# The Cost of Capital: If Not the CAPM, Then What?

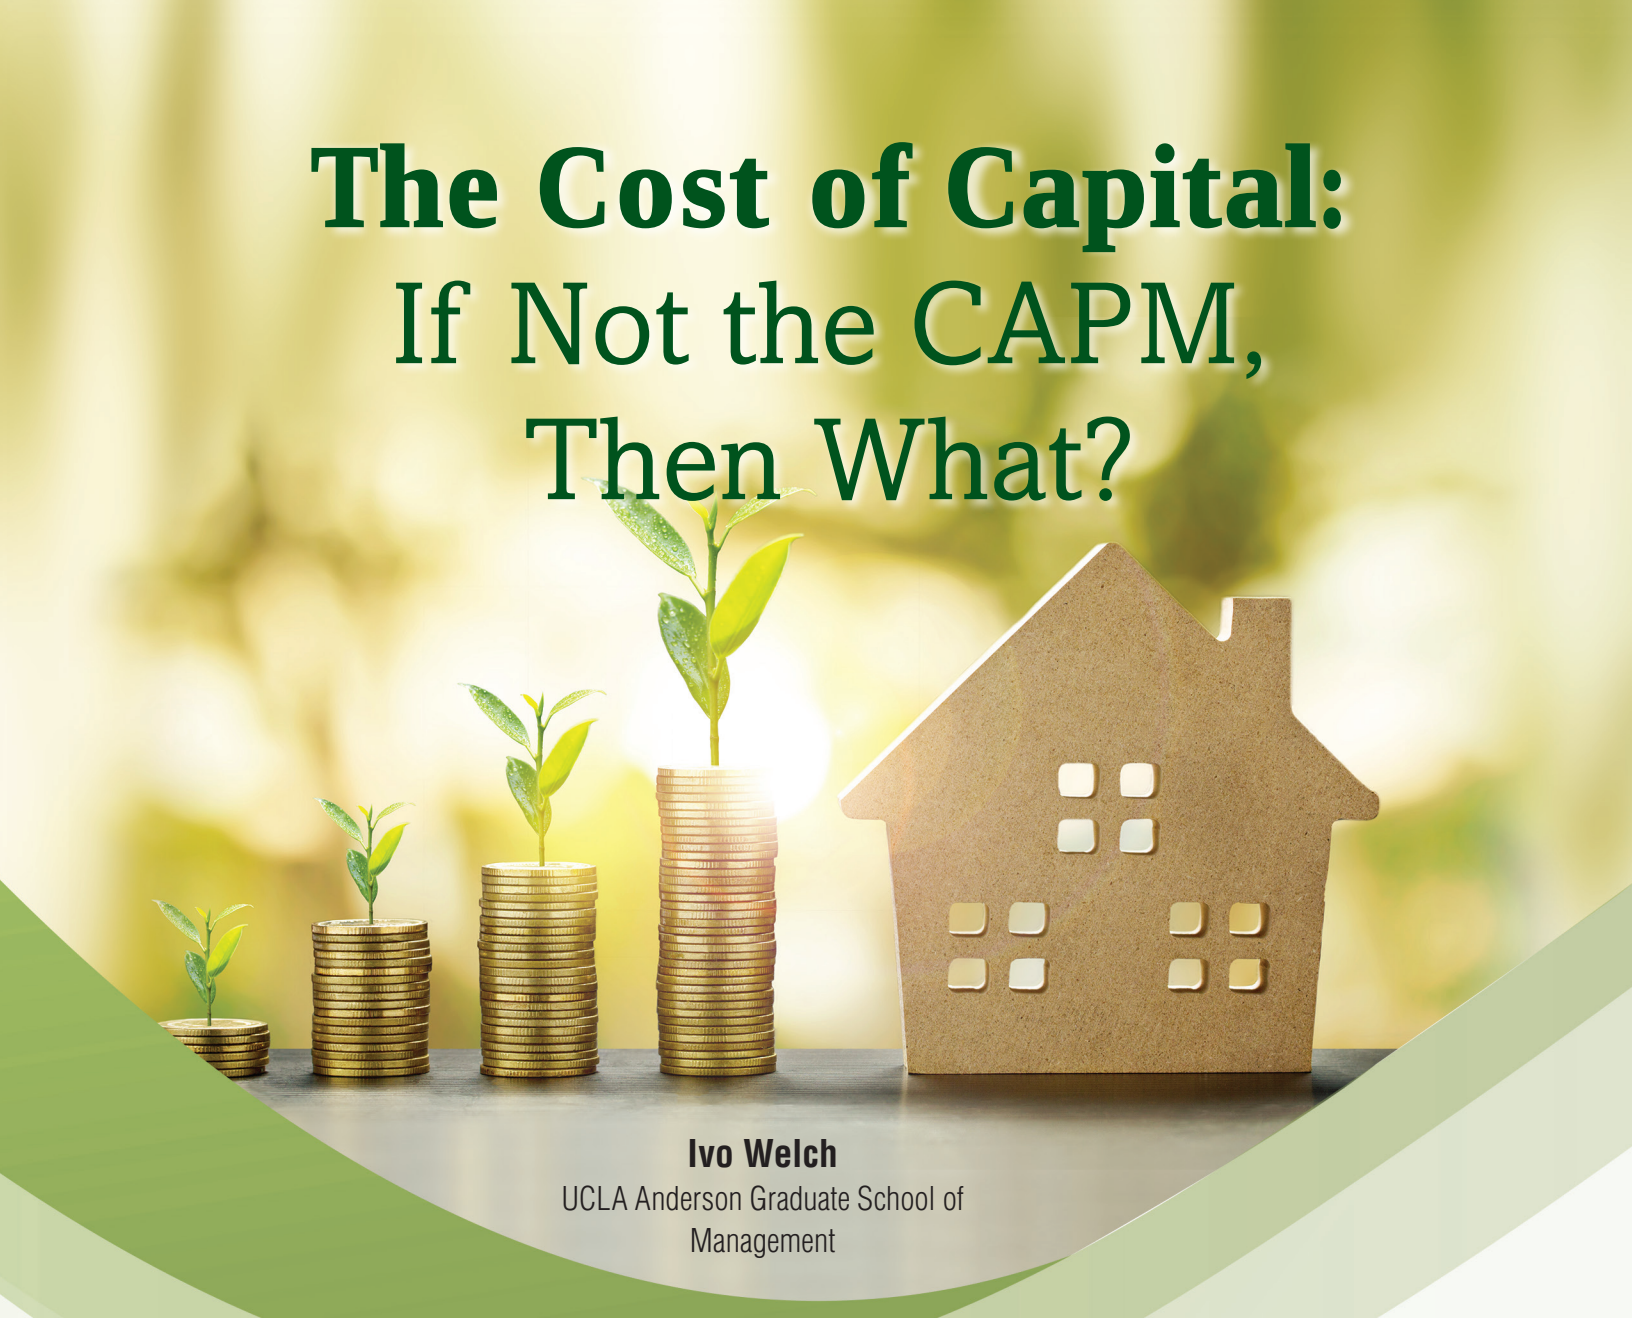

**Ivo Welch**
UCLA Anderson Graduate School of Management

*Twenty years ago, it would have been considered heresy to doubt the usefulness of the capital asset pricing model (CAPM) in assessing the cost of capital. Ivo Welch argues that, today, the CAPM should not just be doubted—it should be discarded.*

## The Cost of Capital and the CAPM Are Ubiquitous

Few numbers in business are more important than the cost of capital (CC). It is not possible to make intelligent investment choices without a good estimate of the CC. It determines which projects should be taken on and which should not. If a project has a CC higher than its expected rate of return, then it costs more than it is worth. And the CC is pervasive. For example, the cost of steel may influence construction projects, but it won't influence, say, a sales campaign. The CC, on the other hand, influences *all* projects.

The workhorse CC model for nearly half a century has been the Capital Asset Pricing Model, or CAPM. It dominates textbooks, teaching, and practice. Over 90 percent of all publicly-traded companies use it. Courts and appraisers also use it. In many contexts, it is even the only accredited model.

Unfortunately—and I write this with a heavy heart—the CAPM is not just imperfect; it is so badly wrong that it is best ignored.

Before I explain why, you are probably thinking: very well, but then why is this model (still) being taught and used so ubiquitously? The answer is disconcerting. It was we academics who committed the original sin. For a long time we were so enamored with the intrinsic beauty of our model that we simply ignored the evidence. It took a long time for us to come to grips with reality. In the meanwhile, the CAPM had taken on a life of its own. These days, the servant has become the master; the model has to be taught *because* it is so widely used. Knowing how to use the model is important because others are using it and they expect everyone else to know and to use it. In a world of Ptolemaists, it is unimportant whether the sun actually revolves around the earth. What's important is to know how Ptolemaists calculate epicycles.

