

Uncertainty, Real Options, and Industrial Energy Efficiency Decisions

Steve Kihm and Claire Cowan, Energy Center of Wisconsin

ABSTRACT

Why do industrial firms require quick paybacks on energy efficiency measures? Real options pricing models can shed light on this question. The discounted cash flow (DCF) model that energy analysts typically use to assess energy efficiency projects fails to consider the dynamic setting in which firms operate. Requiring quick paybacks may be an *ad hoc* way for firms to consider the flexibility and timing considerations that are absent from the DCF approach. In other words, maybe requiring quick paybacks on energy efficiency projects is economically rational after all.

Introduction

As the title for this conference suggests, timing is a key consideration in energy efficiency investment decisions. Energy markets, and energy decisions, are dynamic. As market conditions change, a good energy efficiency investment today might turn out to be an even better investment if we wait until next year to make it (*e.g.*, if technology costs drop substantially). Alternatively, if energy costs decline significantly after a potential energy efficiency project is identified, we might later be glad we did not make the investment. The DCF model typically used to analyze potential energy efficiency investments is a static model that is not well suited to address the dynamic aspects of such decisions.

Consider the following series of events. An energy manager conducts a DCF analysis on various energy efficiency improvement projects applicable to his or her firm. The firm's corporate finance chief agrees that the analyst's input figures are correct, but decides that the company will fund only those projects with a payback of two years or less.

Most energy efficiency experts suggest that, as long as capital is available, all energy efficiency projects with a positive net present value (NPV) should be implemented. When firms pass up efficiency projects with positive NPVs, some efficiency experts suggest that such behavior is economically irrational. However, this conclusion is inconsistent with the principles taught in business schools.

In a dynamic setting, a positive DCF-based NPV result is not a sufficient basis to proceed with a project. Consider the following quote from a leading corporate finance text (Brealy, Myers, & Allen 2006):

The fact that a project has a positive NPV does not mean that you should go ahead today. *It may be better to wait and see how the market develops.* (Emphasis added.)

The DCF model makes the implicit assumption that there are only two possible decisions regarding a potential energy efficiency project: either we implement the project or we reject it. In many real world applications, there are actually three possible decisions: (1) implement the project today; (2) reject the project today; and (3) defer the decision to a later date. To fully evaluate these alternatives we must move beyond a static DCF-based analysis to one that models

dynamic decision processes. Options pricing models accomplish this objective. Options-based models behave quite differently from DCF models, and often provide surprising insights as to the investment decisions of industrial firms.

Shortcomings of the DCF Approach in a Dynamic World

The limitations of the standard DCF model in a dynamic capital budgeting process were first identified in the 1970s by Stewart Myers, a financial economist at the Massachusetts Institute of Technology (MIT). He noted that the DCF model fares poorly in its ability to estimate the economic value of projects if there is a high degree of uncertainty and if the decision to invest can be deferred to a later date (Myers 2003). Myers termed the firm's ability to change its decisions based on changes in future events a "real option." Options provide the holder with the *right*, but not the *obligation*, to take a particular action at some future date.

Some corporate finance departments have moved away from standard DCF analysis and have turned instead to options pricing models. The more risky the business, the more likely it is that they will use options-based models in lieu of traditional DCF models. Consider the quote from the Chief Financial Officer (CFO) of Merck & Co. (Nichols 1994):

Merck's experience with R&D has given us a database of information that allow us to value risk or the volatility of our research projects, a key piece of information in options analysis. Therefore, if I use *option theory* to analyze the investment, I have a tool to examine uncertainty *and to value it*. (Emphasis added.)

The DCF model treats uncertainty as a curse. Options pricing models view uncertainty as a blessing. This point is highly counterintuitive, yet absolutely true. The option value of a dynamic project increases when uncertainty increases (Damodaran 2003). That is because options-based decisions vary depending on whether uncertainty helps or hurts the firm. Put another way, when we have options, we can not only embrace uncertainty when it delivers windfalls, but we can also cut our losses when it delivers negative results. That is because an option bestows a right and not an obligation on its owner.

