A Real Options Based Decision Support Tool for R&D Investment: Application to CO_2 Recycling Technology

P. Deeney[a,*], M. Cummins[a], K. Heintz[b], M. Pryce[b]

[a] *DCU Business School, Dublin City University, Glasnevin, Dublin 9, Ireland*
[b] *School of Chemical Sciences, Dublin City University, Glasnevin, Dublin 9, Ireland*

Abstract

We propose a practice relevant real options based decision support tool to aid in the practical evaluation of R&D investments in technology. Using a Poisson process to simulate the discrete progress typical of advancements in R&D, we take explicit account of the technical risk of the technology development, while market risk exposure and the effect of learning-by-doing through operating the technology is also explicitly modelled. We present a compound real option design, where a European real option structure is used to model the fixed length term typical of early phase research, which is exercisable into an American real option structure to model a subsequent phase R&D. In this latter phase, a successful outcome is acted upon immediately to operationalise the technology. We propose a simulation approach, which models R&D progress in a stylised logistic function or 'S-shape' form, capturing the typically slow rate of R&D progress at the start of the early phase, through to more rapid improvement as the R&D advances, which then slows again as the limitations of the R&D are approached. We propose a business appropriate and workable economic meaning to this progress in the R&D process. We demonstrate the decision support tool with an application to evaluating the R&D investment potential in CO_2 recycling technology, where an energy commodity is produced.

Keywords: Decision Analysis, OR in Research and Development, Real Options, Compound Option Structure, CO_2 Recycling Technology

*Corresponding author
Email addresses: peter.deeney@dcu.ie (P. Deeney), mark.cummins@dcu.ie (M. Cummins), katarina.heintz@dcu.ie (K. Heintz), mary.pryce@dcu.ie (M. Pryce)

1. Introduction

Research and development (R&D) has driven human progress and improved living standards for centuries (Pinker, 2018). This effort consumed approximately US$ 1 trillion in 2010 (The Royal Society, 2010) and approximately US$ 1.7 trillion in 2018 (UNESCO, 2018). The purpose of R&D is to push the boundaries of knowledge and discover the unknown. Such uncertainty makes it difficult to establish the value of an investment in R&D. Real Options Analysis (ROA) has been used in the literature to help make investment choices under uncertainty for some time (Trigeorgis, 1996; Childs et al., 1998; Childs and Triantis, 1999; Schwartz, 1997; Schwartz and Moon, 2000, 2001; Schwartz and Zozaya-Gorostiza, 2003). Trigeorgis and Tsekrekos (2018) provide a comprehensive and contemporary review of the literature on real options applied in operational research. The advantage of ROA over the net present value (NPV) technique is that it takes into account the value of the flexibility which may be present in, or which the firm may engineer into, investment choices (Dixit and Pindyck, 1994). For example, the flexibility to delay the timing of an investment makes that investment more valuable than a similar one which does not give this option. This is because information which comes to hand during the delay may make an abandonment of the investment preferable to an early commitment. ROA is ideally suited to an R&D setting as the purpose of R&D is commonly to garner new information that will aid the firm in deciding whether or not to commit further, larger investment later. In this study, we tailor a real options based decision support tool of practice relevance that allows a firm to evaluate R&D investment in a technology context. The decision support tool is designed to have particular features of relevance to practitioners, with the modelling approach we adopt being motivated by three specific strands of literature, which we synthesise for internal firm investment appraisal purposes.

Firstly, we wish to explicitly model technical risk around the technology under development, where technical risk is defined as the likelihood that an R&D project may not yield a successful outcome (e.g. the commercialisation of the technology). We model the technical risk of an R&D project over time as a random sequence of breakthroughs governed by a Poisson distribution, where we define a breakthrough as an incremental advancement in the technology that leads to some economically measurable improvement. The discrete nature of R&D progress in practice justifies the use of the Poisson distribution. The Poisson process has been used to model the arrival of technological innovation in game theoretic models of technology adoption for a single competitive firm or for multiple competing firms in an economy (Farzin et al., 1998; Hagspiel et al., 2015; Huisman and Kort, 2004). These studies have progressively transitioned from the assumption of constant to the assumption of time-varying technology arrival rates. These models consider the technological innovation as exogenous to the firms in competition, with the decision being whether to adopt a technology early or wait for future technological innovations to come to the market. Motivated by this literature, we too use the Poisson process to model the evolution of technological innovation over time. Importantly though, we differ in considering the technological innovation to be endogenous to the firm.

In an R&D context, only a few authors deal explicitly with technical risk. Such studies include Huchzermeier and Loch (2001), who use a binomial distribution to model changes in product performance as R&D proceeds; one branch allowing for improvement in performance, and the other, deterioration in performance. Schwartz and Zozaya-Gorostiza (2003) consider technical risk in the context of an information technology R&D project with a stochastic cash flow model, describing this uncertainty by a continuously changing Wiener process. Hsu and Schwartz (2008) consider the quality or efficacy of the research output within the context of studying pharmaceutical under-investment in R&D for vaccines to treat diseases affecting the developing regions of the world. The authors use a Beta distribution with support over $[0, 1]$ to model research quality, which can be interpreted in the context of a vaccine as the percentage of a population for which immunisation is effective. Jang et al. (2013) suggest a binomial model of technical risk in renewable energy R&D, whereby a successful (unsuccessful) outcome leads to a constant linear reduction (no reduction) in the unit cost of the generated electricity. Nishihara (2018) uses a discrete measure of technological risk in R&D, allowing the level of success of the R&D and the duration of the R&D to be random. While the approach proposed is generalisable to any number of scenarios, Nishihara (2018) considers three equally likely scenarios for the technology at the end of the R&D period, whereby the project value is either zero, as expected, or twice that expected. The model of Nishihara (2018) is focused on the uncertainty concerning the probability of R&D success. Schwartz and Moon (2000), Miltersen and Schwartz (2004), Schwartz (2004), Cassimon et al. (2011) and Pennings and Sereno (2011) use the Poisson process to model technical risk, but do so as a catastrophic event leading to abandonment of the whole investment, which is a natural consideration in the pharmaceutical field, where a new treatment may be ineffective or unsafe and therefore unusable. In commodity research, such as our example of carbon dioxide recycling, the problem is more nuanced and a matter of degree. Hence, our use of the Poisson model differs from these studies.

Secondly, we have defined an R&D breakthrough as an incremental advancement in the technology that leads to some economically measurable improvement. This is important so that the decision support tool is accessible to practitioners. The modelling of breakthroughs via the Poisson process is somewhat abstract in the sense that it is not possible to define exactly what a breakthrough means, as this will vary considerably depending on the stage of technological development and, indeed, the particular technology in question that is under development. We therefore need some form of functional mapping to give this progress in the R&D process an economic meaning. For our purposes, we specifically consider the problem of evaluating R&D investment in a technology that produces a tradable commodity. This is exactly how carbon dioxide (CO_2) recycling technology is designed, allowing for the conversion of CO_2 emissions into a marketable commodity, such as methane, carbon monoxide, and so on. For the decision support tool, we propose the number of units of the commodity produced by the technology per unit of cost (i.e. units per cost or UPC) as the economic measure of R&D progress, which we suggest be inverted, where appropriate, to the cost per unit of commodity produced by the technology (cost per unit or CPU). With a typical R&D project, one would expect a variable improvement rate in the UPC. We propose a UPC profile that transitions from a slow rate of increase in the early stages of the R&D process, through to a higher rate of increase later as knowledge and expertise accumulates, with the rate of increase eventually slowing as the physical, technical and/or economic limitations of the R&D are reached. Towards specifying a particular functional mapping, we appeal to and align with the literature on the technology life cycle. In such studies, the logistic function, or S-shape model, is used to describe the typical stages of the technology life cycle (Pan and Kohler, 2007): laboratory-invention; decisive demonstration; explosive take-off and turbulent growth; continued high growth; slowdown; and maturity. The use of the logistic function in this way is well established, from Griliches (1957) through to Grübler et al. (1999), Grübler (2003), and Jaffe et al. (2003). In the context of a dynamic input-output economic model, Pan (2006) suggests a stylised S-shape for the unit costs of sector production in response to technical change. Lee et al. (2005) use a logistic curve to measure the way firms use knowledge and find that this measure is closely correlated with R&D expenditure. Pan and Kohler (2007) use a logistic curve to model the technological progress of wind power in the UK measured by the cost of production of a unit of electrical energy. Huchzermeier and Loch (2001) illustrate a proposed link between market payoff variability and requirement variability using an S-bend, but do not specify the curve. Haupt et al. (2007) point to evidence of an S-shape to the evolution of patent applications. Motivated by this literature, we consider the logistic function as a way to capture the UPC profile described above. We map the accumulation of technical breakthroughs in the technology under development to increases in the UPC as described by the S-shape curve, which can then be inverted to give a CPU profile. The use of the logistic function in this way, while stylised, allows for a business-appropriate and economically meaningful way to measure R&D progress.

Thirdly, we wish to design the decision support tool to capture learning-by-doing in the operational phase of the deployed technology. Drawing on the literature pertaining to learning (or experience) curves, we take into account a learning-by-doing effect, whereby the experience accumulated by the firm in handling the deployed technology leads to production efficiencies that further reduce production costs. Learning-by-doing is a concept commonly considered in industrial processes. The concept was brought to academic attention by Wright (1936) and has recently received expositions in Pan and Kohler (2007); Nemet (2009); Yu et al. (2011); Anzanello and Fogliatto (2011); Yeh and Rubin (2012). The learning-by-doing effect provides a further increase (decrease) in the UPC (CPU) during the operational phase, which is unrelated to R&D but which results in further improvements to the cash flows. The specification we follow aligns with Klaassen et al. (2005), Pan and Kohler (2007) and Yu et al. (2011).

The contribution we make, therefore, is in the practice relevance of the real options based decision support tool we propose, which combines the use of a Poisson process to model progress in the technological innovation process, a logistic function to give an economic interpretation to this progress in terms of cost efficiencies, and a learning-by-doing dynamic in the operational phase to model further cost efficiencies post-deployment of the technology. While the Poisson distributed breakthroughs are assumed to arrive independently between consecutive periods, the logistic functional mapping imposes memory in the process. This is consistent with our assumption that the technological innovation is endogenous to the firm. Our work is in the spirit of Berk et al. (2004) in modelling technical risk, cash flow risk, competition and obsolescence risk, and learning-by-doing effects.

The paper proceeds as follows. In Section 2, we set out the R&D investment opportunity for the firm and discuss the investment flexibility that the firm engineers into the process and how these translate into real option structures. The investment opportunity is segmented into alternative phases, comprising staged periods in the R&D process and an operational period, which assumes the developed technology is deployed commercially. Section 3 is dedicated to the Poisson modelling of technical risk and the evolution of progress within the R&D process, along with a description of the logistic functional mapping of this progress to the UPC measure and, its reciprocal, the CPU measure. Section 4 describes the valuation of the real option structures underlying the staged R&D phases, along with the valuation of the cash flows (both inflows and outflows) over the operational phase. The required discretisation of these value formulations for the simulation framework is also presented. Section 5 presents the specific static and stochastic variable settings used in the CO_2 recycling technology simulation. The simulation results are discussed in Section 6, with a range of benchmarking. Section 7 concludes.

2. R&D Investment Opportunity Setup

The objective of our study is to present a decision support tool with practitioner appeal that allows for the evaluation of R&D investments in a setting where a firm seeks to develop a new technology designed to produce an existing commodity. In our application, we consider the development of a new technology designed to recycle $\mathrm{CO_2}$, resulting from some production process, into a tradable commodity, such as, for example, natural gas (methane). We consider an investment process engineered by the firm with the decision flexibility outlined in Figure 1. We assume that the firm seeks to invest in an early phase of R&D that would span a defined period of time, with the objective being to gauge progress with the new technology development. The firm wishes to then make a decision at the end of this period as to whether to proceed to a further phase of R&D (branch (a)) if progress is deemed sufficiently successful or to abandon the R&D in the case of insufficient progress (branch (b)). We consider a fixed period for the early phase of R&D on the assumption that the firm sets a target cost per unit of production at the end the period, to be achieved through the R&D process. If the R&D leads to a reduction in the cost per unit of production below this target level then the firm decides to proceed with a subsequent late phase of R&D, and if not then the firm decides to abandon the R&D process.

The purpose of the late phase of R&D, if progressed, is to advance the new technology development further to a state that would make commercialisation feasible. Given the uncertainty over this progress in the R&D process, the firm engineers in flexibility to commercialise the technology at any point over the period of the late phase of R&D that is optimal to do so (branch (c)). The maximum duration of the late phase of R&D is assumed to be dictated by some secured period of intellectual property or patent protection relating to the new technology, possibly emanating from the early phase of R&D. Such a period of intellectual property or patent protection can be quite accurately predicted. If the R&D fails to progress sufficiently to make the new technology commercially viable by the end of this period then it is assumed that the R&D process fails and terminates with no further investment (branch (e)). Upon deciding to commercialise the technology, there is a period of construction whereby the infrastructure is put in place to deploy the technology. The operational phase then begins and progresses unimpeded for the duration of the operational phase, with this duration reflecting the projected life cycle of the technology (branch (g)).