The CAPM is the cozy bedtime story that tells students and practitioners that the world is in good order and that they have learned something which will allow them to understand it.

So even today, most business schools still teach the CAPM as their main model—even though nearly all finance professors know perfectly well that the model fails all evidentiary standards. Remarkably, we finance professors do not disagree about the evidence. We do, however, disagree about what we should teach instead. Most of us remain more comfortable teaching a beautiful toy model that we fully understand than teaching ad hoc prescriptions of which we understand only bits and pieces. The CAPM is the cozy bedtime story that tells students and practitioners that the world is in good order and that they have learned something which will allow them to understand it. But the real world isn't like that.

### *A Brief Primer on the CAPM*

For readers who have forgotten what the CAPM was all about, or haven't gotten that far yet, let me first briefly explain it, and why it is so beautiful, simple, and useful. The CAPM is a model with only three inputs:

1. the (economy-wide) risk-free rate of interest;
2. the (economy-wide) expected risky rate of return (usually on the stock market);
3. and a "market-beta" (or just "beta," for short) that measures diversifiability.

*Beta risk* is not the own volatility (standard deviation) risk of a project. Instead, it is a measure of how the rate of return on a project correlates with the overall rate of return of the stock market *on average*. A beta of 2 means that for every 1 percent increase or decrease in the stock market, the project tends to go up or down by 2 percent. A beta of –1 means that when the market goes up by 1 percent, the project tends to go down by 1 percent, and vice-versa. A beta of 1 means the project typically moves (noisily) in

**Figure A:** The Security Market Line

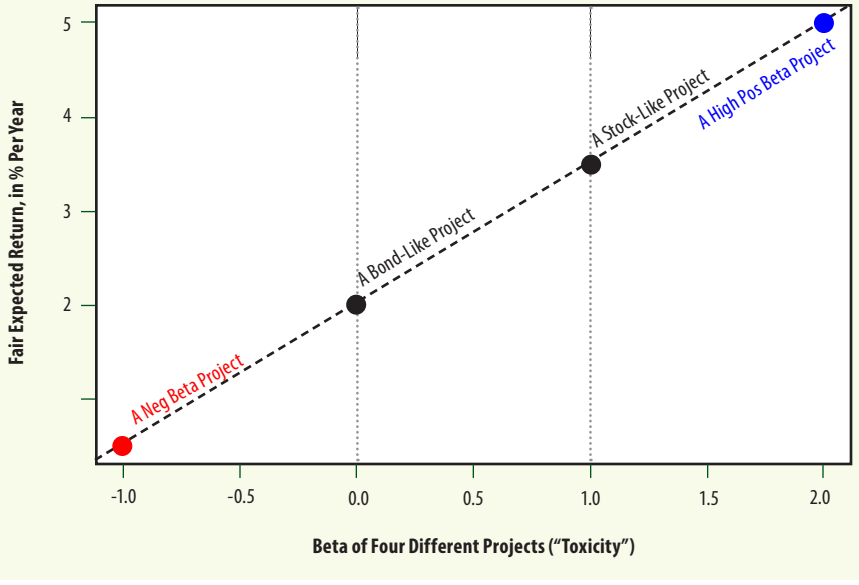

line with the stock market. And a beta of 0 means the project typically does not move together with the market. Figure A shows the security market line, which is the graphic analog of the CAPM.

In a heavily diversified market-like investment portfolio, a good way to think of the market-beta of a project is as a measure of its *toxicity*. Holding the expected rate of return equal, an investor would prefer a project if it rescues her in the next market crash, i.e., if it tends to go up when the rest of her portfolio goes down. In extremis, investors can even be happy with projects which have negative expected rates of return, just as long as their betas are negative enough.

If this seems absurd, think about insurance. Most of the time, the insured pay the premium but get nothing back. The expected rate of return on insurance is usually negative. But customers still want to purchase insurance, because the upside appears in the worst eventuality (e.g., when the house burns down). Likewise, a project with the right beta pays off in the worst stock market eventuality (e.g. the Covid 19 epidemic)

The CAPM says that after investors have done all that smart diversification that allows them to reduce their risk *for free*, what remains is a trade-off between beta risk and reward. If a project has too much reward for its beta risk, too many investors will rush in, drive the price up, and thereby drive the reward down. In the perfect world of the CAPM, such stampedes happen so instantly as to be barely perceptible.

