MODELING ENERGY TECHNOLOGY CHOICES: WHICH FINANCE TOOLS ARE APPROPRIATE?

EMF WP 13.3

Blake Johnson

February 1993

Energy Modeling Forum
Terman Engineering Center
Stanford University
Stanford, California

Modelling Energy Technology Choices: Which Finance Tools are Appropriate?

Abstract

A variety of tools from modern investment theory have been put forward as holding promise for unraveling energy technology investment decisions that often appear anomalous when analyzed using traditional investment analysis methods. This paper reviews the important insights and assumptions of the investment theories most commonly suggested as candidates for explaining the apparent "energy technology investment paradox". The applicability of each theory to energy technology investment decisions is considered in light of important aspects of energy technology investment problems, such as sunk costs, uncertainty, and imperfect information. The theories addressed include the Capital Asset Pricing Model, the Arbitrage Pricing Theory, and the theory of irreversible investment. Enhanced net present value methods are also considered. The relevance of the Capital Asset Pricing Model and the Arbitrage Pricing Theory to energy technology decisions, given the special characteristics of energy technology investments, appears limited to the value of their conceptual insights. The theory of irreversible investment and enhanced net present value methods are found to provide more useful frameworks for modeling and analyzing energy technology decisions, as well as providing additional conceptual insights. An appendix illustrates the methods of irreversible investment theory and enhanced net present value analysis by applying them to a simple attic insulation problem.

Blake Johnson
Energy Modeling Forum

February, 1993

Introduction

The apparent under-investment in cost-effective energy conservation technologies attracts continuing attention in the debate over energy conservation policy. First, whether or not "under-investment" actually occurs, and to what extent, is unclear. Second, if there is underinvestment, why does it occur? Answering the second question is particularly important when considering policy action. If significant under-investment appears to result from market failures, and cost effective responses exist for addressing these failures, appropriate policy action is likely to be welfare improving. If not, enacted policies are likely to be disappointing.

A wide range of approaches from the theory of investment and finance have been proposed as holding promise for unraveling observed energy technology choices. Initially, NPV analyses, as employed in many engineering estimates of new technology cost effectiveness, prevailed. Later, these approaches were criticized for neglecting important aspects of an individual consumer's actual problem, such as installation specific costs and operating efficiencies, and important uncertainties such as the future level of fuel prices. In other words, costly analysis is necessary for a consumer to achieve cost and performance data as reliable as that available to design engineers. Even with accurate cost and performance data, uncertain fuel prices make the value of future energy savings highly uncertain. Finally, most energy technologies must be analyzed as part of a larger energy use system. For instance, a new heating system may be more cost-effective in a well insulated house than otherwise.

These criticisms point out deficiencies in early engineering estimates of cost-effectiveness, especially with regard to their treatment of information costs and the effects of uncertainty. Fortunately, much of the recent work in investment theory addresses these issues. To date, no comprehensive framework for investment under uncertainty with costly information exists, but important insights and useful tools have been developed. However, many of the theories make strong assumptions about the nature of available investment opportunities, about the decision maker's state of information, and about the type of uncertainty faced. As a result, their conclusions should be applied to energy technology investment decisions with caution, and only after a careful review of their relevance to the unique characteristics of energy technology choices.

In order to identify the most promising models of energy technology investment, and to avoid the inappropriate application of popular investment theories, the most influential investment theories in finance and economics are reviewed below. The theories addressed include the Capital Asset Pricing Model, the Arbitrage Pricing Theory, the theory of irreversible investment, and net present value analysis extended to address issues of uncertainty and costly information. For each theory, key insights and assumptions are stressed, and previous applications of the theory to energy technology decisions are noted where available. Some common themes emerge. First, much of modern investment theory was developed for analysis of investments in tradeable securities, as opposed to investment in capital goods, and therefore neglects important aspects of the problem of investing in fixed assets. Second, the decision maker's state of information is still generally treated in a simplistic and unrealistic way to achieve analytical tractability. Even with these simplifications, the solution methods require advanced mathematical methods. In contrast the last method addressed, enhanced net present value analysis, permits more realistic assumptions but generally allows only computational, rather than analytical solutions. The methods of the theory of irreversible investment and enhanced net present value analysis are illustrated in the Appendix, where they are applied to a simple energy efficiency improvement decision.