In order to reflect reality better, however, this base case setting can be modified to account for the potential impact of competition and regulatory risk that may impact the R&D investment process, including the operational phase. Such events are assumed exogenous to the firm and can occur randomly at any point in time.[1] We assume that the firm engineers in the potential for a competition or regulatory event during the late phase of R&D that has an extreme impact and forces the firm to abandon the R&D entirely, with no further investment (branch (d)). The competition event could be, for example, the release of a new technology in the market place that is superior to the proposed technology being developed by the firm. This superiority may come in the form of greater production output of the commodity in question, greater production cost efficiency, or some other form of technological dominance. The regulatory event could be, for example, the introduction of some new regulatory directive that would deem the new technology unusable. The assumption that the firm abandons the R&D process is, of course, extreme but is assumed for illustrative purposes. A competition or regulatory event during the late phase of R&D is assumed to occur with some assumed probability (x%, see branch(d)), which may be informed by the level of competition or regulatory oversight in the sector in question. This assumption is easily built into our simulation framework, which we describe in Section 4.

We also assume that the firm engineers in the potential for a regulatory event during the operational phase that has the impact of reducing the level of cash flows from the deployed technology by some assumed factor (branch (g)). The regulatory event could be the introduction of some new regulatory directive pertaining to the produced commodity that would force the firm to reduce the production rate of the deployed technology. The regulatory event is assumed to occur with some assumed probability (y%, see branch (g)) and at some random time during the operational phase, after which the level of cash flows is reduced by a fixed amount until the end of the operational phase. For our CO2 recycling technology example considered later, we assume

[1] We thank an anonymous referee for suggesting this modelling of competition and regulatory risk as an additional layer to the decision support tool.

that cash flows are reduced by 50% in such a case. We do not consider a competition effect in the operational phase under the assumption that once the technology is deployed then it cannot be readily replaced with a superior technology that enters the market place.

As a final consideration, we assume the firm wishes to incorporate the potential for cost efficiencies gained over the operational phase from working with the newly deployed technology. That is, the firm recognises that there is potential for non-R&D related learning-by-doing benefits as discussed earlier.

The proposed investment structure above can be generalised in a number of directions. Firstly, the early phase of R&D assumes a single decision node to abandon or proceed with the R&D investment. This could readily be extended to a series of decision nodes throughout the early phase period where the firm may consider abandonment of the investment. This would require the firm to set a series of targets for the cost per unit of production that are progressively more ambitious. Brandão et al. (2018), for instance, consider such multistaging in a pharamaceutical R&D setting. Secondly, the late phase of R&D does not incorporate the option to abandon the R&D, except of course at the end of the period. The decision is therefore to deploy the technology or continue with the R&D. Sequenced abandonment options could be readily incorporated into the design of the investment opportunity through a similar series of targets for the cost per unit of production. Thirdly, we propose an extreme impact in the event of a competition or regulatory event during the last phase of R&D, i.e. we assume that such events force the firm to abandon the R&D in its entirety. This could be relaxed, for example, through allowing the R&D to proceed despite the competition or regulatory event but with an option to expand the investment to meet a higher threshold for commercialisation. Finally, we assume that in the event of a regulatory event during the operational phase, the level of cash flows is reduced by a fixed amount for the remainder of the period. This could be relaxed to assume, for instance, a staged decline in the level of cash flows reflecting the gradual impact of the regulatory event. There is of course a multitude of other investment flexibilities that could be considered but the investment structure we assume is deemed sufficient to illustrate the decision support tool in our CO2 recycling technology context.

With the desired investment flexibilities now embedded in the investment process, we move next to translate these to appropriate real option structures. The decision support tool we propose allows decision makers to model this real optionality and ascertain the value of the R&D investment opportunity. Beginning with the early phase of R&D, the objective is to test the viability of the new technology, where the degree of progress made, as measured at the end of the period, informs the firm's decision as to whether to continue into a subsequent phase of R&D or to abandon the project entirely. The early phase of R&D is assumed to start at time $t = 0$ and end at a specified future time $t = T_E$, a fixed period set by the firm. The early phase of R&D is, therefore, in effect a call option on the next phase of R&D and so may be appropriately modelled as a European real option structure. The early phase of R&D is assumed to have a proposed

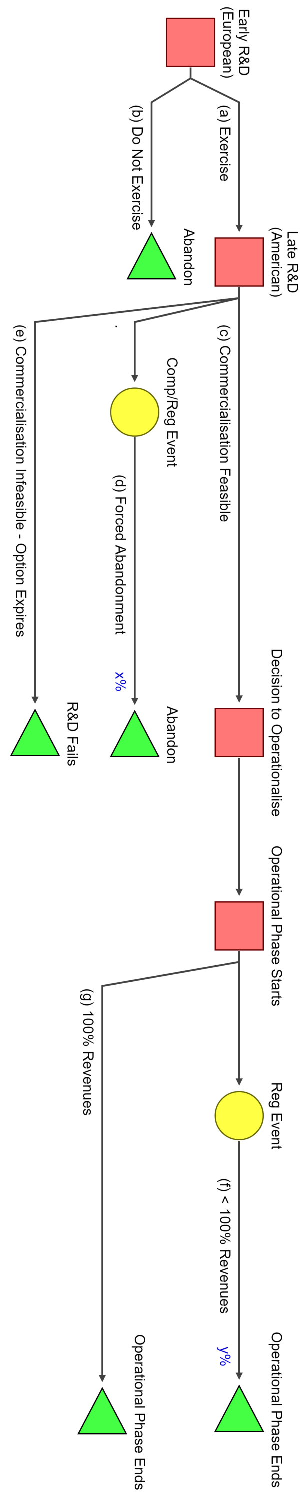

Figure 1: Assumed Decision Tree for R&D Investment in New Technology

investment budget of I_E. The decision support tool we propose allows a financial manager to value the early phase of R&D for comparison against this proposed investment budget I_E. The real option value of this early phase of R&D should be greater than the proposed budget for the project to commence (Mazzucato, 2013; UNEP, 2017). Deciding not to exercise the European real option amounts to the firm exercising its option to abandon the R&D investment.

The objective of the late phase of R&D is to bring the technology to a state of commercial viability, which is marked by the decision of the firm to deploy the technology. It is assumed that the late phase of R&D starts immediately after the early phase of R&D, i.e. at time period T_E+1, and continues until an unknown future time period, T_C, when the decision to commercialise and deploy the technology is made. Since this decision is made as soon as the technology is deemed commercially viable, we model this phase of the R&D process with an American real option structure. The late phase of R&D is, therefore, in effect a call option on the operational phase, with the flexibility to call at any point once commercially viable. We assume that the decision to deploy the technology requires the construction of dedicated infrastructure, which is assumed to cost an amount D. The decision to deploy the technology must be made before some predetermined date T_F, which is assumed to define a period of intellectual property or patent protection $(T_F - T_E)$, after which time the R&D is deemed to have failed if not commercially viable and the investment ceases. Thus, we have that the deployment decision date $T_C \leq T_F$. We use T_F as the expiry date for the American real option structure.

The time required to build the infrastructure to deploy the technology is assumed to last for B time periods, during which time there is no further R&D investment, and after which we enter the operational phase at time $T_C + B + 1$. We assume a time limit of L time periods on the technology, such that the operational phase ends at time $T_C + B + 1 + L$. As already described, during the lifetime of the plant, it is expected that there will be improvements in the efficiency of the production process due to learning-by-doing.

The above mapping confirms that the R&D investment opportunity engineered by the firm is a compound real option structure, whereby the early phase of R&D represents a European real option structure that is exercisable into the late phase of R&D, which represents an American real option structure that is exercisable into the operational phase. Compound option structures have been considered in the literature both for general R&D and specifically for pharmaceutical R&D. For example, Herath and Park (2002) use compound real options to separate four sequential stages of an R&D project: R&D investment, new product introduction, first expansion, and second expansion. Cassimon et al. (2011) model pharmaceutical research stages as a series of nested compound options.

The purpose of the decision support tool is to appraise the R&D investment opportunity just described, and value the embedded real optionality. Our valuation approach works recursively through the phases, transitioning backwards from the operational phase. An essential element of the valuation approach is how we model R&D progress over time and how we map this R&D progress to give it a meaningful economic interpretation to support the decision making process.

We begin with this modelling and mapping exercise (Section 3), and then proceed to discuss in sequence the valuation approach for each of the respective phases (Section 4).

CO_2 Recycling Technology Application Context

With the general R&D investment opportunity set out as above, we take the opportunity to introduce our specific application context. We focus on CO_2 recycling technology, designed to transform CO_2 emissions into alternative products that may be utilised or sold into the market place. In light of climate change, a number of approaches have been advocated to address the rising levels of CO_2 in the atmosphere. The first approach involves direct efforts to decrease the amount of CO_2 produced, by switching, for example, from fossil fuel sources of energy to renewable resources (Bockris, 1977). This may be incentivised by placing a price on the right to emit greenhouse gas (GHG) emissions, such as done under the EU Emissions Trading Scheme (EU ETS). Another approach is to prevent CO_2 emissions from entering the atmosphere by means of carbon capture and storage (CCS) solutions, whereby captured CO_2 is stored under pressure in underground caverns or depleted oil and gas fields, and where it may combine with the surrounding rock (Centi and Perathoner, 2009). However, there are many unanswered questions regarding CCS technology, which is often expensive. A potentially more cost effective and lucrative approach is CO_2 recycling, with emissions being used as a feedstock for synthetic transformations in both the chemical and fuel industries. CO_2 is already used in the synthesis of a wide range of chemical products, including urea, methanol and some polymers. In terms of carbon balance, CO_2 recycling processes can be viewed as carbon neutral, and can thus provide an additional approach to mitigating CO_2 emissions (Balzani et al., 2008),Aresta2010. The recycling of CO_2 is most relevant for those industrial activities producing large volumes of CO_2, for instance, power generation or cement manufacturing (Metz et al., 2005). However, converting CO_2 into a useful commodity is not straightforward as it requires an energy input due to the stability of the CO_2 molecule. Much of the current research in the area of CO_2 recycling focuses on the use of renewable sources of energy to drive the process, with research groups exploring electrocatalytic, photocatalytic or photoelectrocatatlyic routes (Centi and Perathoner, 2009).

Research into CO_2 recycling has led to numerous processes at different technological stages, ranging from projects very much at nascent stages (fundamental research such as, for example, photocatalytic reduction of CO_2, as in Frayne et al. (2018), to others at more established, pilot stages of deployment (such as power to fuel technologies). We apply our new R&D investment decision support tool to an existing procedure for producing methane, which is not economically viable at present and so is a candidate for further R&D endeavour. Specifically, we examine a chemical route known as the Sabatier reaction, which was discovered in the early 20th century by Paul Sabatier.[2] In brief, the Sabatier reaction combines carbon dioxide (CO_2) with hydrogen

[2]For a biological method the interested reader is directed to Prof Jerry Murphy and Dr Magdalena Czyrnek-Deletre of University College Cork https://www.ucc.ie/en/eri/projects/lifecycleanalysisofrenewablegaseousfuelandbiologicalmethanationsystems/

(H_2), to make methane (CH_4), i.e. natural gas, and water (H_2O). The reaction can be facilitated by a nickel catalyst, although other catalysts, such as ruthenium on aluminum oxide, have been suggested. The technical details of this CO_2 recycling technology process are provided in Appendix A.1.

3. Modelling and Mapping of R&D Progress

We wish to model progress in the R&D process as incremental advancements in the proposed technology that lead to economically measurable improvements. We synthesise approaches from the technology adoption literature and the technology life cycle literature. In the former literature, the Poisson process has been advocated for the modelling of technological innovation arrival in game theoretic models of technology adoption (Farzin et al., 1998; Hagspiel et al., 2015; Huisman and Kort, 2004). In the latter case, the logistic function has been advocated as a model of the technology life cycle (Griliches, 1957; Grübler et al., 1999; Grübler, 2003; Jaffe et al., 2003; Pan, 2006; Pan and Kohler, 2007). We leverage these ideas to design a firm-level decision support tool that allows a financial manager to appraise the R&D investment opportunity presented in the previous section.

