and reinvestment. Neither of these issues even arises in our two-period model. Instead, NPV is shown as superior because its use is consistent with maximizing consumer satisfaction, whereas there are situations in which the higher IRR choice is not similarly consistent.

NPV is shown as superior because its use is consistent with maximizing consumer satisfaction, whereas there are situations in which the higher IRR choice is not similarly consistent.

A number of related topics are omitted from our discussion, both because of space constraints and because their inclusion would greatly increase the level of complexity. Prominent among these topics are *taxes* and *risk*. It would also be useful to examine links between the state preference model and the standard coverage of portfolio analysis, with its focus on risk and return. These questions may be addressed in future issues. ■

Notes

1. Someone with little year-1 consumption (at a northwestern point along an indifference curve) would have a high discount rate, giving up a lot of year-2 dollars to gain even a small quantity of extra year-1 dollars. Someone with few year-2 dollars (at a southeastern point along a curve) would have a low discount rate, giving up some scarce year-2 dollars only in return for many extra year-1 dollars.
2. We know that any point on a higher indifference curve is preferred to any point on a lower one, and that any point along a given curve is equivalent in terms of satisfaction to any other point along that particular curve. If two indifference curves were to intersect, consumption represented by the point of intersection would be both equivalent *and* inferior to certain points on each of the intersecting curves.
3. If there were an infinite number of tiny investments that we could rank from better to worse, the required rate of return would equal the marginal IRR, which would equal the interest rate. Thus, the problem here is related to our assumption that there are only two possible investments.