Capital Asset Pricing Model

The Capital Asset Pricing Model (CAPM) has been the most influential model in the theory and practice of finance since its introduction in 1964. It was motivated by earlier work on the role of diversification in maximizing the return of an investment portfolio for a given level of risk. With the benefit of diversification taken as given, the CAPM provided the additional insight that what should determine the return of an asset is the portion of its total risk that cannot be diversified away. This broad concept was made specific and useful by the second insight that the portion of the variance of each individual asset that cannot be diversified away is determined by the risk and return characteristics of the complete portfolio of available assets. The CAPM is operationalized by terming the risk of the complete portfolio of assets that cannot be diversified away "market risk", or synonymously non-diversifiable, or systematic risk. The total variance of an individual asset is then divided into its covariance with market risk and its remaining diversifiable, or "idiosyncratic" risk. Since its idiosyncratic risk can be diversified away, the asset is priced to compensate its holder only for its market risk, as reflected by the covariance term.

The CAPM results depend on the assumption of an idealized, frictionless investment environment where everyone has the same information. In particular, all investors agree on the expected returns and covariances of the assets, all assets are freely tradeable in any amounts, and there are no transaction costs of any kind. Since the CAPM was developed for securities markets, where these assumptions are arguably roughly appropriate, their importance is not always emphasized. However, they are clearly not appropriate for capital investment, where transactions costs are often significant, investment must occur in discrete and often significant amounts, and sunk costs are common. When the idealized market assumptions of the CAPM do not hold, diversification becomes costly or impossible, rather than costless, and the theory fails. When diversification is costly, it is often heuristically said that asset returns should be "somewhere above" what the CAPM would predict since a greater percentage of an asset's overall risk must be born by investors.

Despite its questionable applicability, the CAPM is commonly cited in non-securities market settings both because its requisite assumptions are often overlooked, and also because the powerful insights it provides, if not its specific pricing formulas, have wide ranging relevance. In particular, in the same way that the recommendations to diversify that preceded the CAPM provided a useful conceptual insight, the CAPM's suggestion that only an asset's relation to "systemic" sources of risk should be considered by prospective investors and included in pricing decisions has general conceptual appeal. Rather than the securities market risk which the actual, or "securities market" CAPM identifies, the appropriate "systemic risk" may be something entirely different for "conceptual" applications of the CAPM to markets for other assets. For instance, with regard to energy technologies, energy prices might be the fundamental risk investors seek to address if their goal is to meet future energy service demand at a minimum cost.

Sutherland has directly proposed CAPM-style analysis of the energy technology investment decisions of individuals and corporations. Many other authors allude to the diversification and hedging benefits conservation investments can provide in the energy technology portfolios of utilities and in those of the nation as a whole. Because the returns of conservation investments are generally positively correlated with the level of energy prices, the CAPM suggests that low returns should be expected for conservation investments due to the cost hedging function they provide in technology portfolios designed to meet particular energy service needs. However, a heuristic "correction" upward of this prediction is necessary to account for the inability of most individuals to meaningfully diversify their energy conservation investments. The net result of this approximate process appears a long way from providing a solid analytical framework for modelling consumer energy technology investment decisions.

Arbitrage Pricing Theory

The Arbitrage Pricing Theory (APT) is probably the second most important theory in investment and finance theory and practice. It was developed in the 1970's, after the CAPM, and sought both to generalize the CAPM's insights and to weaken its required assumptions. Like the CAPM, it was built around the distinction between diversifiable and non-diversifiable risk. However, where the CAPM simply defined non-diversifiable risk as the residual risk of the "market portfolio" of all available securities, the APT tried to explain non-diversifiable risk as the product of certain fundamental (but unspecified) "explanatory factors".

After assuming that non-diversifiable risk can be represented by a set of component factors, the APT follows the CAPM and calculates the expected return of an investment as the sum of the investment's covariance with each explanatory factor multiplied by the factor's expected return. To do this, the expected returns of the factors must be estimated, as well as the covariance of each asset with each of the factors. In initial securities market applications of the APT, the explanatory factors were interpreted as things like interest rates, inflation, GNP and energy prices. Later, many potential factors were tested, using historical data and regression-like identification procedures.

The main contribution of the APT was its attempt to identify, through its "risk factors", the sources of what the CAPM simply called market risk. While finding the "right" factors to do this successfully may be practically impossible, working on the problem focused attention on important fundamental risk factors in the economy. Perhaps more importantly, these efforts greatly illuminated the risk profile of individual assets. Whereas under the CAPM an asset is completely represented by its covariance with the market portfolio, under the APT an asset is represented by a set of covariances, each representing its relationship with one of the (presumably) meaningful and observable APT factors.

The more detailed risk profile of an asset provided by the APT has many potential uses. While it may be reasonable to assume that many securities market investors are primarily concerned with "market risk" as identified by the CAPM, many investors considering investments in physical assets may not be. If their individual investment activities lead to a concentration of risk associated with a particular factor, it is this type of risk that they are likely to be most interested in hedging. For those primarily involved in energy production or use, energy prices are likely to be this risk. Using energy as an APT risk factor, appropriate "factor risk pricing" can theoretically by calculated for this risk. An appropriate expected return level can then be calculated, and offsetting risk-hedging transactions in other energy related securities can be made if so desired.