And beyond its beautiful and intuitive logic and graphical representation, the CAPM comes with a beautiful and simple quantifying formula:

**Expected Return of any Investment = Risk-Free Rate**
**+ (Expected Return on the Market – Risk-Free Rate) × Investment's Market Beta**

The formula makes sense. All three inputs are of first-order importance even if the CAPM is not true: (1) The risk-free rate is about whether it is better to consume or better to save and invest. (2) The market premium (also called the equity premium) is about whether it is better to invest in risky projects or non-risky projects. Because risky stocks should earn more than risk-free bonds, it should be positive. (3) *The market-beta is the best measure of the undiversifiable market risk of an investment project*. It is important to any investor who holds mostly the market portfolio. This is true even if the CAPM is wrong.

Even better, by quantifying generic intuition, the CAPM formula gives managers a concrete capital-budgeting input to use in their spreadsheets. It is a common (but not necessarily correct) practice to use historical rates of return for the equity premium. Over the last fifty years, short-term government bills, long-term government bonds, and (long-term) U.S. stocks have produced a geometric average pre-tax rate of return of about 5 percent, 8 percent, and 10 percent per year, respectively. (Inflation ran just under 4 percent.) *If* history is a good guide, a good forward-looking market-premium estimate today would therefore be about 2-3 percent per year.[1] With some finesse, it is easy to estimate a good predictive measure of market-beta. (Warning: Be aware that the beta measures posted on popular websites are generally not very good.[2])

Of course, even if the CAPM were a good model, it would still require well-reasoned inputs and correct usage. Jacobs and Shivdasani's 2012 article[3] provides an excellent overview of common mistakes made in the application of the CAPM. Even in the most capable hands, the CAPM has only ever been considered applicable to large companies held by well-diversified investors. It has often been used incorrectly in situations in which the owners are not well diversified and the available capital markets are imperfect.

## So What is Wrong?

All models are wrong—they are only models, after all. So why be so harsh to the CAPM? Because the CAPM is worse than *just a little* wrong. The data proves that the CAPM is worse than useless. The primary disagreement which remains among finance professors is whether it is merely worse than useless or *statistically significantly* worse than useless.

I could write a treatise on the theoretical and empirical nuances of the CAPM—befitting a PhD practitioner of the Aristotelean art of defending the CAPM—but the truth is much simpler and, once exhibited, difficult to forget.

Ultimately, the CAPM provides one basic prediction: high-beta stocks should outperform low-beta stocks on average, because high-beta stocks are riskier. Unfortunately, the data say the opposite. Even over long periods, average rates of return have been higher for stocks with low betas than stocks with high betas—the opposite of what the CAPM claims. The CAPM prescribes high expected returns for exactly those stocks and industries that have shown low average returns, and vice-versa.

Figure B shows this effect. It plots the historical rates of return for two portfolios, one with all publicly traded stocks in the lowest beta tercile, the other with all stocks in the highest beta tercile. (All portfolio assignments were determined at the start of each year.) The plot shows no evidence that high-beta stocks offered higher average rates of return. Simply put, the high-beta stocks were doubly bad deals for investors who mostly held the overall stock market. They had higher risk and lower average rates of