Options Pricing Concepts on the Home Front

Before we consider an industrial energy efficiency example, let us consider how options-based thinking comes into play in our every day lives. For example, even though you may believe that you would benefit from the purchase of a hybrid-powered automobile today, you might decide to defer the purchase decision until next year due to the uncertainty associated with hybrid vehicle prices. Such an approach could be economically rational. The possibility of getting the vehicle at a substantially lower cost next year may be worth taking the risk that the market prices move in the opposite direction.

Note that this approach does not assume we know which direction vehicles prices will go. It assumes only that those prices are volatile. When uncertainty is high, the option value of waiting (deferring the decision) is high. On the other hand, if future prices are fairly certain, the option value of waiting is likely to be small.

Under our hybrid vehicle example, there could be uncertainty associated not only with future hybrid vehicle prices, but with other aspects of the purchase decision as well. Future gasoline prices, which determine your avoided costs, could be noticeably lower or higher than they are today. Furthermore, future prices of substitute vehicles, such as plug-in all-electric vehicles, could be noticeably higher or lower than they are today.

All of these uncertainties could lead to a very different benefit-cost ratio one year from now. The key is not merely that there is uncertainty—it is that *we can change the decision depending on how the uncertainty plays out*. If hybrid vehicle prices decline substantially, for example, and if we buy next year, as long as other factors remain about the same, we then get a much better deal on the vehicle. On the other hand, if plug-in vehicles drop in price while hybrid prices increase, we can walk away from the hybrid vehicle and purchase the plug-in. If we purchase the hybrid vehicle today, we cannot play the uncertainty game to our benefit. For better or worse, we are stuck with our original decision.

In simple terms, when the future is fairly certain, if a project or purchase produces net economic benefits today, it generally pays to act now. If the future is highly uncertain, however, the rational economic choice for such an investment is often to wait.

An Industrial Example

Pursuing a similar line of reasoning in the context of an industrial manufacturing facility, we analyze a natural gas efficiency measure with a ten-year life. The upfront cost is \$1,000. The vendor is willing to guarantee a level of energy savings of 25 therms and lock in those savings at current natural gas prices (\$6.00 per MMBtu), which translates into energy cost savings of \$150 per year.

If the firm rejects the offer this year, it expects to be able to accept it next year under revised conditions, assuming those conditions are favorable. It can also reject the offer if conditions deteriorate. While this flexibility has value, it comes at the price of increased risk. Conversely, if the firm accepts the offer now, it represents a risk-free deal for the customer. We obtain a NPV of \$158 for this project. Since there is no risk, it is as if we have offered the customer \$158 for free. Why would anyone turn down this deal?

Answering this question requires that we consider risk in its proper context. Finance principles suggest that there is nothing inherently bad about taking on risk, although the concept has a bad connotation outside of finance. As Merck's CFO mentions, firms make money by taking on risk, not by avoiding it. She suggests that "the route to success is to put *more* money at risk, not less" (Nichols 1994). That is not to say that firms and individuals should take on risk without regard for the consequences, as occurred in the recent sub-prime crisis. Proper risk management is about taking on calculated risks that have a better-than-average chance of paying off in terms of increased returns, and that in many cases have limited downside potential.

In this context, a firm might turn down a risk-free project if it faced capital restrictions and if more attractive projects were available, even if those other projects involved taking on risk. If capital is scarce, the firm might prefer to finance projects with shorter paybacks. Note that the payback for the project described in this example is fairly long:

$$\text{payback} = \frac{\$1,000}{\$150/\text{yr}} = 6.7 \text{ years}$$

Many firms require project paybacks be two to three years, or less. In some cases this requirement is due to capital restrictions. However, capital constraints are not the only factor that can drive firms to short paybacks.


What if the firm had ample idle capital? Should it accept the deal? If we apply DCF-based thinking, the answer is yes. Using options-based thinking, however, provides a more complex and somewhat counterintuitive answer. If the future is quite certain, then we should accept the offer today. On the other hand, if the future is highly uncertain, we might want to wait until next year and reconsider the deal at that point. Note that by deferring, we would not merely be putting off the *implementation* for one year. We would be putting off the *decision* for one year. If we defer the decision, we can either accept or reject the offer based future conditions. So by deferring, we pass up the sure-thing investment in the hopes of realizing something better, while recognizing that the worst we can do is to miss one year of savings from the sure-thing investment.