In light of the above discussion, it appears that the APT may have more to bring to an energy conservation investment analysis than the CAPM. However, the APT makes the same assumptions as the CAPM regarding the absence of transaction costs, the infinite divisibility of assets, the absence of sunk costs, and the full availability of mutually agreed on information. Any application of the APT to energy technology investments will therefore have to be quite heuristic as well.

Theory of Irreversible Investment

In recent years, traditional investment models in the economics literature, like those in finance, have been extended to address uncertainty and the effects of limited information. The branch of this work most relevant to energy technology choices has come to be called the theory of irreversible investment. This theory uses recent advances in stochastic optimization and control theory to solve a class of investment problems that involve sunk costs and occur in imperfectly competitive markets. Because energy technology investment decisions often involve sunk costs and imperfectly competitive markets, "irreversibility theory" has been proposed as a method of analyzing energy technology choices. Unfortunately, the theory's complicated analytical methods and other required assumptions leave its real world usefulness in question. Consideration of why these methods and assumptions are required to reach an analytical solution provides a natural introduction to the discussion of enhanced NPV methods below, where a simpler and less restrictive approach leads instead to a computational solution.

In contrast to the CAPM and APT, irreversibility theory was developed to address investment in real, or physical assets, rather than securities. Where the assumptions of the CAPM and APT of finance theory were appropriate for securities and not real assets, the opposite is true for irreversibility theory's assumptions of sunk costs and imperfectly competitive markets. The importance of irreversibility theory in the recent economics literature stems from the fact that many, if not most, investments in physical assets possess these two qualities, and because previous work had not considered them in combination in the presence of uncertainty.

Many energy technology investment problems fit naturally into the irreversibility theory framework. For instance, installation of a more energy-efficient industrial process, or of a solar water heater in a home present the uncertainty of their future effectiveness, as well as of the future level of energy prices. Each requires a sunk cost, since it is unlikely that a significant fraction of the necessary investment could be recovered if the decision was later to be reversed. Finally, the market for these investments is imperfectly competitive - in fact almost "perfectly uncompetitive" - since it is unlikely that anyone else could take advantage of potentially profitable energy conservation investments in someone else's business or home.

The approach irreversibility theory takes to problems of this type is to ask whether the investment should be undertaken now, or postponed for subsequent reconsideration. If postponed, the investment is assumed to be repeatedly reconsidered in the future until the project is either eventually undertaken, or dropped from consideration entirely. The analysis changes over time because the theory assumes that the associated uncertainty is fully or partially resolved with the passage of time. As a result, each time the investment is reconsidered, more is known about its future cashflows. To make the analysis easier, two other important assumptions about the investment opportunity in question are made. First, it is assumed that the uncertainty is resolved "for free", simply by the passage of time. There are no alternative active (and costly) information gathering activities to be considered by the decision maker. Second, the investment opportunity is assumed to be static; it will be available, unchanged, for as long as the decision maker wishes to consider it.

The later two assumptions assure that by simply postponing their decision, the decision maker will be able to reconsider the investment, in unchanged form, with better information in hand. In reality, it is more likely that the decision maker must either search for better information now, or face a possibly different opportunity later. Under its assumptions, irreversibility theory's main conclusion is that due to the presence of the sunk cost, and to the absence of competition for the investment opportunity, it may make sense to wait, even if the project has a *significantly positive* expected NPV now. In particular, the investment should only be undertaken when the value of the incremental information that would be gained by postponing it further is less than the profit that would be earned if the investment was in place over the same period. The value of additional informational when the investment already looks promising is that it *may* reveal that it actually won't be so promising. In other words, waiting is still valuable if it allows you to make the long odds of a bad outcome even longer, and it doesn't force you to forego substantial profits in the mean time. (See the Appendix for a simple example)

The insight which irreversibility theory provides is clearly valuable. A positive NPV is not the appropriate criteria for accepting an investment if you also possess the option to learn more about the investment, at no cost, before you undertake it. Unfortunately, the assumption that waiting for time to pass is the only way to learn about the investment opportunity is generally unrealistic, as is the assumption that the opportunity itself will not change over time. These assumptions are an outgrowth of the complicated analytical methods which irreversibility theory requires which entail the solution of a stochastic partial differential equation. Attempting to consider other decision alternatives such as information gathering or research activities would make the problems analytically intractable, as would assuming that changes in the investment environment are possible. Jaffe and Stavins overcome some of these difficulties by embedding important aspects of irreversibility theory in a larger diffusion theory model of energy technology decisions. In contrast, the enhanced NPV methods described below rely on computer based "number crunching" to solve individual decision problems where information gathering alternatives are considered directly.