**Figure B:** The Performance of High- and Low-Beta Stocks

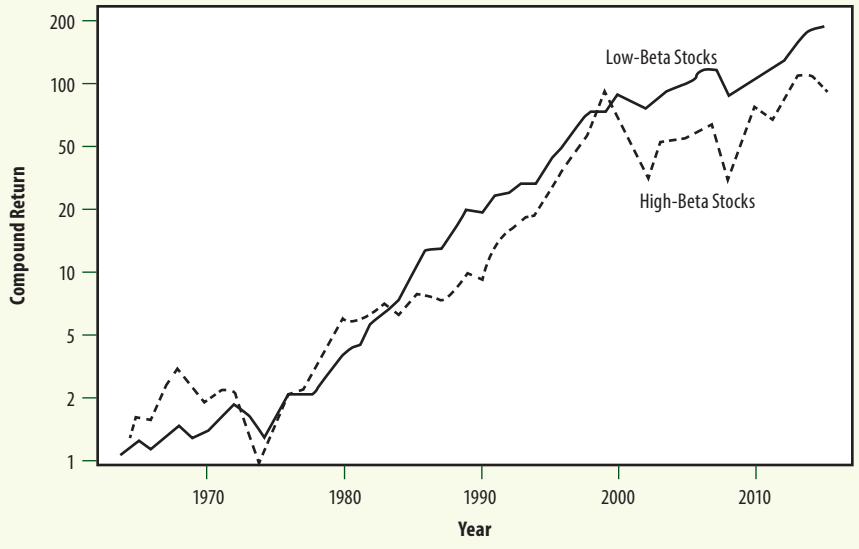


return. In a CAPM world, this should not be the case.

Who would want to put their faith into a rocket that, historically, has exploded half the time?

Of course, this historical return pattern may just be bad sampling luck. If high-beta stocks truly had higher expected returns than low-beta stocks, and Figure B *merely* describes how it happened to turn out during a bad stretch, then maybe the CAPM was just unlucky. This could be true. But then who would want to put their faith into a rocket that, historically, has exploded half the time? Yes, the rocket may be well designed and the explosions merely a run of bad luck. But with this track record, would you want to climb aboard?

Recalling your business school days, with finance professors teaching arbitrage conditions left and right, most notably in options pricing, doesn't the failure of the CAPM violate some sacred natural arbitrage condition? *No*. The CAPM relies on so many strong assumptions that are violated in practice, that in retrospect it would have been a miracle if it had held true.

Nevertheless, that stocks with higher market-betas had lower average rates of return is an uncomfortable truth even in a non-CAPM world. Stocks with low market-betas have always been less risky (to heavily diversified investors) than stocks with high market-betas. Why, then, have these low-beta stocks offered average rates of return that were just as good? We don't really know. It seems perverse, but what diversified investor would not prefer low-beta stocks over high-beta stocks? The evidence does suggest some practical investment advice: as an investor, you should tilt your portfolio towards low-beta stocks. Such portfolios tend to suffer less overall risk for the same average rate of return than un-tilted market portfolios. Just remember, past performance is no guarantee of future performance. Still, and unsurprisingly, the low-beta portfolio strategy also performed very well in the COVID 19 market crash of March 2020.

## Practical Managerial Advice

The failure of the CAPM may present investors with a good opportunity, but it creates a dilemma for corporate managers. How should they calculate the rate of return that diffuse public investors will demand from their projects? If they should not use the CAPM to estimate their hurdle rates, that is, their opportunity costs of capital for similar projects, what model should they use?

Finance researchers have some useful advice. Yet, unfortunately, even professors who dedicate their lives to exploring this subject do not understand the full picture. The real world of managers will never be as neat and nice as the CAPM was.

Assuming such perfectly equal CCs seems to be at least as good a method as using CCs suggested by the CAPM, even for large *publicly-traded firms.*

One running joke among finance professors[4] is that the CC for equity is always the same (maybe 8-10 percent per year arithmetic, 6-8 percent per year geometric). These numbers are equivalent to using the CAPM with (badly estimated) identical market-betas on each and every project equal to 1.0. Given the evidence that high-beta stocks offer lower average rates of return than low-beta stocks, assuming such perfectly equal CCs seems to be at least as good a method as using CCs suggested by the CAPM, even for large *publicly-traded firms.*

None of this means that either beta or risk does not matter. Beta

still measures the diversification effectiveness of projects for most investors. Outcome probabilities still have an immediate correlation to expected cash flows. It also doesn't mean that all projects should actually have the same CC. Furthermore, this equal CC prescription is only appropriate for the equity component and not the debt component. (For small non-diffusely held firms, the CAPM was never appropriate to begin with.) Moreover, there are other important considerations that have or should have played a role in good practice even if the CAPM had, by and large, held. Again, Jacobs and Shivdasani cover a good number of potential CAPM misapplications, many of which still exist in this brave, and without a CAPM, messier, new world.

## Ten Capital-Budgeting and Cost-of-Capital Suggestions

So how can we do better? Unfortunately, in the absence of a rigorous universal model to replace the CAPM, there are no hard rules or formulas. We can only fall back on situational prescriptions widely believed to be solid (at least among finance professors). These require more judgment than the CAPM did, with best practice changing from project to project. In many cases, they require customer-specific estimates and assessments. Of course, project specific estimates have always been important to the net present value (NPV) numerator for expected cash flow estimates. Now they are also important in the NPV denominator, the CC estimates.