Simulation of a Deferred Decision

To analyze the option value of deferring the decision, we begin with the assumption that there is no uncertainty. It makes intuitive sense that under this condition, deferring the decision provides no net benefit. When there is no uncertainty, the NPV associated with pushing off the decision for a positive-NPV project by one year will always be lower than the NPV associated with implementing the project today. If we commit to the project today the NPV is \$158. When there is no uncertainty, deferring the decision reduces the NPV to about \$150. No other analysis is necessary or appropriate here. Under this condition, if capital is available, the firm should accept the offer.

The situation changes when we introduce uncertainty. Let us assume that if we wait a year, the current \$1,000 equipment cost could range anywhere from \$900 to \$1,100. We also assume that volatile natural gas prices could range anywhere from \$4.00 to \$8.00 per MMBtu. While that is quite a wide span of prices, Figure 1 shows that such a range is consistent with conditions experienced in recent times.

Figure 1. Average Price of Natural Gas at the Wellhead (US), 2003-2008



Source: U.S. Department of Energy, Energy Information Administration (2009)

If we simply used the uncertainty ranges to conduct sensitivity analysis, we could still be working within the standard DCF construct. Though DCF-based sensitivity analysis helps us to decide whether or not to accept the offer today, we have to use another approach if we want to see how uncertainty might *change* the decision if we wait a year.

There are three basic approaches that can be used to model contingency-based decision making: (1) partial differential equations (*i.e.*, the Black-Scholes options pricing model); (2) binomial lattice structures (*i.e.*, decision trees); and (3) simulation. The first choice requires complex mathematical calculations that are not well-suited to provide a clear explanation for a broad audience. While the second approach is the most transparent, presenting such an analysis here would exceed the space limitations of this paper. For these reasons, we use a simulation approach to demonstrate the concept.

If we defer a decision to the future, we cannot know with certainty today what that decision will be. This presents an analytic challenge, but it is not a new one. Those conversant with dynamic programming will be on familiar turf in this respect. One year from today, if the NPV is positive, we will commit to the project. If the NPV is zero or negative, we will not. Whether the NPV is positive, negative, or zero depends on how uncertainty plays out over the coming year. To conduct the analysis, therefore, we must work backward in time from the future to the present, which is the construct used in dynamic programming.

We conducted 1,000 simulations, allowing the equipment cost and natural gas price to vary within the ranges identified above.¹ We applied dynamic decision-making within the simulation. We also considered timing differences in energy savings—if we defer the decision by one year, we push the energy savings stream one year into the future, thereby missing the first year’s savings.

For analytical clarity, we used a simplified simulation approach. As modeled for the purposes of this analysis, the alternatives are to invest now or wait a year to decide. At that point we will make a final decision—either accept the project or reject it permanently. In reality, we could continue to treat the project in a dynamic fashion. That is, we could build an elaborate decision model that allows us to apply the options-based approach over a multi-year span. However, our purpose here is to present a concept, not to develop a full-scale options pricing model. Using this simplified analytical structure, we assume that one year hence, we will reject any offer that has an NPV of zero or less. If the NPV is positive one year hence, we will accept the offer. We then calculated the NPVs based on the simulated results. The average NPV for accepting the deal today (entitled “invest now”) and for the deferred decision (entitled “wait”) are shown in Figure 2.

¹ A proper Monte Carlo simulation would involve millions of runs with carefully selected probability distributions. However, our analysis is meant to present the idea of options-based decision making and does not require such an extensive simulation approach.