NPV Analysis With Uncertainty and Costly Information

Despite the impact of the theories described above, net present value analysis is still the most commonly used investment analysis method for project analysis and for investment in physical assets. Because NPV analysis is easy to understand and to apply, it is convincing and practical, and easily supports discussion, even among those with little background in investment analysis. As noted, as it has been applied historically, it also has weaknesses, especially in its treatment of information and uncertainty. These weaknesses result because deterministic cashflows are usually assumed for simplicity, leaving the discount rate to incorporate both the time value of money and to account for uncertainty about the projected cashflows. While these simplifying assumptions are now taken almost as part of the definition of NPV analysis due to their prevalence, they are not necessary, and can be weakened to improve the accuracy and expand the content of the results. In particular, the following extensions to the traditional NPV approach are possible:

1. Uncertainty about the amount of future cashflows can be represented explicitly
2. With cashflow uncertainty represented explicitly, the discount rate can be used to isolate the time value of money
3. The "state of information" of the decision maker can be represented explicitly, and the value of changing this state of information by gathering information or hiring experts can be calculated in monetary terms
4. Project stages, including intermediate decision points, can be represented (e.g. research phase, pilot program, full scale implementation)

Each of these extensions requires conceptual extensions to the standard theory as well as more advanced computation methods. The conceptual basis of the four extensions listed is described briefly below. A simple example which embodies the main ideas presented is provided in the Appendix.

<u>Uncertainty About Future Cashflows</u>

Very few investment opportunities provide certain future period cashflows. Rather, future cashflows usually depend on the outcome of uncertain future events. As a result, beforehand only the distribution of future cashflows, conditional on alternative possible future events, is known. In addition, most decision makers don't have access to the best available estimates of an investment's future cashflows. As a result, they operate with an additional level of uncertainty about the future cashflows in question, which depends on the quality of their own subjective assessment of these cashflows. In traditional NPV analysis, the decision maker's beliefs about the distribution of future cashflow in each period is summarized in one number, a "best guess" or expected value. When the cashflow distribution is condensed in this way, much information about the extent of the decision-maker's uncertainty, such as the range of potential outcomes they believe are possible, is lost.

In addition to this loss of potentially useful information, there are two other problems with the traditional "point estimate of cashflow" simplification. First, it implies that the decision maker is risk-neutral, since they will reach the same conclusion for any two cashflow distributions with the same mean, regardless of their dispersion. Second, when there is more than one source of uncertainty, using a "point estimate" requires the implicit assumption that the two uncertain variables are independent, since any covariance is ignored. Traditionally, the hope has been that these effects aren't large, or will somehow be accounted for elsewhere, such as in the discount rate. Because of the complicated and often unexpected effects of uncertainty, this is likely to be unrealistic.

Full use of information about uncertain variables can be correctly made by calculating the investment's NPV over the range of cashflow values given by the joint distribution of the uncertain variables. This will give a distribution of possible NPV's for the investment, from which the likelihood of the NPV falling in any given range can be read, and best and worst possible outcomes noted. In contrast to providing one best guess number, this method allows the full range of possible levels of profitability of an investment, together with their relative likelihood, to be considered directly. The cost of this method is that the number of calculation required make a computation solution necessary.

A simpler alternative approach also exists if stronger assumptions about the decision maker's preferences are made. If the decision maker feels comfortable with the assumptions of expected utility theory, their risk profile can be represented by a utility function, and the utility function can be used to value uncertain future payoffs. This value is the expected utility of the distribution of possible cashflows at each period, which through the properties of the utility function accounts for the risk tolerance of the decision maker. The discounted value of these values then summarizes the desirability of the investment to the decision-maker in question.

Discount Rate Reflecting Only the Time Value of Money

Since in traditional NPV analyses uncertain future cashflows are summarized with a point estimate, the discount rate is forced to account for both the time value of money and the uncertainty of future cashlows. As a result, discount rates used have been set equal (in an approximate way) to the sum of a risk-free rate to account for time value, and a risk premium believed to be appropriate given the level of risk of the investment. Clearly, this indirect approach to cashflow uncertainty makes only limited use of the decision-maker's knowledge of the nature and sources of the investment's risk, once again summarizing it with a single number.