The uncertainty surrounding the R&D investment opportunity emanates from the fact that, *a priori*, it is not known how successful the R&D will be and from the fact that technological breakthroughs occur randomly over time. We therefore require an appropriate way to model this uncertainty. We use a Poisson counting process to model the accumulation of technological breakthroughs in the R&D process. The Poisson process models increments in the technological breakthrough count as a Poisson random variable with an assumed constant intensity parameter of λ. We expect the rate of arrival of technological breakthroughs in the early phase of the R&D process to be more rapid than the late phase of the R&D process, with advancements in the latter phase less frequent as the more obvious solutions to technological problems are exhausted. As a simple model, we propose that $\lambda = \lambda_{Early}$ during the early phase of R&D ($t \leqslant T_E$), while during the late phase of R&D ($T_E + 1 \leq t \leqslant T_C \leqslant T_F$) the value of $\lambda = \lambda_{Late}$. We impose the condition that $\lambda_{Early} \geqslant \lambda_{Late}$. We define the cumulative number of technological breakthroughs, d_t, up to time t as

$$d_t = \sum_{k=1}^{t} n_k$$

where n_k is drawn from a Poisson distribution at time period k.

To support the decision making process, we need to give our abstract modelling of R&D progress over time a workable economic interpretation. We propose a tractable functional mapping of the Poisson counting process for this purpose. As we are considering a technology that produces a tradable commodity, the functional form we suggest maps the cumulative number

and (Ahern et al., 2015)

of technological breakthroughs, d_t, to the number of units of the commodity produced by the technology per unit of cost (i.e. units of production per unit of cost or UPC). We consider a logistic functional mapping: (i) that is monotonically increasing in the number of technological breakthroughs, reflecting a practical assumption that incremental advancements in the proposed technology lead to increased production of the commodity per unit of cost; and (ii) that allows the rate of production of the commodity per unit of cost to vary over time with a stylised dynamic that allows for lower levels of economic impact (per technological breakthrough) at earlier stages of the R&D timeline, which increases to higher levels of economic impact (per technological breakthrough) in the middle stages of the R&D timeline, fuelled by increased knowledge and know-how around the technology, and, finally, the reversion to lower levels of economic impact (per technological breakthrough) as the limitations of the technology are approached at latter stages of the R&D timeline. This stylised dynamic implies that technological breakthroughs have a varying impact on the UPC measure over time.

We now formally define our logistic functional mapping. We propose a generic logistic function $L(x) = \frac{1}{1+e^{-x}}$, illustrated in Figure 2(a). The function has a limit to the left of zero, and a limit to the right of one. $L(x)$ is not immediately applicable to measure improvements in the UPC as the R&D progresses, as we would like to have a curve similar to that in Figure 2(b) to follow the pattern proposed in our phased R&D model. We introduce a function, $UPC(d_t)$, which generalizes $L(x)$ and provides values of UPC for a given number of technological breakthroughs, d_t:

$$UPC(d_t) = \frac{c}{1 + e^{-(ad_t+b)}} \tag{1}$$

In defining the parameters, a, b and c, of this logistic function generalisation, we wish $UPC(d_t)$ to have three desirable characteristics, which describe the relationship between the number of breakthroughs and the UPC. First, we wish to initialise UPC with $d_0 = 0$ at the start of the R&D process. Second, we wish to impose an upper limit, UPC_{Max}, to which the UPC tends with continued advancements in the R&D. Thirdly, we wish to set a scale for $UPC(d_t)$ to represent how far the technology is from commercialisation.

To achieve the first of these goals, we set $UPC(0) = UPC_{Min}$, where UPC_{Min} is the value of UPC using the existing technology before beginning the R&D and may be considered a lower limit of UPC. This value may be small as the existing technology may be experimental and reside in a laboratory setting.[3] For the second goal, we set $c = UPC_{Max}$, where this upper limit may be unknown and require estimation. There is a degree of subjective judgment involved in estimating UPC_{Max}, which may be informed by internal targets for the project, physical limits

[3] We preclude the lower limit UPC_{Min} value from being zero as it would make $UPC\,(d_t)$ undefined, as will become apparent later.

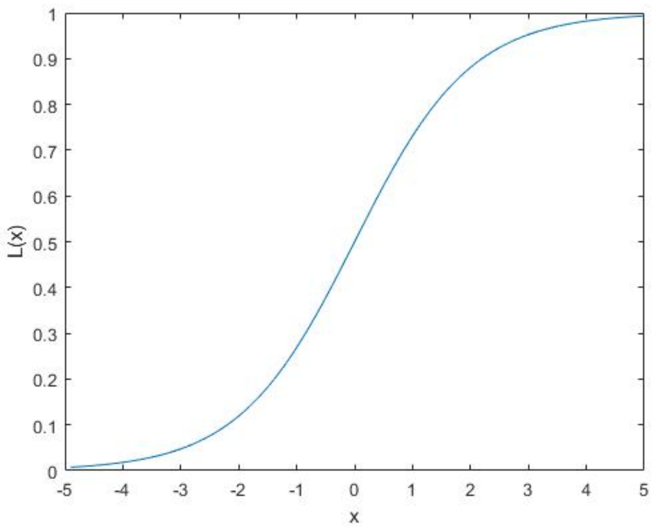

(a)

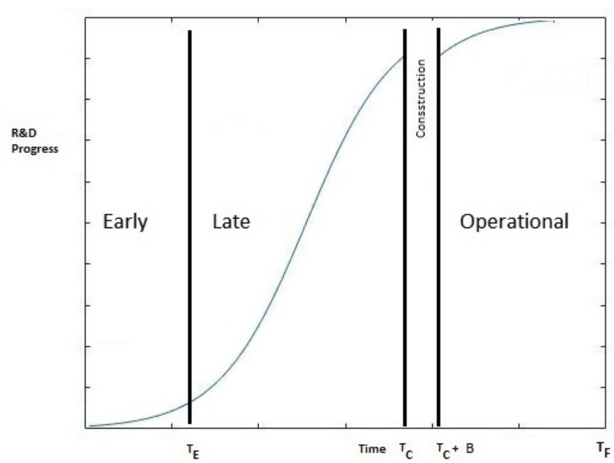

(b)

Figure (a) illustrates the generic logistic function $L(x) = \frac{1}{1+e^{-x}}$, which has a limit to the left of zero and a limit to the right of one. Figure (b) presents the logistic function representation of the units of production per unit of cost (UPC) measure given in Equation (2), which is a specific case of the generalisation in Equation (1). The diagram illustrates the cumulative progress in R&D as measured by UPC. We see that the rate of increase in the UPC varies over time with a stylised dynamic that describes lower levels of economic impact (per technological breakthrough) in the early phase of R&D (up to the end date of this period,T_E), which increases to higher levels of economic impact (per technological breakthrough) in the late phase of R&D, fuelled by increased knowledge and know-how around the technology, which then reverts to lower levels of economic impact (per technological breakthrough) as the limitations of the technology are approached at the end of the late phase of R&D. This is observed between the start of the late phase of R&D, time $T_E + 1$, and the technology deployment decision date T_C. During the operational period the UPC continues to increase due to learning-by-doing, but at a slower rate of increase. This is observed from time $T_C + B + 1$, the start date of the operational phase.

Figure 2: Logistic Functional Mapping

of the proposed technology, or regulatory requirements. For the last goal, we define a parameter d_R to be the value of d_t associated with the maximum rate of increase of the UPC, i.e. the point of inflection of the logistic curve. We select d_R to reflect the present state of the technology, be it a nascent technology that is still in a laboratory, an advancing technology that is some way off commercialisation, or an emerging technology that is near commercialisation. The value of d_R is necessarily different for technologies at different stages of maturity and determines how many more technological breakthroughs need to be made in order to reach the maximum rate of increase in the UPC.

We assume for the sake of simplicity that there are a nominal 100 technological breakthroughs within the R&D process. Thus, our estimation of d_R is a a measure of how much progress has to be made to reach the maximum rate of increase of the UPC. For a well advanced, emerging technology, which is close to commercialisation, we may consider a lower value, for example, $d_R = 25$. In this setting, we would see that the UPC more rapidly approaches its upper limit with continued technological breakthroughs. If the R&D concerns an advancing technology, which is not yet ready for commercialisation, then we may select, for example, $d_R = 50$. This positions the maximum rate of increase of the UPC at the middle of the logistic curve, with the rate of change being symmetrical either side of the inflection point but of different sign. For a nascent technology, where the R&D is significantly away from a state of commercialisation then we may pick, for example, $d_R = 75$. This pushes the maximum rate of increase of the UPC further down the logistic curve, leading to more gradual improvements in the UPC and requiring more technological breakthroughs to approach the upper limit. These choices are illustratively contrasted in Figure 3, where the shift in the inflection point of the logistic curve is evident. The d_R estimate is of course subjective but the previous discussion gives some guidance based on how advanced is the technology.

Using our estimates for UPC_{Min}, UPC_{Max} and d_R, we can determine the values of a and b in Equation (1) as follows:

$$b = -\ln\left(\frac{UPC_{Max}}{UPC_{Min}} - 1\right)$$

$$a = -\frac{b}{d_R};\ \ UPC_{Max} = \lim_{d_t \to \infty} UPC(d_t)$$

Substituting these into Equation (1), we get

$$UPC(d_t) = \frac{UPC_{Max}}{1 + \left(\frac{UPC_{Max}}{UPC_{Min}} - 1\right)^{\left(1 - \frac{d_t}{d_R}\right)}} \tag{2}$$

In order to facilitate our monetary based calculations later, we invert our measure of progress in

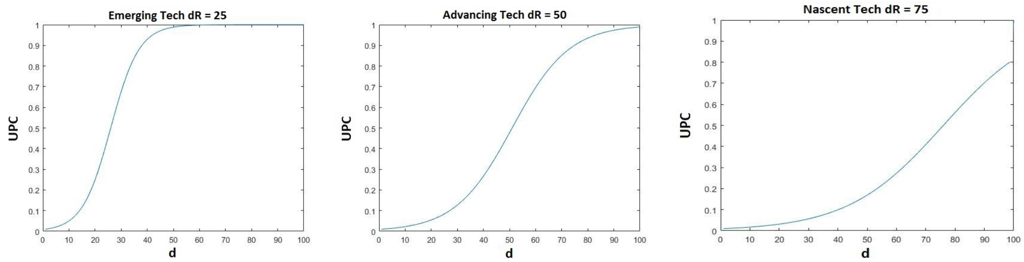

The plots show the stylised logistic curves for technology that is close to commercialisation and emergence onto the market (emerging technology; left plot), technology that is advancing in terms of R&D but still some way off commercialisation (advancing technology; middle plot), and experimental technology that is at an early stage of the R&D process (nascent technology; right plot). The logistic curve provides a formal link between advances in the R&D process, as measured by the technological breakthroughs variable d_t, and improvements in the technology output, as measures by the number of units of production per unit of cost (UPC). We assume for the sake of illustration that there are a nominal 100 technological breakthroughs possible. Defining d_R to be the value of d_t associated with the maximum rate of increase of the UPC, i.e. the point of inflection of the logistic curve, the plots consider three alternative values of d_R, reflecting the maturity of the technology in question. For the emerging technology, we set $d_R = 25$; for the advancing technology, we set $d_R = 50$; and for the nascent technology, we set $d_R = 75$. We scale the examples, such that the lower limit $UPC_{Min} = 0.01$ and the upper limit $UPC_{Max} = 1$.

Figure 3: Stylised Logistic Curves Measured for Alternative Technology Maturities

the R&D process from the number of units of production per unit of cost, to the number of units of cost per unit of production, i.e. we change to cost per unit (CPU). We rearrange Equation (2) as follows:

$$CPU(d_t) = \frac{1}{UPC(d_t)} = CPU_{Min}\left[1 + \left(\frac{CPU_{Max}}{CPU_{Min}} - 1\right)^{\left(1 - \frac{d_t}{d_R}\right)}\right] \tag{3}$$

where CPU_{Min} is the minimum estimated value of CPU as $d_t \to \infty$ (which is the inverse of UPC_{Max}) and CPU_{Max} is the starting value of CPU (which is the inverse of UPC_{Min}).

Remark on Logistic Curve Estimation

We consider next how to estimate the logistic curve. We suggest, for example, that the firm could use past observations on the cost reductions achieved over time on previous technology R&D investments. Let us assume that the firm has past discrete observations over time of the reductions achieved in the cost per unit of production of a given commodity for an earlier technology developed by the firm through an internal R&D process. In this scenario, the firm will know the maximum cost per unit of production (CPU_{Max}) at the beginning of the process, while it will also know the minimum cost per unit of production (CPU_{Miin}) achieved by the end of the R&D process or, potentially, during the operational phase of the technology. Fitting the logistic function ($CPU(d_t)$) to this data would then allow the firm to infer λ parameter estimates by means of calculating the average number of breakthroughs per period (e.g. month or year) between any two points on the curve. For example, if the firm is seeking an estimate for the λ

parameter for, say, a two-year early phase of R&D as per our illustrative example, then the firm could use the fitted logistic function for the earlier technology to infer estimates for $d_0(=0)$ and d_2. It can be readily determined from this that, per month, we expect $(d_2-d_0)/24$ breakthroughs. While this assumes comparability between the R&D processes for the proposed technology and the earlier technology, the firm can easily modify this estimate.