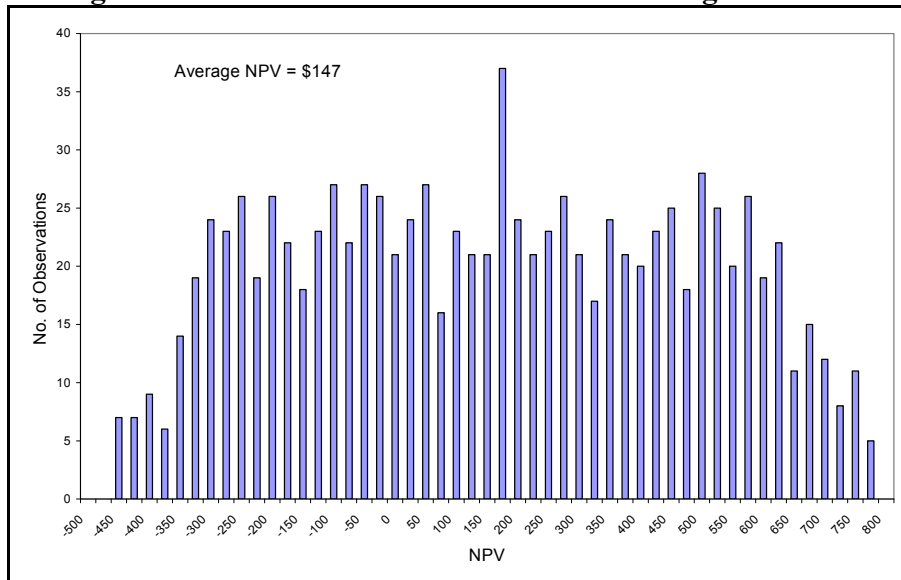
Figure 2. Average NPV Results Based on Simulations



The optimal decision is to defer the decision, even though the project has a positive NPV today. The higher NPV associated with waiting a year suggests that we will be better off by deferring the decision.

But how can that be? We have a positive-NPV project today. We are merely putting off the decision one year. The uncertainty is symmetric, making it just as likely that the situation would deteriorate as it would improve. Yet, the NPV result is noticeably better if we wait. The key is that we treat downside news in a different fashion than we treat upside news, thereby making risk work in our favor instead of against us. With the passive traditional DCF approach the only decision is *whether* we commit to the invest today, and not whether we might make a *different* decision in one year than we would today. Figure 3 shows the NPV values for the 1,000 simulations of the situation one year from today.

Figure 3. Distribution of NPV Based on Waiting One Year

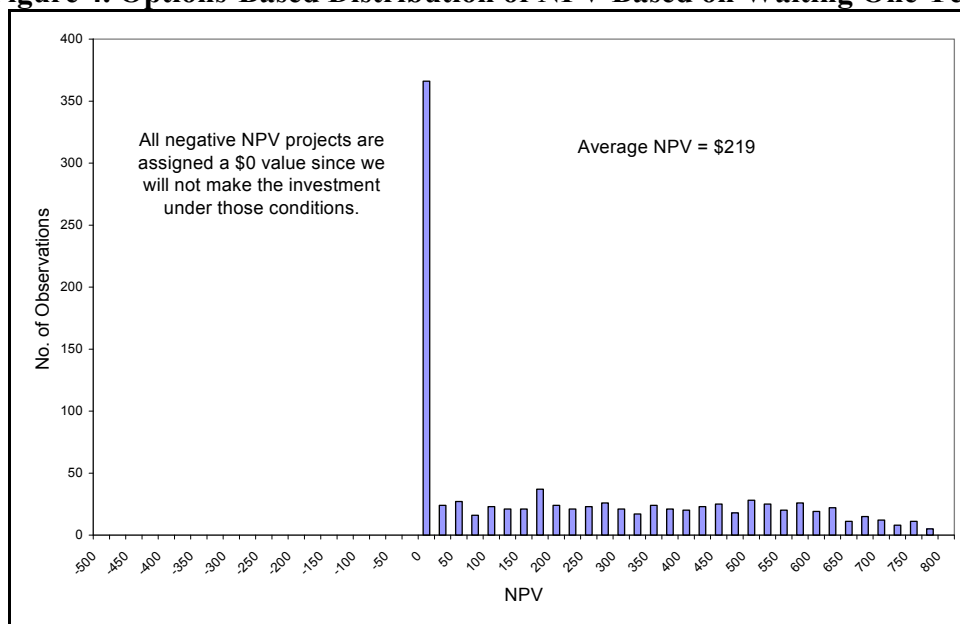


This result is essentially a sensitivity analysis, not an options-based analysis. As we suggested previously, the average NPV of waiting (\$147) is a bit less than the NPV of accepting the offer today (\$158). If we apply DCF-type thinking, all we would do by waiting a year is give up the sure thing in favor of taking on considerable risk. This is a losing proposition, on average.