When uncertainty about future cashflows is addressed directly as described in the section above, the discount rate no longer needs to attempt to account for the effects of the "riskiness" of an investment. Instead, the prevailing treasury market rates for the term of the cashflows being analyzed can be used to adjust for timing effects, and cashflow uncertainty can then be considered directly after being collapsed to one point in time. If the decision-maker prefers to think in terms of risk adjusted rates of return rather than NPV's, the distribution of rates of return associated with the distribution of NPV's described above can of course be provided.

Value of Information Analysis

Real world decision makers rarely approach an investment decision in possession of the best return information available about the investment. Better information is, however, usually available at some cost. Using the framework for modelling cashflow uncertainty described above, the monetary value of additional information about the uncertain variables in question can be calculated. The calculations rely on Bayes rule, and allow the decision-maker's existing beliefs, or state of information, to be augmented by additional information or expert opinion. The end result is that the decision-maker's original state of information, as reflected by their initial probability distribution for the uncertain variables, is replaced with a more accurate distribution gained through some costly process of information gathering or learning. The value of this change can be calculated by considering the circumstances in which the decision-maker's choices change in light of the new information. In particular, some "accept" decisions may change to "reject", avoiding a likely loss, and other "reject" decisions may change to "accept", leading to a likely positive return. The expected value of these changes is the expected value of the improved information.

It is important to note that in most cases, the information received will only be better, and not perfect. The full representation of uncertainty being employed allows the investor's "before" and "after" levels of uncertainty to be compared carefully, and the benefits of their marginal change to be quantified. This same procedure works equally well when applied to any one or any subset of the complete set of uncertain variables, and can be applied repeatedly for repeated refinements to the level of information available.

Sequential Decisions

Many investments are undertaken in stages, with the option to continue or to forego the remainder of the project available at each stage. The development of a new product is an obvious example, with a R&D phase, a testing phase, and finally full scale production. There is clearly an option to continue after each of these stages, but there may also be options to "sell out" the patent or production rights. More generally, problems which require exploratory investments of time or resources, or which allow "scaling up" can be analyzed as sequential decisions.

Using the methodology described above for one stage projects, the return distribution for each stage of a sequential project can be calculated by working backward from the last stage to the first. The process is similar to the one employed in dynamic programming; first, assume you are at the last stage of the project, and that one (of the many possible) sequences of outcomes from the earlier stages has occurred. Apply the one period analysis to the final stage, from this particular starting point, and calculate the investment's return distribution from that point, as determined by the uncertainties of that stage. Use this return distribution to decide whether to undertake the final stage, if faced with it from the starting point in question. Finally, record either "abandon" or the calculated return distribution for the final stage from that starting point based on the decision. Repeat this calculation for the other possible "starting points" of the final stage. These results constitute the "state dependent" decision rule for the final stage.

Next, move to the second to last stage, and repeat the same procedure, using the possible outcomes of the preceding stage as the set of starting points. The "payoffs" to use for calculating the return of the second to last stage are the results of the calculation of the optimal action in the final stage; given a possible outcome for the second to last stage, either the project will be abandoned, or will provide one of the return distributions calculated for the final stage. Finally, calculate the decision rule for each possible starting point of the second to last stage, and record either "abandon" or the return distribution for that starting point.

If the investment being considered has many stages, or more than a few possible outcomes at each stage, the task described above can clearly become substantial. However, since each step is relatively simple, and the process repeats itself, available software can facilitate the required calculations.

Conclusions

The aim of this paper has been to review the relevance of the most important current theories of investment to decisions about energy technology investment. While not exhaustive, together the theories chosen span the important insights and computational methods of modern investment theory. Summarizing, the CAPM suggests risk should be separated into its diversifiable and non-diversifiable components, and that only the later should be considered in pricing calculations. The APT suggests that non-diversifiable risk can and should be thought of as driven by fundamental risk factors. In some cases, an investment's exposure to a particular risk factor, such as energy prices in the case of energy investments, may be the decision maker's primary concern. Because the CAPM and APT depend on assumptions tailored to securities market investments and inappropriate for investment in real assets, only the heuristic value of their insights, and not their precise pricing formulas, appear relevant to energy technology decisions. In contrast, the theory of irreversible investment directly addresses problems of investment in real assets such as energy technologies. It shows that without the threat of competition it pays to postpone an investment decision until foregone profits begin to outweigh the benefits of any further reduction in uncertainty that might result from waiting. In order to get an analytical solution, irreversibility theory is forced to make the restrictive assumption that no information gathering opportunities exist, and that the investment opportunity being considered remains unchanged over time. Finally, various extensions to traditional NPV methods were shown to be flexible and conceptually simple, yet capable of presenting a more complete representation of an investment's return characteristics and of the decision maker's information gathering alternatives. The price of these benefits is that a computational solution is required.