3.1. Dynamic Production/Production Cost Profile

A limitation of the model thus far is that once the logistic curve estimate of the UPC profile is set, it remains static throughout. This does not allow an opportunity for the firm to respond to the information revealed from the progression of the R&D process and to revise its UPC logistic curve estimate accordingly. We therefore modify the model to allow the firm to respond in a couple of ways that would be reflective of how the decision support tool would be applied on a real-time ongoing basis in practice. A firm has to estimate two key parameters when applying the UPC logistic curve in practice. These are (i) the maximum UPC (UPC_{Max}) estimate to which the UPC logistic curve tends to as the R&D progresses, and (ii) the shape of the UPC logistic curve as determined by the point of inflection parameter estimate, d_R, and which is set to reflect the maturity of technology, i.e. whether nascent, advancing or emerging as described in stylised form in Figure 3.

We first discuss the UPC_{Max} estimate. In practice, the firm is likely to learn more about the UPC_{Max} estimate as the R&D progresses and, hence, be positioned to forecast this more accurately. So the firm may wish to revise the UPC_{Max} estimate depending on whether the initial estimate was too optimistic or too pessimistic. If the initial UPC_{Max} estimate is found to have been too high then the estimate can be revised down, which serves to pull down the entire UPC logistic curve. Conversely, if the initial UPC_{Max} estimate is found to have been too low then the estimate can be revised up, which serves to pull up the entire UPC logistic curve. By adjusting the UPC_{Max} estimate in this way, we therefore allow for production cost estimates to be revised upwards or downwards in response to new information visible internally to the firm around the R&D process. Furthermore, as time progresses the firm's uncertainty over the UPC_{Max} estimate is likely to reduce. This can be readily modelled via a time varying standard deviation around the UPC_{Max} estimate. This modelling of variable production costs estimates with declining uncertainty is in the spirit of Weitzman et al. (1981).

While the above allows for the production cost estimates to be revised as the R&D progresses, the initial estimate for the shape of the UPC logistic curve remains unchanged throughout. However, in practice, the shape of this curve may need to be revised as well. Indeed, the process of R&D is designed to advance the underlying technology, which means the shape of the UPC logistic curve will inherently change. With reference to Figure 3, the purpose of R&D is to progress, for instance, a technology at the nascent stage through to the advancing stage and on to the emerging stage. But, of course, there may be instances when the technology is revealed to be less advanced than anticipated and the UPC logistic curve would need to be revised accordingly. Such revisions can be practically applied by adjusting the point of inflection estimate,

d_R. Decreasing (increasing) d_R means that the UPC logistic curve adjusts to reflect advancements (regressions) in the technology relative to initial and ongoing estimates. The implication of this for production costs is that a decrease (increase) in d_R leads to a decrease (increase) in the CPU profile.

Subsequent to presenting the simulation framework in the next section, we outline the specific settings of the above modifications in Section 5.3, as applied for our CO_2 recycling technology case. For comparison purposes, we keep the dynamic UPC_{Max} modification separate from the dynamic UPC logistic curve modification. However, the two dynamic effects could be readily combined.

4. Simulation and Valuation

4.1. Operational Phase: Cash Flow Valuation

In the operational phase, we assume that no further R&D is conducted by the firm. However, we take account of cost efficiencies brought about by learning-by-doing (Pan and Kohler, 2007; Nemet, 2009; Yu et al., 2011; Anzanello and Fogliatto, 2011; Yeh and Rubin, 2012). The learning-by-doing results in further decreases in the CPU, which are unrelated to R&D but result from the knowledge built up from handling the deployed technology. Following Klaassen et al. (2005), Pan and Kohler (2007) and Yu et al. (2011) we propose a discrete model to account for the learning-by-doing effect:

$$CPU_t = A \cdot CP_t^{-\eta} \tag{4}$$

The time index t denotes any time period during the lifetime of the deployed technology, i.e. $T_C + B + 1 \leqslant t \leqslant T_C + B + L$, where, as detailed in Section 2, T_C is the date on which the decision is made to deploy the technology, the construction time for the associated infrastructure is of length B, and the useful lifespan of the deployed technology is of length L. CP_t is notation introduced to denote the cumulative production of the commodity at the end of period t, where we assume that the rate of production of the commodity per unit of time, N_P, is constant. The cumulative production at the end of period t is therefore $CP_t = N_P(t - (T_C + B + 1) + 1)$. A is constant and is defined next, and $\eta > 0$ is the learning-by-doing parameter.

In setting the constant A, we note that the CPU value at the start of the operational phase, when $t = T_C + B + 1$, is the same as the CPU value at the end of the late phase of R&D, when $t = T_C$, i.e. $CPU_{T_C+B+1} = CPU_{T_C}$. This reflects the assumption that during the construction phase, no further R&D is performed by the firm. At the beginning of the operational phase, when $t = T_C + B + 1$, we have that

$$\begin{aligned} CPU_{T_C+B+1} = CPU_{T_C} &= A \cdot N_P^{-\eta} \\ \Rightarrow A &= CPU_{T_C} N_P^{\eta} \end{aligned}$$

Therefore, Equation (4) can be modified to give the following representation of the CPU at any date t during the lifetime of the deployed technology ($T_C + B + 1 \leqslant t \leqslant T_C + B + L$):

$$CPU_t = CPU_{T_C} \left(t - (T_C + B + 1) + 1\right)^{-\eta}. \tag{5}$$

This describes a CPU that declines with the passage of time during the operational phase at a rate dictated by the learning-by-doing parameter η.

With the dynamic of the CPU under learning-by-doing established, we may proceed to set out the cash flows from the deployment of the technology. At each point during the late phase of R&D, we wish to estimate the cash flows from the deployed technology, assuming that the decision to deploy is taken at that point in time. The present value (at time T_C, the technology deployment decision date) of the cash flows from the deployed technology, denoted CF_{T_C}, is therefore

$$CF_{T_C} = \sum_{t=T_C+B+1}^{T_C+B+L} \left(R_t - VC\left(CPU_t\right) - FC_t\right) e^{-r(t-(T_C+B+1))}$$

where R_t is revenues, $VC\left(CPU_t\right)$ is variable costs as a function of production costs CPU_t, and FC_t is fixed costs. In choosing the discount rate, r, we appeal to (Schwartz, 2013) who sets out three real option risk-neutral valuation cases. Under the first two cases, the risk-neutral distribution is either known or unknown but can be inferred from futures prices or other traded assets. The third case is appropriate in our situation as we are dealing with an R&D investment context, where the risk-neutral distribution is unknown and cannot be fully inferred from futures prices or other traded assets. (Schwartz, 2013) suggests that in this case the risk-neutral distribution can be inferred from an equilibrium model, such as the CAPM. For convenience, we choose r to be the risk free rate of interest and, in so doing, assume a zero beta for the CAPM specification. For practical applications, the firm must, of course, consider a non-zero beta, chosen to reflect the risk of the technology been developed, whether this is comparable to the average risk of the firm or is an above- or below-average level of risk.

CO_2 Recycling Technology Application Context

We consider now how the above cash flow representation should be specified under our CO2 recycling technology application. On the revenue side, income is earned by selling the produced commodity into the market place at prevailing market prices. This is the only source of revenue considered under the project. Total revenues for period t are determined as the product of the constant production rate N_P and the commodity price for the period P_t, i.e. $N_P \times P_t$. On the expenditure side, total costs for period t are assumed to comprise both variable and fixed costs. Production costs are the main component of the former, which are determined as the product of the constant production rate N_P and the cost per unit of production CPU_t, i.e. $N_P \times CPU_t$. We take account, however, of the fact that there are cost savings in respect of the emissions allowances saved through the recycling process. Assuming that we have a constant rate of CO_2

savings, N_E, and a prevailing emissions price for period t, E_t, then the value of the savings is given by $N_E \times E_t$. There is a cost though to the carbon capture stage of the CO_2 recycling process. We assume this cost is effectively fixed but that there is a learning-by-doing effect here as well. We adapt Equation (5), giving a formulation for the cost of carbon capture CC_t at any time $t > T_C + B$:

$$CC_t = CC_{T_c+B}(t - (T_C + B))^{-\eta} \tag{6}$$

where the same learning-by-doing parameter, η, as earlier is used for convenience. Other fixed costs, F, associated with the deployed technology are assumed estimated for the first period of the operational phase $T_C + B + 1$, after which we allow for a constant growth at the rate f per period. This growth rate could relate, for instance, to the case of an equipment lease contract, which is a fixed cost with no link to the variability of production but which may increase over time due to payment schedule changes. However, in practice, f could be readily set to zero. It therefore follows that for the CO2 recycling technology context:

$$\begin{aligned} CF_{T_C} = \sum_{t=T_C+B+1}^{T_C+B+L} \Big(& N_P \times P - N_P \times CPU_{T_C} \left(t - T_C - B\right)^{-\eta} + N_E \times E_t \\ & - N_E \times CC_t - F e^{f(t-(T_C+B+1))} \Big) e^{-r(t-(T_C+B+1))} \end{aligned} \tag{7}$$

Within our simulation framework, the cash flow calculations require us to simulate the underlying commodity and associated emissions prices over the full duration of the operational phase. The technical details of this stochastic price modelling are provided in Section 5.

4.2. Late Phase of R&D: American Real Option Valuation

The late phase of R&D is modelled as an American real option structure. We consider the underlying asset of this real options structure to be CF_{T_C}, the cash flows from the deployment of the technology, defined in the previous section. The strike price of the American real option structure is the construction cost of the deployed technology, D, assumed to be constant. As noted earlier, the expiry date T_F of the American real option structure is assumed to be dictated by some period of intellectual property or patent protection, after which time the R&D is deemed to have failed if the technology is not commercially viable and the investment ceases. The value of the late phase of R&D at any time $T_E \leq t \leq T_F$ can therefore be expressed as follows:

$$C_t^L = \sup_{T_E \leq t \leq T_C \leq T_F} E_t \left[\max\left(CF_{T_C} - D,\, 0\right)\right] e^{-r(T_C - t)}$$

where the value of C_t^L depends on the information available at time t. For implementation purposes, we follow the random forest based American option valuation method of Glasserman (2003). We use the mean of the high and low estimators proposed by Glasserman (2003) as the best estimate of the American real option values we require.

We simulate over the length of the American optionality period $(T_F - T_E)$, using s time periods of length δt months. We set the number of branches at each time step of the simulation to be $b_1, b_2, ..., b_s$, such that at each i^{th} stage there are b_i new branches. This use of s periods effectively imposes a Bermudan real option structure since it is only exercisable on a limited number of dates. This is appropriate in our setting as in practice the firm considers the decision to deploy the technology on a discrete basis, rather than on a continuous basis. The use of a discrete number of possible exercise dates is supported by Schwartz (2004).

As we do not know the number of breakthroughs that will be achieved under the early phase of R&D, we wish to run multiple Monte Carlo simulations to value, at time T_E, the American real option structure for each of a range of possible d_{T_E} values, i.e the breakthrough count at the end of the early phase of R&D. As this requires considerable computing time, we wish to optimize the choice of d_{T_E} values to use. The breakthrough count d_{T_E} is assumed to be described by a $Poisson(d_{T_E}, \lambda_{Early})$ distribution, where λ_{Early} is the constant rate of breakthrough arrivals during the early phase of R&D. In order to increase computational efficiency, we concentrate on the range from the 0.5% (d_{lower}) to the 99.5% (d_{upper}) percentiles of the distribution. That is, d_{T_E} lies in the range $d_{lower} \leqslant d_{T_E} \leqslant d_{upper}$, for which the cumulative probability is $\sum_{d_{T_E}=d_{lower}}^{d_{upper}} p_{d_{T_E}} \geqslant 0.99$. We therefore select d_{T_E} values in the range $d_{lower} \leqslant d_{T_E} \leqslant d_{upper}$ and produce an estimate of the value of the American real option structure for each of these breakthrough levels, i.e. $C^L_{T_E}(d_{T_E})$. Thus, we may calculate the value of the American real option structure, i.e. the late phase of R&D, at time $t = T_E$, as follows:

$$C^L_{T_E} = \sum_{d_{T_E}=0}^{\infty} p_{d_{T_E}} C^L_{T_E}(d_{T_E}) \simeq \sum_{d_{T_E}=d_{lower}}^{d_{upper}} p_{d_{T_E}} C^L_{T_E}(d_{T_E}) \tag{8}$$

where $C^L_{T_E}(d_{T_E})$ is the value of the American real option structure at time T_E that begins with a discovery value of d_{T_E}, and $p_{d_{T_E}} = Poisson(d_{T_E}, \lambda_{Early})$ is the probability of achieving a breakthrough count of d_{T_E} according to the Poisson probability density function for the early phase of R&D.