Now let us apply options thinking to these data. If equipment prices decline and natural gas prices go up, we likely will be better off waiting a year to enter the deal (examine the upper-end of the distribution). Note that if equipment prices increase and if natural gas prices go down, however, we would be worse off, but *not to the extent suggested by the prior distribution*.

As Myers noted, a standard DCF-based sensitivity analysis assumes that managers are automatons, which they are not. Even though the NPV one year hence might be negative under these conditions, no one can force the manager to invest in a project with a negative NPV. He or she simply walks away from the project under those conditions, thereby resulting in a dynamic NPV of \$0 for all such negative outcomes. If we recast the distribution in this dynamic context, the revised histogram looks markedly different. This approach changes not only the shape of the distribution, but its average result, as well. Under dynamic decision making, the average NPV associated with waiting is much higher than it is if we invest in the project today. Figure 4 presents the distribution under an options-based approach.

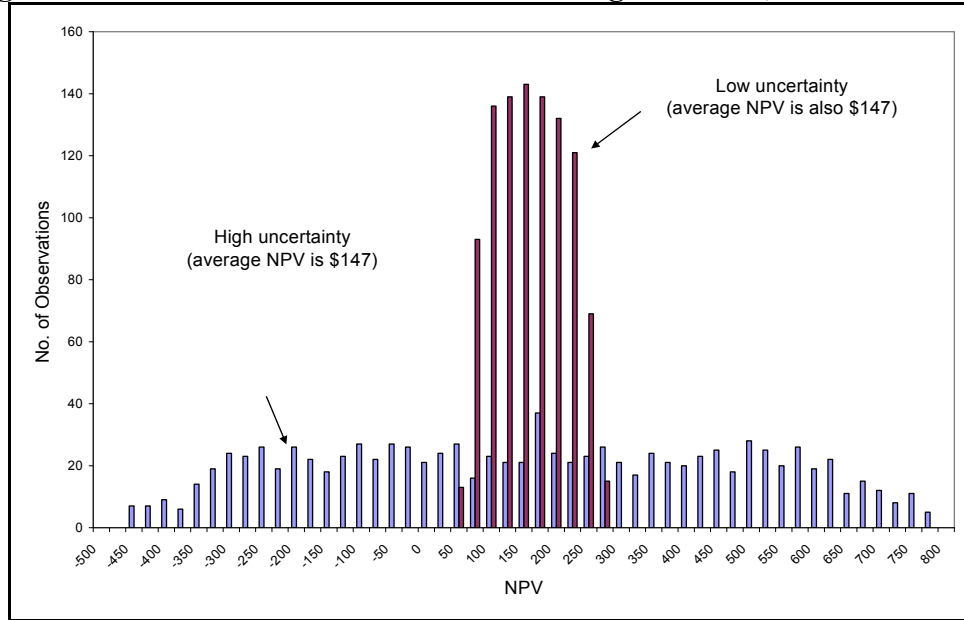
Figure 4. Options-Based Distribution of NPV Based on Waiting One Year



The decision is thus truly dynamic. If uncertainty works in our favor over the coming year, we capture the gain (agree to the deal). If uncertainty works against us, we walk away from the project. *This is how businesses make money—they position themselves so that risk works in their favor!* This is precisely how Merck uses options pricing to benefit from risk.