The sample calculations in the appendix which demonstrate the theory of irreversible investment and extended NPV methods suggest these approaches may offer promise for improving our understanding of energy technology investment decisions. With further work, similar models may be able to provide the insights about the determinants of energy technology investment necessary for the effective design of energy conservation policy.

References

1. Dixit, A., "Investment and Hysteresis", Journal of Economic Perspectives, 1992, vol. 6, no. 1 pp. 107-132.

2. Jaffe, A. and R. Stavins, "The Energy Paradox and the Diffusion of Conservation Technology - Modelling the Effectiveness of Economic Incentives and Direct Regulation", The National Bureau of Economic Research Conference on Economics of the Environment, December 1991.

3. Pindyck, R., "Irreversibility, Uncertainty, and Investment", Journal of Economic Literature, 1991, vol. XXIX, pp. 1110-1148.

4. Ross, S., "The Arbitrage Theory of Capital Asset Pricing", Journal of Economic Theory, 1976, vol. 13 pp. 341-360.

5. Sharpe, W. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk", Journal of Finance, 1964, vol. 19 pp. 425-442.

6. Sutherland, R., "Market Barriers to Energy-Efficiency Investments", The Energy Journal, vol. 12, no. 3, 1991, pp. 15-34.

Appendix

This appendix illustrates the methods of the theory of irreversible investment and NPV with uncertainty through their application to a simple problem. The problem is that of a homeowner deciding whether or not (and when) to add attic insulation to their home. The several decision alternatives considered were constructed to be as simple as possible while still capable of demonstrating the important features of both approaches.

Imagine a homeowner (Jim) who just received a promotional mailing from a supplier of attic insulation. Jim's home, constructed in the 1950's, currently has single-paned windows, limited weather-proofing, and no attic insulation. The mailer advertises that their attic insulation program leads to a 30% reduction in heating costs "for the average home", at a cost of approximately one dollar per square foot. An asterisk references a footnote that cautions that actual savings depend on the design of the house, as well as on its weather-proofing and heating system. Jim's home is a single storey, 2,000 square foot house, with an old gas heating system.

Jim knows that the attic insulation mailer came from a local company, and therefore that the mailer's savings estimates are probably appropriate for local weather conditions. The company appears reputable, but he doesn't know much about it, and he knows even less about how his heating and insulation systems compare to local averages. All in all, he decides that the actual reduction in his heating bill could just as likely be 10% as 25 or 35%. Also, he typically only uses the heating system six months of the year, and doesn't have an air conditioner. His average heating bill during these months is $150. Finally, he knows all his neighbors received the same mailing, and he's fairly sure that at least one of them will have the insulation added this year.

A. Summary of Initial Information:

Expected Cost of Insulation = $1 per square foot × 2,000 square feet = $2,000
Annual heating bill = Monthly heating bill × 6 heating months
= $150 × 6 = $900
Annual heating bill savings:
At 10%: = $900 × 10% = $90
At 25%: = $900 × 25% = $225
At 35%: = $900 × 35% = $315
Assumed discount rate: 10%

B. Irreversibility Theory Approach

Fortunately, Jim knows more about investment analysis than insulation. A simple calculation lets him compare the possibility of insulating his attic now or waiting until after a neighbor has tried it. The problem fits the irreversibility theory framework because Jim has the exclusive right to invest in attic insulation for his home (since no one else could undertake and profit from the project) and because installing the insulation is an irreversible investment (since the insulation can't subsequently be removed and returned or resold). Further, there is the possibility of learning costlessly from a neighbor's experience. Jim expects that if a neighbor installs the insulation and says after their first year's experience that "its working out well", he can revise his beliefs and expect either a 25 and 35% savings with equal likelihood. If the neighbor reports after a year that the insulation isn't working out all that well, then Jim will expect either 10 or 25% savings with equal likelihood. For the one year delay not to detract from the investment opportunity, the horizon for the savings is assumed to be arbitrarily long.

Graphically, Jim's alternatives are:

Time:	0	1	2	...	...	...	∞
Insulate Now:							
		$90	$90	...	...	...	∞
	-$2,000	$225	$225	...	...	...	∞
		$310	$310	...	...	...	∞

$$NPV = -\$2{,}000 + 1/3 \times \left[\sum_{i=1}^{\infty} (1/1.1)^i \times (\$90 + \$225 + \$315) \right] = \$310$$

Wait for Neighbor:						
		$225	$225	...	...	∞
Good News	-$2,000					
		$315	$315	...	...	∞

$$NPV = -\ \$2{,}000 + 1/2 \times \left[\sum_{i=1}^{\infty} (1/1.1)^i \times (\$225 + \$315) \right] = \$970$$