Remark on Glasserman (2003)

Before proceeding to the valuation of the early phase of R&D, we comment more deeply on the use of the random tree method of Glasserman (2003) for the late phase of R&D. The approach solves the full optimal stopping problem for American options in a manner that offers notable advantages. While parametric methods rely on insights into the form of a good stopping rule, the random tree method requires only the ability to simulate the evolution of the state space variables. The random tree method is flexible in that it can be applied to complex state space systems, while being easily implemented and reliable. With only minimal conditions imposed, the method produces two consistent estimators - one biased high and one biased low. Both converge to the true value. The drawback of the method is that it is computationally intensive. However, in practice such computational constraints relating to execution speed, memory demands and storage capacity are readily overcome with a combination of faster programming languages, parallel processing and distributed computing, and cloud based storage. Glasserman (2003) also highlights the advantages of the random tree method with reference to some of the limitations of competing methods. In particular, parametric approximation methods only offer a relatively simple way to obtain a rough estimate of the value of an American option when there is good information about early exercise. State space partitioning use pre-defined states rather than random states. This creates a problem in the selection of these states with higher dimensional problems. The method is not applicable generally with the selection process needing repeating, depending on the American option of focus. Finally, while stochastic mesh methods are good for higher dimensional problems they require initial solutions to problems which may not exist, and experimentation to find useful basis functions.

In the context of a practice-relevant decision support tool, the Glasserman (2003) random tree approach allows us to save the simulations, which can then be interrogated to understand the state space interactions. The simulations, as alternative possible state space scenarios, offer the benefit of providing insights for the firm in respect of best-practice scenario analysis and due diligence around any investment appraisal exercise.

4.3. Early Phase of R&D: European Compound Real Option Valuation

The early phase of R&D is modelled as a European compound real options structure that has as its underlying asset, the value of the American real option structure. As noted earlier, we consider a fixed period for the early phase of R&D on the assumption that the firm sets a target cost per unit of production at the end the period, to be achieved through the R&D process. If the R&D leads to a reduction in the cost per unit of production below this target level then the firm decides to proceed with a subsequent late phase of R&D, and if not then the firm decides to abandon the R&D process. The European compound real option structure can, therefore, be considered an option on the differential between the American real option structure at time T_E and the value of the American real options structure that corresponds to the target cost per unit of production (CPU^{target}) as defined by the firm at the outset. CPU^{target} is defined as some percentage of the initial maximum CPU estimate, i.e. $CPU^{target} = \xi \times CPU_{Max}$. From CPU^{target}, the corresponding target breakthrough count, $d_{T_E}^{target}$, can be inferred. We therefore have the following formulation for the value of the European compound real options structure, i.e. the early phase of R&D, at time $t = 0$:

$$C_0^E = E_0 \left[\max \left(C_{T_E}^L - C_{T_E}^{L,d_{T_E}^{target}}, 0 \right) \right] e^{-rT_E} \tag{9}$$

where the strike price $C_{T_E}^{L,d_{T_E}^{target}}$ is the value of the American real option structure corresponding to $d_{T_E}^{target}$.

For the simulations, we can readily discretise the above valuation under the same Poisson distribution. It follows that:

$$C_0^E \simeq \left[\sum_{d_{T_E}=d_{lower}}^{d_{upper}} p_{d_{T_E}} \max\left(C_{T_E}^L(d_{T_E}) - C_{T_E}^{L,d_{T_E}^{target}}),\, 0 \right) \right] e^{-rT_E} \tag{10}$$

where, again, $C_{T_E}^L(d_{T_E})$ is the value of the late phase American real option which begins with a breakthrough value of d_{T_E}, and $p_{d_{T_E}} = Poisson(d_{T_E}, \lambda_{Early})$ is the probability of achieving a breakthrough value of d_{T_E} according to the Poisson probability density function for the early phase of R&D.

5. Simulation Settings

In this section, we set out the specific settings used for our simulation exercise involving the CO_2 recycling technology case study. Section 5.1 considers the static parameters and Section 5.2 considers the stochastic processes used in our price modelling. Section 5.3 sets out the settings for the modifications made to the modelling of production cost estimates, as per Section 3.1.

5.1. Static Parameters

We begin by explaining the static parameters used in the modelling and mapping of the R&D progress, and describe how the parameter values given in Table 1 are chosen. We assume that the early phase of R&D lasts for $T_E = 24$ months, followed by a period of up to 3 years during which the late phase R&D may be carried out, such that $T_F - T_E = 36$ months. The rate at which breakthroughs occur during the early phase of R&D, as modelled by our proposed Poisson process, has a mean value of $\lambda_{Early} = 1$ per month. For the late phase of R&D we assume a lower rate of $\lambda_{Late} = 0.5$ per month. These values suggest that it becomes increasingly difficult to make new breakthroughs as one transitions from the early phase to the late phase of the R&D process. We assume that the plant costs $D =$ €100 million to build and costs €1 million per month in fixed operating costs, which rise at a rate of $f = 1\%$ per month, calculated continuously. The plant takes $B = 12$ months to build and will operate for its lifetime of $L = 120$ months (10 years). With reference to Table A.1, we assume that the plant produces $N_P = 16,000$ metric tonnes of methane per month. We find that, given the cost of the hydrogen input to the Sabatier process, a reasonable estimate for the starting CPU_{Max} is €40,000 per 16 metric tonnes of methane produced (Table A.1), with an optimistic long-term cost efficiency aim of CPU_{Min} of €1,636 per 16 metric tonnes of methane produced (Table A.2). The number of CO_2 emissions required to be captured as part of the recycling process is $N_E = 44,000$ metric tonnes per month (Table A.1).

We select a learning-by-doing parameter of $\eta = 0.01$ for both the the methane production cost and the carbon capture cost (Section 4.1). This would permit the cost per unit of production to decrease by 0.7% per doubling of cumulative production. This is much less ambitious than the choice of Klaassen et al. (2005) where the CPU decreased by 5.4% for every doubling of the cumulative production. We use a very conservative estimate for η, as an overestimate will exaggerate the size of the cash flows, and hence the value of the R&D. The risk free rate of $r = 0.2568\%$ per month is based on the mean ECB lending rate from the introduction of the currency to the beginning of quantitative easing in 2015.

We set the target cost per unit of production at 50% of the initial cost per unit of production assumed. This production cost reduction target of $\xi = 50\%$ is somewhat arbitrary, but in our setting corresponds to a CPU level of €20,000, which maps to a cumulative number of breakthroughs of 11.68. This threshold lies just above the lower bound of $d_{lower} = 11$ breakthroughs as set in the previous section. A firm may wish to be more conservative than this and so may set the cost reduction target higher. In practice, the firm should set the target level based on its own requirements. And, of course, a firm should perform a sensitivity analysis around this choice as part of best-practice scenario analysis and due diligence efforts.

For the simulations underlying the random tree based valuations, we set the branch parameters $b_1 = b_2 = b_3 = \cdots = b_s = 16$. This branch level choice corresponds to a total of 16,777,216 simulations. The size of the time steps is $\delta t = 6$ months over the late phase of R&D, such that the number of steps over the full three-year period is $s = 6$. For the American real option structure valuation, we therefore use 16,777,216 simulations for each of the breakthrough count levels, d_{T_E}, in the $[d_{lower}, d_{upper}]$ range. This is particularly computationally intensive. We were unable to go any higher with the branch levels due to MATLAB memory and space constraints and the design of the code to store all of the simulations. Our decision to store all of the simulations is made to allow for the interrogation of the simulations as required to better appraise the decision support tool. In practice, such computational constraints relating to execution speed, memory demands and storage capacity are readily overcome with a combination of faster programming languages, parallel processing and distributed computing, and cloud based storage. In Section 6, we include a convergence analysis, where we increment from branch levels of 4 up to branch levels of 16. The convergence is evident. In order to make best use of the computational time, we note the breakthrough count range $d_{lower} \leqslant d_{T_E} \leqslant d_{upper}$ (Section 4.2) and choose the lower and upper limits so as to correspond to the 0.5 and 99.5 percentiles of the distribution of d_{T_E}, which is distributed $Poisson(24, \lambda_{Early} = 1)$. It follows, for our illustrative example, that $d_{lower} = 11$ and $d_{upper} = 37$.

Finally, as proposed in Section 4.1, we assume a fixed cost of capturing the CO_2 but allow for a gradual decrease in this cost due to learning-by-doing. We set the initial cost to be €40 per tonne, which is taken from Covington (2017). This is a more conservative estimate than that of Hanak and Manovic (2018), who estimate a price of £29 sterling per tonne.

5.2. Stochastic Variables

The commodities of interest are the natural gas (methane) produced by the technology and the emission allowances which are saved by preventing the CO_2 emissions entering the atmosphere. Both of these contribute to the positive value of the plant. We use a mean reverting model to capture the price of the natural gas, while we use a geometric Brownian motion (GBM) model to describe the evolution of the EUA emissions price under the EU ETS. Mean reversion models are used widely in the commodities literature, particularly for the modelling of natural gas (Schwartz, 1997; Wong and Lo, 2009; and Jang et al., 2013). We apply a GBM model with a drift to model

Static Parameters Based on Market Data and Model Estimates		
Early Phase Length (months), T_E	$T_E = 24$	user defined
Maximum Late Phase Length (months), $T_F - T_E$	$T_F - T_E = 36$	user defined
Poisson Rate per month	$\lambda_{Early} = 1$, $\lambda_{Late} = 0.5$	user defined
Production Cost Reduction Target	$\xi = 50\%$	user defined
Plant Construction Cost, D (Euro)	$D = 100,000,000$	user defined
Fixed Costs per month, F	$F = 1,000,000$	user defined
Monthly Rate of F Growth, f	$f = 0.01$	user defined
Building Time (months), B	$B = 12$	user defined
Lifecycle of plant (months), L	$L = 120$	user defined
Output Methane Units* per month, N_P	$N_P = 16,000$	user defined
Captured CO_2 Units* per month, N_E	$N_E = 44,000$	user defined
Initial CPU_{Max} (Euro) **	$CPU_{Max} = 40,000$	market estimate
Limit CPU_{Min} (Euro) **	$CPU_{Min} = 1,636$	market estimate
Max Rate of Progress, d_R	$d_R = 50$	user defined
Learning-By-Doing Parameter, η	$\eta = 0.01$	user defined
Monthly Risk Free rate, r	$r = 0.002568$	Average ECB rate ***
Parameters for Monte Carlo Simulations		
Number of branches at other simulations, b_i	$b_2 = b_3 = ... = 16$	user defined
Length of simulation step, δt	$\delta t = 6$	user defined
Number of simulation steps, s	$s = 6$	user defined
Lower limit of d_{T_E} for calculation	$d_{lower} = 11$	0.5 perc, Poisson($T_E \lambda_{Early}$)
Upper limit of d_{T_E} for calculation	$d_{upper} = 37$	99.5 perc, Poisson($T_E \lambda_{Early}$)

The table shows the static parameters for the Monte Carlo simulations. Time is measured in months and currency in Euro.

* A mass unit of production in our application case is based on Table A.1, and is 16 metric tonnes of methane and 44 metric tonnes of CO_2. We assume 1,000 mass units are produced per month.

** Given on a cost per mass unit of methane production, i.e. per 16 metric tonnes of methane produced. Implies that 40,000/16 = 2,500 (1636/16 = 102.25) is the maximum (minimum) cost per metric tonne of methane produced.

*** This is the average ECB lending rate since the introduction of the Euro until quantitative easing in 2015.

Table 1: Static Parameters for Monte Carlo Simulations of the Sabatier Process R&D

emission allowance prices following Daskalakis et al. (2009), Kalayci et al. (2014) and Gurler et al. (2016).

Modelling of Commodity Prices

We wish to produce simulated prices for natural gas, which is our proxy for the income from the sale of the methane to be produced by the CO_2 recycling technology. In order to produce a mean reverting series of natural gas prices, G, for the Monte Carlo simulations, we assume an Ornstein-Uhlenbeck model following Glasserman (2003), Wong and Lo (2009), and we use the estimation method of Marin Sanchez and Palacio (2013):

$$dg(t) = \alpha(\nu - g(t))dt + \sigma dX(t)$$

where $g(t) = \ln G(t)$ is the log price of natural gas at time t, α is the speed of mean reversion and $\nu = \ln \mu$ is the log of the mean level to which the price reverts, and where $X(t)$ is a standard Brownian motion and σ is a constant volatility coefficient. When the mean is a constant, μ, following Glasserman (2003) we have

$$G(t_{i+1}) = G(t_i)e^{-\alpha\delta t} + \mu(1 - e^{-\alpha\delta t}) + \sigma\sqrt{\frac{1}{2\alpha}(1 - e^{-2\alpha\delta t})}X_i$$

with X_i drawn from $N(0, 1)$ and the time step $\delta t = t_{i+1} - t_i$ constant. We use the method of Marin Sanchez and Palacio (2013), which is a maximum likelihood technique using the Euler-Maruyma scheme, to estimate the long term mean $\mu = 50.58$, volatility $\sigma = 1.158$ and the speed of mean reversion $\alpha = 0.099$ using monthly NBP prompt month futures natural gas prices from July 2008 to June 2018.