### Project-Specific

My first two recommendations pertain generically to specific project valuation:

1. **Comparables: Avoid NPV/IRR (internal rate of return) capital budgeting decisions *if* you can.** If there is a competitive and liquid market for similar projects, chances are that there are good comps, that is comparables, in the market. Rely on them! Good comps tend to be better than NPV analyses. Not only does NPV analysis require estimating the CC, it also requires estimating expected cash flows. If dozens of similar goods have recently been sold, why not instead learn from the value assessments of their buyers and sellers?
   For example, if buildings have sold at twenty times the rental rate, chances are that a similar building is a bargain if the seller asks for fifteen times the rental rate and a dud if the seller asks for twenty-five times. Although an NPV and sensitivity analysis can still help you to better understand the economics of buildings, chances are that ad hoc situation-adjusted comps will be more accurate in pricing buildings than even the best NPV analyses.
2. **Cost-of-Capital Weighting: The overall CC remains a weighted average of debt and equity CC.** WACC (the weighted average cost of capital on debt and equity) works just as well without a CAPM. Debt often provides cheaper project financing than equity, especially for firms that have use for the corporate income tax shelter that debt provides. For example, a building may be financed by a conforming mortgage covering 80 percent of the cost at 6 percent a year, or a jumbo mortgage covering 90 percent at 7 percent a year. If the cost of equity capital remains approximately 10 percent a year regardless of capital structure, the CC is 6.8 percent with the conforming mortgage and 7.3 percent with the jumbo. For a firm in a 60 percent corporate income tax bracket, the WACC is 4.88 percent for the conforming and 4.78 percent for the jumbo.

   Using the same CC for equity regardless of leverage creates a puzzle: Why would the cost of equity not be higher in the more levered capital structure? The residual equity is much riskier in the jumbo capital structure than in the conforming one. Yet, although the investors *should* care, the empirical evidence for our publicly traded stocks suggests that they do not. They have not demanded or received sufficiently better terms to compensate them for the higher risk of the more levered equity. No one understands why.

### Factor Adjustments

Although market-beta does not predict expected *equity* returns, there are other analogous methods that do. From a zoo of similar regularities I have chosen four empirical regularities that [1] make sense to me and [2] have persisted for a long time. These are good candidates for judging the informal CC of equity adjustments.

3. **Market Cap Adjustment: Smaller firms have to offer higher average returns.**
   Market cap may be used as a proxy for a lot of different attributes, perhaps the most important being access to perfect capital markets. Smaller firms can often obtain capital only on worse terms than larger ones. From 1995 to 2018, the average publicly-traded firm in the bottom 30 percent by market cap (averaging about $300 million in 2018) returned a geometric 10 percent per year. Firms in the top 70 percent returned only about 9 percent per year. So the market cap spread was about 1 percent.[5] Although we do not have systematic data from which to estimate the historical spread for small non-publicly traded firms, it is

almost surely even higher than 1 percent.

4. **Project Duration Adjustment: Longer-term assets have to offer higher average returns.** Projects with longer lifespans require higher CCs. A prevailing "term premium" can be estimated from yield curves found on the web. As I write this, for example, high-quality corporate bonds yield about 2 percent for one-year bonds, 3 percent for ten-year bonds, and 4 percent for fifty-year bonds. A rental commitment of one year should thus be discounted at about a 1 percent higher CC than a ten year commitment.
5. **Liquidity Adjustment: Less liquid assets have to offer higher average returns.** Illiquid assets are those that have high round trip transaction costs—and illiquidity can increase quickly in a financial crises. For example, in the 2009 crisis, only treasuries were supremely liquid and even GNMA bonds (backed by the U.S. government) were so difficult to resell that they commanded a 1 percent higher spread (lower price). Many other assets, especially bonds of various kinds, were far less liquid, and some briefly sold for cents on the dollar. Even in ordinary times, a half-finished building project may have comparatively high spreads, because half-finished buildings can be sold only slowly or at fire-sale prices. The courts typically apply liquidity discounts of about 25 percent to privately held assets to reflect lack of marketability even in good times. This discount seems much too high, but then courts are composed of legal rather than financial experts.
6. **Asset-Class Adjustment: Riskier classes of assets have to offer higher average returns.** We do not understand why riskier stocks have not offered higher average returns, but it seems to be mostly a phenomenon internal to the stock market.[6] Over the last fifty years, stocks overall have outperformed risk-free government bonds by about 2-3 percent a year, and bills by 4-6 percent.[7] Projects that look more like debt investments (and financing that is more debt-like) should thus be assumed to have a lower CC than equivalent projects that look like equity.