		$90	$90	...	...	∞
Bad News	-$2,000					
		$225	$225	...	...	∞

$$NPV = -\ \$2{,}000 + 1/2 \times \left[\sum_{i=1}^{\infty} (1/1.1)^i \times (\$90 + \$225) \right] = -\$267.50$$

If Jim goes ahead now, he can expect a $310 positive NPV. However, if he waits until next year and learns from the experience of his neighbor, he can avoid the possible downside he faces now (a meager 10% heating bill reduction), and invest only if the neighbor has "good news". Assuming the neighbor is as likely to have good news as bad, Jim's expected NPV if he waits is $1/2 \times \$970 = \$485$, greater than the NPV of going ahead today. It therefore pays to wait, even though the investment's NPV is significantly positive now.

C. NPV Analysis with Uncertainty

The analysis above can be extended and made more realistic by incorporating additional information gathering opportunities, and by considering the impact of the finite number of years Jim is actually likely to own the house. This requires the methods of NPV with uncertainty, as illustrated below.

First, Jim knows he won't remain in his house forever, as he was forced to assume above so that the investment opportunity would remain unchanged despite its one year postponement. Instead, he would like to consider the insulation investment's payoff if he stays in the house for only 10 more years, as well as if he stays 20 more years. He estimates that a future buyer is likely to compensate him for only one-half the cost of the insulation, or $1,000.

Jim would also like to address two other complications. First, when he called the insulation company to ask about the offer, he was told that for $100 one of the company's representatives would carefully examine his home in order to provide a more accurate assessment of the benefits that attic insulation would provide him. If he later went ahead with the insulation, the $100 would be credited against the $2,000 cost. Second, at a cost of a weekend and about $500 in materials, Jim believes he can do a lot of the other basic weatherproofing that his house needs, and at the same time get a better idea of the benefit the attic insulation might provide. He estimates that the weatherproofing would reduce his heating bill by 5%, independent of the benefit of the attic insulation. His own lost weekend would be offset by any positive NPV of this alternative, and by the knowledge that he had "done the right thing". In addition, weatherstripping the front door would allow him to finally resolve whether the living room was always cold in the winter because of the gaps around the front door or because the kids left the door open so often.

First, Jim reconsiders the basic alternatives considered above, but now takes into consideration the alternative 10 or 20 year term, and the $1,000 "salvage value":

<u>Insulate Now</u>:

Time:	0	1	2	...	...	...	10 or 20	11 or 21
		$90	$90	...	...	...	$90	$1,000
	-$2,000	$225	$225	...	...	...	$225	$1,000
		$310	$310	...	...	...	$315	$1,000

Given a 10 year horizon,

$$NPV = -\$2{,}000 + 1/3 \times \left[\sum_{i=1}^{10} (1/1.1)^i \times (\$90 + \$225 + \$315)\right] + (1/1.1)^{11} \times \$1{,}000 = -\$359.15$$

Given a 20 year horizon,

$$NPV = -\$2,000 + 1/3 \times [\sum_{i=1}^{20} (1/1.1)^i \times (\$90 + \$225 + \$315)] + (1/1.1)^{21} \times \$1,000 = -\$77.02$$

Therefore, Jim won't go ahead without further information under either the 10 or 20 year planning horizon. Note that this result conflicts with the \$310 *positive* NPV for the "insulate now" alternative when an infinite horizon was assumed.

<u>Wait for Neighbor</u>:

• Case 1: Neighbor has Good News

Time:	0	1	2	...	...	...	10 or 20	11 or 21
			\$225	\$225	...	...	\$225	\$1,000
Cashflows:	-\$2,000							
			\$315	\$315	...	...	\$315	\$1,000

Given a 9 year remaining horizon,

$$NPV = -\$2,000 + 1/2 \times [\sum_{i=1}^{9} (1/1.1)^i \times (\$225 + \$315)] + (1/1.1)^{10} \times \$1,000 = -\$59.52$$

Given a 19 year remaining horizon,

$$NPV = -\$2,000 + 1/2 \times [\sum_{i=1}^{19} (1/1.1)^i \times (\$225 + \$315)] + (1/1.1)^{20} \times \$1,000 = \$407.17$$

• Case 2: Neighbor has Bad News

Time:	0	1	2	...	...	...	10 or 20	11 or 21
			\$90	\$90	...	...	\$90	\$1,000
Cashflows:	-\$2,000							
			\$225	\$225	...	...	\$225	\$1,000

Given a 9 year remaining horizon,

$$NPV = -\$2,000 + 1/2 \times [\sum_{i=1}^{9} (1/1.1)^i \times (\$90 + \$225)] + (1/1.1)^{10} \times \$1,000 = -\$707.41$$