Modelling of EUA Prices

We model the price of EUAs using a GBM with drift. We therefore have for the log returns of the EUA price, $m(t)$,

$$dm(t) = \beta dt + \varsigma dY(t)$$

such that in a discretised setting

$$M(t_{i+1}) = M(t_i)e^{(\beta - \frac{1}{2}\varsigma^2)\delta t + \varsigma\sqrt{\delta t}Y_i)}$$

where $M(t)$ is the price of emissions at time t, ς is the standard deviation of EUA log returns and β is the mean of the monthly log returns; Y_i is a series of random variables chosen from $N(0, 1)$; and the time step is again a constant $\delta t = t_{i+1} - t_i$ (Glasserman, 2003). We calibrate the model using monthly EUA futures data during the present phase of the EU ETS (Phase III, 2013 to 2020), taken on 9th July 2018 when the EUA futures price was €16.00, following Campbell et al. (1997). When we use monthly data we have $\beta_{month} = 0.0196$ and $\varsigma_{month} = 0.1488$, which suggests an EUA price in ten years of $16e^{(120\times 0.0196)}$ = €168. Since the supply of EUAs in the

EU Emissions Trading Scheme is set to drop by 2.2% per year in Phase IV of the EU ETS, it seems reasonable to expect price growth. However, as we are using the example of carbon dioxide to illustrate the real options valuation model we have developed, rather than building a specific predictive model for EUA prices, we will use the estimated price of EUA in ten years to be €100, which is a reasonable expectation in order to ensure that we have a conservative estimate of EUA price growth. Therefore, we pick $\beta = 0.0157$, which we find separately to be consistent with the same calculation using daily data.

5.3. Dynamic Production/Production Cost Settings

As noted in Section 3.1, we can extend the modelling framework to move from a static to a dynamic UPC profile to reflect that the firm implementing the decision support tool may use information revealed as the R&D progresses to adjust its initial and ongoing estimates of production rates and, hence, production costs. We propose this through (i) the dynamic UPC_{Max} estimate modification and (ii) the dynamic UPC logistic curve estimate modification as outlined in Section 3.1. We present the specific settings chosen for the CO_2 recycling technology case.

We begin with the dynamic UPC_{Max} estimate modification. For practical purposes, we work with the production cost CPU profile. We allow CPU_{Min} to vary during the late phase period of R&D. Specifically, we model $CPU_{Min} \sim N(\mu_t, \sigma_t)$, such that the minimum cost per unit of production is assumed to be drawn from a normal distribution with a time-varying mean, μ_t, and time-varying standard deviation, σ_t. For each of the six dates that the late phase simulations are sampled, the time varying mean at period t adjusts to the previous CPU_{Min} estimate at period $(t-1)$, where we set the mean for the first six monthly period to be our original CPU_{Min} estimate of €1636. In this way, we mimic in our simulations a process whereby the firm periodically adjusts its estimate of CPU_{Min}. In respect of the uncertainty around the CPU_{Min} estimate, we let the time-varying standard deviation σ_t take the values $S, (5/6)S, (4/6)S, (3/6)S, (2/6)S, (1/6)S$, applied respectively to each of the six dates on which the late phase simulations are sampled. S is some assumed initial standard deviation level. This allows for the uncertainty in the CPU_{Min} estimate to decline in a simple linear fashion as the late phase R&D progresses. We show results for alternative implementations with initial standard deviations levels of $S = 100, 200$ and 300. Note that while we use, for convenience, the Normal distribution to model upward and downward revisions in the CPU_{Min} estimates, an asymmetric distribution could be readily applied to weight more towards downward revisions in production costs rather than upwards revisions.

Turning to the dynamic UPC logistic curve estimate modification, we wish to update our simulation framework in a way that allows for the shape of the UPC logistic curve estimate to be adjusted in response to the natural advancement of the technology, but also the firm's view of the advancement of the technology. While there are many ways this time varying UPC logistic curve could be modelled, we proceed as follows in our implementation. As outlined in Section 5.1, we begin with an initial estimate of $d_R = 50$ given the use of the Sabatier process for our CO_2 recycling technology case. We allow this to randomly change every six months within the late phase of R&D as follows, with each having equal probability:

- Decrease by a level of 0.5;
- Decrease by a level of 0.25;
- No change;
- Increase by a level of 0.25.

In this way, we model a greater likelihood for advancement in the technology, while accounting for the potential that the technology will not advance or that it will regress relative to previous estimates. The settings allow for conservative advancement in the technology given the established nature of the Sabatier procedure. As noted earlier, decreases (increases) in d_R mean that the UPC logistic curve adjusts to reflect advancements (regressions) in the technology relative to initial and ongoing estimates. The implication of this for production costs is that a decreased (increased) d_R leads to a decreased (increased) CPU profile.

6. Simulation Results

We present the results of the decision support tool applied to the R&D investment opportunity around the Sabatier based CO_2 recycling technology as described in Section 2. The simulation exercise returns the resulting valuations for the American real options structure and the European compound real options structure, which respectively model the late phase and early phase of the R&D process. Our main results are reported in Table 2. The early phase of R&D has a reported value of €45.54m, which indicates that the R&D should proceed if the proposed budget, I_E, is less than this value. Convergence to this valuation can be seen in Table 3, where we increment from branch levels of 4 up to branch levels of 16. Of particular use for practitioners is the breakdown that the decision support tool provides around the valuation of the late phase of R&D. The value of the American real option structure at each possible d_{T_E} level, and corresponding $CPU\,(d_{T_T})$ level, is reported in the third column of Table 2, with associated probabilities in the fourth column. We see that the value of the late phase of R&D increases considerably as the CPU (d_{T_E}) achieved over the early phase of R&D decreases. This is because the early progress means a better starting position for the late phase of R&D. This breakdown provides insights for the firm in setting future production cost targets and planning future budgets in respect of the late phase of R&D.

As noted in Section 4.2, the simulation framework of Glasserman (2003) is ideally suited to support scenario analysis and due diligence around the R&D investment opportunity. In particular, the underlying simulations can be interrogated to better understand how key variables feed into decision making. For instance, we consider here the exercise region of the American real options structure embedded in the late phase of R&D. We are able to to identify from the simulations that there are scenarios where early exercise into the operational phase is optimal and that these occur when, in particular, sufficiently high emissions prices are combined with sufficiently high (low) levels of R&D progress (production costs). Figure 4 illustrates this by means of charting the evolution of the (early) exercise region through the first, second and third years of the late phase of R&D. A clear pattern is identifiable from this evolution. As the cost of production lowers through technological breakthroughs then early exercise can occur at lower levels of emission prices, while early exercise can be observed at higher cost of production levels once the emissions prices is sufficiently high. The interrogation of the simulations also reveals that the price of natural gas does not play a role in terms of early exercise, reflecting the mean reverting nature of the price dynamic assumed for natural gas. Relatedly, we can also investigate when the R&D fails entirely. This is represented by the terminal region of the late phase of R&D where the firms decides not to operationalise the technology. This is due to insufficient progress and/or unfavourable emissions market conditions - see Figure 5.

As a final comment, we emphasise the importance of the embedded optionality as engineered into the R&D investment opportunity (Section 2). We do this through benchmarking the real option based valuation against a naive net present value (NPV) that fails to consider the abandonment and early exercise optionality. Specifically, to align with our five year term spanning the early and late phases of R&D, we consider an operational phase that commences five years hence. This operational phase comprises the same one-year period for constructing the required infrastructure, followed by the same ten-year period over which the CO2 recycling technology is operated. We determine the NPV for each of a range of starting breakthrough levels, ranging from 1 breakthrough through to 100 breakthroughs. We then probability weight these using a Poisson distribution with mean 42, chosen to align with the Poisson intensity parameter choices used in the real options analysis; namely, one breakthrough per month for the first two years and 0.5 breakthroughs per month for the subsequent three years. The probability weighted average of the NPV values then serves as the benchmark NPV value. This is highly negative at approximately -€73m, emphasising the importance of the abandonment optionality in particular.

Competition and Regulatory Risk

Extending the analysis to include competition and regulatory risk, as described in Section 2, shows a considerable reduction in the R&D investment value (Table 4, Panel B). For this analysis, we assume a 5% probability of a competition/regulatory event during the late phase of R&D, a realisation of which forces the company to abandon the R&D entirely, while we assume a 5% probability of a regulatory event during the operational phase of the deployed technology, which results in a 50% reduction in the cash flows from the project. The reduction in the value of the early phase of R&D is pronounced under these settings, dropping by almost 19%. While the settings are conservative, the analysis emphasises the importance of considering competition and regulatory risk in the R&D investment problem.

Benchmarking

As a form of benchmarking, we consider two alternative specifications of the decision support tool. These comprise (i) a deterministic rate of R&D breakthroughs (Table 4, Panel C) to replace the Poisson driven rate of R&D breakthroughs, and (ii) a linear functional mapping to replace the logistic functional mapping (Table 4, Panel D). Assuming a deterministic rate of R&D breakthroughs is a constraint that has a dramatic valuation effect, failing to capture much of the value relative to the base case. For this implementation, we simply use the arrival rates that we have assumed for the Poisson modelling across the early and late phases of R&D. For the early phase of R&D, we use the monthly rate of $\lambda_{Early} = 1$, such that at the end of the two-year period we are certain to have accumulated 24 breakthroughs. For the late phase of R&D, we use the monthly rate of $\lambda_{Late} = 0.5$, such that at the end of this additional three-year period we are certain to have accumulated a total of 42 breakthroughs, i.e. a further 18 breakthroughs. Given the certainty over the R&D progress, the decision to exercise the American real option structure in our simulation framework is determined entirely from the joint evolution of the natural gas and emissions prices.

Replacing the logistic functional mapping with the straight line functional mapping, as a way to give economic meaning to the R&D progress modelled, offers considerable new insights. The straight line function assumes an equal impact on the production rate, which contrast to the logistic function assumption of a varying, albeit stylised, impact. The linear function overestimates progress in UPC improvements relative to the logistic function at both the left tail and right tail of the cumulative breakthrough (d) spectrum. This overestimation is particularly pronounced for the right tail, corresponding to high values of the cumulative breakthrough level. Our results show that, relative to the logistic function case, the value of the American real option structure is higher (by approximately 14%). On the other hand, the value of the European real option structure is lower (by approximately 31%). This latter observation reflects the fact that the strike price of the European real option structure (i.e. the value of the American real option structure that corresponds to the target CPU level set by the firm) is higher under the linear function case than the logistic function case.

Discount Rate Sensitivity

We now conduct a sensitivity analysis of the real option valuations with respect to the discount rate. As outlined in Section 4.1, we choose, for convenience, the discount rate to be the risk free rate of interest and, in so doing, assume a zero beta for the CAPM specification. For practical applications, the firm must, of course, consider a non-zero beta, chosen to reflect the risk of the technology been developed, whether this is comparable to the average risk of the firm or is an above-average level of risk. For the sensitivity analysis, we move from our base case of an approximate 3% p.a. discount rate and consider both 5% p.a. and 7% p.a. discount rates. We implement the decision support tool under both scenarios. Table 4, Panel E, presents the results, where the heavier discounting is evident through considerable valuation reductions. The real options structures are particularly sensitive to the discount rate choice and so the firm needs to carefully choose this parameter.

Dynamic Production Costs

We conclude with some insights into our modelling extension that allows for dynamic production rate and production cost estimates. We generalise the modelling approach by moving from a static CPU profile to a dynamic one, through building in (i) a dynamic CPU_{Min} estimate, or (ii) a dynamic CPU curve estimate, as outlined in Section 3.1. The exact settings used in each case are given in Section 5.3. Table 4, Panel F, shows the results for the dynamic CPU_{Min} estimate modification. It can be seen that allowing for variation in the production cost estimates through CPU_{Min} revisions leads to a moderate increase in the value of the R&D investment relative to the base case. Indeed, the higher the initial variation that is set around the CPU_{Min} estimate then the higher the increase in value. Panel G of Table 4 shows the results for the case of the dynamic CPU production cost curve estimate, which, as per Section 5.3, is modelled through adjustments to the point of inflection of the UPC production rate curve. Allowing for variation in the production cost estimates in this manner, can again be seen to moderately increase the value of the R&D investment. This analysis emphasises that the uncertainty around the estimation of production costs is an important factor within the R&D investment process that the decision support tool needs to capture.