## Investor-Specific

Models like the CAPM assume perfect capital markets: All investors are alike and compete, so only the project characteristics matter. This view is often too simplistic. Instead, it seems that the CC depends on both project supply (their future cash flows) and project demand (*available* investor capital).

7. **Relevant Exposures: The appropriate CC depends on the investor.** For example, outside capital may not be competitively available to entrepreneurs whose entire wealth is in their firms. For them, it doesn't even make sense to think of their CC in terms of covariation (beta) with the stock market. Instead, what matters to them is their firm-specific risk exposure, measured by the variance or volatility. Investor heterogeneity can also be caused by preferences or tax status. Tax-exempt investors should have a lower CC and therefore offer capital at lower prices.

Whether investors differ because of risk aversion, preferences, or tax bracket, the appropriate CC must reflect not just the project, but the specific situation of the owners. Valuation is far more difficult when a single project may have a CC of 5 percent for one investor and a CC of 10 percent for another.

## Analysis-Specific

Along with these specific technical recommendations, I want to close with some general analysis advice. By necessity, this is even more vague than my preceding suggestions.

8. **Conflicts of Interest: Know others.** Many bad capital-budgeting inputs and outputs occur because someone (you?) wants to arrive at a particular answer. If your input estimates come from someone who has a conflict of interest, do not trust them. Estimates from employees, lawyers, investment bankers, and so forth are rarely unbiased. Think about what's in it for them. Solicit advice from various parties, ideally ones with opposite motives. Consider hiring a devil's advocate consultant, tasked with talking you out of whatever course of action and cost of capital you prefer. And be warned: You may not like what you hear.[8]
9. **Judgment Errors: Know thyself.** We are all overly optimistic, but admitting we have a problem is the first step to recovery. The most important factor, and perhaps the most obvious, is forgetting the realistic probability of utter failure and so incorrectly judging the most likely outcomes. For each 1 percent probability that a pandemic, earthquake, fire, plane crash, or employee death will wipe out entire projects, the internal rate of return for the project will be 1 percent lower. Recognizing common human biases can sometimes justify conservatively increasing project investment hurdle rates to exceed the CC.

Many managers, especially pessimistic ones, like scenario analyses which help them incorporate potential failures into their estimates. I am skeptical about

the usefulness of these analyses for CC assessments, though not for expected cash flow assessments. I do not know of a way to effectively use bad scenarios in a formal decision process.

10. **Humility: CC assessments for non-trivial projects have low accuracy.**
Except for fixed-income investments, the most important prescription for assessing the CC may well be humility—a quality that does not come easily to either practitioners or finance professors. Even the best CC estimates are rough.[9] It is cold comfort that the CC estimates in the NPV denominator are still usually better than the expected cash flow estimates in the NPV numerator when you need both to be sound. Sensitivity analyses can help you understand the project better, but we don't really know how to incorporate them appropriately into our decision processes, either. It may be better to adopt a pessimistic view than the average one.

While students often believe that theory is more difficult than practice, unfortunately, the opposite is true.

To many managers who already have an intuitive understanding of their own ignorance, the flaws in the CAPM may not be news. Such managers tend to err on the side of caution.[10] Jagannathan et al showed that firms, on average, reported using *twice* their own estimated CC as their hurdle rates! So even firms with abundant access to capital may prefer to forgo many profitable projects.

I close with the observation that, while students often believe that theory is more difficult than practice, unfortunately, the opposite is true.

## A Speculative Capital-Structure Theory with Capital-Budgeting Implications

In a Modigliani-Miller (M&M) perfect capital market, the overall WACC remains the same regardless of capital structure. Mathematically, a capital structure with more leverage has a higher cost of debt and a higher cost of equity but tilts the weighting from higher-cost equity towards lower-cost debt. Of course, the Modigliani-Miller world is primarily a thought experiment.