Given a 19 year remaining horizon,

$$NPV = -\$2,000 + 1/2 \times [\sum_{i=1}^{19} (1/1.1)^i \times (\$90 + \$225)] + (1/1.1)^{20} \times \$1,000 = -\$533.88$$

The analysis reveals that waiting to learn from the neighbor's experience creates a positive NPV for Jim's project if he assumes a 20 year horizon (which leaves 19 years after the one year wait). In that case, he will go ahead with the project if the news is good, which occurs with probability .5, and do nothing if the news is bad. The NPV *today* in this case is $(.5 \times \$407.17)/1.1 = \$185.08$. If he assumes a 10 year horizon, he won't insulate regardless of what his neighbor says, since the investment has a negative NPV over the remaining 9 years of his residency even if the news is good.

<u>$100 Inspection by Insulation Company Representative</u>

Jim believes the insulation company's inspection will be helpful. In particular, he wants to know whether the relatively poor quality of his home's weatherproofing and wall insulation will increase or decrease the impact of the attic insulation. If the inspector suggests a decrease will result and sounds discouraging, Jim will assume his savings will be 10% for sure. If the inspector suggests an increase will result and sounds encouraging, Jim will assume that 25% or 35% savings will result, with equal likelihood. This inspection alternative can be represented as follows:

• Get Inspection:

	Time: 0	1	2	...	...	...	10 or 20	11 or 21
Bad News:	1/3 × -\$100	0	0				0	0
		\$225	\$225	...	...	...	\$225	\$1,000
Good News:	2/3 × -\$2000							
		\$315	\$315	...	...	...	\$315	\$1,000

Given a 10 year horizon,

$$NPV = 1/3 \times -\$100 + 2/3 \times [-\$2,000 + 1/2 \times [\sum_{i=1}^{10} (1/1.1)^i \times (\$225 + \$315)] + (1/1.1)^{11} \times \$1,000] = -\$73.01$$

Given a 20 year horizon,

$$NPV = 1/3 \times -\$100 + 2/3 \times [-\$2,000 + 1/2 \times [\sum_{i=1}^{20} (1/1.1)^i \times (\$225 + \$315)] + (1/1.1)^{21} \times \$1,000] = \$238.12$$

As was true for the "wait for the neighbor" alternative, it makes sense to proceed with the inspection if the planning horizon is 20 years, but not 10. Also, the NPV in this case, $238.12, is larger than the expected NPV for waiting given a 20 year horizon, which was $185.08. Jim would be wiser to spend $100 for this active information gathering alternative than to wait a year for the free, but less reliable, information from his neighbor.

<u>$500 Weekend Weatherizing Project</u>

By investing $500 and a weekend Jim believes he can learn as much as he would from the insulation company inspection, as well as reduce his future heating bill by 5% regardless of whether he ends up adding the attic insulation. Jim expects that a future buyer of his home would compensate him for one-half the cost of the weatherproofing, as was assumed for the cost of the attic insulation. This weekend "do it yourself" alternative can be represented as follows:

Do Weather-proofing:

	Time: 0	1	2	...	...	...	10 or 20	11 or 21
Looks Bad	1/3 × -$500	$45	$45	...	...	...	$45	$250
Looks Good	2/3 × -$2500	$270	$270	...	...	...	$270	$1,250
		$360	$360	...	...	...	$360	$1,250

Given a 10 year horizon,

$$NPV = 1/3 \times [\, -\$500 + [\sum_{i=1}^{10} (1/1.1)^i \times \$45] + (1/1.1)^{11} \times \$250]$$

$$+\ 2/3 \times [\, -\$2,500 + 1/2 \times [\sum_{i=1}^{10} (1/1.1)^i \times (\$270 + \$360)] + (1/1.1)^{11} \times \$1,250] = -\$129.52$$

Given a 20 year horizon,

$$NPV = 1/3 \times [\, -\$500 + [\sum_{i=1}^{20} (1/1.1)^i \times \$45] + (1/1.1)^{21} \times \$250]$$

$$+\ 2/3 \times [\, -\$2,500 + 1/2 \times [\sum_{i=1}^{20} (1/1.1)^i \times (\$270 + \$360)] + (1/1.1)^{21} \times \$1,250] = \$206.08$$

As was true for Jim's alternatives to wait and to have the insulation company inspection, this alternative only has a positive NPV over a 20 year horizon. For the 20 year horizon, Jim is better off doing the weatherstripping than simply waiting for his neighbor's results (NPV of $206.08 vs. $185.08), but would be even better off if he chose the insulation company's $100 inspection, yielding a $238.12 NPV.