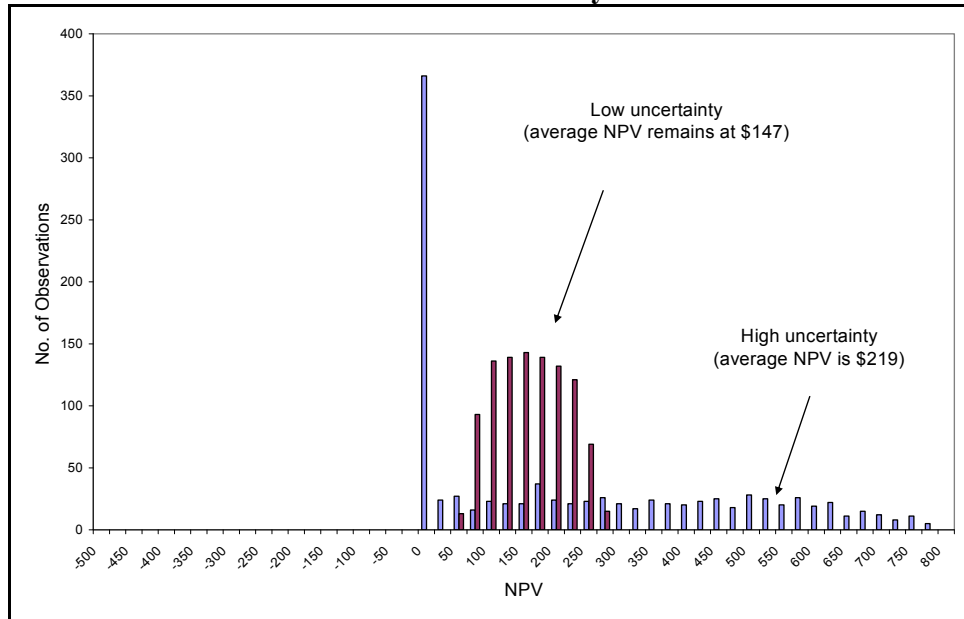
Next, we evaluate how results change when uncertainty is reduced. Let us say that equipment prices next year will likely range between \$975 and \$1,025, and natural gas prices will likely fluctuate between \$5.50 and \$6.50 per MMBtu. Figures 5 and 6 present the distributions under these parameters.

Figure 5. Distribution of NPV Based on Waiting One Year, Low Uncertainty



Note that the distribution of future NPV results is much tighter under the low uncertainty scenario. Note also that the average NPV is the same under each scenario. The most interesting observation is that there are no negative NPV results for the low-uncertainty scenario. This means that there is not likely much option value associated with waiting. It is by taking on the upside risk and dynamically avoiding those negative NPVs that the firm makes money. As we noted above, when option value exists, the overall distribution does not reflect the economics of the project. It is the upside that we care about. This concept is illustrated in the following chart.

Figure 6. Options-Based Distribution of NPV Based on Waiting One Year, Low Uncertainty



When uncertainty is low, so is the option value of waiting. Under such conditions it would be better to implement the project today. On the other hand, when uncertainty is high, the potential gains are huge and the downside is limited if we wait to make the decision. The greater the uncertainty, the longer the upper tail of the distribution and the higher the average NPV.

Uncertainty and Energy Efficiency

If uncertainty increases, the value of waiting to make energy efficiency investments increases. That is, the more uncertain the world becomes, the less energy efficiency we should see today because waiting for uncertainty to work its way through the market is on average a winning proposition for industrial firms.

One must be careful not to confuse energy price *levels* with energy price *uncertainty*. Higher price levels increase the demand for energy efficiency investments. Increased uncertainty, on the other hand, reduces the demand for those investments. If uncertainty is high, rational managers should defer some energy efficiency investments even though they have a positive DCF-based NPV. The higher the uncertainty, the more projects should be deferred. Put another way, using the NPV estimate from a DCF analysis of energy efficiency investments overstates the amount of cost-effective investment that should be taking place today. When uncertainty is high, the value of waiting to invest at a later date can dwarf the DCF-based NPV results.

Implementing Real Options Thinking in Capital Budgeting

Many firms apply the real options concept, even if they do not estimate real option value directly. While firms such as Merck employ options pricing models, due to the complex nature of these models many firms apply options pricing *concepts* instead. Other firms are led by intuition to behave in a manner consistent with options pricing principles, even if they are not familiar with the formal concept.