When the capital markets are not perfect, firms can minimize their WACC by choosing the best capital structure—the one that minimizes their tax obligation, moral hazard and agency conflicts, adverse information disclosure, transaction costs, etc. Nevertheless, the M&M indifference prescription remains surprisingly accurate *as long as* debt is less than, say, half of the firm's financing. This works because WACC tends to be very insensitive to modest levels of leverage. Put differently, it matters little whether a firm chooses a capital structure of 10 percent debt or 30 percent debt; the WACC typically remains about the same.

Fine-tuning their optimal choice of leverage really matters only for firms that are high-leverage (say, 80 percent or more), such as financial services firms, firms near financial distress, or firms in leveraged buyouts.

As shown in Figure B above, the empirical evidence for large publicly traded firms suggests that the expected rate of return on equity does not increase with market-beta and leverage. What does this imply to an enlightened manager about optimal capital-structure and WACC?

Of the three effects of leverage (higher CC on debt, higher CC on equity, more weight on the debt component) only the first and last remain. The WACC then decreases as long as the expected rate of return on marginal debt remains below the expected rate of return on equity. Managers can thus obtain the lowest WACC with a capital structure in which the expected rate of return on debt is equal to the (roughly constant) expected rate of return on equity *on the margin*. (The average cost of capital on debt should be lower.)

To the extent that managers care only about equity returns (if only because debt default could get them fired), and if the debt comes from external capital providers (and not from the equity holders themselves), they can follow an even simpler rule. They can compare the *quoted* interest rate on debt to the expected rate of return on equity, and use debt financing until the two become equal. If they pursue this capital structure policy, then all three rates are the same: the CC for capital budgeting purposes, the quoted rate of return on debt, and the expected rate of return on equity. This equal cost of equity capital structure theory has clear flaws, but it may be more realistic and useful than its competitors. The evidence will tell. ■

**Ivo Welch** is the J. Fred Weston Professor of Finance at the UCLA Anderson Graduate School of Management, a 2015 Humboldt fellow, editor of the Critical Finance Review, and author of a free corporate finance textbook. His published work has been highly influential both in finance and economics. His current research agenda focuses on the estimation and uses of market-beta.

## Endnotes

1. The shorter bill rate should be used for short-term projects in both the risk-free rate and the equity premium, and the long-term bonds rate for long-term projects. A more detailed discussion (with information about data) appears in Part II of my free corporate finance textbook at book.ivo-welch.info.
2. The best-known estimator of market-beta is also one of the easiest. It is explained in Welch (2019) available on SSRN, 3371240.
3. Jacobs, Michael T., and Anil Shivdasani. "Do You Know Your Cost of Capital?" Harvard Business Review. July/August 2012. https://hbr.org/2012/07/do-you-know-your-cost-of-capital.
4. Finance professors usually lack a sense of humor.
5. From 1962 to 2018, small stocks offered an even higher 1.5 percent per year bonus spread over large stocks. The compound rate of return reflects not only the arithmetic mean, but also the volatility. For example, projects with a volatility of 20 percent offer 2 percent lower rates of return than projects with 0 percent volatility.
6. Interestingly, within stocks, there is some evidence that riskier stocks have not outbut underperformed. (Besides, anyone who can predict equity performance should quit corporate practice and work for a hedge fund. This sort of predictive capability has exceeded that of even the most sophisticated finance professors.)
7. My quoted costs of capital are context specific, but they are also lower than those often quoted. This is primarily because a geometric equity premium of 1-3 percent per annum with respect to long-term bonds over long horizons is a reasonable estimate, equivalent to an arithmetic premium of about 3-5 percent.
8. In short, my advice is to thoroughly understand the problem. I do not advise you to succumb to inertia through endless analyses and committees.
9. A manifestation of this lack of humility is the frequent occurrence of equity CC estimates quoted as percents with several digits after the decimal point. Outside of fixed-income cash flows, such pseudo-precision is laughable. We have nowhere near the ability to assess CC at better than 1% intervals.
10. Jagannathan, Ravi, David A. Matsa, Iwan Meier, and Vefa Tarhan. "Why do Firms Use High Discount Rates" Journal of Financial Economics. 120, no. 3 (2016) https://econpapers.repec.org/article/eeejfinec/v_3a120_3ay_3a2016_3ai_3a3_3ap_3a445-463.htm