7. Conclusions

The practice relevant decision support tool that we propose allows firms, that wish to appraise R&D investment in new technology, to generate a rational estimation of the value of the R&D based on real options analysis. The main novelty of the decision support tool lies in how it combines the use of a Poisson process to model progress in the technological innovation process, a logistic function to give an economic interpretation to this progress in terms of production cost efficiencies, and a learning-by-doing dynamic in the operational phase to model further cost efficiencies post-deployment of the technology. Our modelling approach accounts for technical risk, cash flow risk, competition and obsolescence risk, and learning-by-doing effects (Berk et al., 2004). A number of messages emerge from our work.

The decision support tool is readily adapted to alternative forms of investment flexibility. The investment problem we devise for illustrative purposes, incorporates the flexibility for the firm to (i) abandon the R&D at the end of a defined early phase of R&D, if a predetermined production cost efficiency target is not achieved, and (ii) deploy the technology as early as possible, once commercially viable, during a subsequent phase of R&D. The decision support tool can be readily extended to explicitly model competition and regulatory risk. Accounting for these sources of risk within the R&D process, and indeed within the operational phase, has important valuation implications. Realisations of such risk exposures can lead to significant reductions in value, as demonstrated in our simulation exercise. Furthermore, the logistic functional mapping we propose has the advantage of alignment with technology lifecycle theory. Indeed, an alternative linear functional mapping reveals considerable under-estimation (over-estimation) in the early (late) R&D phase valuation relative to the logistic functional mapping. The valuations resulting from the decision support tool are shown to be sensitive to the choice of discount rate. In this R&D investment context, where the risk-neutral distribution is unknown and cannot be fully inferred from futures prices or other traded assets, we recommend users follow the advice of Schwartz (2013), who suggests that in this case the risk-neutral distribution can be inferred from an equilibrium model, such as the CAPM. The firm must consider the beta choice carefully to reflect the risk of the technology been developed, whether this reflects the average risk of the firm or some higher or lower level of risk. Finally, we show how the modelling approach can be generalised to allow for a dynamic product cost profile throughout the R&D process, reflecting periodic revisions of production cost estimates based on information that reveals itself to the firm.

d_{T_E}	$CPU_{d_{T_E}}$	American Real Option Value	$p_{d_{T_E}}$	E[American]	European Payoff	E[European]
		A	p	pA	E	pE
11	20,800	39,439,469	0.001	56,759	-	-
12	19,628	43,395,716	0.003	124,904	1,265,999	3,644
13	18,528	46,639,212	0.005	247,827	4,509,495	23,962
14	17,495	50,013,371	0.009	455,583	7,883,654	71,814
15	16,526	53,498,373	0.015	779,726	11,368,656	165,695
16	15,615	57,046,231	0.022	1,247,153	14,916,514	326,107
17	14,760	60,748,620	0.031	1,874,958	18,618,903	574,658
18	13,958	64,454,511	0.041	2,652,449	22,324,794	918,716
19	13,204	68,181,745	0.052	3,544,211	26,052,028	1,354,232
20	12,497	72,044,666	0.062	4,494,015	29,914,949	1,866,040
21	11,833	75,919,628	0.071	5,412,260	33,789,911	2,408,860
22	11,209	79,940,564	0.078	6,216,993	37,810,847	2,940,557
23	10,624	84,033,183	0.081	6,819,419	41,903,467	3,400,529
24	10,074	88,075,734	0.081	7,147,478	45,946,017	3,728,588
25	9,558	92,203,244	0.078	7,183,135	50,073,527	3,901,000
26	9,074	96,258,433	0.072	6,922,205	54,128,716	3,892,543
27	8,619	100,322,568	0.064	6,412,861	58,192,851	3,719,828
28	8,192	104,346,462	0.055	5,717,210	62,216,745	3,408,895
29	7,791	108,277,218	0.045	4,909,720	66,147,501	2,999,391
30	7,415	112,238,504	0.036	4,071,472	70,108,788	2,543,209
31	7,061	116,070,331	0.028	3,259,721	73,940,614	2,076,549
32	6,730	119,840,515	0.021	2,524,202	77,710,798	1,636,823
33	6,418	123,526,150	0.015	1,892,242	81,396,433	1,246,876
34	6,126	127,018,559	0.011	1,373,464	84,888,842	917,911
35	5,851	130,575,858	0.007	968,180	88,446,142	655,801
36	5,594	133,994,764	0.005	662,353	91,865,047	454,101
37	5,352	137,371,823	0.003	440,463	95,242,106	305,380
				€ 87,410,961		*€45,541,709*

The number of breakthroughs at the end of the early phase of R&D, d_{T_E}, are presented with the Monte Carlo simulation results. CPU and d_{T_E} are related by Equation (2) with parameters $UPC_{Max} = 1/1636$, $UPC_{Min} = 1/40000$ and $d_R = 50$. The values of the early and late phases are calculated using Equations (10) and (8), using a branching parameter of 16 for each of the six periods of six months during the late phase of R&D. Thus the sample size for the Monte Carlo simulation is 16^6=16,777,216. Results are presented, discounted to the start of the R&D process, i.e. when $t = 0$. Column 5 calculates the expected value of the American real options structure based on the Poison distribution of d_{T_E}. The strike price of the European real options structure is €$42,129,717$, which is the value of the American real options structure that corresponds to the target cost per unit of production (CPU^{target}) as defined by the firm at the outset. CPU^{target} is defined as some percentage of the initial CPU, before the R&D is conducted, i.e. $CPU^{target} = \xi \times CPU_{Max}$. From CPU^{target}, the corresponding target breakthrough count, $d_{T_E}^{target}$, can be inferred. For operational purposes, we have set the target cost per unit of production at 50% of the initial cost per unit of production, i.e. $\xi = 0.5$. We calculate the payoff of the European compound real option structure in column 6 and its expected value in column 7. We present results for $11 \leqslant d_{T_E} \leqslant 37$ which corresponds to the 0.5 and 99.5 percentiles of the distribution of the sum of 24 independent draws from a Poisson($\lambda_{Early} = 1$) distribution. We note that the contributions to value of the American and European real option structures in columns 5 and 7 are very small at either end of this selection, and therefore supports our choice of $d_{lower} = 11$ and $d_{upper} = 37$.

Table 2: R&D Phase Valuations

Branch Level	# Branches	American Real Option Value	European Real Option Value
4	4,096	-	59,249,050
5	15,625	-	52,250,550
6	46,656	84,661,781	42,202,864
7	117,649	102,164,775	50,314,009
8	262,144	89,389,338	48,373,720
9	531,441	94,695,242	50,232,072
10	1,000,000	92,564,254	46,970,627
11	1,771,561	97,778,152	48,250,435
12	2,985,984	96,522,465	49,884,446
13	4,826,809	102,955,902	50,570,243
14	7,529,536	91,939,177	46,781,534
15	11,390,625	86,754,542	44,001,455
16	16,777,216	87,410,961	45,541,709

Table presents the American and European real options structure valuations based on implementing the decision support tool for alternative branch levels.

Table 3: Convergence Analysis

	American Real Option Value	European Real Option Value	American Value Difference	European Value Difference
	Panel A: Poisson Breakthrough Rate, Logistic Functional Mapping			
	87,410,961	45,541,709		
	Panel B: As per Panel A, Competition and Regulatory Risk Included			
	73,940,075	36,865,913	-15.41%	-19.05%
	Panel C: Deterministic Breakthrough Rate, Logistic Functional Mapping			
	39,818,239	22,329,309	-54.45%	-50.97%
	Panel D: Poisson Breakthrough Rate, Linear Functional Mapping			
	99,354,542	31,388,266	13.66%	-31.08%
	Panel E: As per Panel A, Alternative Discount Rates			
5% p.a.	69,055,500	36,531,861	-21.00%	-19.78%
7% p.a.	50,040,471	26,290,697	-42.75%	-42.27%
	Panel F: As Per Panel A, Dynamic CPU_{Min}			
100	88,110,249	45,792,227	0.8%	0.6%
200	90,033,290	46,736,838	3.0%	2.1%
300	91,956,331	48,407,939	5.2%	3.6%
	Panel G: As per Panel A, Dynamic d_R			
	89,945,879	49,279,288	2.9%	1.8%

Table presents the American and European real options structure valuations based on implementing the decision support tool for alternative specifications. Panel A presents the main valuations under the assumption of a Poisson breakthrough rate and the logistic functional mapping. Panel B is as per Panel A but with competition and regulatory risk included in the valuations. Panel C presents valuations whereby the Poisson breakthrough rate is replaced with a deterministic rate. Panel D presents valuations whereby the logistic functional mapping is replaced with a linear functional mapping. Panel E is as per Panel A but with valuations derived using alternative discount rates to the main analysis. Panel F is as per Panel A but with dynamic CPU_{Min} values, implemented for initial standard deviations of 100, 200 and 300 (see Section 5.3). Panel G is as per Panel A but with dynamic point of inflection, d_R, for the logistic function specification of the UPC production profile (see Section 5.3). Value differences are relative to the main valuations reported in Panel A.

Table 4: Alternative Specifications of Decision Support Tool

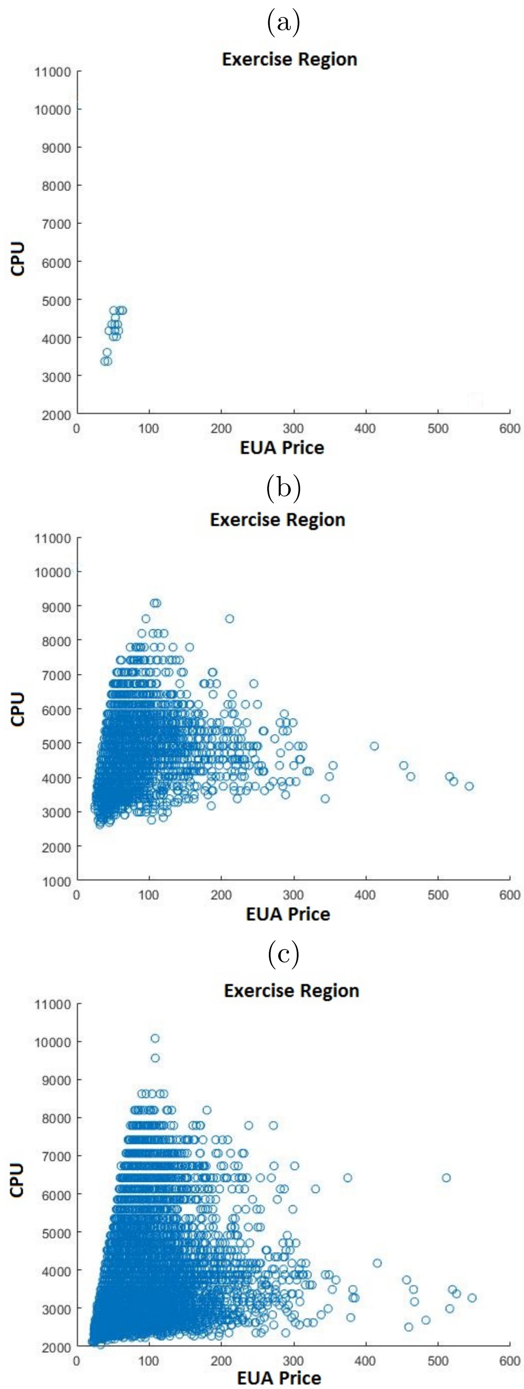

The plots show the region where the firm decides to exercise into the operational phase, as recoreded on the producfion cost-emissions price plane for a sample of simulations. Plots (a) and (b) show production cost-emissions price combinations that, respectively, record early exercise one year and two years into the late phase of R&D. Plot (c) shows production cost-emissions price combinations that record exercise in the final year of the late phase of R&D.

Figure 4: Operational Phase Excercise Region: Production Cost-Emissions Price Plane

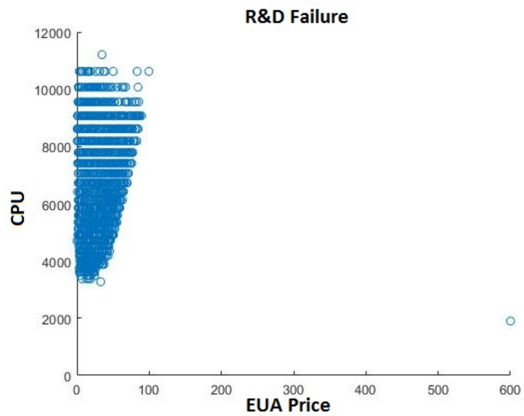

The plot shows the region where the R&D fails in the final year of the late phase of R&D. Production cost-emissions price combinations for a sample of simulations are given, where the firm decides not to exercise into the operational phase, and instead abandons the R&D.