An imperfect, but pragmatic way of employing the real options concept is to implement only those projects with short paybacks. The firm implements the most cost-effective investments and waits to see what happens to the other potential projects with positive DCF-based NPVs. It commits to those projects in the future only when uncertainty works in the firm's favor (*i.e.*, when the paybacks for those projects are driven to acceptable levels).

Our recent analysis of achievable energy efficiency potential for a Midwest state suggests that 90 percent of the cost-effective measures (based on DCF analysis) have associated paybacks of 6.2 years or less. Options pricing analysis suggests that in today's uncertain energy markets, it is likely that many of those longer-payback projects should probably be deferred rather than implemented today. In other words, using options-based thinking, the upper limit on paybacks should be significantly less than six years.

Empirical data from industrial energy efficiency programs supports this contention. Analysis of the U.S. Department of Energy's Industrial Assessment Centers (IAC) Database, which lists data on tens of thousands of actual energy efficiency measures implemented over the past 25 years or so, shows that 90 percent of the IAC recommendations that firms actually implemented have paybacks of 2.3 years or less. This could be due in part to capital rationing

within firms, but also could be the result of dynamic options-based thinking. Under the options-based approach the question is not only *whether* a measure might be implemented, but *when* it might be.

A capital budgeting protocol that invests only in projects with short paybacks therefore may implicitly reflect sound capital budgeting practice as suggested by finance theory, even if capital is freely available. This approach is especially likely when uncertainty is great because under such conditions, the option value of deferring the decision is at its maximum.

Is Uncertainty Ever Resolved?

This analysis begs a fundamental question. What if uncertainty is never resolved? This can create a situation referred to as hysteresis, which is a form of inertia. Under such conditions, the firm never makes the energy efficiency investment. It keeps waiting for uncertainty to be resolved.

While some areas of uncertainty may never be resolved, uncertainty clearly decreases in other areas. For example, when innovative new technologies are introduced, future costs can be difficult to predict. Over time, however, such costs become more predictable as the technology transitions from emerging to established. As a case in point, it seems likely that LED lighting costs will be more stable in five years than they are today. As long as the uncertainty associated with some of the aspects affecting the cost-effectiveness of an energy efficiency investment are resolved over time, the real options approach can be useful.

Conclusion

The purpose of this paper is to shed some light on why industrial firms behave the way they do in making energy efficiency investment decisions. We know that such firms typically require short paybacks. This paper shows that it may be economically rational to defer some energy efficiency investments, even if the projects have positive NPVs today. When uncertainty is high, the option value of waiting is high. An options-based approach helps to explain why industrial firms invest in only the most cost-effective projects, while taking a wait-and-see approach with others.

Understanding how uncertainty affects industrial decision-making processes is essential from a policy perspective, as well as from an energy efficiency program design perspective. At the policy level, historical results have shown that a consistent, long-term commitment to funding energy efficiency initiatives is critically important. Regulatory uncertainty can also lead to deferred investments. A recent survey of energy management decision-makers conducted by Johnson Controls suggests that present uncertainty about future climate change legislation is likely contributing to decisions to defer potential energy efficiency investments until that uncertainty is resolved (Johnson Controls 2009).

References

Brealey, Myers, and Allen. 2006. *Principles of Corporate Finance*. McGraw Hill-Irwin.

Damodaran. 2003. "The Promise of Real Options." *The Revolution in Corporate Finance*. Basil Blackwell.

- Johnson Controls. 2009. *New Research From Johnson Controls Indicates Energy Efficiency Is Still Important to Business Leaders, but Investment Lags*. <http://www.johnsoncontrols.com/publish/us/en/news.html>. Press Release. May 6.
- Myers. 2003. "Finance Theory and Financial Strategy." *The Revolution in Corporate Finance*, Basil Blackwell.
- Nichols. 1994. "Scientific Management at Merck." *Harvard Business Review* (Jan.-Feb).
- U.S. Department of Energy, Energy Information Administration. 2009. U.S. Natural Gas Navigator. http://tonto.eia.doe.gov/dnav/ng/ng_pri_sum_dcu_nus_m.htm. Release Date: April 29.
- U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy, Industrial Technologies Program. *Industrial Assessment Centers (IAC) Database*. <http://www.iac.rutgers.edu/database/>.