Figure 5: R&D Failure Region: Production Cost-Emissions Price Plane

References

Ahern, E.P., Deane, P., Persson, T., Gallachï¿œeir, B.., Murphy, J.D., 2015. A perspective on the potential role of renewable gas in a smart energy island system. Renewable Energy 78, 648 – 656. URL: `http://www.sciencedirect.com/science/article/pii/S096014811500066X`, doi:https://doi.org/10.1016/j.renene.2015.01.048.

Anzanello, M.J., Fogliatto, F.S., 2011. Learning curve models and applications: Literature review and research directions. International Journal of Industrial Ergonomics 41, 573 – 583. doi:https://doi.org/10.1016/j.ergon.2011.05.001.

Balzani, V., Credi, A., Venturi, M., 2008. Photochemical conversion of solar energy. ChemSusChem 1, 26–58.

Berk, J.B., Green, R.C., Naik, V., 2004. Valuation and return dynamics of new ventures. Review of Financial Studies 17, 1 – 35.

Bockris, J., 1977. Environmental chemistry, in: Environmental Chemistry. Springer, Boston, MA, pp. 1–18.

Brandão, L.E., Fernandes, G., Dyer, J.S., 2018. Valuing multistage investment projects in the pharmaceutical industry. European Journal of Operational Research 271, 720–732.

Campbell, J., Lo, W., MacKinlay, A., 1997. The Econometrics of Financial Markets. Princeton University Press, Princeton New Jersey.

Cassimon, D., Backer, M.D., Engelen, P., Wouwe, M.V., Yordanov, V., 2011. Incorporating technical risk in compound real option models to value a pharmaceutical R&D licensing opportunity. Research Policy 40, 1200 – 1216.

URL: http://www.sciencedirect.com/science/article/pii/S0048733311001016, doi:https://doi.org/10.1016/j.respol.2011.05.020.

Centi, G., Perathoner, S., 2009. Opportunities and prospects in the chemical recycling of carbon dioxide to fuels. Catalysis Today 148, 191 – 205. URL: http://www.sciencedirect.com/science/article/pii/S0920586109004416, doi:https://doi.org/10.1016/j.cattod.2009.07.075. special Issue of the 10th International Conference on CO2 Utilization, Tianijn, China, May 17-21, 2009.

Childs, P.D., Ott, S.H., Triantis, A.J., 1998. Capital budgeting for interrelated projects: A real options approach. Journal of Financial and Quantitative Analysis 33, 305 – 334.

Childs, P.D., Triantis, A.J., 1999. Dynamic R&D investment policies. Management Science 45, 1359 – 1377.

Covington, H., 2017. Investment consequences of the Paris Climate Agreement. Journal of Sustainable Finance & Investment 7, 54–63. URL: https://doi.org/10.1080/20430795.2016.1196556, doi:10.1080/20430795.2016.1196556, arXiv:https://doi.org/10.1080/20430795.2016.1196556.

Daskalakis, G., Psychoyios, D., Markellos, R.N., 2009. Modeling CO2 emission allowance prices and derivatives: Evidence from the European trading scheme. Journal of Banking & Finance 33, 1230 – 1241. URL: http://www.sciencedirect.com/science/article/pii/S037842660900003X, doi:https://doi.org/10.1016/j.jbankfin.2009.01.001.

Dixit, A., Pindyck, R., 1994. Investment under Uncertainty. Princeton University Press.

Farzin, Y.H., Huisman, K.J., Kort, P.M., 1998. Optimal timing of technology adoption. Journal of Economic Dynamics and Control 22, 779–799.

Frayne, L., Das, N., Paul, A., Amirjalayer, S., Buma, W.J., Woutersen, S., Long, C., Vos, J.G., Pryce, M.T., 2018. Photo and electrochemical properties of a CO2 reducing ruthenium rhenium quaterpyridine based catalyst. ChemPhotoChem 2, 323–331. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/cptc.201700197, doi:10.1002/cptc.201700197, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/cptc.201700197

Glasserman, P., 2003. Monte Carlo methods in financial engineering. Springer, New York.

Griliches, Z., 1957. Hybrid corn: An exploration in the economics of technological change. Econometrica, Journal of the Econometric Society , 501–522.

Grübler, A., 2003. Technology and global change. Cambridge University Press.

Grübler, A., Nakićenović, N., Victor, D.G., 1999. Dynamics of energy technologies and global change. Energy policy 27, 247–280.

Gurler, l., Yenigï¿œen, D., ï¿œaglar, M., Berk, E., 2016. On the modeling of CO2 EUA and CER prices of EU ETS for the 2008 to 2012 period. Applied Stochastic Models in Business & Industry 32, 375 – 395.

Hagspiel, V., Huisman, K.J., Nunes, C., 2015. Optimal technology adoption when the arrival rate of new technologies changes. European Journal of Operational Research 243, 897–911.

Hanak, D.P., Manovic, V., 2018. Techno economic feasibility assessment of CO2 capture from coal-fired power plants using molecularly imprinted polymer. Fuel 214, 512 – 520. URL: http://www.sciencedirect.com/science/article/pii/S0016236117313650, doi:https://doi.org/10.1016/j.fuel.2017.10.107.

Haupt, R., Kloyer, M., Lange, M., 2007. Patent indicators for the technology life cycle development. Research Policy 36, 387–398.

Herath, H.S., Park, C.S., 2002. Multi-stage capital investment opportunities as compound real options. The Engineering Economist 47, 1–27.

Hsu, J.C., Schwartz, E.S., 2008. A model of R&D valuation and the design of research incentives. Insurance: Mathematics and Economics 43, 350 – 367. URL: http://www.sciencedirect.com/science/article/pii/S0167668708000577, doi:https://doi.org/10.1016/j.insmatheco.2008.05.003.

Huchzermeier, A., Loch, C., 2001. Project management under risk: Using the real options approach to evaluate flexibility in R&D. Management Science 47, 85–101. doi:10.1287/mnsc.47.1.85.10661.

Huisman, K.J., Kort, P.M., 2004. Strategic technology adoption taking into account future technological improvements: A real options approach. European Journal of Operational Research 159, 705–728.

Jaffe, A.B., Newell, R.G., Stavins, R.N., 2003. Technological change and the environment, in: Handbook of environmental economics. Elsevier. volume 1, pp. 461–516.

Jang, Y.S., Lee, D.J., Oh, H.S., 2013. Evaluation of new and renewable energy technologies in Korea using real options. International Journal of Energy Research 37, 1645–1656. URL: http://dx.doi.org/10.1002/er.2985, doi:10.1002/er.2985.

Kalayci, E., Gaygisiz, E., Weber, G.W., 2014. A multi-period stochastic portfolio optimization model applied for an airline company in the EU ETS. Optimization 63, 1817–1835. URL: https://doi.org/10.1080/02331934.2014.895900, doi:10.1080/02331934.2014.895900, arXiv:https://doi.org/10.1080/02331934.2014.895900.

Klaassen, G., Miketa, A., Larsen, K., Sundqvist, T., 2005. The impact of R&D on innovation for wind energy in Denmark, Germany and the United Kingdom. Ecological Economics 54, 227 – 240. URL: http://www.sciencedirect.com/science/article/pii/S0921800905000340, doi:https://doi.org/10.1016/j.ecolecon.2005.01.008. technological Change and the Environment.

Lee, K.C., Lee, S., Kang, I.W., 2005. KMPI: Measuring knowledge management performance. Information & Management 42, 469 – 482. URL: http://www.sciencedirect.com/science/article/pii/S0378720604000473, doi:https://doi.org/10.1016/j.im.2004.02.003.

Marin Sanchez, F.H., Palacio, J.S., 2013. Gaussian estimation of one-factor mean reversion processes. Journal of Probability & Statistics , 1 – 10.

Mazzucato, M., 2013. The entrepreneurial state: Debunking public vs. private sector myths. Anthem Press.

Metz, B., Davidson, O., de Coninck M., Loos, M., Meyer, L., 2005. Carbon dioxide capture and storage: special report of the Intergovernmental Panel on Climate Change. Cambridge University Press.

Miltersen, K., Schwartz, E., 2004. R&D investments with competitive interactions. European Finance Review 8, 1–47.

Nemet, G.F., 2009. Interim monitoring of cost dynamics for publicly supported energy technologies. Energy Policy 37, 825 – 835. URL: http://www.sciencedirect.com/science/article/pii/S0301421508005910, doi:https://doi.org/10.1016/j.enpol.2008.10.031.

Nishihara, M., 2018. Valuation of R&D investment under technological, market, and rival preemption uncertainty. Managerial and Decision Economics 39, 200–212. URL: http://dx.doi.org/10.1002/mde.2896, doi:10.1002/mde.2896. mDE-16-0076.R1.

Pan, H., 2006. Dynamic and endogenous change of input–output structure with specific layers of technology. Structural Change and Economic Dynamics 17, 200–223.

Pan, H., Kohler, J., 2007. Technological change in energy systems: Learning curves, logistic curves and input-output coefficients. Ecological Economics 63, 749 – 758. URL: http://www.sciencedirect.com/science/article/pii/S0921800907000912, doi:https://doi.org/10.1016/j.ecolecon.2007.01.013. sustainability and Cost-Benefit Analysis.

Pennings, E., Sereno, L., 2011. Evaluating pharmaceutical R&D under technical and economic uncertainty. European Journal of Operational Research 212, 374 – 385. URL: http://www.sciencedirect.com/science/article/pii/S0377221711001093, doi:https://doi.org/10.1016/j.ejor.2011.01.055.

Pinker, S., 2018. Enlightenment now: The case for reason, science, humanism and progress. Allen Lane.

Schwartz, E., 2013. The real options approach to valuation: Challenges and opportunities. Latin American Journal of Economics 50, 163–177.

Schwartz, E.S., 1997. The stochastic behavior of commodity prices: Implications for valuation and hedging. The Journal of Finance 52, 923–973. URL: http://www.jstor.org/stable/2329512.

Schwartz, E.S., 2004. Patents and R&D as real options. Economic Notes 33, 23 – 54. doi:https://doi-org.dcu.idm.oclc.org/10.1111/j.0391-5026.2004.00124.x.

Schwartz, E.S., Moon, M., 2000. Rational pricing of internet companies. Financial Analysts Journal 56, 62–75. URL: https://search.proquest.com/docview/219213897?accountid=15753. copyright - Copyright Association for Investment Management and Research May/Jun 2000; Last updated - 2014-09-13; CODEN - FIAJA4; SubjectsTermNotLitGenreText - United States; US.

Schwartz, E.S., Moon, M., 2001. Rational pricing of internet companies revisited. Financial Review 36, 7–26.

Schwartz, E.S., Zozaya-Gorostiza, C., 2003. Investment under uncertainty in information technology: Acquisition and development projects. Management Science 49, 57 – 70. doi:doi.org/10.1287/mnsc.49.1.57.12753.

The Royal Society, 2010. URL: royalsociety.org. accesssed on 1st November 2018.

Trigeorgis, L., 1996. Real Options Managerial Flexibility and Strategy in Resource Allocation. MIT Press.

Trigeorgis, L., Tsekrekos, A.E., 2018. Real options in operations research: A review. European Journal of Operational Research 270, 1–24.

UNEP, 2017. Global trends in renewable energy investment 2017. URL: http://fs-unep-centre.org/sites/default/files/publications/globaltrendsinrenewableenergyinve accessed on 1st November 2018.

UNESCO, 2018. URL: http://uis.unesco.org/apps/visualisations/research-and-development-spending accessed on 1st November 2018.

Weitzman, M.L., Newey, W., Rabin, M., 1981. Sequential r&d strategy for synfuels. The Bell Journal of Economics , 574–590.

Wong, H.Y., Lo, Y.W., 2009. Option pricing with mean reversion and stochastic volatility. European Journal of Operational Research 197, 179 – 187. URL: http://www.sciencedirect.com/science/article/pii/S0377221708004244, doi:https://doi.org/10.1016/j.ejor.2008.05.014.

Wright, T.P., 1936. Factors affecting the cost of airplanes. Journal of the Aeronautical Sciences 3, 122 – 128. URL: http://www.uvm.edu/pdodds/research/papers/others/1936/wright1936a.pdf.

Yeh, S., Rubin, E.S., 2012. A review of uncertaities in technology experience curves. Energy Economics 34, 762 – 771. URL: http://www.sciencedirect.com/science/article/pii/S0140988311002805, doi:https://doi.org/10.1016/j.eneco.2011.11.006.

Yu, C., van Sark, W., Alsema, E., 2011. Unraveling the photovoltaic technology learning curve by incorporation of input price changes and scale effects. Renewable and Sustainable Energy Reviews 15, 324 – 337. URL: http://www.sciencedirect.com/science/article/pii/S1364032110002881, doi:https://doi.org/10.1016/j.rser.2010.09.001.